PSY 190: STATISTICS FOR THE BEHAVIORAL SCIENCES



Kevin R. Smith
Rio Hondo College



Psychological Statistics Remix

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (https://LibreTexts.org) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden born from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

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.



LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by NICE CXone Expert and was supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, 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 adaptions contact info@LibreTexts.org or visit our main website at https://LibreTexts.org.

This text was compiled on 04/19/2026

TABLE OF CONTENTS

Licensing

Note about this book/version

1: Introduction

- 1.1: What are statistics?
- 1.2: Why do we study statistics?
- 1.3: Types of Data and How to Collect Them
- 1.4: Collecting Data
- 1.5: Type of Research Designs
- 1.6: Types of Statistical Analyses
- 1.7: Mathematical Notation
- 1.E: Introduction (Exercises)

2: Describing Data using Distributions and Graphs

- 2.1: Graphing Qualitative Variables
- 2.2: Graphing Quantitative Variables
- 2.E: Describing Data using Distributions and Graphs (Exercises)

3: Measures of Central Tendency and Spread

- 3.1: What is Central Tendency?
- 3.2: Measures of Central Tendency
- 3.3: Spread and Variability
- 3.E: Measures of Central Tendency and Spread (Exercises)

4: Basic Concepts of Probability: The math

- 4.1: Sample Spaces, Events, and Their Probabilities
- 4.2: Complements, Intersections, and Unions
- 4.3: Conditional Probability and Independent Events
- 4.E: Basic Concepts of Probability (Exercises)

5: Z-scores and the Standard Normal Distribution

- 5.1: Normal Distributions
- 5.2: Z-scores
- 5.3: Z-scores and the Area under the Curve
- 5.E: Z-scores and the Standard Normal Distribution (Exercises)

6: Probability

- 6.1: What is Probability
- 6.2: Probability in Graphs and Distributions
- 6.3: The Bigger Picture
- 6.E: Probability (Exercises)

7: Sampling Distributions

- 7.1: People, Samples, and Populations
- 7.2: The Sampling Distribution of Sample Means
- 7.3: Using Standard Error for Probability
- 7.4: Sampling Distribution, Probability and Inference
- 7.E: Sampling Distributions (Exercises)

8: Introduction to Hypothesis Testing

- 8.1: Logic and Purpose of Hypothesis Testing
- 8.2: The Probability Value
- 8.3: The Null Hypothesis
- 8.4: The Alternative Hypothesis
- 8.5: Critical values, p-values, and significance level
- 8.6: Steps of the Hypothesis Testing Process
- 8.7: Movie Popcorn
- 8.8: Effect Size
- 8.9: Office Temperature
- 8.10: Different Significance Level
- 8.11: Other Considerations in Hypothesis Testing
- 8.E: Introduction to Hypothesis Testing (Exercises)

9: Introduction to t-tests

- 9.1: The t-statistic
- 9.2: Hypothesis Testing with t
- 9.3: Confidence Intervals
- 9.E: Introduction to t-tests (Exercises)

10: Repeated Measures

- 10.1: Change and Differences
- 10.2: Hypotheses of Change and Differences
- 10.3: Increasing Satisfaction at Work
- 10.4: Bad Press
- 10.E: Repeated Measures (Exercises)

11: Independent Samples

- 11.1: Difference of Means
- 11.2: Research Questions about Independent Means
- 11.3: Hypotheses and Decision Criteria
- 11.4: Independent Samples t-statistic
- 11.5: Standard Error and Pooled Variance
- 11.6: Movies and Mood
- 11.7: Effect Sizes and Confidence Intervals
- 11.8: Homogeneity of Variance
- 11.E: Independent Samples (Exercises)

12: Analysis of Variance

- 12.1: Observing and Interpreting Variability
- 12.2: Sources of Variance
- 12.3: ANOVA Table
- 12.4: ANOVA and Type I Error
- 12.5: Hypotheses in ANOVA
- 12.6: Scores on Job Application Tests
- 12.7: Variance Explained
- 12.8: Post Hoc Tests
- 12.9: Other ANOVA Designs
- 12.10: Analysis of Variance (Exercises)

13: Correlations

- 13.1: Variability and Covariance
- 13.2: Visualizing Relations
- 13.3: Three Characteristics
- 13.4: Pearson's r
- 13.5: Anxiety and Depression
- 13.6: Effect Size
- 13.7: Correlation versus Causation
- 13.8: Final Considerations
- 13.E: Correlations (Exercises)

14: Chi-square

- 14.1: Categories and Frequency Tables
- 14.2: Goodness-of-Fit
- 14.3: χ^2 Statistic
- 14.4: Pineapple on Pizza
- 14.5: Contingency Tables for Two Variables
- 14.6: Test for Independence
- 14.7: College Sports
- 14.E: Chi-square (Exercises)

Index

Glossary

Detailed Licensing

Licensing

*A detailed breakdown of this resource's licensing can be found in **Back Matter/Detailed Licensing**.*

Note about this book/version

CHAPTER OVERVIEW

1: Introduction

1.1: What are statistics?
1.2: Why do we study statistics?
1.3: Types of Data and How to Collect Them
1.4: Collecting Data
1.5: Type of Research Designs
1.6: Types of Statistical Analyses
1.7: Mathematical Notation
1.E: Introduction (Exercises)

This page titled 1: Introduction is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

1.1: What are statistics?

Statistics include numerical facts and figures. For instance:

- The largest earthquake measured 9.2 on the Richter scale.
- Men are at least 10 times more likely than women to commit murder.
- One in every 8 South Africans is HIV positive.
- By the year 2020, there will be 15 people aged 65 and over for every new baby born.

The study of statistics involves math and relies upon calculations of numbers. But it also relies heavily on how the numbers are chosen and how the statistics are interpreted. For example, consider the following three scenarios and the interpretations based upon the presented statistics. You will find that the numbers may be right, but the interpretation may be wrong. Try to identify a major flaw with each interpretation before we describe it.

1. A new advertisement for Ben and Jerry's ice cream introduced in late May of last year resulted in a 30% increase in ice cream sales for the following three months. Thus, the advertisement was effective. A major flaw is that ice cream consumption generally increases in the months of June, July, and August regardless of advertisements. This effect is called a history effect and leads people to interpret outcomes as the result of one variable when another variable (in this case, one having to do with the passage of time) is actually responsible.
2. The more churches in a city, the more crime there is. Thus, churches lead to crime. A major flaw is that both increased churches and increased crime rates can be explained by larger populations. In bigger cities, there are both more churches and more crime. This problem, which we will discuss in more detail later in the text, refers to the third-variable problem. Namely, a third variable can cause both situations; however, people erroneously believe that there is a causal relationship between the two primary variables rather than recognize that a third variable can cause both.
3. 75% more interracial marriages are occurring this year than 25 years ago. Thus, our society accepts interracial marriages. A major flaw is that we don't have the information that we need. What is the rate at which marriages are occurring? Suppose only 1% of marriages 25 years ago were interracial and so now 1.75% of marriages are interracial (1.75 is 75% higher than 1). But this latter number is hardly evidence suggesting the acceptability of interracial marriages. In addition, the statistic provided does not rule out the possibility that the number of interracial marriages has seen dramatic fluctuations over the years and this year is not the highest. Again, there is simply not enough information to understand fully the impact of the statistics.

As a whole, these examples show that statistics are not only facts and figures; they are something more than that. In the broadest sense, “statistics” refers to a range of techniques and procedures for analyzing, interpreting, displaying, and making decisions based on data.

Statistics is the language of science and data. The ability to understand and communicate using statistics enables researchers from different labs, different languages, and different fields articulate to one another exactly what they have found in their work. It is an objective, precise, and powerful tool in science and in everyday life.

What statistics are not

Many psychology students dread the idea of taking a statistics course, and more than a few have changed majors upon learning that it is a requirement. That is because many students view statistics as a math class, which is actually not true. While many of you will not believe this or agree with it, statistics isn’t math. Although math is a central component of it, statistics is a broader way of organizing, interpreting, and communicating information in an objective manner. Indeed, great care has been taken to eliminate as much math from this course as possible (students who do not believe this are welcome to ask the professor what matrix algebra is). Statistics is a way of viewing reality as it exists around us in a way that we otherwise could not.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 1.1: What are statistics? is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.1: What are statistics?** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

1.2: Why do we study statistics?

Virtually every student of the behavioral sciences takes some form of statistics class. This is because statistics is how we communicate in science. It serves as the link between a research idea and usable conclusions. Without statistics, we would be unable to interpret the massive amounts of information contained in data. Even small datasets contain hundreds – if not thousands – of numbers, each representing a specific observation we made. Without a way to organize these numbers into a more interpretable form, we would be lost, having wasted the time and money of our participants, ourselves, and the communities we serve.

Beyond its use in science, however, there is a more personal reason to study statistics. Like most people, you probably feel that it is important to "take control of your life." But what does this mean? Partly, it means being able to properly evaluate the data and claims that bombard you every day. If you cannot distinguish good from faulty reasoning, then you are vulnerable to manipulation and to decisions that are not in your best interest. Statistics provides tools that you need in order to react intelligently to information you hear or read. In this sense, statistics is one of the most important things that you can study.

To be more specific, here are some claims that we have heard on several occasions. (We are not saying that each one of these claims is true!)

- 4 out of 5 dentists recommend Dentine.
- Almost 85% of lung cancers in men and 45% in women are tobacco-related.
- Condoms are effective 94% of the time.
- People tend to be more persuasive when they look others directly in the eye and speak loudly and quickly.
- Women make 75 cents to every dollar a man makes when they work the same job.
- A surprising new study shows that eating egg whites can increase one's life span.
- People predict that it is very unlikely there will ever be another baseball player with a batting average over 400.
- There is an 80% chance that in a room full of 30 people that at least two people will share the same birthday.
- 79.48% of all statistics are made up on the spot.

All of these claims are statistical in character. We suspect that some of them sound familiar; if not, we bet that you have heard other claims like them. Notice how diverse the examples are. They come from psychology, health, law, sports, business, etc. Indeed, data and data interpretation show up in discourse from virtually every facet of contemporary life.

Statistics are often presented in an effort to add credibility to an argument or advice. You can see this by paying attention to television advertisements. Many of the numbers thrown about in this way do not represent careful statistical analysis. They can be misleading and push you into decisions that you might find cause to regret. For these reasons, learning about statistics is a long step towards taking control of your life. (It is not, of course, the only step needed for this purpose.) The purpose of this course, beyond preparing you for a career in psychology, is to help you learn statistical essentials. It will make you into an intelligent consumer of statistical claims.

You can take the first step right away. To be an intelligent consumer of statistics, your first reflex must be to question the statistics that you encounter. The British Prime Minister Benjamin Disraeli is quoted by Mark Twain as having said, "There are three kinds of lies -- lies, damned lies, and statistics." This quote reminds us why it is so important to understand statistics. So let us invite you to reform your statistical habits from now on. No longer will you blindly accept numbers or findings. Instead, you will begin to think about the numbers, their sources, and most importantly, the procedures used to generate them.

The above section puts an emphasis on defending ourselves against fraudulent claims wrapped up as statistics, but let us look at a more positive note. Just as important as detecting the deceptive use of statistics is the appreciation of the proper use of statistics. You must also learn to recognize statistical evidence that supports a stated conclusion. Statistics are all around you, sometimes used well, sometimes not. We must learn how to distinguish the two cases. In doing so, statistics will likely be the course you use most in your day to day life, even if you do not ever run a formal analysis again.

This page titled 1.2: Why do we study statistics? is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.2: Why do we study statistics?** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

1.3: Types of Data and How to Collect Them

In order to use statistics, we need data to analyze. Data come in an amazingly diverse range of formats, and each type gives us a unique type of information. In virtually any form, data represent the measured value of variables. A variable is simply a characteristic or feature of the thing we are interested in understanding. In psychology, we are interested in people, so we might get a group of people together and measure their levels of stress (one variable), anxiety (a second variable), and their physical health (a third variable). Once we have data on these three variables, we can use statistics to understand if and how they are related. Before we do so, we need to understand the nature of our data: what they represent and where they came from.

Types of Variables

When conducting research, experimenters often manipulate variables. For example, an experimenter might compare the effectiveness of four types of antidepressants. In this case, the variable is "type of antidepressant." When a variable is manipulated by an experimenter, it is called an independent variable. The experiment seeks to determine the effect of the independent variable on relief from depression. In this example, relief from depression is called a dependent variable. In general, the independent variable is manipulated by the experimenter and its effects on the dependent variable are measured.

Example 1.3.1

Can blueberries slow down aging? A study indicates that antioxidants found in blueberries may slow down the process of aging. In this study, 19-month-old rats (equivalent to 60-year-old humans) were fed either their standard diet or a diet supplemented by either blueberry, strawberry, or spinach powder. After eight weeks, the rats were given memory and motor skills tests. Although all supplemented rats showed improvement, those supplemented with blueberry powder showed the most notable improvement.

a. What is the independent variable? (dietary supplement: none, blueberry, strawberry, and spinach)
b. What are the dependent variables? (memory test and motor skills test)

Example 1.3.2

Does beta-carotene protect against cancer? Beta-carotene supplements have been thought to protect against cancer. However, a study published in the Journal of the National Cancer Institute suggests this is false. The study was conducted with 39,000 women aged 45 and up. These women were randomly assigned to receive a beta-carotene supplement or a placebo, and their health was studied over their lifetime. Cancer rates for women taking the betacarotene supplement did not differ systematically from the cancer rates of those women taking the placebo.

a. What is the independent variable? (supplements: beta-carotene or placebo)
b. What is the dependent variable? (occurrence of cancer)

Example 1.3.3

How bright is right? An automobile manufacturer wants to know how bright brake lights should be in order to minimize the time required for the driver of a following car to realize that the car in front is stopping and to hit the brakes.

a. What is the independent variable? (brightness of brake lights)
b. What is the dependent variable? (time to hit brakes)

Levels of an Independent Variable

If an experiment compares an experimental treatment with a control treatment, then the independent variable (type of treatment) has two levels: experimental and control. If an experiment were comparing five types of diets, then the independent variable (type of diet) would have 5 levels. In general, the number of levels of an independent variable is the number of experimental conditions.

Qualitative and Quantitative Variables

An important distinction between variables is between qualitative variables and quantitative variables. Qualitative variables are those that express a qualitative attribute such as hair color, eye color, religion, favorite movie, gender, and so on. The values of a qualitative variable do not imply a numerical ordering. Values of the variable “religion” differ qualitatively; no ordering of religions is implied. Qualitative variables are sometimes referred to as categorical variables. Quantitative variables are those variables that are measured in terms of numbers. Some examples of quantitative variables are height, weight, and shoe size.

In the study on the effect of diet discussed previously, the independent variable was type of supplement: none, strawberry, blueberry, and spinach. The variable “type of supplement” is a qualitative variable; there is nothing quantitative about it. In contrast, the dependent variable “memory test” is a quantitative variable since memory performance was measured on a quantitative scale (number correct).

Discrete and Continuous Variables

Variables such as number of children in a household are called discrete variables since the possible scores are discrete points on the scale. For example, a household could have three children or six children, but not **4.53** children. Other variables such as “time to respond to a question” are continuous variables since the scale is continuous and not made up of discrete steps. The response time could be **1.64** seconds, or it could be **1.64237123922121** seconds. Of course, the practicalities of measurement preclude most measured variables from being truly continuous.

Levels of Measurement

Before we can conduct a statistical analysis, we need to measure our dependent variable. Exactly how the measurement is carried out depends on the type of variable involved in the analysis. Different types are measured differently. To measure the time taken to respond to a stimulus, you might use a stop watch. Stop watches are of no use, of course, when it comes to measuring someone's attitude towards a political candidate. A rating scale is more appropriate in this case (with labels like “very favorable,” “somewhat favorable,” etc.). For a dependent variable such as “favorite color,” you can simply note the color-word (like “red”) that the subject offers.

Although procedures for measurement differ in many ways, they can be classified using a few fundamental categories. In a given category, all of the procedures share some properties that are important for you to know about. The categories are called “scale types,” or just “scales,” and are described in this section.

Nominal scales

When measuring using a nominal scale, one simply names or categorizes responses. Gender, handedness, favorite color, and religion are examples of variables measured on a nominal scale. The essential point about nominal scales is that they do not imply any ordering among the responses. For example, when classifying people according to their favorite color, there is no sense in which green is placed “ahead of” blue. Responses are merely categorized. Nominal scales embody the lowest level of measurement.

Ordinal scales

A researcher wishing to measure consumers' satisfaction with their microwave ovens might ask them to specify their feelings as either “very dissatisfied,” “somewhat dissatisfied,” “somewhat satisfied,” or “very satisfied.” The items in this scale are ordered, ranging from least to most satisfied. This is what distinguishes ordinal from nominal scales. Unlike nominal scales, ordinal scales allow comparisons of the degree to which two subjects possess the dependent variable. For example, our satisfaction ordering makes it meaningful to assert that one person is more satisfied than another with their microwave ovens. Such an assertion reflects the first person's use of a verbal label that comes later in the list than the label chosen by the second person.

On the other hand, ordinal scales fail to capture important information that will be present in the other scales we examine. In particular, the difference between two levels of an ordinal scale cannot be assumed to be the same as the difference between two other levels. In our satisfaction scale, for example, the difference between the responses “very dissatisfied” and “somewhat dissatisfied” is probably not equivalent to the difference between “somewhat dissatisfied” and “somewhat satisfied.” Nothing in our measurement procedure allows us to determine whether the two differences reflect the same difference in psychological

satisfaction. Statisticians express this point by saying that the differences between adjacent scale values do not necessarily represent equal intervals on the underlying scale giving rise to the measurements. (In our case, the underlying scale is the true feeling of satisfaction, which we are trying to measure.)

What if the researcher had measured satisfaction by asking consumers to indicate their level of satisfaction by choosing a number from one to four? Would the difference between the responses of one and two necessarily reflect the same difference in satisfaction as the difference between the responses two and three? The answer is No. Changing the response format to numbers does not change the meaning of the scale. We still are in no position to assert that the mental step from 1 to 2 (for example) is the same as the mental step from 3 to 4.

Interval scales

Interval scales are numerical scales in which intervals have the same interpretation throughout. As an example, consider the Fahrenheit scale of temperature. The difference between 30 degrees and 40 degrees represents the same temperature difference as the difference between 80 degrees and 90 degrees. This is because each 10-degree interval has the same physical meaning (in terms of the kinetic energy of molecules).

Interval scales are not perfect, however. In particular, they do not have a true zero point even if one of the scaled values happens to carry the name "zero." The Fahrenheit scale illustrates the issue. Zero degrees Fahrenheit does not represent the complete absence of temperature (the absence of any molecular kinetic energy). In reality, the label "zero" is applied to its temperature for quite accidental reasons connected to the history of temperature measurement. Since an interval scale has no true zero point, it does not make sense to compute ratios of temperatures. For example, there is no sense in which the ratio of 40 to 20 degrees Fahrenheit is the same as the ratio of 100 to 50 degrees; no interesting physical property is preserved across the two ratios. After all, if the "zero" label were applied at the temperature that Fahrenheit happens to label as 10 degrees, the two ratios would instead be 30 to 10 and 90 to 40, no longer the same! For this reason, it does not make sense to say that 80 degrees is "twice as hot" as 40 degrees. Such a claim would depend on an arbitrary decision about where to "start" the temperature scale, namely, what temperature to call zero (whereas the claim is intended to make a more fundamental assertion about the underlying physical reality).

Ratio scales

The ratio scale of measurement is the most informative scale. It is an interval scale with the additional property that its zero position indicates the absence of the quantity being measured. You can think of a ratio scale as the three earlier scales rolled up in one. Like a nominal scale, it provides a name or category for each object (the numbers serve as labels). Like an ordinal scale, the objects are ordered (in terms of the ordering of the numbers). Like an interval scale, the same difference at two places on the scale has the same meaning. And in addition, the same ratio at two places on the scale also carries the same meaning.

The Fahrenheit scale for temperature has an arbitrary zero point and is therefore not a ratio scale. However, zero on the Kelvin scale is absolute zero. This makes the Kelvin scale a ratio scale. For example, if one temperature is twice as high as another as measured on the Kelvin scale, then it has twice the kinetic energy of the other temperature.

Another example of a ratio scale is the amount of money you have in your pocket right now (25 cents, 55 cents, etc.). Money is measured on a ratio scale because, in addition to having the properties of an interval scale, it has a true zero point: if you have zero money, this implies the absence of money. Since money has a true zero point, it makes sense to say that someone with 50 cents has twice as much money as someone with 25 cents (or that Bill Gates has a million times more money than you do).

What level of measurement is used for psychological variables?

Rating scales are used frequently in psychological research. For example, experimental subjects may be asked to rate their level of pain, how much they like a consumer product, their attitudes about capital punishment, their confidence in an answer to a test question. Typically these ratings are made on a 5-point or a 7-point scale. These scales are ordinal scales since there is no assurance that a given difference represents the same thing across the range of the scale. For example, there is no way to be sure that a treatment that reduces pain from a rated pain level of 3 to a rated pain level of 2 represents the same level of relief as a treatment that reduces pain from a rated pain level of 7 to a rated pain level of 6.

In memory experiments, the dependent variable is often the number of items correctly recalled. What scale of measurement is this? You could reasonably argue that it is a ratio scale. First, there is a true zero point; some subjects may get no items correct at all. Moreover, a difference of one represents a difference of one item recalled across the entire scale. It is certainly valid to say that someone who recalled 12 items recalled twice as many items as someone who recalled only 6 items.

But number-of-items recalled is a more complicated case than it appears at first. Consider the following example in which subjects are asked to remember as many items as possible from a list of 10. Assume that (a) there are 5 easy items and 5 difficult items, (b) half of the subjects are able to recall all the easy items and different numbers of difficult items, while (c) the other half of the subjects are unable to recall any of the difficult items but they do remember different numbers of easy items. Some sample data are shown below.

Table **1.3.1**: Sample Data

Subject	Easy Items					Difficult Items					Score
A	0	0	1	1	0	0	0	0	0	0	2
B	1	0	1	1	0	0	0	0	0	0	3
C	1	1	1	1	1	1	1	0	0	0	7
D	1	1	1	1	1	0	1	1	0	1	8

Let's compare (i) the difference between Subject A's score of 2 and Subject B's score of 3 and (ii) the difference between Subject C's score of 7 and Subject D's score of 8. The former difference is a difference of one easy item; the latter difference is a difference of one difficult item. Do these two differences necessarily signify the same difference in memory? We are inclined to respond “No” to this question since only a little more memory may be needed to retain the additional easy item whereas a lot more memory may be needed to retain the additional hard item. The general point is that it is often inappropriate to consider psychological measurement scales as either interval or ratio.

Consequences of level of measurement

Why are we so interested in the type of scale that measures a dependent variable? The crux of the matter is the relationship between the variable's level of measurement and the statistics that can be meaningfully computed with that variable. For example, consider a hypothetical study in which 5 children are asked to choose their favorite color from blue, red, yellow, green, and purple. The researcher codes the results as follows:

Table **1.3.2**: Favorite color data code

Color	Code
Blue	1
Red	2
Yellow	3
Green	4
Purple	5

This means that if a child said her favorite color was “Red,” then the choice was coded as “2,” if the child said her favorite color was “Purple,” then the response was coded as 5, and so forth. Consider the following hypothetical data

Table **1.3.3**: Favorite color data

Subject	Color	Code
1	Blue	1
2	Blue	1
3	Green	4
4	Green	4
5	Purple	5

Each code is a number, so nothing prevents us from computing the average code assigned to the children. The average happens to be 3, but you can see that it would be senseless to conclude that the average favorite color is yellow (the color with a code of 3). Such nonsense arises because favorite color is a nominal scale, and taking the average of its numerical labels is like counting the number of letters in the name of a snake to see how long the beast is.

Does it make sense to compute the mean of numbers measured on an ordinal scale? This is a difficult question, one that statisticians have debated for decades. The prevailing (but by no means unanimous) opinion of statisticians is that for almost all practical situations, the mean of an ordinally-measured variable is a meaningful statistic. However, there are extreme situations in which computing the mean of an ordinally-measured variable can be very misleading.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 1.3: Types of Data and How to Collect Them is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.3: Types of Data and How to Collect Them** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

1.4: Collecting Data

We are usually interested in understanding a specific group of people. This group is known as the population of interest, or simply the population. The population is the collection of all people who have some characteristic in common; it can be as broad as “all people” if we have a very general research question about human psychology, or it can be extremely narrow, such as “all freshmen psychology majors at Midwestern public universities” if we have a specific group in mind.

Populations and samples

In statistics, we often rely on a sample --- that is, a small subset of a larger set of data --- to draw inferences about the larger set. The larger set is known as the population from which the sample is drawn.

Example 1.4.1

You have been hired by the National Election Commission to examine how the American people feel about the fairness of the voting procedures in the U.S. Who will you ask?

Solution

It is not practical to ask every single American how he or she feels about the fairness of the voting procedures. Instead, we query a relatively small number of Americans, and draw inferences about the entire country from their responses. The Americans actually queried constitute our sample of the larger population of all Americans.

A sample is typically a small subset of the population. In the case of voting attitudes, we would sample a few thousand Americans drawn from the hundreds of millions that make up the country. In choosing a sample, it is therefore crucial that it not over-represent one kind of citizen at the expense of others. For example, something would be wrong with our sample if it happened to be made up entirely of Florida residents. If the sample held only Floridians, it could not be used to infer the attitudes of other Americans. The same problem would arise if the sample were comprised only of Republicans. Inferences from statistics are based on the assumption that sampling is representative of the population. If the sample is not representative, then the possibility of sampling bias occurs. Sampling bias means that our conclusions apply only to our sample and are not generalizable to the full population.

Example 1.4.2

We are interested in examining how many math classes have been taken on average by current graduating seniors at American colleges and universities during their four years in school.

Solution

Whereas our population in the last example included all US citizens, now it involves just the graduating seniors throughout the country. This is still a large set since there are thousands of colleges and universities, each enrolling many students. (New York University, for example, enrolls 48,000 students.) It would be prohibitively costly to examine the transcript of every college senior. We therefore take a sample of college seniors and then make inferences to the entire population based on what we find. To make the sample, we might first choose some public and private colleges and universities across the United States. Then we might sample 50 students from each of these institutions. Suppose that the average number of math classes taken by the people in our sample were 3.2. Then we might speculate that 3.2 approximates the number we would find if we had the resources to examine every senior in the entire population. But we must be careful about the possibility that our sample is non-representative of the population. Perhaps we chose an overabundance of math majors, or chose too many technical institutions that have heavy math requirements. Such bad sampling makes our sample unrepresentative of the population of all seniors.

To solidify your understanding of sampling bias, consider the following example. Try to identify the population and the sample, and then reflect on whether the sample is likely to yield the information desired.

Example 1.4.3

A substitute teacher wants to know how students in the class did on their last test. The teacher asks the 10 students sitting in the front row to state their latest test score. He concludes from their report that the class did extremely well. What is the sample? What is the population? Can you identify any problems with choosing the sample in the way that the teacher did?

Solution

The population consists of all students in the class. The sample is made up of just the 10 students sitting in the front row. The sample is not likely to be representative of the population. Those who sit in the front row tend to be more interested in the class and tend to perform higher on tests. Hence, the sample may perform at a higher level than the population.

Example 1.4.4

A coach is interested in how many cartwheels the average college freshmen at his university can do. Eight volunteers from the freshman class step forward. After observing their performance, the coach concludes that college freshmen can do an average of 16 cartwheels in a row without stopping.

Solution

The population is the class of all freshmen at the coach's university. The sample is composed of the 8 volunteers. The sample is poorly chosen because volunteers are more likely to be able to do cartwheels than the average freshman; people who can't do cartwheels probably did not volunteer! In the example, we are also not told of the gender of the volunteers. Were they all women, for example? That might affect the outcome, contributing to the non-representative nature of the sample (if the school is co-ed).

Simple Random Sampling

Researchers adopt a variety of sampling strategies. The most straightforward is simple random sampling. Such sampling requires every member of the population to have an equal chance of being selected into the sample. In addition, the selection of one member must be independent of the selection of every other member. That is, picking one member from the population must not increase or decrease the probability of picking any other member (relative to the others). In this sense, we can say that simple random sampling chooses a sample by pure chance. To check your understanding of simple random sampling, consider the following example. What is the population? What is the sample? Was the sample picked by simple random sampling? Is it biased?

Example 1.4.5

A research scientist is interested in studying the experiences of twins raised together versus those raised apart. She obtains a list of twins from the National Twin Registry, and selects two subsets of individuals for her study. First, she chooses all those in the registry whose last name begins with Z. Then she turns to all those whose last name begins with B. Because there are so many names that start with B, however, our researcher decides to incorporate only every other name into her sample. Finally, she mails out a survey and compares characteristics of twins raised apart versus together.

Solution

The population consists of all twins recorded in the National Twin Registry. It is important that the researcher only make statistical generalizations to the twins on this list, not to all twins in the nation or world. That is, the National Twin Registry may not be representative of all twins. Even if inferences are limited to the Registry, a number of problems affect the sampling procedure we described. For instance, choosing only twins whose last names begin with Z does not give every individual an equal chance of being selected into the sample. Moreover, such a procedure risks over-representing ethnic groups with many surnames that begin with Z. There are other reasons why choosing just the Z's may bias the sample. Perhaps such people are more patient than average because they often find themselves at the end of the line! The same problem occurs with choosing twins whose last name begins with B. An additional problem for the B's is that the “every-other-one” procedure disallowed adjacent names on the B part of the list from being both selected. Just this defect alone means the sample was not formed through simple random sampling.

Sample size matters

Recall that the definition of a random sample is a sample in which every member of the population has an equal chance of being selected. This means that the sampling procedure rather than the results of the procedure define what it means for a sample to be random. Random samples, especially if the sample size is small, are not necessarily representative of the entire population. For example, if a random sample of 20 subjects were taken from a population with an equal number of males and females, there would be a nontrivial probability (0.06) that 70% or more of the sample would be female. Such a sample would not be representative, although it would be drawn randomly. Only a large sample size makes it likely that our sample is close to representative of the population. For this reason, inferential statistics take into account the sample size when generalizing results from samples to populations. In later chapters, you'll see what kinds of mathematical techniques ensure this sensitivity to sample size.

More complex sampling

Sometimes it is not feasible to build a sample using simple random sampling. To see the problem, consider the fact that both Dallas and Houston are competing to be hosts of the 2012 Olympics. Imagine that you are hired to assess whether most Texans prefer Houston to Dallas as the host, or the reverse. Given the impracticality of obtaining the opinion of every single Texan, you must construct a sample of the Texas population. But now notice how difficult it would be to proceed by simple random sampling. For example, how will you contact those individuals who don't vote and don't have a phone? Even among people you find in the telephone book, how can you identify those who have just relocated to California (and had no reason to inform you of their move)? What do you do about the fact that since the beginning of the study, an additional 4,212 people took up residence in the state of Texas? As you can see, it is sometimes very difficult to develop a truly random procedure. For this reason, other kinds of sampling techniques have been devised. We now discuss two of them.

Stratified Sampling

Since simple random sampling often does not ensure a representative sample, a sampling method called stratified random sampling is sometimes used to make the sample more representative of the population. This method can be used if the population has a number of distinct "strata" or groups. In stratified sampling, you first identify members of your sample who belong to each group. Then you randomly sample from each of those subgroups in such a way that the sizes of the subgroups in the sample are proportional to their sizes in the population.

Let's take an example: Suppose you were interested in views of capital punishment at an urban university. You have the time and resources to interview 200 students. The student body is diverse with respect to age; many older people work during the day and enroll in night courses (average age is 39), while younger students generally enroll in day classes (average age of 19). It is possible that night students have different views about capital punishment than day students. If 70% of the students were day students, it makes sense to ensure that 70% of the sample consisted of day students. Thus, your sample of 200 students would consist of 140 day students and 60 night students. The proportion of day students in the sample and in the population (the entire university) would be the same. Inferences to the entire population of students at the university would therefore be more secure.

Convenience Sampling

Not all sampling methods are perfect, and sometimes that's okay. For example, if we are beginning research into a completely unstudied area, we may sometimes take some shortcuts to quickly gather data and get a general idea of how things work before fully investing a lot of time and money into well-designed research projects with proper sampling. This is known as convenience sampling, named for its ease of use. In limited cases, such as the one just described, convenience sampling is okay because we intend to follow up with a representative sample. Unfortunately, sometimes convenience sampling is used due only to its convenience without the intent of improving on it in future work.

This page titled 1.4: Collecting Data is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.4: Collecting Data** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

1.5: Type of Research Designs

Research studies come in many forms, and, just like with the different types of data we have, different types of studies tell us different things. The choice of research design is determined by the research question and the logistics involved. Though a complete understanding of different research designs is the subject for at least one full class, if not more, a basic understanding of the principles is useful here. There are three types of research designs we will discuss: experimental, quasiexperimental, and non-experimental.

Experimental Designs

If we want to know if a change in one variable causes a change in another variable, we must use a true experiment. An experiment is defined by the use of random assignment to treatment conditions and manipulation of the independent variable. To understand what this means, let's look at an example:

A clinical researcher wants to know if a newly developed drug is effective in treating the flu. Working with collaborators at several local hospitals, she randomly samples 40 flu patients and randomly assigns each one to one of two conditions: Group A receives the new drug and Group B received a placebo. She measures the symptoms of all participants after 1 week to see if there is a difference in symptoms between the groups.

In the example, the independent variable is the drug treatment; we manipulate it into 2 levels: new drug or placebo. Without the researcher administering the drug (i.e. manipulating the independent variable), there would be no difference between the groups. Each person, after being randomly sampled to be in the research, was then randomly assigned to one of the 2 groups. That is, random sampling and random assignment are not the same thing and cannot be used interchangeably. For research to be a true experiment, random assignment must be used. For research to be representative of the population, random sampling must be used. The use of both techniques helps ensure that there are no systematic differences between the groups, thus eliminating the potential for sampling bias.

The dependent variable in the example is flu symptoms. Barring any other intervention, we would assume that people in both groups, on average, get better at roughly the same rate. Because there are no systematic differences between the 2 groups, if the researcher does find a difference in symptoms, she can confidently attribute it to the effectiveness of the new drug.

Quasi-Experimental Designs

Quasi-experimental research involves getting as close as possible to the conditions of a true experiment when we cannot meet all requirements. Specifically, a quasiexperiment involves manipulating the independent variable but not randomly assigning people to groups. There are several reasons this might be used. First, it may be unethical to deny potential treatment to someone if there is good reason to believe it will be effective and that the person would unduly suffer if they did not receive it. Alternatively, it may be impossible to randomly assign people to groups. Consider the following example:

A professor wants to test out a new teaching method to see if it improves student learning. Because he is teaching two sections of the same course, he decides to teach one section the traditional way and the other section using the new method. At the end of the semester, he compares the grades on the final for each class to see if there is a difference.

In this example, the professor has manipulated his teaching method, which is the independent variable, hoping to find a difference in student performance, the dependent variable. However, because students enroll in courses, he cannot randomly assign the students to a particular group, thus precluding using a true experiment to answer his research question. Because of this, we cannot know for sure that there are no systematic differences between the classes other than teaching style and therefore cannot determine causality.

Non-Experimental Designs

Finally, non-experimental research (sometimes called correlational research) involves observing things as they occur naturally and recording our observations as data. Consider this example:

A data scientist wants to know if there is a relation between how conscientious a person is and whether that person is a good employee. She hopes to use this information to predict the job performance of future employees by measuring their personality when they are still job applicants. She randomly samples volunteer employees from several different companies, measuring their conscientiousness and having their bosses rate their performance on the job. She analyzes this data to find a relation.

Here, it is not possible to manipulate conscientious, so the researcher must gather data from employees as they are in order to find a relation between her variables.

Although this technique cannot establish causality, it can still be quite useful. If the relation between conscientiousness and job performance is consistent, then it doesn't necessarily matter if conscientiousness causes good performance or if they are both caused by something else – she can still measure conscientiousness to predict future performance. Additionally, these studies have the benefit of reflecting reality as it actually exists since we as researchers do not change anything.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 1.5: Type of Research Designs is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.5: Type of Research Designs** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

1.6: Types of Statistical Analyses

Now that we understand the nature of our data, let's turn to the types of statistics we can use to interpret them. There are 2 types of statistics: descriptive and inferential.

Descriptive Statistics

Descriptive statistics are numbers that are used to summarize and describe data. The word "data" refers to the information that has been collected from an experiment, a survey, an historical record, etc. (By the way, "data" is plural. One piece of information is called a "datum.") If we are analyzing birth certificates, for example, a descriptive statistic might be the percentage of certificates issued in New York State, or the average age of the mother. Any other number we choose to compute also counts as a descriptive statistic for the data from which the statistic is computed. Several descriptive statistics are often used at one time to give a full picture of the data. Descriptive statistics are just descriptive. They do not involve generalizing beyond the data at hand. Generalizing from our data to another set of cases is the business of inferential statistics, which you'll be studying in another section. Here we focus on (mere) descriptive statistics. Some descriptive statistics are shown in Table **1.6.1**. The table shows the average salaries for various occupations in the United States in 1999.

Table **1.6.1**: Average salaries for various occupations in 1999.

Salary	Occupation
$112,760	pediatricians
$106,130	dentists
$100,090	podiatrists
$76,140	physicists
$53,410	architects
$49,720	school, clinical, and counseling psychologists
$47,910	flight attendants
$39,560	elementary school teachers
$38,710	police officers
$18,980	floral designers

Descriptive statistics like these offer insight into American society. It is interesting to note, for example, that we pay the people who educate our children and who protect our citizens a great deal less than we pay people who take care of our feet or our teeth.

For more descriptive statistics, consider Table **1.6.2**. It shows the number of unmarried men per 100 unmarried women in U.S. Metro Areas in 1990. From this table we see that men outnumber women most in Jacksonville, NC, and women outnumber men most in Sarasota, FL. You can see that descriptive statistics can be useful if we are looking for an opposite-sex partner! (These data come from the Information Please Almanac.)

Table **1.6.2**: Number of unmarried men per 100 unmarried women in U.S. Metro Areas in 1990.

Cities with mostly men	Men per 100 Women	Cities with mostly women	Men per 100 Women
1. Jacksonville, NC	224	1. Sarasota, FL	66
2. Killeen-Temple, TX	123	2. Bradenton, FL	68
3. Fayetteville, NC	118	3. Altoona, PA	69
4. Brazoria, TX	117	4. Springfield, IL	70
5. Lawton, OK	116	5. Jacksonville, TN	70
6. State College, PA	113	6. Gadsden, AL	70
7. ClarksvilleHopkinsville, TN-KY	113	7. Wheeling, WV	70
8. Anchorage, Alaska	112	8. Charleston, WV	71
9. Salinas-SeasideMonterey, CA	112	9. St. Joseph, MO	71
10. Bryan-College Station, TX	111	10. Lynchburg, VA	71

NOTE: Unmarried includes never-married, widowed, and divorced persons, 15 years or older.

These descriptive statistics may make us ponder why the numbers are so disparate in these cities. One potential explanation, for instance, as to why there are more women in Florida than men may involve the fact that elderly individuals tend to move down to the Sarasota region and that women tend to outlive men. Thus, more women might live in Sarasota than men. However, in the absence of proper data, this is only speculation.

You probably know that descriptive statistics are central to the world of sports. Every sporting event produces numerous statistics such as the shooting percentage of players on a basketball team. For the Olympic marathon (a foot race of 26.2 miles), we possess data that cover more than a century of competition. (The first modern Olympics took place in 1896.) The following table shows the winning times for both men and women (the latter have only been allowed to compete since 1984).

Table **1.6.3**: Winning Olympic marathon times.

Year	Winner	Country	Time
Women			
1984	Joan Benoit	USA	2:24:52
1988	Rosa Mota	POR	2:25:40
1992	Valentina Yegorova	UT	2:32:41
1996	Fatuma Roba	ETH	2:26:05
2000	Naoko Takahashi	JPN	2:23:14
2004	Mizuki Noguchi	JPN	2:26:20
Men			
1896	Spiridon Louis	GRE	2:58:50
1900	Michel Theato	FRA	2:59:45
1904	Thomas Hicks	USA	3:28:53
1906	Billy Sherring	CAN	2:51:23
1908	Johnny Hayes	USA	2:55:18
1912	Kenneth McArthur	S. Afr.	2:36:54
1920	Hannes Kolehmainen	FIN	2:32:35
1924	Albin Stenroos	FIN	2:41:22
1928	Boughra El Ouafi	FRA	2:32:57
1932	Juan Carlos Zabala	ARG	2:31:36
1936	Sohn Kee-Chung	JPN	2:29:19
1948	Delfo Cabrera	ARG	2:34:51
1952	Emil Ztopek	CZE	2:23:03
1956	Alain Mimoun	FRA	2:25:00
1960	Abebe Bikila	ETH	2:15:16
1964	Abebe Bikila	ETH	2:12:11
1968	Mamo Wolde	ETH	2:20:26
1972	Frank Shorter	USA	2:12:19
1976	Waldemar Cierpinski	E.Ger	2:09:55

Year	Winner	Country	Time
1980	Waldemar Cierpinski	E.Ger	2:11:03
1984	Carlos Lopes	POR	2:09:21
1988	Gelindo Bordin	ITA	2:10:32
1992	Hwang Young-Cho	S. Kor	2:13:23
1996	Josia Thugwane	S. Afr.	2:12:36
2000	Gezahenge Abera	ETH	2:10.10
2004	Stefano Baldini	ITA	2:10:55

There are many descriptive statistics that we can compute from the data in the table. To gain insight into the improvement in speed over the years, let us divide the men's times into two pieces, namely, the first 13 races (up to 1952) and the second 13 (starting from 1956). The mean winning time for the first 13 races is 2 hours, 44 minutes, and 22 seconds (written 2:44:22). The mean winning time for the second 13 races is 2:13:18. This is quite a difference (over half an hour). Does this prove that the fastest men are running faster? Or is the difference just due to chance, no more than what often emerges from chance differences in performance from year to year? We can't answer this question with descriptive statistics alone. All we can affirm is that the two means are "suggestive."

Examining Table **1.6.3** leads to many other questions. We note that Takahashi (the lead female runner in 2000) would have beaten the male runner in 1956 and all male runners in the first 12 marathons. This fact leads us to ask whether the gender gap will close or remain constant. When we look at the times within each gender, we also wonder how far they will decrease (if at all) in the next century of the Olympics. Might we one day witness a sub-2 hour marathon? The study of statistics can help you make reasonable guesses about the answers to these questions.

It is also important to differentiate what we use to describe populations vs what we use to describe samples. A population is described by a parameter; the parameter is the true value of the descriptive in the population, but one that we can never know for sure. For example, the Bureau of Labor Statistics reports that the average hourly wage of chefs is $23.87. However, even if this number was computed using information from every single chef in the United States (making it a parameter), it would quickly become slightly off as one chef retires and a new chef enters the job market. Additionally, as noted above, there is virtually no way to collect data from every single person in a population. In order to understand a variable, we estimate the population parameter using a sample statistic. Here, the term "statistic" refers to the specific number we compute from the data (e.g. the average), not the field of statistics. A sample statistic is an estimate of the true population parameter, and if our sample is representative of the population, then the statistic is considered to be a good estimator of the parameter.

Even the best sample will be somewhat off from the full population, earlier referred to as sampling bias, and as a result, there will always be a tiny discrepancy between the parameter and the statistic we use to estimate it. This difference is known as sampling error, and, as we will see throughout the course, understanding sampling error is the key to understanding statistics. Every observation we make about a variable, be it a full research study or observing an individual's behavior, is incapable of being completely representative of all possibilities for that variable. Knowing where to draw the line between an unusual observation and a true difference is what statistics is all about.

Inferential Statistics

Descriptive statistics are wonderful at telling us what our data look like. However, what we often want to understand is how our data behave. What variables are related to other variables? Under what conditions will the value of a variable change? Are two groups different from each other, and if so, are people within each group different or similar? These are the questions answered by inferential statistics, and inferential statistics are how we generalize from our sample back up to our population. Units 2 and 3 are all about inferential statistics, the formal analyses and tests we run to make conclusions about our data.

For example, we will learn how to use a t statistic to determine whether people change over time when enrolled in an intervention. We will also use an F statistic to determine if we can predict future values on a variable based on current known values of a variable. There are many types of inferential statistics, each allowing us insight into a different behavior of the data we collect. This course will only touch on a small subset (or a sample) of them, but the principles we learn along the way will make it easier to learn new tests, as most inferential statistics follow the same structure and format.

This page titled 1.6: Types of Statistical Analyses is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.6: Types of Statistical Analyses** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

1.7: Mathematical Notation

As noted above, statistics is not math. It does, however, use math as a tool. Many statistical formulas involve summing numbers. Fortunately there is a convenient notation for expressing summation. This section covers the basics of this summation notation.

Let's say we have a variable **X** that represents the weights (in grams) of 4 grapes:

Table **1.7.1**

Grape	X
1	4.6
2	5.1
3	4.9
4	4.4

We label Grape 1's weight X_1, Grape 2's weight X_2, etc. The following formula means to sum up the weights of the four grapes:

$$\sum_{i=1}^{4} X_i \tag{1.7.1}$$

The Greek letter Σ indicates summation. The "i = 1" at the bottom indicates that the summation is to start with X_1 and the 4 at the top indicates that the summation will end with X_4. The "X_i" indicates that **X** is the variable to be summed as i goes from 1 to 4. Therefore,

$$\sum_{i=1}^{4} X_i = X_1 + X_2 + X_3 + X_4 = 4.6 + 5.1 + 4.9 + 4.4 = 19$$

The symbol

$$\sum_{i=1}^{3} X_i$$

indicates that only the first 3 scores are to be summed. The index variable i goes from 1 to 3.

When all the scores of a variable (such as **X**) are to be summed, it is often convenient to use the following abbreviated notation:

$$\sum \mathrm{X}$$

Thus, when no values of i are shown, it means to sum all the values of **X**.

Many formulas involve squaring numbers before they are summed. This is indicated as

$$\begin{aligned}\sum X^2 &= 4.6^2 + 5.1^2 + 4.9^2 + 4.4^2 \\ &= 21.16 + 26.01 + 24.01 + 19.36 = 90.54\end{aligned}$$

Notice that:

$$\left(\sum \mathrm{X}\right)^2 \neq \sum \mathrm{X}^2 \tag{1.7.2}$$

because the expression on the left means to sum up all the values of **X** and then square the sum ($19^2 = 361$), whereas the expression on the right means to square the numbers and then sum the squares (90.54, as shown).

Some formulas involve the sum of cross products. Below are the data for variables **X** and **Y**. The cross products (**XY**) are shown in the third column. The sum of the cross products is 3 + 4 + 21 = 28.

Table **1.7.2**

X	Y	XY
1	3	3
2	2	4
3	7	21

In summation notation, this is written as:

$$\sum XY = 28$$

This page titled 1.7: Mathematical Notation is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.7: Mathematical Notation** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

1.E: Introduction (Exercises)

1. In your own words, describe why we study statistics.

Answer:

Your answer could take many forms but should include information about objectively interpreting information and/or communicating results and research conclusions

2. For each of the following, determine if the variable is continuous or discrete:
   a. Time taken to read a book chapter
   b. Favorite food
   c. Cognitive ability
   d. Temperature
   e. Letter grade received in a class
3. For each of the following, determine the level of measurement:
   a. T-shirt size
   b. Time taken to run 100 meter race
   c. First, second, and third place in 100 meter race
   d. Birthplace
   e. Temperature in Celsius

Answer:

   a. Ordinal
   b. Ratio
   c. Ordinal
   d. Nominal
   e. Interval

4. What is the difference between a population and a sample? Which is described by a parameter and which is described by a statistic?
5. What is sampling bias? What is sampling error?

Answer:

Sampling bias is the difference in demographic characteristics between a sample and the population it should represent. Sampling error is the difference between a population parameter and sample statistic that is caused by random chance due to sampling bias.

6. What is the difference between a simple random sample and a stratified random sample?
7. What are the two key characteristics of a true experimental design?

Answer:

Random assignment to treatment conditions and manipulation of the independent variable 9

8. When would we use a quasi-experimental design?
9. Use the following dataset for the computations below:

X	Y
2	8
3	8
7	4
5	1
9	4

a. $\sum X$
b. $\sum Y^2$
c. $\sum XY$
d. $(\sum Y)^2$

Answer:

a. 26
b. 161
c. 109
d. 625

10. What are the most common measures of central tendency and spread?

This page titled 1.E: Introduction (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **1.E: Introduction (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

2: Describing Data using Distributions and Graphs

2.1: Graphing Qualitative Variables

2.2: Graphing Quantitative Variables

2.E: Describing Data using Distributions and Graphs (Exercises)

This page titled 2: Describing Data using Distributions and Graphs is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

2.1: Graphing Qualitative Variables

When Apple Computer introduced the iMac computer in August 1998, the company wanted to learn whether the iMac was expanding Apple's market share. Was the iMac just attracting previous Macintosh owners? Or was it purchased by newcomers to the computer market and by previous Windows users who were switching over? To find out, 500 iMac customers were interviewed. Each customer was categorized as a previous Macintosh owner, a previous Windows owner, or a new computer purchaser.

This section examines graphical methods for displaying the results of the interviews. We'll learn some general lessons about how to graph data that fall into a small number of categories. A later section will consider how to graph numerical data in which each observation is represented by a number in some range. The key point about the qualitative data that occupy us in the present section is that they do not come with a pre-established ordering (the way numbers are ordered). For example, there is no natural sense in which the category of previous Windows users comes before or after the category of previous Macintosh users. This situation may be contrasted with quantitative data, such as a person's weight. People of one weight are naturally ordered with respect to people of a different weight.

Frequency Tables

All of the graphical methods shown in this section are derived from frequency tables. Table 1 shows a frequency table for the results of the iMac study; it shows the frequencies of the various response categories. It also shows the relative frequencies, which are the proportion of responses in each category. For example, the relative frequency for "none" of 0.17 = 85/500.

Table **2.1.1**: Frequency Table for the iMac Data.

Previous Ownership	Frequency	Relative Frequency
None	85	0.17
Windows	60	0.12
Macintosh	355	0.71
Total	500	1

Pie Charts

The pie chart in Figure **2.1.1** shows the results of the iMac study. In a pie chart, each category is represented by a slice of the pie. The area of the slice is proportional to the percentage of responses in the category. This is simply the relative frequency multiplied by 100. Although most iMac purchasers were Macintosh owners, Apple was encouraged by the 12% of purchasers who were former Windows users, and by the 17% of purchasers who were buying a computer for the first time.

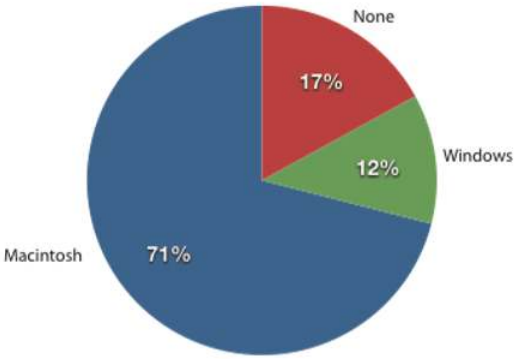


Figure **2.1.1**: Pie chart of iMac purchases illustrating frequencies of previous computer ownership.

Pie charts are effective for displaying the relative frequencies of a small number of categories. They are not recommended, however, when you have a large number of categories. Pie charts can also be confusing when they are used to compare the

outcomes of two different surveys or experiments. In an influential book on the use of graphs, Edward Tufte asserted "The only worse design than a pie chart is several of them." Here is another important point about pie charts. If they are based on a small number of observations, it can be misleading to label the pie slices with percentages. For example, if just 5 people had been interviewed by Apple Computers, and 3 were former Windows users, it would be misleading to display a pie chart with the Windows slice showing 60%. With so few people interviewed, such a large percentage of Windows users might easily have occurred since chance can cause large errors with small samples. In this case, it is better to alert the user of the pie chart to the actual numbers involved. The slices should therefore be labeled with the actual frequencies observed (e.g., 3) instead of with percentages.

Bar charts

Bar charts can also be used to represent frequencies of different categories. A bar chart of the iMac purchases is shown in Figure **2.1.2**. Frequencies are shown on the Y-axis and the type of computer previously owned is shown on the X-axis. Typically, the Y-axis shows the number of observations in each category rather than the percentage of observations in each category as is typical in pie charts.

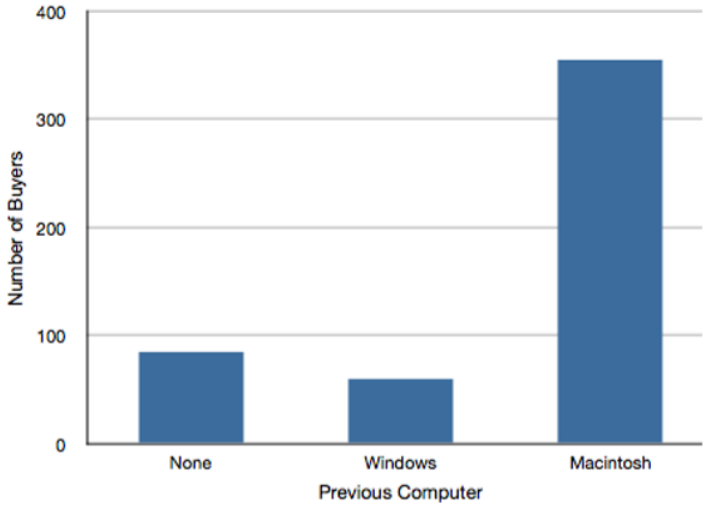


Figure **2.1.2**: Bar chart of iMac purchases as a function of previous computer ownership.

Comparing Distributions

Often we need to compare the results of different surveys, or of different conditions within the same overall survey. In this case, we are comparing the "distributions" of responses between the surveys or conditions. Bar charts are often excellent for illustrating differences between two distributions. Figure **2.1.3** shows the number of people playing card games at the Yahoo web site on a Sunday and on a Wednesday in the spring of 2001. We see that there were more players overall on Wednesday compared to Sunday. The number of people playing Pinochle was nonetheless the same on these two days. In contrast, there were about twice as many people playing hearts on Wednesday as on Sunday. Facts like these emerge clearly from a well-designed bar chart.

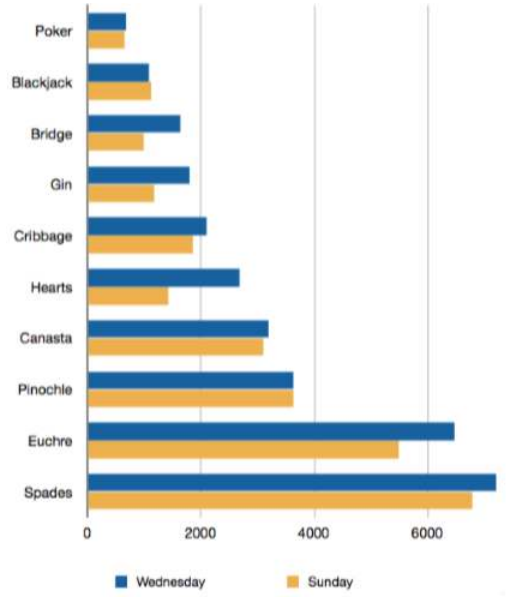


Figure **2.1.3**: A bar chart of the number of people playing different card games on Sunday and Wednesday.

The bars in Figure **2.1.3** are oriented horizontally rather than vertically. The horizontal format is useful when you have many categories because there is more room for the category labels. We'll have more to say about bar charts when we consider numerical quantities later in this chapter.

Some graphical mistakes to avoid

Don't get fancy! People sometimes add features to graphs that don't help to convey their information. For example, 3-dimensional bar charts such as the one shown in Figure **2.1.4** are usually not as effective as their two-dimensional counterparts.

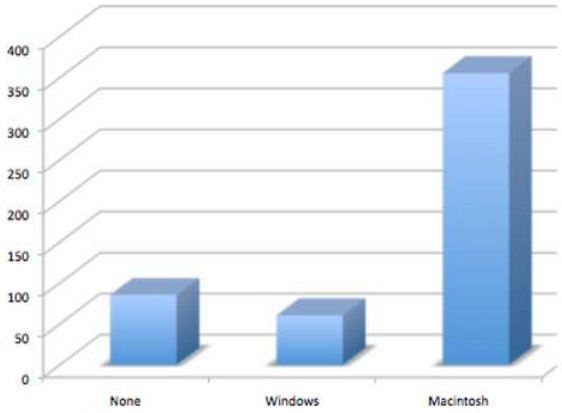


Figure **2.1.4**: A three-dimensional version of Figure **2.1.2**

Here is another way that fanciness can lead to trouble. Instead of plain bars, it is tempting to substitute meaningful images. For example, Figure **2.1.5** presents the iMac data using pictures of computers. The heights of the pictures accurately represent the number of buyers, yet Figure **2.1.5** is misleading because the viewer's attention will be captured by areas. The areas can exaggerate the size differences between the groups. In terms of percentages, the ratio of previous Macintosh owners to previous Windows owners is about 6 to 1. But the ratio of the two areas in Figure **2.1.5** is about 35 to 1. A biased person wishing to hide the fact that many Windows owners purchased iMacs would be tempted to use Figure **2.1.5** instead of Figure **2.1.2**! Edward Tufte coined the term "lie factor" to refer to the ratio of the size of the effect shown in a graph to the size of the effect shown in the data. He suggests that lie factors greater than 1.05 or less than 0.95 produce unacceptable distortion.

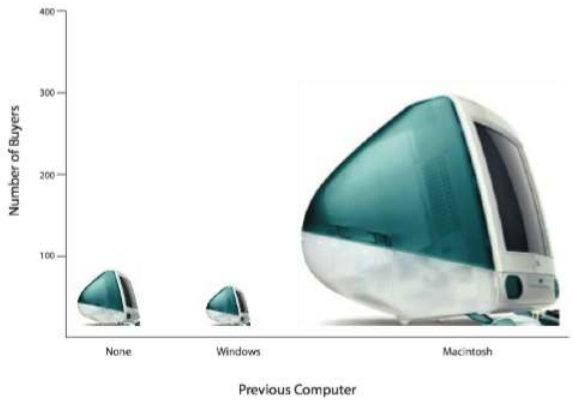


Figure 2.1.5: A redrawing of Figure 2.1.2 with a lie factor greater than 8.

Another distortion in bar charts results from setting the baseline to a value other than zero. The baseline is the bottom of the Y-axis, representing the least number of cases that could have occurred in a category. Normally, but not always, this number should be zero. Figure 6 shows the iMac data with a baseline of 50. Once again, the differences in areas suggests a different story than the true differences in percentages. The number of Windows-switchers seems minuscule compared to its true value of 12%.

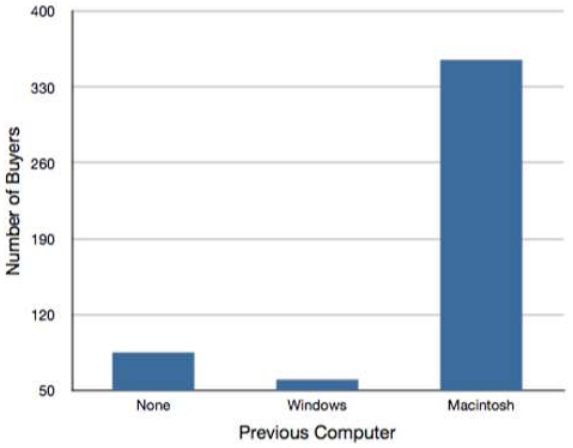


Figure 2.1.6: A redrawing of Figure 2.1.2 with a baseline of 50.

Finally, we note that it is a serious mistake to use a line graph when the X-axis contains merely qualitative variables. A line graph is essentially a bar graph with the tops of the bars represented by points joined by lines (the rest of the bar is suppressed). Figure 2.1.7 inappropriately shows a line graph of the card game data from Yahoo. The drawback to Figure 7 is that it gives the false impression that the games are naturally ordered in a numerical way when, in fact, they are ordered alphabetically.

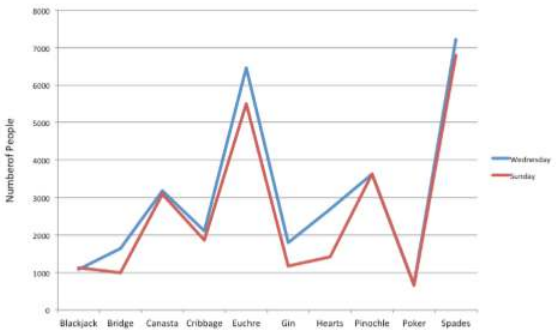


Figure 2.1.7: A line graph used inappropriately to depict the number of people playing different card games on Sunday and Wednesday.

Summary

Pie charts and bar charts can both be effective methods of portraying qualitative data. Bar charts are better when there are more than just a few categories and for comparing two or more distributions. Be careful to avoid creating misleading graphs.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 2.1: Graphing Qualitative Variables is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **2.1: Graphing Qualitative Variables** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

2.2: Graphing Quantitative Variables

As discussed in the section on variables in Chapter 1, quantitative variables are variables measured on a numeric scale. Height, weight, response time, subjective rating of pain, temperature, and score on an exam are all examples of quantitative variables. Quantitative variables are distinguished from categorical (sometimes called qualitative) variables such as favorite color, religion, city of birth, favorite sport in which there is no ordering or measuring involved.

There are many types of graphs that can be used to portray distributions of quantitative variables. The upcoming sections cover the following types of graphs:

1. stem and leaf displays
2. histograms
3. frequency polygons
4. box plots
5. bar charts
6. line graphs
7. dot plots
8. scatter plots (discussed in a different chapter)

Some graph types such as stem and leaf displays are best-suited for small to moderate amounts of data, whereas others such as histograms are best suited for large amounts of data. Graph types such as box plots are good at depicting differences between distributions. Scatter plots are used to show the relationship between two variables.

Stem and Leaf Displays

A stem and leaf display is a graphical method of displaying data. It is particularly useful when your data are not too numerous. In this section, we will explain how to construct and interpret this kind of graph.

As usual, we will start with an example. Consider Figure **2.2.1** that shows the number of touchdown passes (TD passes) thrown by each of the 31 teams in the National Football League in the 2000 season.

37, 33, 33, 32, 29, 28,
28, 23, 22, 22, 22, 21,
21, 21, 20, 20, 19, 19,
18, 18, 18, 18, 16, 15,
14, 14, 14, 12, 12, 9, 6

Figure **2.2.1**: Number of touchdown passes.

A stem and leaf display of the data is shown in Figure **2.2.2**. The left portion of Figure **2.2.2** contains the stems. They are the numbers 3, 2, 1, and 0, arranged as a column to the left of the bars. Think of these numbers as 10's digits. A stem of 3, for example, can be used to represent the 10's digit in any of the numbers from 30 to 39. The numbers to the right of the bar are leaves, and they represent the 1's digits. Every leaf in the graph therefore stands for the result of adding the leaf to 10 times its stem.

3|2337
2|001112223889
1|2244456888899
0|69

Figure **2.2.2**: Stem and leaf display of the number of touchdown passes.

To make this clear, let us examine Figure **2.2.2** more closely. In the top row, the four leaves to the right of stem 3 are 2, 3, 3, and 7. Combined with the stem, these leaves represent the numbers 32, 33, 33, and 37, which are the numbers of TD passes for the first four teams in Figure **2.2.1**. The next row has a stem of 2 and 12 leaves. Together, they represent 12 data points, namely, two occurrences of 20 TD passes, three occurrences of 21 TD passes, three occurrences of 22 TD passes, one occurrence of 23 TD passes, two occurrences of 28 TD passes, and one occurrence of 29 TD passes. We leave it to you to figure out what the third row represents. The fourth row has a stem of 0 and two leaves. It stands for the last two entries in Figure **2.2.1**, namely 9 TD passes and 6 TD passes. (The latter two numbers may be thought of as 09 and 06.)

One purpose of a stem and leaf display is to clarify the shape of the distribution. You can see many facts about TD passes more easily in Figure **2.2.2** than in Figure **2.2.1**. For example, by looking at the stems and the shape of the plot, you can tell that most of the teams had between 10 and 29 passing TD's, with a few having more and a few having less. The precise numbers of TD passes can be determined by examining the leaves.

We can make our figure even more revealing by splitting each stem into two parts. Figure **2.2.3** shows how to do this. The top row is reserved for numbers from 35 to 39 and holds only the 37 TD passes made by the first team in Figure **2.2.1**. The second row is reserved for the numbers from 30 to 34 and holds the 32, 33, and 33 TD passes made by the next three teams in the table. You can see for yourself what the other rows represent.

```
3|7
3|233
2|889
2|001112223
1|56888899
1|22444
0|69
```

Figure **2.2.3**: Stem and leaf display with the stems split in two.

Figure **2.2.3** is more revealing than Figure **2.2.2** because the latter figure lumps too many values into a single row. Whether you should split stems in a display depends on the exact form of your data. If rows get too long with single stems, you might try splitting them into two or more parts.

There is a variation of stem and leaf displays that is useful for comparing distributions. The two distributions are placed back to back along a common column of stems. The result is a “back-to-back stem and leaf display.” Figure **2.2.4** shows such a graph. It compares the numbers of TD passes in the 1998 and 2000 seasons. The stems are in the middle, the leaves to the left are for the 1998 data, and the leaves to the right are for the 2000 data. For example, the second-to-last row shows that in 1998 there were teams with 11, 12, and 13 TD passes, and in 2000 there were two teams with 12 and three teams with 14 TD passes.

11	4	
	3	7
332	3	233
8865	2	889
44331110	2	001112223
987776665	1	56888899
321	1	22444
7	0	69

Figure **2.2.4**: Back-to-back stem and leaf display. The left side shows the 1998 TD data and the right side shows the 2000 TD data.

Figure **2.2.4** helps us see that the two seasons were similar, but that only in 1998 did any teams throw more than 40 TD passes.

There are two things about the football data that make them easy to graph with stems and leaves. First, the data are limited to whole numbers that can be represented with a one-digit stem and a one-digit leaf. Second, all the numbers are positive. If the data include numbers with three or more digits, or contain decimals, they can be rounded to two-digit accuracy. Negative values are also easily handled. Let us look at another example.

Figure **2.2.5** shows data from the case study Weapons and Aggression. Each value is the mean difference over a series of trials between the times it took an experimental subject to name aggressive words (like “punch”) under two conditions. In one condition, the words were preceded by a non-weapon word such as “bug.” In the second condition, the same words were preceded by a weapon word such as “gun” or “knife.” The issue addressed by the experiment was whether a preceding weapon word would speed up (or prime) pronunciation of the aggressive word compared to a non-weapon priming word. A positive difference implies greater priming of the aggressive word by the weapon word. Negative differences imply that the priming by the weapon word was less than for a neutral word.

```
43.2, 42.9, 35.6, 25.6, 25.4, 23.6, 20.5, 19.9, 14.4, 12.7, 11.3,
10.2, 10.0, 9.1, 7.5, 5.4, 4.7, 3.8, 2.1, 1.2, -0.2, -6.3, -6.7,
-8.8, -10.4, -10.5, -14.9, -14.9, -15.0, -18.5, -27.4
```

Figure **2.2.5**: The effects of priming (thousandths of a second).

You see that the numbers range from 43.2 to -27.4. The first value indicates that one subject was 43.2 milliseconds faster pronouncing aggressive words when they were preceded by weapon words than when preceded by neutral words. The value -27.4 indicates that another subject was 27.4 milliseconds slower pronouncing aggressive words when they were preceded by weapon words.

The data are displayed with stems and leaves in Figure **2.2.6**. Since stem and leaf displays can only portray two whole digits (one for the stem and one for the leaf) the numbers are first rounded. Thus, the value 43.2 is rounded to 43 and represented with a stem of 4 and a leaf of 3. Similarly, 42.9 is rounded to 43. To represent negative numbers, we simply use negative stems. For example, the bottom row of the figure represents the number –27. The second-to-last row represents the numbers -10, -10, -15, etc. Once again, we have rounded the original values from Figure **2.2.5**.

```
 4|33
 3|6
 2|00456
 1|00134
 0|1245589
-0|0679
-1|005559
-2|7
```

Figure **2.2.6**: Stem and leaf display with negative numbers and rounding.

Observe that the figure contains a row headed by “0” and another headed by “-0.” The stem of 0 is for numbers between 0 and 9, whereas the stem of -0 is for numbers between 0 and -9. For example, the fifth row of the table holds the numbers 1, 2, 4, 5, 5, 8, 9 and the sixth row holds 0, -6, -7, and -9. Values that are exactly 0 before rounding should be split as evenly as possible between the “0” and “-0” rows. In Figure **2.2.5**, none of the values are 0 before rounding. The “0” that appears in the “-0” row comes from the original value of -0.2 in the table.

Although stem and leaf displays are unwieldy for large data sets, they are often useful for data sets with up to 200 observations. Figure **2.2.7** portrays the distribution of populations of 185 US cities in 1998. To be included, a city had to have between 100,000 and 500,000 residents.

```
4|899
4|6
4|4455
4|333
4|01
3|99
3|677777
3|55
3|223
3|111
2|8899
2|666667
2|444455
2|22333
2|000000
1|88888888888899999999999
1|666666777777
1|444444444445555555555
1|22222222222222222333333333
1|00000000000000011111111111111111111111111
```

Figure **2.2.7**: Stem and leaf display of populations of 185 US cities with populations between 100,000 and 500,000 in 1988.

Since a stem and leaf plot shows only two-place accuracy, we had to round the numbers to the nearest 10,000. For example the largest number (493,559) was rounded to 490,000 and then plotted with a stem of 4 and a leaf of 9. The fourth highest number (463,201) was rounded to 460,000 and plotted with a stem of 4 and a leaf of 6. Thus, the stems represent units of 100,000 and the leaves represent units of 10,000. Notice that each stem value is split into five parts: 0-1, 2-3, 4-5, 67, and 8-9.

Whether your data can be suitably represented by a stem and leaf display depends on whether they can be rounded without loss of important information. Also, their extreme values must fit into two successive digits, as the data in Figure 11 fit into the 10,000 and 100,000 places (for leaves and stems, respectively). Deciding what kind of graph is best suited to displaying your data thus requires good judgment. Statistics is not just recipes!

Histograms

A histogram is a graphical method for displaying the shape of a distribution. It is particularly useful when there are a large number of observations. We begin with an example consisting of the scores of 642 students on a psychology test. The test consists of 197 items each graded as “correct” or “incorrect.” The students' scores ranged from 46 to 167.

The first step is to create a frequency table. Unfortunately, a simple frequency table would be too big, containing over 100 rows. To simplify the table, we group scores together as shown in Table **2.2.1**.

Table **2.2.1**: Grouped Frequency Distribution of Psychology Test Scores

Interval's Lower Limit	Interval's Upper Limit	Class Frequency
39.5	49.5	3
49.5	59.5	10
59.5	69.5	53
69.5	79.5	107
79.5	89.5	147
89.5	99.5	130
99.5	109.5	78
109.5	119.5	59
119.5	129.5	36
129.5	139.5	11
139.5	149.5	6
149.5	159.5	1
159.5	169.5	1

To create this table, the range of scores was broken into intervals, called class intervals. The first interval is from 39.5 to 49.5, the second from 49.5 to 59.5, etc. Next, the number of scores falling into each interval was counted to obtain the class frequencies. There are three scores in the first interval, 10 in the second, etc.

Class intervals of width 10 provide enough detail about the distribution to be revealing without making the graph too “choppy.” More information on choosing the widths of class intervals is presented later in this section. Placing the limits of the class intervals midway between two numbers (e.g., 49.5) ensures that every score will fall in an interval rather than on the boundary between intervals.

In a histogram, the class frequencies are represented by bars. The height of each bar corresponds to its class frequency. A histogram of these data is shown in Figure **2.2.8**.

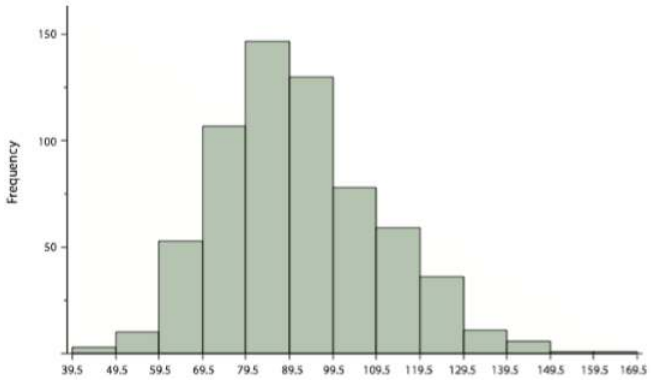


Figure **2.2.8**: Histogram of scores on a psychology test.

The histogram makes it plain that most of the scores are in the middle of the distribution, with fewer scores in the extremes. You can also see that the distribution is not symmetric: the scores extend to the right farther than they do to the left. The distribution is therefore said to be skewed. (We'll have more to say about shapes of distributions in Chapter 3.)

In our example, the observations are whole numbers. Histograms can also be used when the scores are measured on a more continuous scale such as the length of time (in milliseconds) required to perform a task. In this case, there is no need to worry about fence sitters since they are improbable. (It would be quite a coincidence for a task to require exactly 7 seconds, measured to the nearest thousandth of a second.) We are therefore free to choose whole numbers as boundaries for our class intervals, for example, 4000, 5000, etc. The class frequency is then the number of observations that are greater than or equal to the lower bound, and strictly less than the upper bound. For example, one interval might hold times from 4000 to 4999 milliseconds. Using whole numbers as boundaries avoids a cluttered appearance, and is the practice of many computer programs that create histograms. Note also that some computer programs label the middle of each interval rather than the end points.

Histograms can be based on relative frequencies instead of actual frequencies. Histograms based on relative frequencies show the proportion of scores in each interval rather than the number of scores. In this case, the Y-axis runs from 0 to 1 (or somewhere in between if there are no extreme proportions). You can change a histogram based on frequencies to one based on relative frequencies by (a) dividing each class frequency by the total number of observations, and then (b) plotting the quotients on the Y-axis (labeled as proportion).

There is more to be said about the widths of the class intervals, sometimes called bin widths. Your choice of bin width determines the number of class intervals. This decision, along with the choice of starting point for the first interval, affects the shape of the histogram. The best advice is to experiment with different choices of width, and to choose a histogram according to how well it communicates the shape of the distribution.

Frequency Polygons

Frequency polygons are a graphical device for understanding the shapes of distributions. They serve the same purpose as histograms, but are especially helpful for comparing sets of data. Frequency polygons are also a good choice for displaying cumulative frequency distributions.

To create a frequency polygon, start just as for histograms, by choosing a class interval. Then draw an X-axis representing the values of the scores in your data. Mark the middle of each class interval with a tick mark, and label it with the middle value represented by the class. Draw the Y-axis to indicate the frequency of each class. Place a point in the middle of each class interval at the height corresponding to its frequency. Finally, connect the points. You should include one class interval below the lowest value in your data and one above the highest value. The graph will then touch the X-axis on both sides.

A frequency polygon for 642 psychology test scores shown in Figure **2.2.8** was constructed from the frequency table shown in Table **2.2.2**.

Table **2.2.2**: Frequency Distribution of Psychology Test Scores

Lower Limit	Upper Limit	Count	Cumulative Count
29.5	39.5	0	0
39.5	49.5	3	3
49.5	59.5	10	13
59.5	69.5	53	66
69.5	79.5	107	173
79.5	89.5	147	320
89.5	99.5	130	450
99.5	109.5	78	528
109.5	119.5	59	587
119.5	129.5	36	623
129.5	139.5	11	634
139.5	149.5	6	640
149.5	159.5	1	641
159.5	169.5	1	642
169.5	170.5	0	642

The first label on the X-axis is 35. This represents an interval extending from 29.5 to 39.5. Since the lowest test score is 46, this interval has a frequency of 0. The point labeled 45 represents the interval from 39.5 to 49.5. There are three scores in this interval. There are 147 scores in the interval that surrounds 85.

You can easily discern the shape of the distribution from Figure **2.2.9**. Most of the scores are between 65 and 115. It is clear that the distribution is not symmetric inasmuch as good scores (to the right) trail off more gradually than poor scores (to the left). In the terminology of Chapter 3 (where we will study shapes of distributions more systematically), the distribution is skewed.

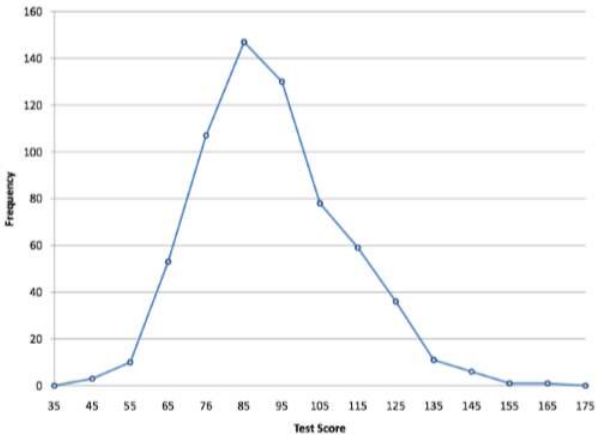


Figure **2.2.9**: Frequency polygon for the psychology test scores.

A cumulative frequency polygon for the same test scores is shown in Figure **2.2.10**. The graph is the same as before except that the Y value for each point is the number of students in the corresponding class interval plus all numbers in lower intervals. For example, there are no scores in the interval labeled “35,” three in the interval “45,” and 10 in the interval “55.” Therefore, the Y value corresponding to “55” is 13. Since 642 students took the test, the cumulative frequency for the last interval is 642.

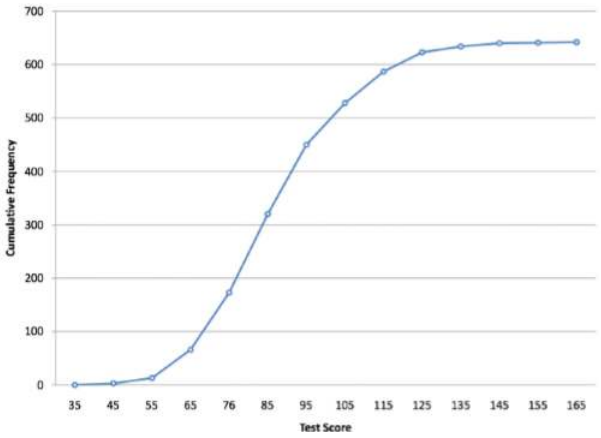


Figure **2.2.10**: Cumulative frequency polygon for the psychology test scores.

Frequency polygons are useful for comparing distributions. This is achieved by overlaying the frequency polygons drawn for different data sets. Figure 2.1.3 provides an example. The data come from a task in which the goal is to move a computer cursor to a target on the screen as fast as possible. On 20 of the trials, the target was a small rectangle; on the other 20, the target was a large rectangle. Time to reach the target was recorded on each trial. The two distributions (one for each target) are plotted together in Figure **2.2.11**. The figure shows that, although there is some overlap in times, it generally took longer to move the cursor to the small target than to the large one.

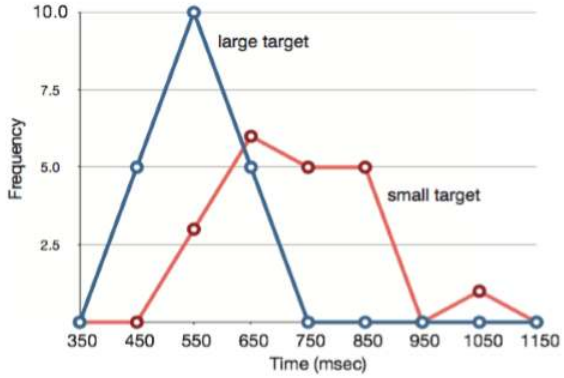


Figure **2.2.11**: Overlaid frequency polygons.

It is also possible to plot two cumulative frequency distributions in the same graph. This is illustrated in Figure **2.2.12** using the same data from the cursor task. The difference in distributions for the two targets is again evident.

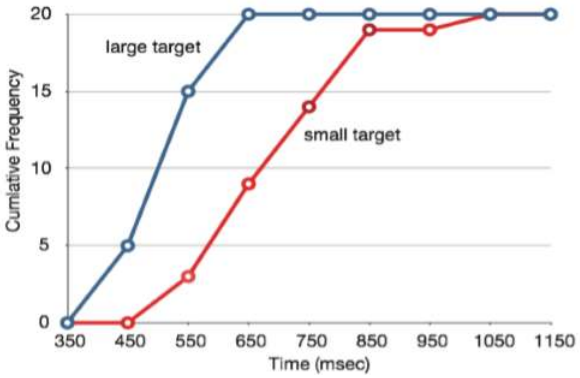


Figure **2.2.12**: Overlaid cumulative frequency polygons.

Box Plots

We have already discussed techniques for visually representing data (see histograms and frequency polygons). In this section we present another important graph, called a box plot. Box plots are useful for identifying outliers and for comparing distributions. We will explain box plots with the help of data from an in-class experiment. Students in Introductory Statistics were presented with a page containing 30 colored rectangles. Their task was to name the colors as quickly as possible. Their times (in seconds) were recorded. We'll compare the scores for the 16 men and 31 women who participated in the experiment by making separate box plots for each gender. Such a display is said to involve parallel box plots.

There are several steps in constructing a box plot. The first relies on the 25th, 50th, and 75th percentiles in the distribution of scores. Figure **2.2.14** shows how these three statistics are used. For each gender we draw a box extending from the 25th percentile to the 75th percentile. The 50th percentile is drawn inside the box. Therefore, the bottom of each box is the 25th percentile, the top is the 75th percentile, and the line in the middle is the 50th percentile. The data for the women in our sample are shown in Figure **2.2.13**.

14, 15, 16, 16, 17, 17, 17, 17, 17, 18, 18, 18, 18, 18, 18, 19, 19, 19
20, 20, 20, 20, 20, 20, 21, 21, 22, 23, 24, 24, 29

Figure **2.2.13**: Women's times.

For these data, the 25th percentile is 17, the 50th percentile is 19, and the 75th percentile is 20. For the men (whose data are not shown), the 25th percentile is 19, the 50th percentile is 22.5, and the 75th percentile is 25.5.

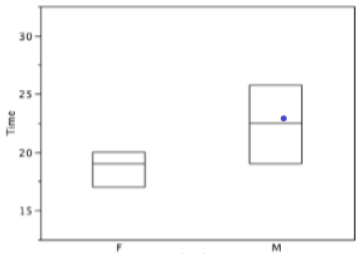


Figure **2.2.14**: The first step in creating box plots.

Before proceeding, the terminology in Table **2.2.3** is helpful.

Table **2.2.3**: Box plot terms and values for women's times.

Name	Formula	Value
Upper Hinge	75th Percentile	20
Lower Hinge	25th Percentile	17
H-Spread	Upper Hinge - Lower Hinge	3
Step	1.5 x H-Spread	4.5
Upper Inner Fence	Upper Hinge + 1 Step	24.5
Lower Inner Fence	Lower Hinge - 1 Step	12.5
Upper Outer Fence	Upper Hinge + 2 Steps	29
Lower Outer Fence	Lower Hinge - 2 Steps	8
Upper Adjacent	Largest value below Upper Inner Fence	24
Lower Adjacent	Smallest value above Lower Inner Fence	14
Outside Value	A value beyond an Inner Fence but not beyond an Outer Fence	29
Far Out Value	A value beyond an Outer Fence	None

Continuing with the box plots, we put “whiskers” above and below each box to give additional information about the spread of data. Whiskers are vertical lines that end in a horizontal stroke. Whiskers are drawn from the upper and lower hinges to the upper and lower adjacent values (24 and 14 for the women's data), as shown in Figure **2.2.15**.

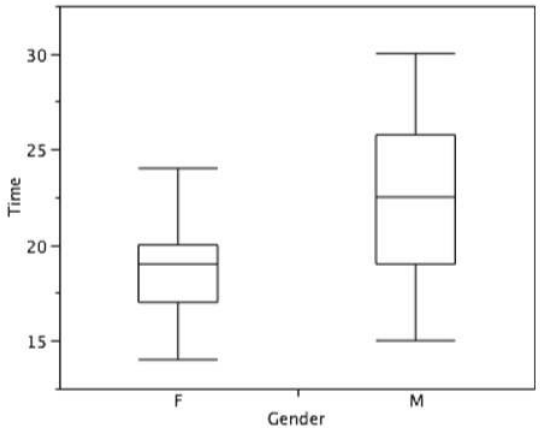


Figure **2.2.15**: The box plots with the whiskers drawn.

Although we don't draw whiskers all the way to outside or far out values, we still wish to represent them in our box plots. This is achieved by adding additional marks beyond the whiskers. Specifically, outside values are indicated by small “o's” and far out values are indicated by asterisks (*). In our data, there are no farout values and just one outside value. This outside value of 29 is for the women and is shown in Figure **2.2.16**.

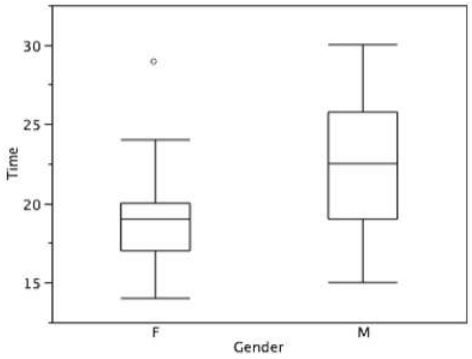


Figure **2.2.16**: The box plots with the outside value shown.

There is one more mark to include in box plots (although sometimes it is omitted). We indicate the mean score for a group by inserting a plus sign. Figure **2.2.17** shows the result of adding means to our box plots.

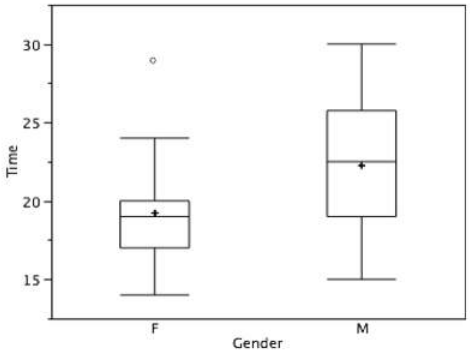


Figure **2.2.17**: The completed box plots.

Figure 2.2.17 provides a revealing summary of the data. Since half the scores in a distribution are between the hinges (recall that the hinges are the 25th and 75th percentiles), we see that half the women's times are between 17 and 20 seconds whereas half the men's times are between 19 and 25.5 seconds. We also see that women generally named the colors faster than the men did, although one woman was slower than almost all of the men. Figure 2.2.18 shows the box plot for the women's data with detailed labels.

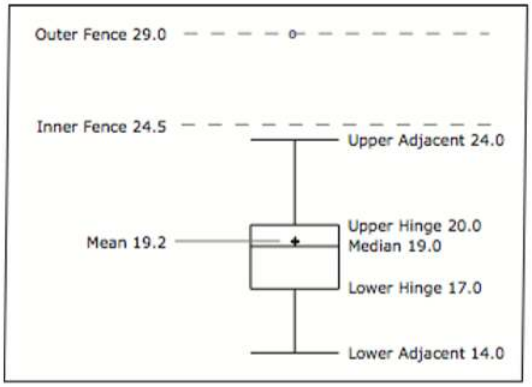


Figure 2.2.18: The box plots for the women's data with detailed labels.

Box plots provide basic information about a distribution. For example, a distribution with a positive skew would have a longer whisker in the positive direction than in the negative direction. A larger mean than median would also indicate a positive skew. Box plots are good at portraying extreme values and are especially good at showing differences between distributions. However, many of the details of a distribution are not revealed in a box plot and to examine these details one should use create a histogram and/or a stem and leaf display.

Bar Charts

In the section on qualitative variables, we saw how bar charts could be used to illustrate the frequencies of different categories. For example, the bar chart shown in Figure 2.2.19 shows how many purchasers of iMac computers were previous Macintosh users, previous Windows users, and new computer purchasers.

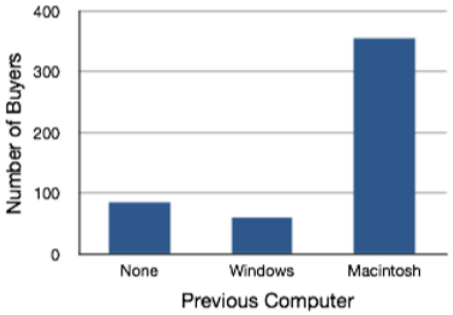


Figure 2.2.19: iMac buyers as a function of previous computer ownership.

In this section we show how bar charts can be used to present other kinds of quantitative information, not just frequency counts. The bar chart in Figure 2.2.20 shows the percent increases in the Dow Jones, Standard and Poor 500 (S & P), and Nasdaq stock indexes from May 24th 2000 to May 24th 2001. Notice that both the S & P and the Nasdaq had “negative increases” which means that they decreased in value. In this bar chart, the Y-axis is not frequency but rather the signed quantity percentage increase.

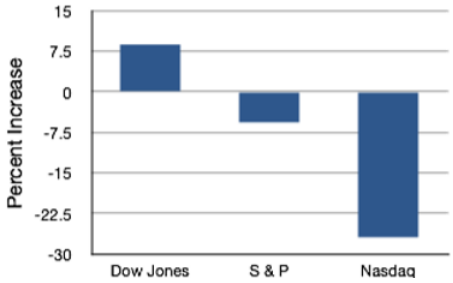


Figure 2.2.20: Percent increase in three stock indexes from May 24th 2000 to May 24th 2001.

Bar charts are particularly effective for showing change over time. Figure 2.2.21, for example, shows the percent increase in the Consumer Price Index (CPI) over four three-month periods. The fluctuation in inflation is apparent in the graph.

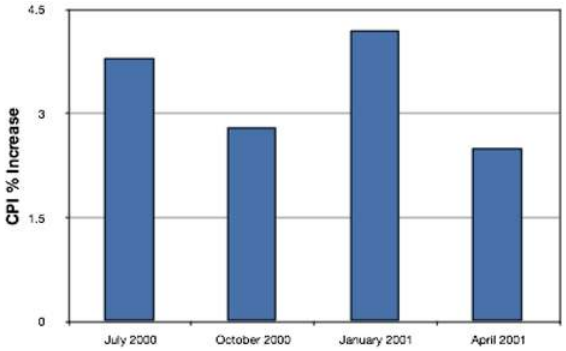


Figure 2.2.21: Percent change in the CPI over time. Each bar represents percent increase for the three months ending at the date indicated.

Bar charts are often used to compare the means of different experimental conditions. Figure 2.1.4 shows the mean time it took one of us (DL) to move the cursor to either a small target or a large target. On average, more time was required for small targets than for large ones.

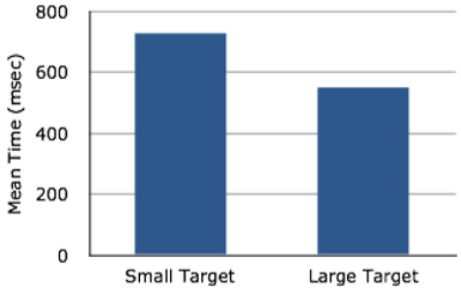


Figure 2.2.22: Bar chart showing the means for the two conditions.

Although bar charts can display means, we do not recommend them for this purpose. Box plots should be used instead since they provide more information than bar charts without taking up more space. For example, a box plot of the cursor-movement data is shown in Figure 2.2.23. You can see that Figure 2.2.23 reveals more about the distribution of movement times than does Figure 2.2.22.

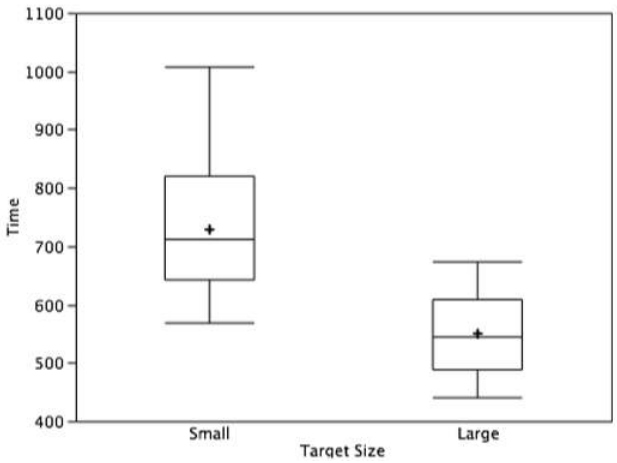


Figure 2.2.23: Box plots of times to move the cursor to the small and large targets.

The section on qualitative variables presented earlier in this chapter discussed the use of bar charts for comparing distributions. Some common graphical mistakes were also noted. The earlier discussion applies equally well to the use of bar charts to display quantitative variables.

Line Graphs

A line graph is a bar graph with the tops of the bars represented by points joined by lines (the rest of the bar is suppressed). For example, Figure 2.2.24 was presented in the section on bar charts and shows changes in the Consumer Price Index (CPI) over time.

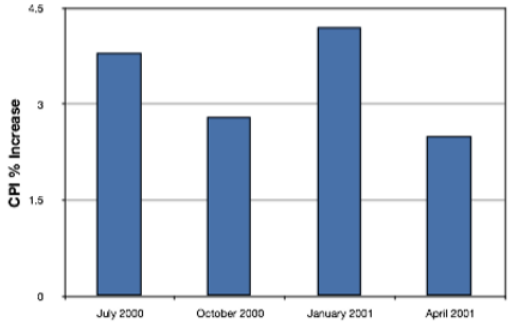


Figure 2.2.24: A bar chart of the percent change in the CPI over time. Each bar represents percent increase for the three months ending at the date indicated.

A line graph of these same data is shown in Figure 2.2.25. Although the figures are similar, the line graph emphasizes the change from period to period.

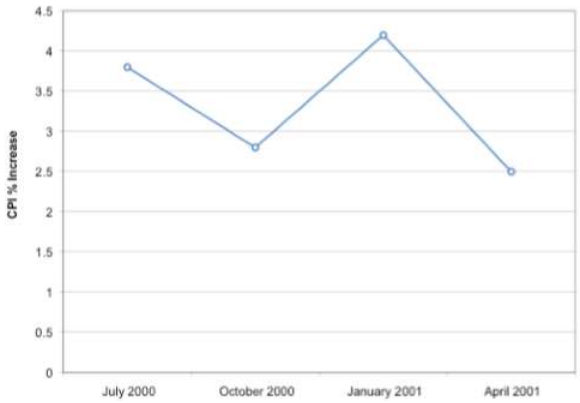


Figure 2.2.25: A line graph of the percent change in the CPI over time. Each point represents percent increase for the three months ending at the date indicated.

Line graphs are appropriate only when both the X- and Y-axes display ordered (rather than qualitative) variables. Although bar charts can also be used in this situation, line graphs are generally better at comparing changes over time. Figure **2.2.26**, for example, shows percent increases and decreases in five components of the CPI. The figure makes it easy to see that medical costs had a steadier progression than the other components. Although you could create an analogous bar chart, its interpretation would not be as easy.

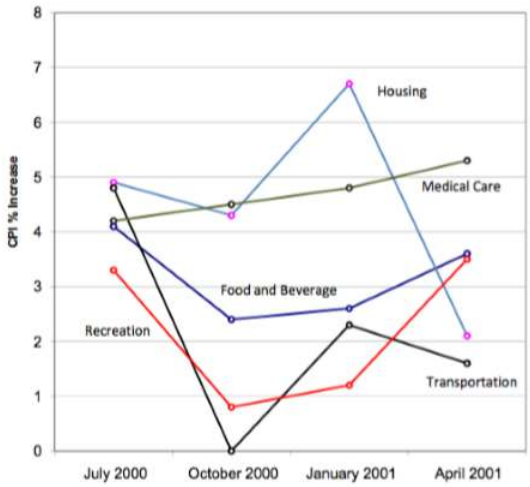


Figure **2.2.26**: A line graph of the percent change in five components of the CPI over time.

Let us stress that it is misleading to use a line graph when the X-axis contains merely qualitative variables. Figure **2.2.27** inappropriately shows a line graph of the card game data from Yahoo, discussed in the section on qualitative variables. The defect in Figure **2.2.27** is that it gives the false impression that the games are naturally ordered in a numerical way.

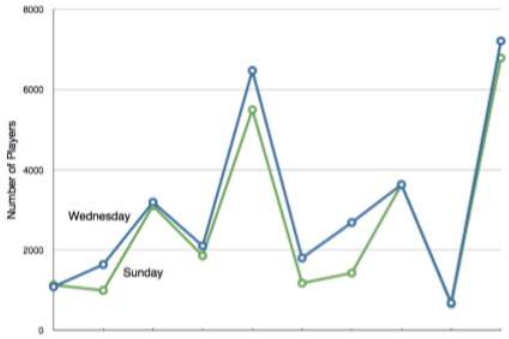


Figure **2.2.27**: A line graph, inappropriately used, depicting the number of people playing different card games on Wednesday and Sunday.

The Shape of Distribution

Finally, it is useful to present discussion on how we describe the shapes of distributions, which we will revisit in the next chapter to learn how different shapes affect our numerical descriptors of data and distributions.

The primary characteristic we are concerned about when assessing the shape of a distribution is whether the distribution is symmetrical or skewed. A symmetrical distribution, as the name suggests, can be cut down the center to form 2 mirror images. Although in practice we will never get a perfectly symmetrical distribution, we would like our data to be as close to symmetrical as possible for reasons we delve into in Chapter 3. Many types of distributions are symmetrical, but by far the most common and pertinent distribution at this point is the normal distribution, shown in Figure **2.2.28**. Notice that although the symmetry is not perfect (for instance, the bar just to the right of the center is taller than the one just to the left), the two sides are roughly the same shape. The normal distribution has a single peak, known as the center, and two tails that extend out equally, forming what is known as a bell shape or bell curve.

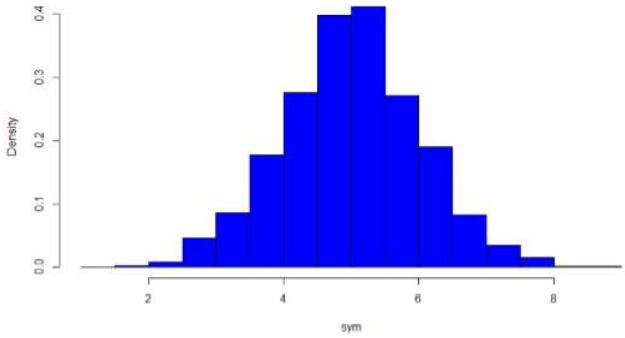


Figure 2.2.28: A symmetrical distribution

Symmetrical distributions can also have multiple peaks. Figure 2.2.29 shows a bimodal distribution, named for the two peaks that lie roughly symmetrically on either side of the center point. As we will see in the next chapter, this is not a particularly desirable characteristic of our data, and, worse, this is a relatively difficult characteristic to detect numerically. Thus, it is important to visualize your data before moving ahead with any formal analyses.

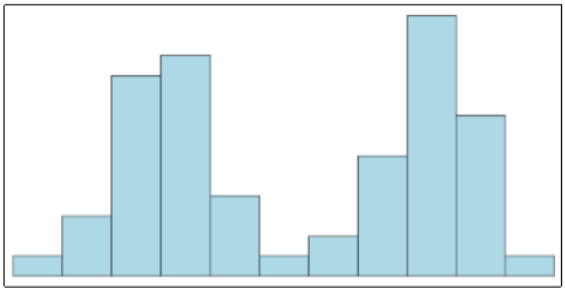


Figure 2.2.29: A bimodal distribution.

Distributions that are not symmetrical also come in many forms, more than can be described here. The most common asymmetry to be encountered is referred to as skew, in which one of the two tails of the distribution is disproportionately longer than the other. This property can affect the value of the averages we use in our analyses and make them an inaccurate representation of our data, which causes many problems.

Skew can either be positive or negative (also known as right or left, respectively), based on which tail is longer. It is very easy to get the two confused at first; many students want to describe the skew by where the bulk of the data (larger portion of the histogram, known as the body) is placed, but the correct determination is based on which tail is longer. You can think of the tail as an arrow: whichever direction the arrow is pointing is the direction of the skew. Figures 2.2.30 and 2.2.31 show positive (right) and negative (left) skew, respectively.

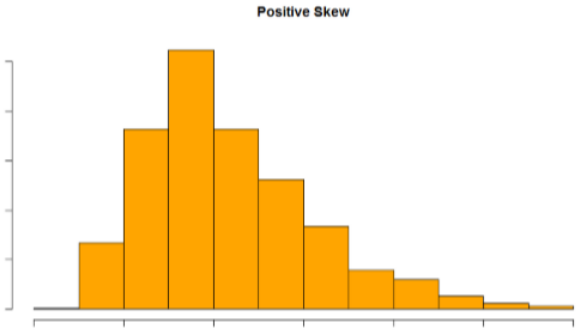


Figure 2.2.30: A positively skewed distribution.

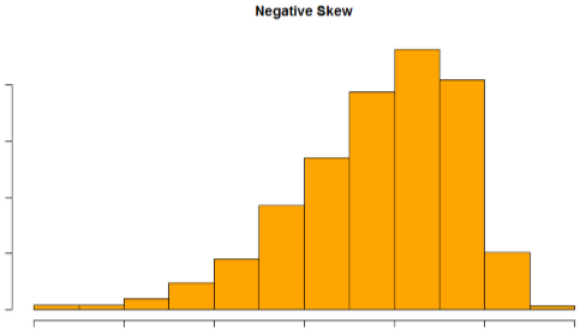


Figure **2.2.31**: A negatively skewed distribution .

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 2.2: Graphing Quantitative Variables is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **2.2: Graphing Quantitative Variables** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

2.E: Describing Data using Distributions and Graphs (Exercises)

1. Name some ways to graph quantitative variables and some ways to graph qualitative variables.

Answer:

Qualitative variables are displayed using pie charts and bar charts. Quantitative variables are displayed as box plots, histograms, etc.

2. Given the following data, construct a pie chart and a bar chart. Which do you think is the more appropriate or useful way to display the data?

Favorite Movie Genre	Frequency
Comedy	14
Horror	9
Romance	8
Action	12

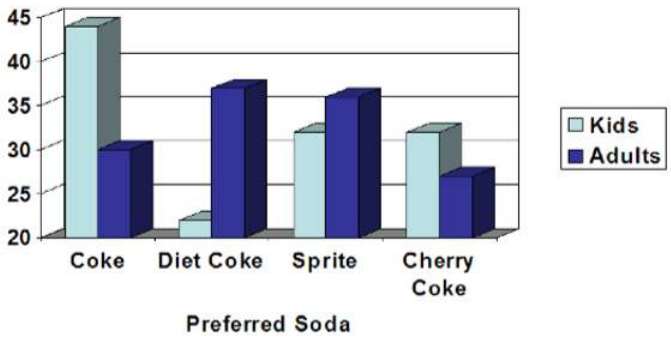
3. Pretend you are constructing a histogram for describing the distribution of salaries for individuals who are 40 years or older, but are not yet retired.
   a. What is on the Y-axis? Explain.
   b. What is on the X-axis? Explain.
   c. What would be the probable shape of the salary distribution? Explain why.

Answer:

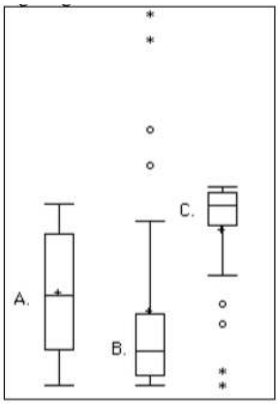
[You do not need to draw the histogram, only describe it below]

   a. The Y-axis would have the frequency or proportion because this is always the case in histograms
   b. The X-axis has income, because this is out quantitative variable of interest
   c. Because most income data are positively skewed, this histogram would likely be skewed positively too

4. A graph appears below showing the number of adults and children who prefer each type of soda. There were 130 adults and kids surveyed. Discuss some ways in which the graph below could be improved.



5. Which of the box plots on the graph has a large positive skew? Which has a large negative skew?



Answer:

Chart B has the positive skew because the outliers (dots and asterisks) are on the upper (higher) end; chart C has the negative skew because the outliers are on the lower end.

6. Create a histogram of the following data representing how many shows children said they watch each day:

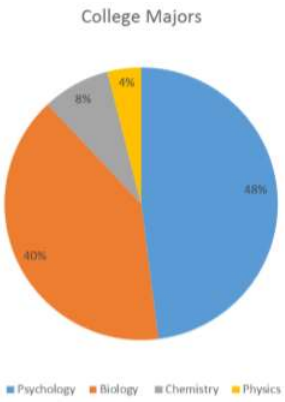
Number of TV Shows	Frequency
0	2
1	18
2	36
3	7
4	3

7. Explain the differences between bar charts and histograms. When would each be used?

Answer:

In bar charts, the bars do not touch; in histograms, the bars do touch. Bar charts are appropriate for qualitative variables, whereas histograms are better for quantitative variables.

8. Draw a histogram of a distribution that is
   a. Negatively skewed
   b. Symmetrical
   c. Positively skewed
9. Based on the pie chart below, which was made from a sample of 300 students, construct a frequency table of college majors.



Answer:

Use the following dataset for the computations below:

Major	Frequency
Psychology	144
Biology	120
Chemistry	24
Physics	12

10. Create a histogram of the following data. Label the tails and body and determine if it is skewed (and direction, if so) or symmetrical.

Hours worked per week	Proportion
0 -10	4
10 -20	8
20 - 30	11
30 - 40	51
40 - 50	12
50 - 60	9
60+	5

This page titled 2.E: Describing Data using Distributions and Graphs (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **2.E: Describing Data using Distributions and Graphs (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

3: Measures of Central Tendency and Spread

3.1: What is Central Tendency?

3.2: Measures of Central Tendency

3.3: Spread and Variability

3.E: Measures of Central Tendency and Spread (Exercises)

This page titled 3: Measures of Central Tendency and Spread is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

3.1: What is Central Tendency?

What is “central tendency,” and why do we want to know the central tendency of a group of scores? Let us first try to answer these questions intuitively. Then we will proceed to a more formal discussion.

Imagine this situation: You are in a class with just four other students, and the five of you took a 5-point pop quiz. Today your instructor is walking around the room, handing back the quizzes. She stops at your desk and hands you your paper. Written in bold black ink on the front is “3/5.” How do you react? Are you happy with your score of 3 or disappointed? How do you decide? You might calculate your percentage correct, realize it is 60%, and be appalled. But it is more likely that when deciding how to react to your performance, you will want additional information. What additional information would you like?

If you are like most students, you will immediately ask your neighbors, “Whad'ja get?” and then ask the instructor, “How did the class do?” In other words, the additional information you want is how your quiz score compares to other students' scores. You therefore understand the importance of comparing your score to the class distribution of scores. Should your score of 3 turn out to be among the higher scores, then you'll be pleased after all. On the other hand, if 3 is among the lower scores in the class, you won't be quite so happy.

This idea of comparing individual scores to a distribution of scores is fundamental to statistics. So let's explore it further, using the same example (the pop quiz you took with your four classmates). Three possible outcomes are shown in Table **3.1.1**. They are labeled “Dataset A,” “Dataset B,” and “Dataset C.” Which of the three datasets would make you happiest? In other words, in comparing your score with your fellow students' scores, in which dataset would your score of 3 be the most impressive?

Table **3.1.1**: Three possible datasets for the 5-point make-up quiz.

Student	Dataset A	Dataset B	Dataset C
You	3	3	3
John's	3	4	2
Maria's	3	4	2
Shareecia's	3	4	2
Luther's	3	5	1

In Dataset A, everyone's score is 3. This puts your score at the exact center of the distribution. You can draw satisfaction from the fact that you did as well as everyone else. But of course it cuts both ways: everyone else did just as well as you.

Now consider the possibility that the scores are described as in Dataset B. This is a depressing outcome even though your score is no different than the one in Dataset A. The problem is that the other four students had higher grades, putting yours below the center of the distribution.

Finally, let's look at Dataset C. This is more like it! All of your classmates score lower than you so your score is above the center of the distribution.

Now let's change the example in order to develop more insight into the center of a distribution. Figure **3.1.1** shows the results of an experiment on memory for chess positions. Subjects were shown a chess position and then asked to reconstruct it on an empty chess board. The number of pieces correctly placed was recorded. This was repeated for two more chess positions. The scores represent the total number of chess pieces correctly placed for the three chess positions. The maximum possible score was 89.

	8	05
	7	156
	6	233
	5	168
330	4	06
9420	3	
622	2	

Figure **3.1.1**: Back-to-back stem and leaf display. The left side shows the memory scores of the non-players. The right side shows the scores of the tournament players.

Two groups are compared. On the left are people who don't play chess. On the right are people who play a great deal (tournament players). It is clear that the location of the center of the distribution for the non-players is much lower than the center of the distribution for the tournament players.

We're sure you get the idea now about the center of a distribution. Here is another way for you to try to grasp the idea of a central tendency.

Balance Scale

One definition of central tendency is the point at which the distribution is in balance. Figure **3.1.2** shows the distribution of the five numbers 2, 3, 4, 9, 16 placed upon a balance scale. If each number weighs one pound, and is placed at its position along the number line, then it would be possible to balance them by placing a fulcrum at 6.8.

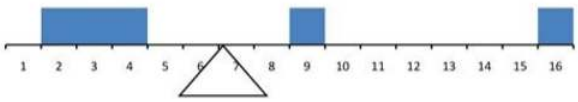


Figure **3.1.2**: A balance scale.

For another example, consider the distribution shown in Figure **3.1.3**. It is balanced by placing the fulcrum in the geometric middle.

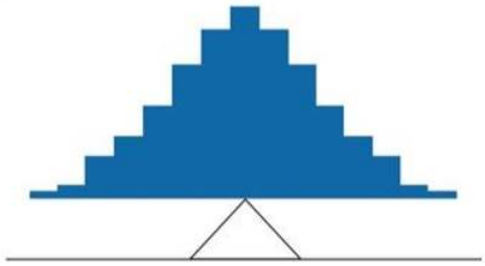


Figure **3.1.3**: A distribution balanced on the tip of a triangle.

Figure **3.1.4** illustrates that the same distribution can't be balanced by placing the fulcrum to the left of center.

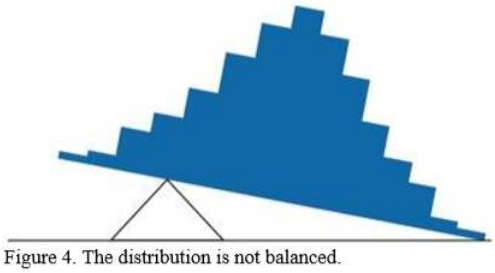


Figure 4. The distribution is not balanced.

Figure **3.1.4**: The distribution is not balanced.

Figure **3.1.5** shows an asymmetric distribution. To balance it, we cannot put the fulcrum halfway between the lowest and highest values (as we did in Figure **3.1.3**). Placing the fulcrum at the “half way” point would cause it to tip towards the left.

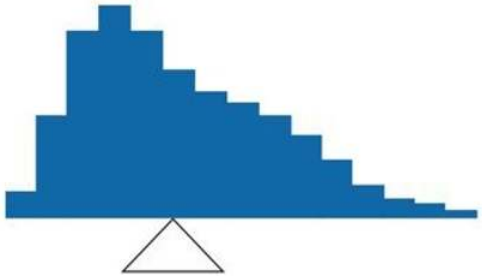


Figure **3.1.5** An asymmetric distribution balanced on the tip of a triangle.

There are other ways to try to define the central tendency, but for our purposes, this will suffice.

Contributors

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 3.1: What is Central Tendency? is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

3.2: Measures of Central Tendency

In the previous section we saw that there are several ways to define central tendency. This section defines the three most common measures of central tendency: the mean, the median, and the mode. The relationships among these measures of central tendency and the definitions given in the previous section will probably not be obvious to you.

This section gives only the basic definitions of the mean, median and mode. A further discussion of the relative merits and proper applications of these statistics is presented in a later section.

Arithmetic Mean

The arithmetic mean is the most common measure of central tendency. It is simply the sum of the numbers divided by the number of numbers. The symbol "μ" (pronounced "mew") is used for the mean of a population. The symbols M and "$\overline{X}$" (pronounced "X-bar") are used for the mean of a sample. The formula for μ is shown below:

$$\mu = \frac{\sum X}{N} \tag{3.2.1}$$

where ΣX is the sum of all the numbers in the population and N is the number of numbers in the population.

The formula for M is essentially identical:

$$M = \frac{\sum X}{N} \tag{3.2.2}$$

where ΣX is the sum of all the numbers in the sample and N is the number of numbers in the sample. The only distinction between these two equations is whether we are referring to the population (in which case we use the parameter μ) or a sample of that population (in which case we use the statistic M or $\overline{X}$).

As an example, the mean of the numbers 1, 2, 3, 6, 8 is 20/5 = 4 regardless of whether the numbers constitute the entire population or just a sample from the population.

Figure 3.2.1 shows the number of touchdown (TD) passes thrown by each of the 31 teams in the National Football League in the 2000 season. The mean number of touchdown passes thrown is 20.45 as shown below.

$$\mu = \frac{\sum X}{N} = \frac{634}{31} = 20.45$$

37, 33, 33, 32, 29, 28,
28, 23, 22, 22, 22, 21,
21, 21, 20, 20, 19, 19,
18, 18, 18, 18, 16, 15,
14, 14, 14, 12, 12, 9, 6

Figure 3.2.1: Number of touchdown passes.

Although the arithmetic mean is not the only "mean" (there is also a geometric mean, a harmonic mean, and many others that are all beyond the scope of this course), it is by far the most commonly used. Therefore, if the term "mean" is used without specifying whether it is the arithmetic mean, the geometric mean, or some other mean, it is assumed to refer to the arithmetic mean.

Median

The median is also a frequently used measure of central tendency. The median is the midpoint of a distribution: the same number of scores is above the median as below it. For the data in Figure 3.2.1, there are 31 scores. The 16th highest score (which equals 20) is the median because there are 15 scores below the 16th score and 15 scores above the 16th score. The median can also be thought of as the 50th percentile.

When there is an odd number of numbers, the median is simply the middle number. For example, the median of 2, 4, and 7 is 4. When there is an even number of numbers, the median is the mean of the two middle numbers. Thus, the median of the numbers 2, 4, 7, 12 is:

$$\frac{4+7}{2} = 5.5$$

When there are numbers with the same values, each appearance of that value gets counted. For example, in the set of numbers 1, 3, 4, 4, 5, 8, and 9, the median is 4 because there are three numbers (1, 3, and 4) below it and three numbers (5, 8, and 9) above it. If we only counted 4 once, the median would incorrectly be calculated at 4.5 (4+5 divided by 2). When in doubt, writing out all of the numbers in order and marking them off one at a time from the top and bottom will always lead you to the correct answer.

Mode

The mode is the most frequently occurring value in the dataset. For the data in Figure **3.2.1**, the mode is 18 since more teams (4) had 18 touchdown passes than any other number of touchdown passes. With continuous data, such as response time measured to many decimals, the frequency of each value is one since no two scores will be exactly the same (see discussion of continuous variables). Therefore the mode of continuous data is normally computed from a grouped frequency distribution. Table **3.2.1** shows a grouped frequency distribution for the target response time data. Since the interval with the highest frequency is 600-700, the mode is the middle of that interval (650). Though the mode is not frequently used for continuous data, it is nevertheless an important measure of central tendency as it is the only measure we can use on qualitative or categorical data.

Table **3.2.1**: Grouped frequency distribution

Range	Frequency
500 - 600	3
600 - 700	6
700 - 800	5
800 - 900	5
900 - 1000	0
1000 - 1100	1

More on the Mean and Median

It is important to note that the center of a distribution could be defined three ways:

1. the point on which a distribution would balance
2. the value whose average absolute deviation from all the other values is minimized
3. the value whose squared difference from all the other values is minimized.

The mean is the point on which a distribution would balance, the median is the value that minimizes the sum of absolute deviations, and the mean is the value that minimizes the sum of the squared deviations.

Table **3.2.2** shows the absolute and squared deviations of the numbers 2, 3, 4, 9, and 16 from their median of 4 and their mean of 6.8. You can see that the sum of absolute deviations from the median (20) is smaller than the sum of absolute deviations from the mean (22.8). On the other hand, the sum of squared deviations from the median (174) is larger than the sum of squared deviations from the mean (134.8).

Table **3.2.2**: Absolute & squared deviations from the median of 4 and the mean of 6.8.

Value	Absolute Deviation from Median	Absolute Deviation from Mean	Squared Deviation from Median	Squared Deviation from Mean
2	2	4.8	4	23.04
3	1	3.8	1	14.44
4	0	2.8	0	7.84
9	5	2.2	25	4.84
16	12	9.2	144	84.64
Total	20	22.8	174	134.8

Figure **3.2.2** shows that the distribution balances at the mean of 6.8 and not at the median of 4. The relative advantages and disadvantages of the mean and median are discussed in the section "Comparing Measures" later in this chapter.



Figure **3.2.2**: The distribution balances at the mean of 6.8 and not at the median of 4.0.

When a distribution is symmetric, then the mean and the median are the same. Consider the following distribution: 1, 3, 4, 5, 6, 7, 9. The mean and median are both 5. The mean, median, and mode are identical in the bell-shaped normal distribution.

Comparing Measures of Central Tendency

How do the various measures of central tendency compare with each other? For symmetric distributions, the mean and median are equal, as is the mode except in bimodal distributions. Differences among the measures occur with skewed distributions. Figure **3.2.3** shows the distribution of 642 scores on an introductory psychology test. Notice this distribution has a slight positive skew.

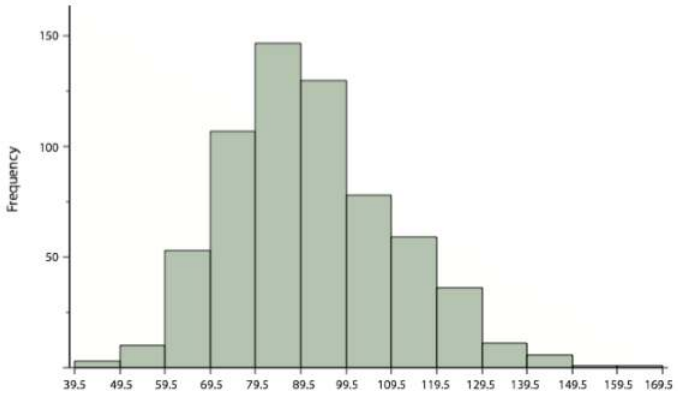


Figure **3.2.3**: A distribution with a positive skew.

Measures of central tendency are shown in Table **3.2.3**. Notice they do not differ greatly, with the exception that the mode is considerably lower than the other measures. When distributions have a positive skew, the mean is typically higher than the median, although it may not be in bimodal distributions. For these data, the mean of 91.58 is higher than the median of 90. This pattern holds true for any skew: the mode will remain at the highest point in the distribution, the median will be pulled slightly out into the skewed tail (the longer end of the distribution), and the mean will be pulled the farthest out. Thus, the mean is more sensitive to skew than the median or mode, and in cases of extreme skew, the mean may no longer be appropriate to use.

Table **3.2.3**: Measures of central tendency for the test scores.

Measure	Value
Mode	84.00
Median	90.00
Mean	91.58

The distribution of baseball salaries (in 1994) shown in Figure **3.2.4** has a much more pronounced skew than the distribution in Figure **3.2.3**.

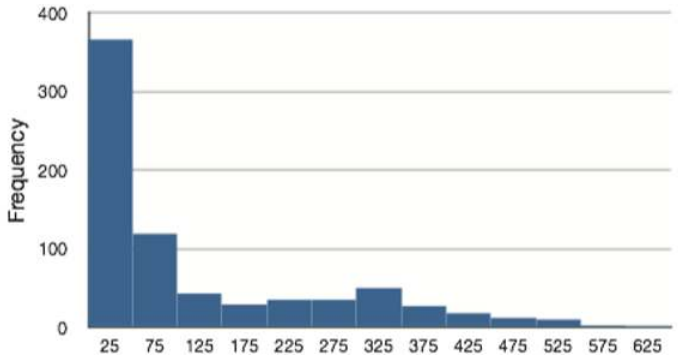


Figure **3.2.3**: A distribution with a very large positive skew. This histogram shows the salaries of major league baseball players (in thousands of dollars).

Table **3.2.4** shows the measures of central tendency for these data. The large skew results in very different values for these measures. No single measure of central tendency is sufficient for data such as these. If you were asked the very general question: "So, what do baseball players make?" and answered with the mean of $1,183,000, you would not have told the whole story since only about one third of baseball players make that much. If you answered with the mode of $250,000 or the median of $500,000, you would not be giving any indication that some players make many millions of dollars. Fortunately, there is no need to summarize a distribution with a single number. When the various measures differ, our opinion is that you should report the mean and median. Sometimes it is worth reporting the mode as well. In the media, the median is usually reported to summarize the center of skewed distributions. You will hear about median salaries and median prices of houses sold, etc. This is better than reporting only the mean, but it would be informative to hear more statistics.

Table **3.2.4**: Measures of central tendency for baseball salaries (in thousands of dollars).

Measure	Value
Mode	250
Median	500
Mean	1,183

Contributors

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 3.2: Measures of Central Tendency is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

3.3: Spread and Variability

Variability refers to how “spread out” a group of scores is. To see what we mean by spread out, consider graphs in Figure **3.3.1**. These graphs represent the scores on two quizzes. The mean score for each quiz is 7.0. Despite the equality of means, you can see that the distributions are quite different. Specifically, the scores on Quiz 1 are more densely packed and those on Quiz 2 are more spread out. The differences among students were much greater on Quiz 2 than on Quiz 1.

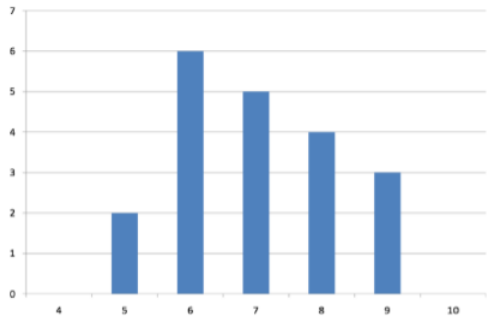


Figure **3.3.1**: Bar chart of quiz one.

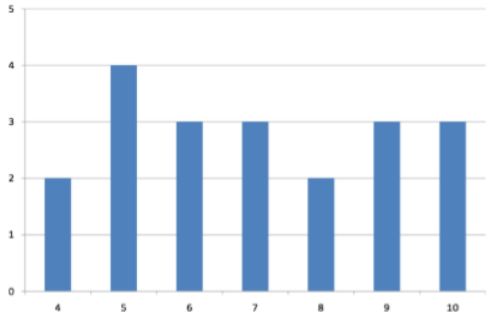


Figure **3.3.2**: Bar chart of quiz two.

The terms variability, spread, and dispersion are synonyms, and refer to how spread out a distribution is. Just as in the section on central tendency where we discussed measures of the center of a distribution of scores, in this chapter we will discuss measures of the variability of a distribution. There are three frequently used measures of variability: range, variance, and standard deviation. In the next few paragraphs, we will look at each of these measures of variability in more detail.

Range

The range is the simplest measure of variability to calculate, and one you have probably encountered many times in your life. The range is simply the highest score minus the lowest score. Let’s take a few examples. What is the range of the following group of numbers: 10, 2, 5, 6, 7, 3, 4? Well, the highest number is 10, and the lowest number is 2, so 10 - 2 = 8. The range is 8. Let’s take another example. Here’s a dataset with 10 numbers: 99, 45, 23, 67, 45, 91, 82, 78, 62, 51. What is the range? The highest number is 99 and the lowest number is 23, so 99 - 23 equals 76; the range is 76. Now consider the two quizzes shown in Figure **3.3.1** and Figure **3.3.2**. On Quiz 1, the lowest score is 5 and the highest score is 9. Therefore, the range is 4. The range on Quiz 2 was larger: the lowest score was 4 and the highest score was 10. Therefore the range is 6.

The problem with using range is that it is extremely sensitive to outliers, and one number far away from the rest of the data will greatly alter the value of the range. For example, in the set of numbers 1, 3, 4, 4, 5, 8, and 9, the range is 8 (9 – 1).

However, if we add a single person whose score is nowhere close to the rest of the scores, say, 20, the range more than doubles from 8 to 19.

Interquartile Range

The interquartile range (IQR) is the range of the middle 50% of the scores in a distribution and is sometimes used to communicate where the bulk of the data in the distribution are located. It is computed as follows:

$$\text{IQR} = 75\text{th percentile} - 25\text{th percentile} \tag{3.3.1}$$

For Quiz 1, the 75th percentile is 8 and the 25th percentile is 6. The interquartile range is therefore 2. For Quiz 2, which has greater spread, the 75th percentile is 9, the 25th percentile is 5, and the interquartile range is 4. Recall that in the discussion of box plots, the 75th percentile was called the upper hinge and the 25th percentile was called the lower hinge. Using this terminology, the interquartile range is referred to as the H-spread.

Sum of Squares

Variability can also be defined in terms of how close the scores in the distribution are to the middle of the distribution. Using the mean as the measure of the middle of the distribution, we can see how far, on average, each data point is from the center. The data from Quiz 1 are shown in Table **3.3.1**. The mean score is 7.0 ($\Sigma X/N = 140/20 = 7$). Therefore, the column "$X - M$" contains deviations (how far each score deviates from the mean), here calculated as the score minus 7. The column "$(X - M)^2$" has the "Squared Deviations" and is simply the previous column squared.

There are a few things to note about how Table **3.3.1** is formatted, as this is the format you will use to calculate variance (and, soon, standard deviation). The raw data scores (**X**) are always placed in the left-most column. This column is then summed at the bottom to facilitate calculating the mean (simply divided this number by the number of scores in the table). Once you have the mean, you can easily work your way down the middle column calculating the deviation scores. This column is also summed and has a very important property: it will always sum to 0 (or close to zero if you have rounding error due to many decimal places). This step is used as a check on your math to make sure you haven't made a mistake. If this column sums to 0, you can move on to filling in the third column of squared deviations. This column is summed as well and has its own name: the **Sum of Squares (abbreviated as** SS and given the formula $\sum (X - M)^2$). As we will see, the Sum of Squares appears again and again in different formulas – it is a very important value, and this table makes it simple to calculate without error.

Table **3.3.1**: Calculation of Variance for Quiz 1 scores.

X	$X - M$	$(X - M)^2$
9	2	4
9	2	4
9	2	4
8	1	1
8	1	1
8	1	1
8	1	1
7	0	0
7	0	0
7	0	0
7	0	0
7	0	0
6	-1	1
6	-1	1
6	-1	1
6	-1	1
6	-1	1
6	-1	1
5	-2	4
5	-2	4
$\Sigma = 140$	$\Sigma = 0$	$\Sigma = 30$

Variance

Now that we have the Sum of Squares calculated, we can use it to compute our formal measure of average distance from the mean, the variance. The variance is defined as the average squared difference of the scores from the mean. We square the deviation scores because, as we saw in the Sum of Squares table, the sum of raw deviations is always 0, and there's nothing we can do mathematically without changing that.

The population parameter for variance is σ^2 ("sigma-squared") and is calculated as:

$$\sigma^2 = \frac{\sum (X - \mu)^2}{N} \tag{3.3.2}$$

Notice that the numerator in that formula is identical to the formula for Sum of Squares presented above with M replaced by μ. Thus, we can use the Sum of Squares table to easily calculate the numerator then simply divide that value by N to get variance. If we assume that the values in Table **3.3.1** represent the full population, then we can take our value of Sum of Squares and divide it by N to get our population variance:

$$\sigma^2 = \frac{30}{20} = 1.5$$

So, on average, scores in this population are 1.5 squared units away from the mean. This measure of spread is much more robust (a term used by statisticians to mean resilient or resistant to) outliers than the range, so it is a much more useful value to compute. Additionally, as we will see in future chapters, variance plays a central role in inferential statistics.

The sample statistic used to estimate the variance is s^2 ("s-squared"):

$$s^2 = \frac{\sum (X - M)^2}{N - 1} \tag{3.3.3}$$

This formula is very similar to the formula for the population variance with one change: we now divide by N–1 instead of N. The value N–1 has a special name: the **degrees of freedom (abbreviated as df)**. You don't need to understand in depth what degrees of freedom are (essentially they account for the fact that we have to use a sample statistic to estimate the mean (M) before we estimate the variance) in order to calculate variance, but knowing that the denominator is called df provides a nice shorthand for the variance formula: SS/df.

Going back to the values in Table **3.3.1** and treating those scores as a sample, we can estimate the sample variance as:

$$s^2 = \frac{30}{20 - 1} = 1.58 \tag{3.3.4}$$

Notice that this value is slightly larger than the one we calculated when we assumed these scores were the full population. This is because our value in the denominator is slightly smaller, making the final value larger. In general, as your sample size N gets bigger, the effect of subtracting 1 becomes less and less. Comparing a sample size of 10 to a sample size of 1000; 10 – 1 = 9, or 90% of the original value, whereas 1000 – 1 = 999, or 99.9% of the original value. Thus, larger sample sizes will bring the estimate of the sample variance closer to that of the population variance. This is a key idea and principle in statistics that we will see over and over again: larger sample sizes better reflect the population.

Standard Deviation

The standard deviation is simply the square root of the variance. This is a useful and interpretable statistic because taking the square root of the variance (recalling that variance is the average squared difference) puts the standard deviation back into the original units of the measure we used. Thus, when reporting descriptive statistics in a study, scientists virtually always report mean and standard deviation. Standard deviation is therefore the most commonly used measure of spread for our purposes.

The population parameter for standard deviation is σ ("sigma"), which, intuitively, is the square root of the variance parameter σ^2 (on occasion, the symbols work out nicely that way). The formula is simply the formula for variance under a square root sign:

$$\sigma = \sqrt{\frac{\sum (X - \mu)^2}{N}} \tag{3.3.5}$$

Back to our earlier example from Table **3.3.1**:

$$\sigma = \sqrt{\frac{30}{20}} = \sqrt{1.5} = 1.22$$

The sample statistic follows the same conventions and is given as s:

$$s = \sqrt{\frac{\sum (X - M)^2}{N - 1}} = \sqrt{\frac{SS}{df}} \tag{3.3.6}$$

The sample standard deviation from Table **3.3.1** is:

$$s = \sqrt{\frac{30}{20 - 1}} = \sqrt{1.58} = 1.26$$

The standard deviation is an especially useful measure of variability when the distribution is normal or approximately normal because the proportion of the distribution within a given number of standard deviations from the mean can be calculated. For example, 68% of the distribution is within one standard deviation (above and below) of the mean and approximately 95% of the distribution is within two standard deviations of the mean. Therefore, if you had a normal distribution with a mean of 50 and a

standard deviation of 10, then 68% of the distribution would be between 50 - 10 = 40 and 50 +10 =60. Similarly, about 95% of the distribution would be between 50 - 2 x 10 = 30 and 50 + 2 x 10 = 70.

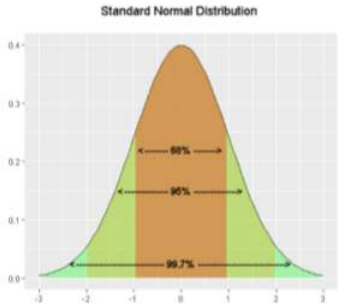


Figure **3.3.3**: Percentages of the normal distribution

Figure **3.3.4** shows two normal distributions. The red distribution has a mean of 40 and a standard deviation of 5; the blue distribution has a mean of 60 and a standard deviation of 10. For the red distribution, 68% of the distribution is between 45 and 55; for the blue distribution, 68% is between 50 and 70. Notice that as the standard deviation gets smaller, the distribution becomes much narrower, regardless of where the center of the distribution (mean) is. Figure **3.3.5** presents several more examples of this effect.

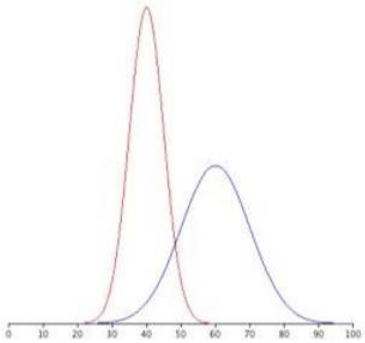


Figure **3.3.4**: Normal distributions with standard deviations of 5 and 10.

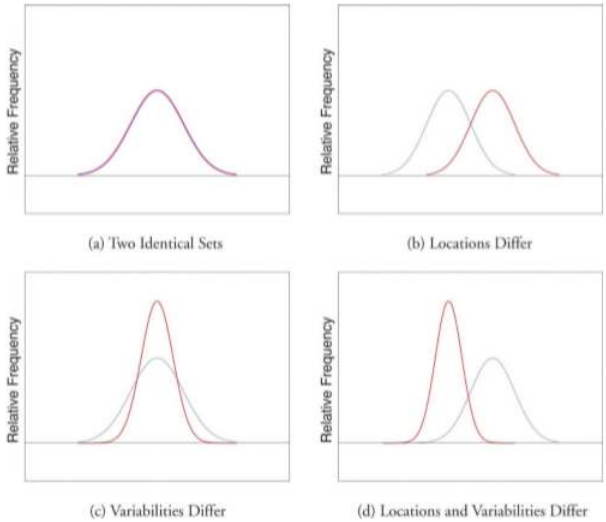


Figure **3.3.5**: Differences between two datasets.

Contributors

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 3.3: Spread and Variability is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

3.E: Measures of Central Tendency and Spread (Exercises)

1. If the mean time to respond to a stimulus is much higher than the median time to respond, what can you say about the shape of the distribution of response times?

Answer:

If the mean is higher, that means it is farther out into the right-hand tail of the distribution. Therefore, we know this distribution is positively skewed.

2. Compare the mean, median, and mode in terms of their sensitivity to extreme scores.
3. Your younger brother comes home one day after taking a science test. He says that some- one at school told him that “60% of the students in the class scored above the median test grade.” What is wrong with this statement? What if he had said “60% of the students scored above the mean?”

Answer:

The median is defined as the value with 50% of scores above it and 50% of scores below it; therefore, 60% of score cannot fall above the median. If 60% of scores fall above the mean, that would indicate that the mean has been pulled down below the value of the median, which means that the distribution is negatively skewed

4. Make up three data sets with 5 numbers each that have:
   a. the same mean but different standard deviations.
   b. the same mean but different medians.
   c. the same median but different means.
5. Compute the population mean and population standard deviation for the following scores (remember to use the Sum of Squares table): 5, 7, 8, 3, 4, 4, 2, 7, 1, 6

Answer:

$\mu = 4.80, \sigma^2 = 2.36$

6. For the following problem, use the following scores: 5, 8, 8, 8, 7, 8, 9, 12, 8, 9, 8, 10, 7, 9, 7, 6, 9, 10, 11, 8
   a. Create a histogram of these data. What is the shape of this histogram?
   b. How do you think the three measures of central tendency will compare to each other in this dataset?
   c. Compute the sample mean, the median, and the mode
   d. Draw and label lines on your histogram for each of the above values. Do your results match your predictions?
7. Compute the range, sample variance, and sample standard deviation for the following scores: 25, 36, 41, 28, 29, 32, 39, 37, 34, 34, 37, 35, 30, 36, 31, 31

Answer:

range = 16, $s^2 = 18.40$, $s = 4.29$

8. Using the same values from problem 7, calculate the range, sample variance, and sample standard deviation, but this time include 65 in the list of values. How did each of the three values change?
9. Two normal distributions have exactly the same mean, but one has a standard deviation of 20 and the other has a standard deviation of 10. How would the shapes of the two distributions compare?

Answer:

If both distributions are normal, then they are both symmetrical, and having the same mean causes them to overlap with one another. The distribution with the standard deviation of 10 will be narrower than the other distribution

10. Compute the sample mean and sample standard deviation for the following scores: -8, -4, -7, -6, -8, -5, -7, -9, -2, 0

This page titled 3.E: Measures of Central Tendency and Spread (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **3.E: Measures of Central Tendency and Spread (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

4: Basic Concepts of Probability: The math

NOTE: The following chapter is from an Introductory Statistics book by Shafer and Zhang. You may notice a slight change of tone/voice in the way it is written.

Suppose a polling organization questions 1,200 voters in order to estimate the proportion of all voters who favor a particular bond issue. We would expect the proportion of the 1,200 voters in the survey who are in favor to be close to the proportion of all voters who are in favor, but this need not be true. There is a degree of randomness associated with the survey result. If the survey result is highly likely to be close to the true proportion, then we have confidence in the survey result. If it is not particularly likely to be close to the population proportion, then we would perhaps not take the survey result too seriously. The likelihood that the survey proportion is close to the population proportion determines our confidence in the survey result. For that reason, we would like to be able to compute that likelihood. The task of computing it belongs to the realm of probability, which we study in this chapter.

4.1: Sample Spaces, Events, and Their Probabilities

4.2: Complements, Intersections, and Unions

4.3: Conditional Probability and Independent Events

4.E: Basic Concepts of Probability (Exercises)

4: Basic Concepts of Probability: The math is shared under a CC BY-SA license and was authored, remixed, and/or curated by LibreTexts.

4.1: Sample Spaces, Events, and Their Probabilities

Learning Objectives

- To learn the concept of the sample space associated with a random experiment.
- To learn the concept of an event associated with a random experiment.
- To learn the concept of the probability of an event.

Sample Spaces and Events

Rolling an ordinary six-sided die is a familiar example of a *random experiment*, an action for which all possible outcomes can be listed, but for which the actual outcome on any given trial of the experiment cannot be predicted with certainty. In such a situation we wish to assign to each outcome, such as rolling a two, a number, called the *probability* of the outcome, that indicates how likely it is that the outcome will occur. Similarly, we would like to assign a probability to any *event*, or collection of outcomes, such as rolling an even number, which indicates how likely it is that the event will occur if the experiment is performed. This section provides a framework for discussing probability problems, using the terms just mentioned.

Definition: random experiment

A *random experiment* is a mechanism that produces a definite outcome that cannot be predicted with certainty. The sample space associated with a random experiment is the set of all possible outcomes. An event is a subset of the sample space.

Definition: Element and Occurrence

An event E is said to occur on a particular trial of the experiment if the outcome observed is an element of the set E.

Example 4.1.1: Sample Space for a single coin

Construct a sample space for the experiment that consists of tossing a single coin.

Solution

The outcomes could be labeled h for heads and t for tails. Then the sample space is the set: $S = \{h, t\}$

Example 4.1.2: Sample Space for a single die

Construct a sample space for the experiment that consists of rolling a single die. Find the events that correspond to the phrases "an even number is rolled" and "a number greater than two is rolled."

Solution:

The outcomes could be labeled according to the number of dots on the top face of the die. Then the sample space is the set $S = \{1, 2, 3, 4, 5, 6\}$

The outcomes that are even are $2, 4,$ and 6, so the event that corresponds to the phrase "an even number is rolled" is the set $\{2, 4, 6\}$, which it is natural to denote by the letter E. We write $E = \{2, 4, 6\}$.

Similarly the event that corresponds to the phrase "a number greater than two is rolled" is the set $T = \{3, 4, 5, 6\}$, which we have denoted T.

A graphical representation of a sample space and events is a *Venn diagram*, as shown in Figure 4.1.1. In general the sample space S is represented by a rectangle, outcomes by points within the rectangle, and events by ovals that enclose the outcomes that compose them.



(a): One Coin Toss (b): One Die Roll

Figure 4.1.1: Venn Diagrams for Two Sample Spaces

Example 4.1.3: Sample Spaces for two coines

A random experiment consists of tossing two coins.

a. Construct a sample space for the situation that the coins are indistinguishable, such as two brand new pennies.
b. Construct a sample space for the situation that the coins are distinguishable, such as one a penny and the other a nickel.

Solution:

a. After the coins are tossed one sees either two heads, which could be labeled $2h$, two tails, which could be labeled $2t$, or coins that differ, which could be labeled d Thus a sample space is $S = \{2h, 2t, d\}$.
b. Since we can tell the coins apart, there are now two ways for the coins to differ: the penny heads and the nickel tails, or the penny tails and the nickel heads. We can label each outcome as a pair of letters, the first of which indicates how the penny landed and the second of which indicates how the nickel landed. A sample space is then $S' = \{hh, ht, th, tt\}$.

A device that can be helpful in identifying all possible outcomes of a random experiment, particularly one that can be viewed as proceeding in stages, is what is called a *tree diagram*. It is described in the following example.

Example 4.1.4: Tree diagram

Construct a sample space that describes all three-child families according to the genders of the children with respect to birth order.

Solution:

Two of the outcomes are "two boys then a girl," which we might denote bbg, and "a girl then two boys," which we would denote gbb.

Clearly there are many outcomes, and when we try to list all of them it could be difficult to be sure that we have found them all unless we proceed systematically. The tree diagram shown in Figure 4.1.2, gives a systematic approach.

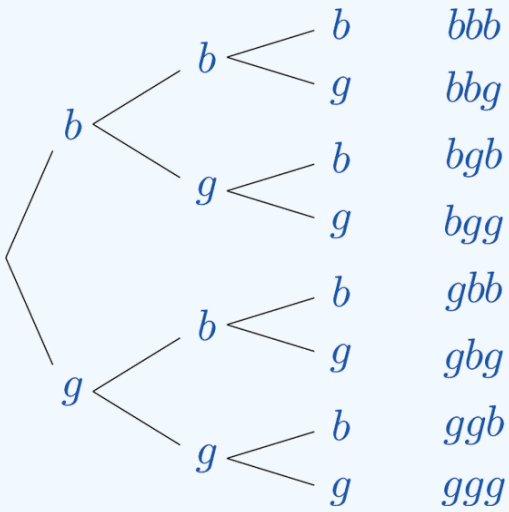


Figure 4.1.2: *Tree Diagram For Three-Child Families*

The diagram was constructed as follows. There are two possibilities for the first child, boy or girl, so we draw two line segments coming out of a starting point, one ending in a b for "boy" and the other ending in a g for "girl." For each of these two possibilities for the first child there are two possibilities for the second child, "boy" or "girl," so from each of the b and g we draw two line segments, one segment ending in a b and one in a g. For each of the four ending points now in the diagram there are two possibilities for the third child, so we repeat the process once more.

The line segments are called **branches** of the tree. The right ending point of each branch is called a **node**. The nodes on the extreme right are the **final nodes**; to each one there corresponds an outcome, as shown in the figure.

From the tree it is easy to read off the eight outcomes of the experiment, so the sample space is, reading from the top to the bottom of the final nodes in the tree,

$$S = \{bbb,\ bbg,\ bgb,\ bgg,\ gbb,\ gbg,\ ggb,\ ggg\} \tag{4.1.1}$$

Probability

Definition: probability

The probability of an outcome e in a sample space S is a number P between 1 and 0 that measures the likelihood that e will occur on a single trial of the corresponding random experiment. The value $P = 0$ corresponds to the outcome e being impossible and the value $P = 1$ corresponds to the outcome e being certain.

Definition: probability of an event

The *probability of an event* A is the sum of the probabilities of the individual outcomes of which it is composed. It is denoted $P(A)$.

The following formula expresses the content of the definition of the probability of an event:

If an event E is $E = \{e_1, e_2, ..., e_k\}$, then

$$P(E) = P(e_1) + P(e_2) + ... + P(e_k) \tag{4.1.2}$$

The following figure expresses the content of the definition of the probability of an event:

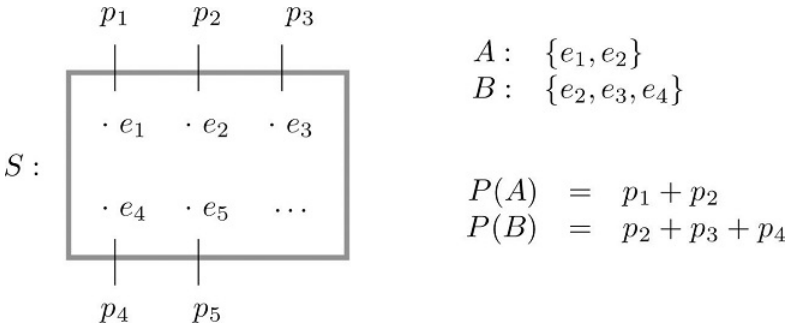


Figure **4.1.3***: Sample Spaces and Probability*

Since the whole sample space S is an event that is certain to occur, the sum of the probabilities of all the outcomes must be the number 1.

In ordinary language probabilities are frequently expressed as percentages. For example, we would say that there is a 70% chance of rain tomorrow, meaning that the probability of rain is 0.70. We will use this practice here, but in all the computational formulas that follow we will use the form 0.70 and not 70%.

Example 4.1.5

A coin is called "balanced" or "fair" if each side is equally likely to land up. Assign a probability to each outcome in the sample space for the experiment that consists of tossing a single fair coin.

Solution:

With the outcomes labeled h for heads and t for tails, the sample space is the set

$$S = \{h, t\} \tag{4.1.3}$$

Since the outcomes have the same probabilities, which must add up to 1, each outcome is assigned probability $1/2$.

Example 4.1.6

A die is called "balanced" or "fair" if each side is equally likely to land on top. Assign a probability to each outcome in the sample space for the experiment that consists of tossing a single fair die. Find the probabilities of the events E: "an even number is rolled" and T: "a number greater than two is rolled."

Solution:

With outcomes labeled according to the number of dots on the top face of the die, the sample space is the set

$$S = \{1, 2, 3, 4, 5, 6\} \tag{4.1.4}$$

Since there are six equally likely outcomes, which must add up to 1, each is assigned probability $1/6$.

Since $E = \{2, 4, 6\}$,

$$P(E) = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = \frac{3}{6} = \frac{1}{2} \tag{4.1.5}$$

Since $T = \{3, 4, 5, 6\}$,

$$P(T) = \frac{4}{6} = \frac{2}{3} \tag{4.1.6}$$

Example 4.1.7

Two fair coins are tossed. Find the probability that the coins match, i.e., either both land heads or both land tails.

Solution:

In Example 4.1.3 we constructed the sample space $S = \{2h, 2t, d\}$ for the situation in which the coins are identical and the sample space $S' = \{hh, ht, th, tt\}$ for the situation in which the two coins can be told apart.

The theory of probability does not tell us how to assign probabilities to the outcomes, only what to do with them once they are assigned. Specifically, using sample space S, matching coins is the event $M = \{2h, 2t\}$ which has probability $P(2h) + P(2t)$. Using sample space S', matching coins is the event $M' = \{hh, tt\}$, which has probability $P(hh) + P(tt)$. In the physical world it should make no difference whether the coins are identical or not, and so we would like to assign probabilities to the outcomes so that the numbers $P(M)$ and $P(M')$ are the same and best match what we observe when actual physical experiments are performed with coins that seem to be fair. Actual experience suggests that the outcomes in S' are equally likely, so we assign to each probability $\frac{1}{4}$, and then...

$$P(M') = P(hh) + P(tt) = \frac{1}{4} + \frac{1}{4} = \frac{1}{2} \tag{4.1.7}$$

Similarly, from experience appropriate choices for the outcomes in S are:

$$P(2h) = \frac{1}{4} \tag{4.1.8}$$

$$P(2t) = \frac{1}{4} \tag{4.1.9}$$

$$P(d) = \frac{1}{2} \tag{4.1.10}$$

The previous three examples illustrate how probabilities can be computed simply by counting when the sample space consists of a finite number of equally likely outcomes. In some situations the individual outcomes of any sample space that represents the experiment are unavoidably unequally likely, in which case probabilities cannot be computed merely by counting, but the computational formula given in the definition of the probability of an event must be used.

Example 4.1.8

The breakdown of the student body in a local high school according to race and ethnicity is 51% white, 27% black, 11% Hispanic, 6% Asian, and 5% for all others. A student is randomly selected from this high school. (To select "randomly" means that every student has the same chance of being selected.) Find the probabilities of the following events:

a. B: the student is black,
b. M: the student is minority (that is, not white),
c. N: the student is not black.

Solution:

The experiment is the action of randomly selecting a student from the student population of the high school. An obvious sample space is $S = \{w, b, h, a, o\}$. Since 51% of the students are white and all students have the same chance of being selected, $P(w) = 0.51$, and similarly for the other outcomes. This information is summarized in the following table:

Outcome	w	b	h	a	o
Probability	0.51	0.27	0.11	0.06	0.05

(4.1.11)

a. Since $B = \{b\}$, $P(B) = P(b) = 0.27$
b. Since $M = \{b, h, a, o\}$, $P(M) = P(b) + P(h) + P(a) + P(o) = 0.27 + 0.11 + 0.06 + 0.05 = 0.49$
c. Since $N = \{w, h, a, o\}$, $P(N) = P(w) + P(h) + P(a) + P(o) = 0.51 + 0.11 + 0.06 + 0.05 = 0.73$

Example 4.1.9

The student body in the high school considered in the last example may be broken down into ten categories as follows: 25% white male, 26% white female, 12% black male, 15% black female, 6% Hispanic male, 5% Hispanic female, 3% Asian male, 3% Asian female, 1% male of other minorities combined, and 4% female of other minorities combined. A student is randomly selected from this high school. Find the probabilities of the following events:

a. B: the student is black
b. MF: the student is a non-white female
c. FN: the student is female and is not black

Solution:

Now the sample space is $S = \{wm, bm, hm, am, om, wf, bf, hf, af, of\}$. The information given in the example can be summarized in the following table, called a two-way contingency table:

Gender	Race / Ethnicity				
	White	Black	Hispanic	Asian	Others
Male	0.25	0.12	0.06	0.03	0.01
Female	0.26	0.15	0.05	0.03	0.04

a. Since $B = \{bm, bf\}$, $P(B) = P(bm) + P(bf) = 0.12 + 0.15 = 0.27$
b. Since $MF = \{bf, hf, af, of\}$, $P(M) = P(bf) + P(hf) + P(af) + P(of) = 0.15 + 0.05 + 0.03 + 0.04 = 0.27$
c. Since $FN = \{wf, hf, af, of\}$, $P(FN) = P(wf) + P(hf) + P(af) + P(of) = 0.26 + 0.05 + 0.03 + 0.04 = 0.38$

Key Takeaway

- The sample space of a random experiment is the collection of all possible outcomes.
- An event associated with a random experiment is a subset of the sample space.
- The probability of any outcome is a number between 0 and 1. The probabilities of all the outcomes add up to 1.
- The probability of any event A is the sum of the probabilities of the outcomes in A.

4.1: Sample Spaces, Events, and Their Probabilities is shared under a CC BY-SA license and was authored, remixed, and/or curated by LibreTexts.

4.2: Complements, Intersections, and Unions

skills to develop

- To learn how some events are naturally expressible in terms of other events.
- To learn how to use special formulas for the probability of an event that is expressed in terms of one or more other events.

Some events can be naturally expressed in terms of other, sometimes simpler, events.

Complements

Definition: complement

The complement of an event A in a sample space S, denoted A^c, is the collection of all outcomes in S that are not elements of the set A. It corresponds to negating any description in words of the event A.

Example 4.2.1

Two events connected with the experiment of rolling a single die are E: ***"the number rolled is even"*** and T: ***"the number rolled is greater than two."*** Find the complement of each.

Solution:

In the sample space $S = \{1, 2, 3, 4, 5, 6\}$ the corresponding sets of outcomes are $E = \{2, 4, 6\}$ and $T = \{3, 4, 5, 6\}$. The complements are $E^c = \{1, 3, 5\}$ and $T^c = \{1, 2\}$.

In words the complements are described by "the number rolled is not even" and "the number rolled is not greater than two." Of course easier descriptions would be "the number rolled is odd" and "the number rolled is less than three."

If there is a 60% chance of rain tomorrow, what is the probability of fair weather? The obvious answer, 40%, is an instance of the following general rule.

Definition: Probability Rule for Complements

The Probability Rule for Complements states that

$$P(A^c) = 1 - P(A) \tag{4.2.1}$$

This formula is particularly useful when finding the probability of an event directly is difficult.

Example 4.2.2

Find the probability that at least one heads will appear in five tosses of a fair coin.

Solution:

Identify outcomes by lists of five hs and ts, such as $tthtt$ and $hhttt$. Although it is tedious to list them all, it is not difficult to count them. Think of using a tree diagram to do so. There are two choices for the first toss. For each of these there are two choices for the second toss, hence $2 \times 2 = 4$ outcomes for two tosses. For each of these four outcomes, there are two possibilities for the third toss, hence $4 \times 2 = 8$ outcomes for three tosses. Similarly, there are $8 \times 2 = 16$ outcomes for four tosses and finally $16 \times 2 = 32$ outcomes for five tosses.

Let O denote the event "at least one heads." There are many ways to obtain at least one heads, but only one way to fail to do so: all tails. Thus although it is difficult to list all the outcomes that form O, it is easy to write $O^c = \{ttttt\}$. Since there are 32 equally likely outcomes, each has probability $\frac{1}{32}$, so $P(O^c) = 1 / 32$, hence $P(O) = 1 - \frac{1}{32} \approx 0.97$ or about a 97% chance.

Intersection of Events

Definition: intersections

The intersection of events A and B, denoted $A \cap B$, is the collection of all outcomes that are elements of both of the sets A and B. It corresponds to combining descriptions of the two events using the word "and."

To say that the event $A \cap B$ occurred means that on a particular trial of the experiment both A and B occurred. A visual representation of the intersection of events A and B in a sample space S is given in Figure 4.2.1. The intersection corresponds to the shaded lens-shaped region that lies within both ovals.

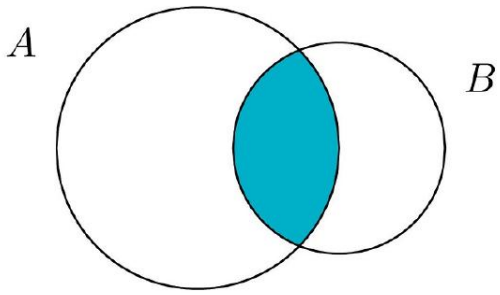


Figure 4.2.1: *The Intersection of Events A and B*

Example 4.2.3

In the experiment of rolling a single die, find the intersection $E \cap T$ of the events E: *"the number rolled is even"* and T: *"the number rolled is greater than two."*

Solution:

The sample space is $S = \{1, 2, 3, 4, 5, 6\}$. Since the outcomes that are common to $E = \{2, 4, 6\}$ and $T = \{3, 4, 5, 6\}$ are 4 and 6, $E \cap T = \{4, 6\}$.

In words the intersection is described by "the number rolled is even and is greater than two." The only numbers between one and six that are both even and greater than two are four and six, corresponding to $E \cap T$ given above.

Example 4.2.4

A single die is rolled.

1. Suppose the die is fair. Find the probability that the number rolled is both even and greater than two.
2. Suppose the die has been "loaded" so that $P(1) = \frac{1}{12}$, $P(6) = \frac{3}{12}$, and the remaining four outcomes are equally likely with one another. Now find the probability that the number rolled is both even and greater than two.

Solution:

In both cases the sample space is $S = \{1, 2, 3, 4, 5, 6\}$ and the event in question is the intersection $E \cap T = \{4, 6\}$ of the previous example.

1. Since the die is fair, all outcomes are equally likely, so by counting we have $P(E \cap T) = \frac{2}{6}$.
2. The information on the probabilities of the six outcomes that we have so far is

Outcome	1	2	3	4	5	6
Probablity	$\frac{1}{12}$	p	p	p	p	$\frac{3}{12}$

(4.2.2)

Since $P(1) + P(6) = \frac{4}{6} = \frac{1}{3}$

$$P(2) + P(3) + P(4) + P(5) = 1 - \frac{1}{3} = \frac{2}{3} \tag{4.2.3}$$

Thus $4p = \frac{2}{3}$, so $p = \frac{1}{6}$. In particular $P(4) = \frac{1}{6}$ therefore:

$$P(E \cap T) = P(4) + P(6) = \frac{1}{6} + \frac{3}{12} = \frac{5}{12} \tag{4.2.4}$$

Definition: mutually exclusive

Events A and B are mutually exclusive (cannot both occur at once) if they have no elements in common.

For A and B to have no outcomes in common means precisely that it is impossible for both A and B to occur on a single trial of the random experiment. This gives the following rule:

Definition: Probability Rule for Mutually Exclusive Events

Events A and B are mutually exclusive if and only if

$$P(A \cap B) = 0 \tag{4.2.5}$$

Any event A and its complement A^c are mutually exclusive, but A and B can be mutually exclusive without being complements.

Example 4.2.5

In the experiment of rolling a single die, find three choices for an event A so that the events A and E: "the number rolled is even" are mutually exclusive.

Solution:

Since $E = \{2, 4, 6\}$ and we want A to have no elements in common with E, any event that does not contain any even number will do. Three choices are $\{1, 3, 5\}$ (the complement E^c, the odds), $\{1, 3\}$, and $\{5\}$.

Union of Events

Definition: Union of Events

The union of events A and B, denoted $A \cup B$, is the collection of all outcomes that are elements of one or the other of the sets A and B, or of both of them. It corresponds to combining descriptions of the two events using the word "or."

To say that the event $A \cup B$ occurred means that on a particular trial of the experiment either A or B occurred (or both did). A visual representation of the union of events A and B in a sample space S is given in Figure 4.2.2. The union corresponds to the shaded region.

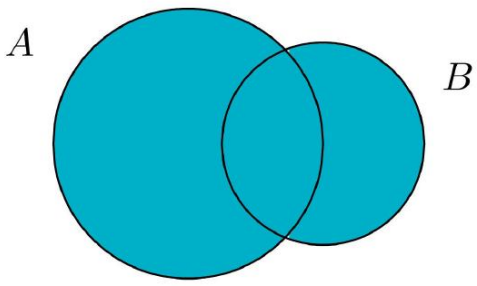


***Figure* 4.2.2**: *The Union of Events A and B*

Example 4.2.6

In the experiment of rolling a single die, find the union of the events E: **"the number rolled is even"** and T: **"the number rolled is greater than two."**

Solution:

Since the outcomes that are in either $E = \{2,4,6\}$ or $T = \{3,4,5,6\}$ (or both) are $2, 3, 4, 5,$ and 6, that means $E \cup T = \{2,3,4,5,6\}$.

Note that an outcome such as 4 that is in both sets is still listed only once (although strictly speaking it is not incorrect to list it twice).

In words the union is described by "the number rolled is even or is greater than two." Every number between one and six except the number one is either even or is greater than two, corresponding to $E \cup T$ given above.

Example 4.2.7

A two-child family is selected at random. Let B denote the event that at least one child is a boy, let D denote the event that the genders of the two children differ, and let M denote the event that the genders of the two children match. Find $B \cup D$ and $B \cup M$.

Solution:

A sample space for this experiment is $S = \{bb, bg, gb, gg\}$, where the first letter denotes the gender of the firstborn child and the second letter denotes the gender of the second child. The events $B, D,$ and M are $B = \{bb, bg, gb\}$, $D = \{bg, gb\}$, $M = \{bb, gg\}$.

Each outcome in D is already in B, so the outcomes that are in at least one or the other of the sets B and D is just the set B itself: $B \cup D = \{bb, bg, gb\} = B$.

Every outcome in the whole sample space S is in at least one or the other of the sets B and M, so $B \cup M = \{bb, bg, gb, gg\} = S$.

Definition: Additive Rule of Probability

A useful property to know is the Additive Rule of Probability, which is

$$P(A \cup B) = P(A) + P(B) - P(A \cap B) \tag{4.2.6}$$

The next example, in which we compute the probability of a union both by counting and by using the formula, shows why the last term in the formula is needed.

Example 4.2.8

Two fair dice are thrown. Find the probabilities of the following events:

1. both dice show a four
2. at least one die shows a four

Solution:

As was the case with tossing two identical coins, actual experience dictates that for the sample space to have equally likely outcomes we should list outcomes as if we could distinguish the two dice. We could imagine that one of them is red and the other is green. Then any outcome can be labeled as a pair of numbers as in the following display, where the first number in the pair is the number of dots on the top face of the green die and the second number in the pair is the number of dots on the top face of the red die.

Illegal pream-token (1)

1. There are 36 equally likely outcomes, of which exactly one corresponds to two fours, so the probability of a pair of fours is $1/36$.
2. From the table we can see that there are 11 pairs that correspond to the event in question: the six pairs in the fourth row (the green die shows a four) plus the additional five pairs other than the pair 44, already counted, in the fourth column (the red die is four), so the answer is $11/36$. To see how the formula gives the same number, let A_G denote the event that the green die is a four and let A_R denote the event that the red die is a four. Then clearly by counting we get: $P(A_G) = 6/36$ and $P(A_R) = 6/36$. Since $A_G \cap A_R = \{44\}$, $P(A_G \cap A_R) = 1/36$. This is the computation from part 1, of course. Thus by the Additive Rule of Probability we get:

$$P(A_G \cap A_R) = P(A_G) + P(A_R) - P(A_G - A_R) = 6/36 + 6/36 - 1/36 = \frac{11}{36} \tag{4.2.7}$$

Example 4.2.9

A tutoring service specializes in preparing adults for high school equivalence tests. Among all the students seeking help from the service, 63% need help in mathematics, 34% need help in English, and 27% need help in both mathematics and English. What is the percentage of students who need help in either mathematics or English?

Solution:

Imagine selecting a student at random, that is, in such a way that every student has the same chance of being selected. Let M denote the event "the student needs help in mathematics" and let E denote the event "the student needs help in English." The information given is that $P(M) = 0.63$, $P(E) = 0.34$ and $P(M \cap E) = 0.27$. Thus the Additive Rule of Probability gives:

$$P(M \cup E) = P(M) + P(E) - P(M \cap E) = 0.63 + 0.34 - 0.27 = 0.70 \tag{4.2.8}$$

Note how the naïve reasoning that if 63% need help in mathematics and 34% need help in English then 63 plus 34 or 97% need help in one or the other gives a number that is too large. The percentage that need help in both subjects must be subtracted off, else the people needing help in both are counted twice, once for needing help in mathematics and once again for needing help in English. The simple sum of the probabilities would work if the events in question were mutually exclusive, for then $P(A \cap B)$ is zero, and makes no difference.

Example 4.2.10

Volunteers for a disaster relief effort were classified according to both specialty (C: construction, E: education, M: medicine) and language ability (S: speaks a single language fluently, T: speaks two or more languages fluently). The results are shown in the following two-way classification table:

Specialty	Language Ability	
	S	T
C	12	1
E	4	3
M	6	2

The first row of numbers means that 12 volunteers whose specialty is construction speak a single language fluently, and 1 volunteer whose specialty is construction speaks at least two languages fluently. Similarly for the other two rows.

A volunteer is selected at random, meaning that each one has an equal chance of being chosen. Find the probability that:

1. his specialty is medicine and he speaks two or more languages;
2. either his specialty is medicine or he speaks two or more languages;
3. his specialty is something other than medicine.

Solution:

When information is presented in a two-way classification table it is typically convenient to adjoin to the table the row and column totals, to produce a new table like this:

Specialty	Language Ability		Total
	S	T	
C	12	1	13
E	4	3	7
M	6	2	8
Total	22	6	28

1. The probability sought is $P(M \cap T)$. The table shows that there are 2 such people, out of 28 in all, hence $P(M \cap T) = 2/28 \approx 0.07$ or about a 7% chance.
2. The probability sought is $P(M \cup T)$. The third row total and the grand total in the sample give $P(M) = 8/28$. The second column total and the grand total give $P(T) = 6/28$. Thus using the result from part (1),

$$P(M \cup T) = P(M) + P(T) - P(M \cap T) = 8/28 + 6/28 - 2/28 = 12/28 \approx 0.43 \tag{4.2.9}$$

   or about a 43% chance.

3. This probability can be computed in two ways. Since the event of interest can be viewed as the event $C \cup E$ and the events C and E are mutually exclusive, the answer is, using the first two row totals,

$$P(C \cup E) = P(C) + P(E) - P(C \cap E) = 13/28 + 7/28 - 0/28 = 20/28 \approx 0.71 \tag{4.2.10}$$

On the other hand, the event of interest can be thought of as the complement M^c of M, hence using the value of $P(M)$computed in part (2),

$$P(M^c) = 1 - P(M) = 1 - 8/28 = 20/28 \approx 0.71 \tag{4.2.11}$$

as before.

Key Takeaway

- The probability of an event that is a complement or union of events of known probability can be computed using formulas.

4.2: Complements, Intersections, and Unions is shared under a CC BY-SA license and was authored, remixed, and/or curated by LibreTexts.

4.3: Conditional Probability and Independent Events

Learning Objectives

- To learn the concept of a conditional probability and how to compute it.
- To learn the concept of independence of events, and how to apply it.

Suppose a fair die has been rolled and you are asked to give the probability that it was a five. There are six equally likely outcomes, so your answer is $1/6$. But suppose that before you give your answer you are given the extra information that the number rolled was odd. Since there are only three odd numbers that are possible, one of which is five, you would certainly revise your estimate of the likelihood that a five was rolled from $1/6$ to $1/3$. In general, the *revised* probability that an event A has occurred, taking into account the additional information that another event B has definitely occurred on this trial of the experiment, is called the *conditional probability of A given B* and is denoted by $P(A \mid B)$. The reasoning employed in this example can be generalized to yield the computational formula in the following definition.

Definition: conditional probability

The *conditional probability* of A given B, denoted $P(A \mid B)$, is the probability that event A has occurred in a trial of a random experiment for which it is known that event B has definitely occurred. It may be computed by means of the following formula:

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)} \tag{4.3.1}$$

Example 4.3.1: Rolling a Die

A fair (unbiased) die is rolled.

a. Find the probability that the number rolled is a five, given that it is odd.
b. Find the probability that the number rolled is odd, given that it is a five.

Solution:

The sample space for this experiment is the set $S = 1, 2, 3, 4, 5, 6$ consisting of six equally likely outcomes. Let F denote the event "a five is rolled" and let O denote the event "an odd number is rolled," so that

$$F = 5 \text{ and } O = 1, 3, 5$$

a. This is the introductory example, so we already know that the answer is $1/3$. To use Equation 4.3.1 to confirm this we must replace A in the formula (the event whose likelihood we seek to estimate) by F and replace B (the event we know for certain has occurred) by O:

$$P(F \mid O) = \frac{P(F \cap O)}{P(O)}$$

Since

$$F \cap O = 5 \cap 1, 3, 5 = 5, \; P(F \cap O) = 1/6 \tag{4.3.2}$$

Since

$$O = 1, 3, 5, \; P(O) = 3/6. \tag{4.3.3}$$

Thus

$$P(F \mid O) = \frac{P(F \cap O)}{P(O)} = \frac{1/6}{3/6} = \frac{1}{3}$$

b. This is the same problem, but with the roles of F and O reversed. Since we are given that the number that was rolled is five, which is odd, the probability in question must be 1. To apply Equation 4.3.1 to this case we must now replace A (the event whose likelihood we seek to estimate) by O and B (the event we know for certain has occurred) by F:

$$P(O \mid F) = \frac{P(O \cap F)}{P(F)}$$

Obviously $P(F) = 1/6$. In part (a) we found that $P(F \mid O) = 1/6$. Thus

$$P(O \mid F) = \frac{P(O \cap F)}{P(F)} = \frac{1/6}{1/6} = 1$$

Just as we did not need the computational formula in this example, we do not need it when the information is presented in a *two-way classification table*, as in the next example.

Example 4.3.2: Marriage and Gender

In a sample of 902 individuals under 40 who were or had previously been married, each person was classified according to gender and age at first marriage. The results are summarized in the following two-way classification table, where the meaning of the labels is:

- M: male
- F: female
- E: a teenager when first married
- W: in one's twenties when first married
- H: in one's thirties when first married

	E	W	H	Total
M	43	293	114	450
F	82	299	71	452
Total	125	592	185	902

The numbers in the first row mean that 43 people in the sample were men who were first married in their teens, 293 were men who were first married in their twenties, 114 men who were first married in their thirties, and a total of 450 people in the sample were men. Similarly for the numbers in the second row. The numbers in the last row mean that, irrespective of gender, 125 people in the sample were married in their teens, 592 in their twenties, 185 in their thirties, and that there were 902 people in the sample in all. Suppose that the proportions in the sample accurately reflect those in the population of all individuals in the population who are under 40 and who are or have previously been married. Suppose such a person is selected at random.

1. Find the probability that the individual selected was a teenager at first marriage.
2. Find the probability that the individual selected was a teenager at first marriage, given that the person is male.

Solution:

It is natural to let E also denote the event that the person selected was a teenager at first marriage and to let M denote the event that the person selected is male.

1. According to the table, the proportion of individuals in the sample who were in their teens at their first marriage is $125/902$. This is the relative frequency of such people in the population, hence $P(E) = 125/902 \approx 0.139$ or about 14%.
2. Since it is known that the person selected is male, all the females may be removed from consideration, so that only the row in the table corresponding to men in the sample applies:

	E	W	H	Total
M	43	293	114	450

The proportion of males in the sample who were in their teens at their first marriage is $43/450$. This is the relative frequency of such people in the population of males, hence $P(E/M) = 43/450 \approx 0.096$ or about 10%.

In the next example, the computational formula in the definition must be used.

Example 4.3.3: Body Weigth and hypertension

Suppose that in an adult population the proportion of people who are both overweight and suffer hypertension is **0.09**; the proportion of people who are not overweight but suffer hypertension is **0.11**; the proportion of people who are overweight but do not suffer hypertension is **0.02**; and the proportion of people who are neither overweight nor suffer hypertension is **0.78**. An adult is randomly selected from this population.

a. Find the probability that the person selected suffers hypertension given that he is overweight.
b. Find the probability that the selected person suffers hypertension given that he is not overweight.
c. Compare the two probabilities just found to give an answer to the question as to whether overweight people tend to suffer from hypertension.

Solution:

Let H denote the event "the person selected suffers hypertension." Let O denote the event "the person selected is overweight." The probability information given in the problem may be organized into the following contingency table:

	O	O^c
H	0.09	0.11
H^c	0.02	0.78

a. Using the formula in the definition of conditional probability (Equation 4.3.1),

$$P(H|O) = \frac{P(H \cap O)}{P(O)} = \frac{0.09}{0.09 + 0.02} = 0.8182 \tag{4.3.4}$$

b. Using the formula in the definition of conditional probability (Equation 4.3.1),

$$P(H|O) = \frac{P(H \cap O^c)}{P(O^c)} = \frac{0.11}{0.11 + 0.78} = 0.1236 \tag{4.3.5}$$

c. $P(H|O) = 0.8182$ is over six times as large as $P(H|O^c) = 0.1236$, which indicates a much higher rate of hypertension among people who are overweight than among people who are not overweight. It might be interesting to note that a direct comparison of $P(H \cap O) = 0.09$ and $P(H \cap O^c) = 0.11$ does not answer the same question.

Independent Events

Although typically we expect the conditional probability $P(A \mid B)$ to be different from the probability $P(A)$ of A, it does not have to be different from $P(A)$. When $P(A \mid B) = P(A)$, the occurrence of B has no effect on the likelihood of A. Whether or not the event A has occurred is independent of the event B.

Using algebra it can be shown that the equality $P(A \mid B) = P(A)$ holds if and only if the equality $P(A \cap B) = P(A) \cdot P(B)$ holds, which in turn is true if and only if $P(B \mid A) = P(B)$. This is the basis for the following definition.

Definition: Independent and Dependent Events

Events A and B are *independent* (i.e., events whose probability of occurring together is the product of their individual probabilities). if

$$P(A \cap B) = P(A) \cdot P(B) \tag{4.3.6}$$

If A and B are not independent then they are *dependent*.

The formula in the definition has two practical but exactly opposite uses:

- In a situation in which we can compute all three probabilities $P(A), P(B)$ and $P(A \cap B)$, it is used to check whether or not the events A and B are independent:

- If $P(A \cap B) = P(A) \cdot P(B)$, then A and B are independent.
- If $P(A \cap B) \neq P(A) \cdot P(B)$, then A and B are not independent.

- In a situation in which each of $P(A)$ and $P(B)$ can be computed and it is known that A and B are independent, then we can compute $P(A \cap B)$ by multiplying together $P(A)$ and $P(B)$: $P(A \cap B) = P(A) \cdot P(B)$.

Example 4.3.4: Rolling a Die again

A single fair die is rolled. Let $A = \{3\}$ and $B = \{1, 3, 5\}$. Are A and B independent?

Solution:

In this example we can compute all three probabilities $P(A) = 1/6$, $P(B) = 1/2$, and $P(A \cap B) = P(\{3\}) = 1/6$. Since the product $P(A) \cdot P(B) = (1/6)(1/2) = 1/12$ is not the same number as $P(A \cap B) = 1/6$, the events A and B are not independent.

Example 4.3.5

The two-way classification of married or previously married adults under 40 according to gender and age at first marriage produced the table

	E	W	H	Total
M	43	293	114	450
F	82	299	71	452
Total	125	592	185	902

Determine whether or not the events F: "female" and E: "was a teenager at first marriage" are independent.

Solution:

The table shows that in the sample of 902 such adults, 452 were female, 125 were teenagers at their first marriage, and 82 were females who were teenagers at their first marriage, so that

$$P(F) = \frac{452}{902},$$
$$P(E) = \frac{125}{902}$$
$$P(F \cap E) = \frac{82}{902}$$

Since

$$P(F) \cdot P(E) = \frac{452}{902} \cdot \frac{125}{902} = 0.069$$

is not the same as

$$P(F \cap E) = \frac{82}{902} = 0.091$$

we conclude that the two events are not independent.

Example 4.3.6

Many diagnostic tests for detecting diseases do not test for the disease directly but for a chemical or biological product of the disease, hence are not perfectly reliable. The *sensitivity* of a test is the probability that the test will be positive when administered to a person who has the disease. The higher the sensitivity, the greater the detection rate and the lower the false negative rate.

Suppose the sensitivity of a diagnostic procedure to test whether a person has a particular disease is 92%. A person who actually has the disease is tested for it using this procedure by two independent laboratories.

a. What is the probability that both test results will be positive?
b. What is the probability that at least one of the two test results will be positive?

Solution:

a. Let A_1 denote the event "the test by the first laboratory is positive" and let A_2 denote the event "the test by the second laboratory is positive." Since A_1 and A_2 are independent,

$$\begin{aligned} P(A_1 \cap A_2) &= P(A_1) \cdot P(A_2) \\ &= 0.92 \times 0.92 \\ &= 0.8464 \end{aligned}$$

b. Using the Additive Rule for Probability and the probability just computed,

$$\begin{aligned} P(A_1 \cup A_2) &= P(A_1) + P(A_2) - P(A_1 \cap A_2) \\ &= 0.92 + 0.92 - 0.8464 \\ &= 0.9936 \end{aligned}$$

Example 4.3.7: specificity of a diagnostic test

The *specificity* of a diagnostic test for a disease is the probability that the test will be negative when administered to a person who does not have the disease. The higher the specificity, the lower the false positive rate. Suppose the specificity of a diagnostic procedure to test whether a person has a particular disease is 89%.

a. A person who does not have the disease is tested for it using this procedure. What is the probability that the test result will be positive?
b. A person who does not have the disease is tested for it by two independent laboratories using this procedure. What is the probability that both test results will be positive?

Solution:

a. Let B denote the event "the test result is positive." The complement of B is that the test result is negative, and has probability the specificity of the test, 0.89. Thus

$$P(B) = 1 - P(B^c) = 1 - 0.89 = 0.11$$

b. Let B_1 denote the event "the test by the first laboratory is positive" and let B_2 denote the event "the test by the second laboratory is positive." Since B_1 and B_2 are independent, by part (a) of the example

$$P(B_1 \cap B_2) = P(B_1) \cdot P(B_2) = 0.11 \times 0.11 = 0.0121$$

The concept of independence applies to any number of events. For example, three events A, B, and C are independent if $P(A \cap B \cap C) = P(A) \cdot P(B) \cdot P(C)$. Note carefully that, as is the case with just two events, this is not a formula that is always valid, but holds precisely when the events in question are independent.

Example 4.3.8: redundancy

The reliability of a system can be enhanced by redundancy, which means building two or more independent devices to do the same job, such as two independent braking systems in an automobile. Suppose a particular species of trained dogs has a 90% chance of detecting contraband in airline luggage. If the luggage is checked three times by three different dogs independently of one another, what is the probability that contraband will be detected?

Solution:

Let D_1 denote the event that the contraband is detected by the first dog, D_2 the event that it is detected by the second dog, and D_3 the event that it is detected by the third. Since each dog has a 90% of detecting the contraband, by the Probability Rule for

Complements it has a 10% chance of failing. In symbols,

$$P(D_1^c) = 0.10, \quad P(D_2^c) = 0.10, \quad P(D_3^c) = 0.10 \tag{4.3.7}$$

Let D denote the event that the contraband is detected. We seek $P(D)$. It is easier to find $P(D^c)$, because although there are several ways for the contraband to be detected, there is only one way for it to go undetected: all three dogs must fail. Thus $D^c = D_1^c \cap D_2^c \cap D_3^c$ and

$$P(D) = 1 - P(D^c) = 1 - P(D_1^c \cap D_2^c \cap D_3^c) \tag{4.3.8}$$

But the events D_1, D_2, and D_3 are independent, which implies that their complements are independent, so

$$P(D_1^c \cap D_2^c \cap D_3^c) = P(D_1^c) \cdot P(D_2^c) \cdot P(D_3^c) = 0.10 \times 0.10 \times 0.10 = 0.001 \tag{4.3.9}$$

Using this number in the previous display we obtain

$$P(D) = 1 - 0.001 = 0.999 \tag{4.3.10}$$

That is, although any one dog has only a 90% chance of detecting the contraband, three dogs working independently have a 99.9% chance of detecting it.

Probabilities on Tree Diagrams

Some probability problems are made much simpler when approached using a tree diagram. The next example illustrates how to place probabilities on a tree diagram and use it to solve a problem.

Example 4.3.9: A jar of Marbles

A jar contains 10 marbles, 7 black and 3 white. Two marbles are drawn without replacement, which means that the first one is not put back before the second one is drawn.

a. What is the probability that both marbles are black?
b. What is the probability that exactly one marble is black?
c. What is the probability that at least one marble is black?

Solution:

A tree diagram for the situation of drawing one marble after the other without replacement is shown in Figure 4.3.1. The circle and rectangle will be explained later, and should be ignored for now.

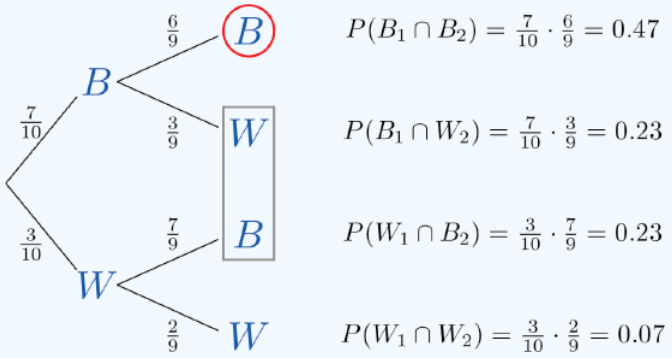


Figure 4.3.1: Tree Diagram for Drawing Two Marbles

The numbers on the two leftmost branches are the probabilities of getting either a black marble, 7 out of 10, or a white marble, 3 out of 10, on the first draw. The number on each remaining branch is the probability of the event corresponding to the node on the right end of the branch occurring, given that the event corresponding to the node on the left end of the branch has occurred. Thus for the top branch, connecting the two Bs, it is $P(B_2 \mid B_1)$, where B_1 denotes the event "the first marble drawn is black" and B_2 denotes the event "the second marble drawn is black." Since after drawing a black marble out there are 9 marbles left, of which 6 are black, this probability is 6/9.

The number to the right of each final node is computed as shown, using the principle that if the formula in the Conditional Rule for Probability is multiplied by $P(B)$, then the result is

$$P(B \cap A) = P(B) \cdot P(A \mid B) \tag{4.3.11}$$

a. The event "both marbles are black" is $B_1 \cap B_2$ and corresponds to the top right node in the tree, which has been circled. Thus as indicated there, it is 0.47.
b. The event "exactly one marble is black" corresponds to the two nodes of the tree enclosed by the rectangle. The events that correspond to these two nodes are mutually exclusive: black followed by white is incompatible with white followed by black. Thus in accordance with the Additive Rule for Probability we merely add the two probabilities next to these nodes, since what would be subtracted from the sum is zero. Thus the probability of drawing exactly one black marble in two tries is $0.23 + 0.23 = 0.46$.
c. The event "at least one marble is black" corresponds to the three nodes of the tree enclosed by either the circle or the rectangle. The events that correspond to these nodes are mutually exclusive, so as in part (b) we merely add the probabilities next to these nodes. Thus the probability of drawing at least one black marble in two tries is $0.47 + 0.23 + 0.23 = 0.93$.

Of course, this answer could have been found more easily using the Probability Law for Complements, simply subtracting the probability of the complementary event, "two white marbles are drawn," from 1 to obtain $1 - 0.07 = 0.93$.

As this example shows, finding the probability for each branch is fairly straightforward, since we compute it knowing everything that has happened in the sequence of steps so far. Two principles that are true in general emerge from this example:

Probabilities on Tree Diagrams

- The probability of the event corresponding to any node on a tree is the product of the numbers on the unique path of branches that leads to that node from the start.
- If an event corresponds to several final nodes, then its probability is obtained by *adding* the numbers next to those nodes.

Key Takeaway

- A conditional probability is the probability that an event has occurred, taking into account additional information about the result of the experiment.
- A conditional probability can always be computed using the formula in the definition. Sometimes it can be computed by discarding part of the sample space.
- Two events A and B are independent if the probability $P(A \cap B)$ of their intersection $A \cap B$ is equal to the product $P(A) \cdot P(B)$ of their individual probabilities.

4.3: Conditional Probability and Independent Events is shared under a CC BY-SA license and was authored, remixed, and/or curated by LibreTexts.

- **3.3: Conditional Probability and Independent Events** by Anonymous is licensed CC BY-NC-SA 3.0. Original source: https://2012books.lardbucket.org/books/beginning-statistics.

4.E: Basic Concepts of Probability (Exercises)

These are homework exercises to accompany the Textmap created for "Introductory Statistics" by Shafer and Zhang.

3.1: Sample Spaces, Events, and Their Probabilities

Basic

Q3.1.1

A box contains 10 white and 10 black marbles. Construct a sample space for the experiment of randomly drawing out, with replacement, two marbles in succession and noting the color each time. (To draw “with replacement” means that the first marble is put back before the second marble is drawn.)

Q3.1.2

A box contains 16 white and 16 black marbles. Construct a sample space for the experiment of randomly drawing out, with replacement, three marbles in succession and noting the color each time. (To draw “with replacement” means that each marble is put back before the next marble is drawn.)

Q3.1.3

A box contains 8 red, 8 yellow, and 8 green marbles. Construct a sample space for the experiment of randomly drawing out, with replacement, two marbles in succession and noting the color each time.

Q3.1.4

A box contains 6 red, 6 yellow, and 6 green marbles. Construct a sample space for the experiment of randomly drawing out, with replacement, three marbles in succession and noting the color each time.

Q3.1.5

In the situation of **Exercise 1**, list the outcomes that comprise each of the following events.

a. At least one marble of each color is drawn.
b. No white marble is drawn.

Q3.1.6

In the situation of **Exercise 2**, list the outcomes that comprise each of the following events.

a. At least one marble of each color is drawn.
b. No white marble is drawn.
c. More black than white marbles are drawn.

Q3.1.7

In the situation of **Exercise 3**, list the outcomes that comprise each of the following events.

a. No yellow marble is drawn.
b. The two marbles drawn have the same color.
c. At least one marble of each color is drawn.

Q3.1.8

In the situation of **Exercise 4**, list the outcomes that comprise each of the following events.

a. No yellow marble is drawn.
b. The three marbles drawn have the same color.
c. At least one marble of each color is drawn.

Q3.1.9

Assuming that each outcome is equally likely, find the probability of each event in **Exercise 5**.

Q3.1.10

Assuming that each outcome is equally likely, find the probability of each event in **Exercise 6**.

Q3.1.11

Assuming that each outcome is equally likely, find the probability of each event in **Exercise 7**.

Q3.1.12

Assuming that each outcome is equally likely, find the probability of each event in **Exercise 8**.

Q3.1.13

A sample space is $S = \{a, b, c, d, e\}$. Identify two events as $U = \{a, b, d\}$ and $V = \{b, c, d\}$. Suppose $P(a)$ and $P(b)$ are each 0.2 and $P(c)$ and $P(d)$ are each 0.1.

a. Determine what $P(e)$ must be.
b. Find $P(U)$.
c. Find $P(V)$

Q3.1.14

A sample space is $S = \{u, v, w, x\}$. Identify two events as $A = \{v, w\}$ and $B = \{u, w, x\}$. Suppose $P(u) = 0.22$, $P(w) = 0.36$, and $P(x) = 0.27$.

a. Determine what $P(v)$ must be.
b. Find $P(A)$.
c. Find $P(B)$.

Q3.1.15

A sample space is $S = \{m, n, q, r, s\}$. Identify two events as $U = \{m, q, s\}$ and $V = \{n, q, r\}$. The probabilities of some of the outcomes are given by the following table:

Outcome	m	n	q	r	s
Probability	0.18	0.16		0.24	0.21

(4.E.1)

a. Determine what $P(q)$ must be.
b. Find $P(U)$.
c. Find $P(V)$.

Q3.1.16

A sample space is $S = \{d, e, f, g, h\}$. Identify two events as $M = \{e, f, g, h\}$ and $N = \{d, g\}$. The probabilities of some of the outcomes are given by the following table:

Outcome	d	e	f	g	h
Probability	0.22	0.13	0.27		0.19

(4.E.2)

a. Determine what $P(g)$ must be.
b. Find $P(M)$.
c. Find $P(N)$.

Applications

Q3.1.17

The sample space that describes all three-child families according to the genders of the children with respect to birth order was constructed in "Example 3.1.4". Identify the outcomes that comprise each of the following events in the experiment of selecting a

three-child family at random.

a. At least one child is a girl.
b. At most one child is a girl.
c. All of the children are girls.
d. Exactly two of the children are girls.
e. The first born is a girl.

Q3.1.18

The sample space that describes three tosses of a coin is the same as the one constructed in "Example 3.1.4" with “boy” replaced by “heads” and “girl” replaced by “tails.” Identify the outcomes that comprise each of the following events in the experiment of tossing a coin three times.

a. The coin lands heads more often than tails.
b. The coin lands heads the same number of times as it lands tails.
c. The coin lands heads at least twice.
d. The coin lands heads on the last toss.

Q3.1.19

Assuming that the outcomes are equally likely, find the probability of each event in **Exercise 17**.

Q3.1.20

Assuming that the outcomes are equally likely, find the probability of each event in **Exercise 18**.

Additional Exercises

Q3.1.21

The following two-way contingency table gives the breakdown of the population in a particular locale according to age and tobacco usage:

Age	Tobacco Use	
	Smoker	Non-smoker
Under 30	0.05	0.20
Over 30	0.20	0.55

A person is selected at random. Find the probability of each of the following events.

a. The person is a smoker.
b. The person is under 30.
c. The person is a smoker who is under 30.

Q3.1.22

The following two-way contingency table gives the breakdown of the population in a particular locale according to party affiliation ($A, B, C,$ or None) and opinion on a bond issue:

Affiliation	Opinion		
	Favors	Opposes	Undecided
A	0.12	0.09	0.07
B	0.16	0.12	0.14
C	0.04	0.03	0.06
None	0.08	0.06	0.03

A person is selected at random. Find the probability of each of the following events.

a. The person is affiliated with party B.
b. The person is affiliated with some party.
c. The person is in favor of the bond issue.
d. The person has no party affiliation and is undecided about the bond issue.

Q3.1.23

The following two-way contingency table gives the breakdown of the population of married or previously married women beyond child-bearing age in a particular locale according to age at first marriage and number of children:

Age	Number of Children		
	0	1 *or* 2	3 or More
Under 20	0.02	0.14	0.08
20 – 29	0.07	0.37	0.11
30 and above	0.10	0.10	0.01

A woman is selected at random. Find the probability of each of the following events.

a. The woman was in her twenties at her first marriage.
b. The woman was 20 or older at her first marriage.
c. The woman had no children.
d. The woman was in her twenties at her first marriage and had at least three children.

Q3.1.24

The following two-way contingency table gives the breakdown of the population of adults in a particular locale according to highest level of education and whether or not the individual regularly takes dietary supplements:

Education	Use of Supplements	
	Takes	Does Not Take
No High School Diploma	0.04	0.06
High School Diploma	0.06	0.44
Undergraduate Degree	0.09	0.28
Graduate Degree	0.01	0.02

An adult is selected at random. Find the probability of each of the following events.

a. The person has a high school diploma and takes dietary supplements regularly.
b. The person has an undergraduate degree and takes dietary supplements regularly.

c. The person takes dietary supplements regularly.
d. The person does not take dietary supplements regularly.

Large Data Set Exercises

Q3.1.25

Large Data Set 4 and Data Set 4A record the results of 500 tosses of a coin. Find the relative frequency of each outcome $1, 2, 3, 4, 5,$ *and* 6. Does the coin appear to be “balanced” or “fair”?

Q3.1.26

Large Data Set 6, Data Set 6A, and Data Set 6B record results of a random survey of 200 voters in each of two regions, in which they were asked to express whether they prefer Candidate A for a U.S. Senate seat or prefer some other candidate.

a. Find the probability that a randomly selected voter among these 400 prefers Candidate A.
b. Find the probability that a randomly selected voter among the 200 who live in Region 1 prefers Candidate A (separately recorded in **Large Data Set 6A**).
c. Find the probability that a randomly selected voter among the 200 who live in Region 2 prefers Candidate A (separately recorded in **Large Data Set 6B**).

Answers

S3.1.1

$S = \{bb, bw, wb, ww\}$

S3.1.3

$S = \{rr, ry, rg, yr, yy, yg, gr, gy, gg\}$

S3.1.5

a. $\{bw, wb\}$
b. $\{bb\}$

S3.1.7

a. $\{rr, rg, gr, gg\}$
b. $\{rr, yy, gg\}$
c. $\varnothing$

S3.1.9

a. $1/4$
b. $2/4$

S3.1.11

a. $4/9$
b. $3/9$
c. 0

S3.1.13

a. 0.4
b. 0.5
c. 0.4

S3.1.15

a. 0.61
b. 0.6
c. 0.21

S3.1.17

a. $\{gbb, gbg, ggb, ggg\}$
b. $\{bgg, gbg, ggb\}$
c. $\{ggg\}$
d. $\{bbb, bbg, bgb, gbb\}$
e. $\{bbg, bgb, bgg, gbb, gbg, ggb, ggg\}$

S3.1.19

a. $4/8$
b. $3/8$
c. $1/8$
d. $4/8$
e. $7/8$

S3.1.21

a. 0.05
b. 0.25
c. 0.25

S3.1.23

a. 0.11
b. 0.19
c. 0.76
d. 0.55

S3.1.25

The relative frequencies for 1 through 6 are $0.16, 0.194, 0.162, 0.164, 0.154$ *and* 0.166. It would appear that the die is not balanced.

3.2: Complements, Intersections and Unions

Basic

1. For the sample space $S = \{a, b, c, d, e\}$ identify the complement of each event given.
   a. $A = \{a, d, e\}$
   b. $B = \{b, c, d, e\}$
   c. S
2. For the sample space $S = \{r, s, t, u, v\}$ identify the complement of each event given.
   a. $R = \{t, u\}$
   b. $T = \{r\}$
   c. $\varnothing$ (the "empty" set that has no elements)
3. The sample space for three tosses of a coin is $S = \{hhh, hht, hth, htt, thh, tht, tth, ttt\}$ Define events

$$\begin{array}{c}\text{H:at least one head is observed} \\ \text{M:more heads than tails are observed}\end{array} \tag{4.E.3}$$

   a. List the outcomes that comprise H and M.
   b. List the outcomes that comprise $H \cap M$, $H \cup M$, and H^c.
   c. Assuming all outcomes are equally likely, find $P(H \cap M)$, $P(H \cup M)$, and $P(H^c)$.
   d. Determine whether or not H^c and M are mutually exclusive. Explain why or why not.
4. For the experiment of rolling a single six-sided die once, define events

$$\begin{array}{c}\text{T:the number rolled is three} \\ \text{G:the number rolled is four or greater}\end{array} \tag{4.E.4}$$

a. List the outcomes that comprise T and G.
b. List the outcomes that comprise $T \cap G$, $T \cup G$, T^c, and $(T \cup G)^c$.
c. Assuming all outcomes are equally likely, find $P(T \cap G)$, $P(T \cup G)$, and $P(T^c)$.
d. Determine whether or not T and G are mutually exclusive. Explain why or why not.

5. A special deck of **16** cards has **4** that are blue, **4** yellow, **4** green, and **4** red. The four cards of each color are numbered from one to four. A single card is drawn at random. Define events

$$\begin{gathered}\text{B:the card is blue} \\ \text{R:the card is red} \\ \text{N:the number on the card is at most two}\end{gathered} \tag{4.E.5}$$

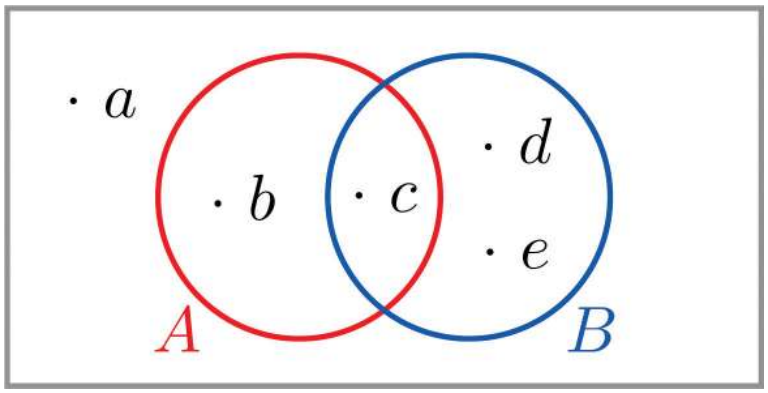
a. List the outcomes that comprise B, R, and N.
b. List the outcomes that comprise $B \cap R$, $B \cup R$, $B \cap N$, $R \cup N$, B^c, and $(B \cup R)^c$.
c. Assuming all outcomes are equally likely, find the probabilities of the events in the previous part.
d. Determine whether or not B and N are mutually exclusive. Explain why or why not.

6. In the context of the previous problem, define events

$$\begin{gathered}\text{Y:the card is yellow} \\ \text{I:the number on the card is not a one} \\ \text{J:the number on the card is a two or a four}\end{gathered} \tag{4.E.6}$$

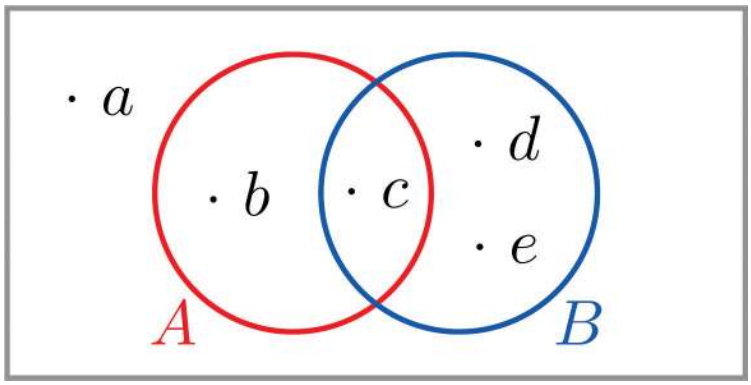
a. List the outcomes that comprise Y, I, and J.
b. List the outcomes that comprise $Y \cap I$, $Y \cup J$, $I \cap J$, I^c, and $(Y \cup J)^c$.
c. Assuming all outcomes are equally likely, find the probabilities of the events in the previous part.
d. Determine whether or not I^c and J are mutually exclusive. Explain why or why not.

7. The Venn diagram provided shows a sample space and two events A and B. Suppose $P(a) = 0.13, P(b) = 0.09, P(c) = 0.27, P(d) = 0.20,$ and $P(e) = 0.31$. Confirm that the probabilities of the outcomes add up to **1**, then compute the following probabilities.



a. $P(A)$.
b. $P(B)$.
c. $P(A^c)$. Two ways: (i) by finding the outcomes in A^c and adding their probabilities, and (ii) using the Probability Rule for Complements.
d. $P(A \cap B)$.
e. $P(A \cup B)$ Two ways: (i) by finding the outcomes in $A \cup B$ and adding their probabilities, and (ii) using the Additive Rule of Probability.

8. The Venn diagram provided shows a sample space and two events A and B. Suppose $P(a) = 0.32, P(b) = 0.17, P(c) = 0.28,$ and $P(d) = 0.23$. Confirm that the probabilities of the outcomes add up to **1**, then compute the following probabilities.



a. $P(A)$.
b. $P(B)$.
c. $P(A^c)$. Two ways: (i) by finding the outcomes in A^c and adding their probabilities, and (ii) using the Probability Rule for Complements.
d. $P(A \cap B)$.
e. $P(A \cup B)$ Two ways: (i) by finding the outcomes in $A \cup B$ and adding their probabilities, and (ii) using the Additive Rule of Probability.

9. Confirm that the probabilities in the two-way contingency table add up to 1, then use it to find the probabilities of the events indicated.

	U	V	W
A	0.15	0.00	0.23
B	0.22	0.30	0.10

a. $P(A), P(B), P(A \cap B)$.
b. $P(U), P(W), P(U \cap W)$.
c. $P(U \cup W)$.
d. $P(V^c)$.
e. Determine whether or not the events A and U are mutually exclusive; the events A and V.

10. Confirm that the probabilities in the two-way contingency table add up to 1, then use it to find the probabilities of the events indicated.

	R	S	T
M	0.09	0.25	0.19
N	0.31	0.16	0.00

a. $P(R), P(S), P(R \cap S)$.
b. $P(M), P(N), P(M \cap N)$.
c. $P(R \cup S)$.
d. $P(R^c)$.
e. Determine whether or not the events N and S are mutually exclusive; the events N and T.

Applications

11. Make a statement in ordinary English that describes the complement of each event (do not simply insert the word “not”).

a. In the roll of a die: “five or more.”
b. In a roll of a die: “an even number.”
c. In two tosses of a coin: “at least one heads.”

d. In the random selection of a college student: “Not a freshman.”

12. Make a statement in ordinary English that describes the complement of each event (do not simply insert the word “not”).
    a. In the roll of a die: “two or less.”
    b. In the roll of a die: “one, three, or four.”
    c. In two tosses of a coin: “at most one heads.”
    d. In the random selection of a college student: “Neither a freshman nor a senior.”
13. The sample space that describes all three-child families according to the genders of the children with respect to birth order is $S = \{bbb, bbg, bgb, bgg, gbb, gbg, ggb, ggg\}$. For each of the following events in the experiment of selecting a three-child family at random, state the complement of the event in the simplest possible terms, then find the outcomes that comprise the event and its complement.
    a. At least one child is a girl.
    b. At most one child is a girl.
    c. All of the children are girls.
    d. Exactly two of the children are girls.
    e. The first born is a girl.
14. The sample space that describes the two-way classification of citizens according to gender and opinion on a political issue is $S = \{mf, ma, mn, ff, fa, fn\}$, where the first letter denotes gender (m: male, f: female) and the second opinion (f: for, a: against, n: neutral). For each of the following events in the experiment of selecting a citizen at random, state the complement of the event in the simplest possible terms, then find the outcomes that comprise the event and its complement.
    a. The person is male.
    b. The person is not in favor.
    c. The person is either male or in favor.
    d. The person is female and neutral.
15. A tourist who speaks English and German but no other language visits a region of Slovenia. If 35% of the residents speak English, 15% speak German, and 3% speak both English and German, what is the probability that the tourist will be able to talk with a randomly encountered resident of the region?
16. In a certain country 43% of all automobiles have airbags, 27% have anti-lock brakes, and 13% have both. What is the probability that a randomly selected vehicle will have both airbags and anti-lock brakes?
17. A manufacturer examines its records over the last year on a component part received from outside suppliers. The breakdown on source (supplier A, supplier B) and quality (H: high, U: usable, D: defective) is shown in the two-way contingency table.

	H	U	D
A	0.6937	0.0049	0.0014
B	0.2982	0.0009	0.0009

The record of a part is selected at random. Find the probability of each of the following events.

    a. The part was defective.
    b. The part was either of high quality or was at least usable, in two ways: (i) by adding numbers in the table, and (ii) using the answer to (a) and the Probability Rule for Complements.
    c. The part was defective and came from supplier B.
    d. The part was defective or came from supplier B, in two ways: by finding the cells in the table that correspond to this event and adding their probabilities, and (ii) using the Additive Rule of Probability.

18. Individuals with a particular medical condition were classified according to the presence (T) or absence (N) of a potential toxin in their blood and the onset of the condition (E: early, M: midrange, L: late). The breakdown according to this classification is shown in the two-way contingency table.

	E	M	L
T	0.012	0.124	0.013
N	0.170	0.638	0.043

One of these individuals is selected at random. Find the probability of each of the following events.

a. The person experienced early onset of the condition.
b. The onset of the condition was either midrange or late, in two ways: (i) by adding numbers in the table, and (ii) using the answer to (a) and the Probability Rule for Complements.
c. The toxin is present in the person's blood.
d. The person experienced early onset of the condition and the toxin is present in the person's blood.
e. The person experienced early onset of the condition or the toxin is present in the person's blood, in two ways: (i) by finding the cells in the table that correspond to this event and adding their probabilities, and (ii) using the Additive Rule of Probability.

19. The breakdown of the students enrolled in a university course by class (**F: freshman, So: sophomore, J: junior, Se: senior**) and academic major (**S: science, mathematics, or engineering, L: liberal arts, O: other**) is shown in the two-way classification table.

Major	Class			
	F	So	J	Se
S	92	42	20	13
L	368	167	80	53
O	460	209	100	67

A student enrolled in the course is selected at random. Adjoin the row and column totals to the table and use the expanded table to find the probability of each of the following events.

a. The student is a freshman.
b. The student is a liberal arts major.
c. The student is a freshman liberal arts major.
d. The student is either a freshman or a liberal arts major.
e. The student is not a liberal arts major.

20. The table relates the response to a fund-raising appeal by a college to its alumni to the number of years since graduation.

Response	Years Since Graduation			
	0 – 5	6 – 20	21 – 35	Over 35
Positive	120	440	210	90
None	1380	3560	3290	910

An alumnus is selected at random. Adjoin the row and column totals to the table and use the expanded table to find the probability of each of the following events.

a. The alumnus responded.
b. The alumnus did not respond.
c. The alumnus graduated at least 21 years ago.
d. The alumnus graduated at least 21 years ago and responded.

Additional Exercises

21. The sample space for tossing three coins is $S = \{hhh, hht, hth, htt, thh, tht, tth, ttt\}$
    a. List the outcomes that correspond to the statement “All the coins are heads.”
    b. List the outcomes that correspond to the statement “Not all the coins are heads.”
    c. List the outcomes that correspond to the statement “All the coins are not heads.”

Answers

1. a. $\{b, c\}$
   b. $\{a\}$
   c. $\varnothing$
2.
3. a. $H = \{hhh, hht, hth, htt, thh, tht, tth\},\ M = \{hhh, hht, hth, thh\}$
   b. $H \cap M = \{hhh, hht, hth, thh\}, H \cup M = H, H^c = \{ttt\}$
   c. $P(H \cap M) = 4/8, P(H \cup M) = 7/8, P(H^c) = 1/8$
   d. Mutually exclusive because they have no elements in common.
4.
5. a. $B = \{b1, b2, b3, b4\},\ R = \{r1, r2, r3, r4\},\ N = \{b1, b2, y1, y2, g1, g2, r1, r2\}$
   b. $B \cap R = \varnothing, B \cup R = \{b1, b2, b3, b4, r1, r2, r3, r4\},\ B \cap N = \{b1, b2\},$
   $R \cup N = \{b1, b2, y1, y2, g1, g2, r1, r2, r3, r4\},$
   $B^c = \{y1, y2, y3, y4, g1, g2, g3, g4, r1, r2, r3, r4\},\ (B \cup R)^c = \{y1, y2, y3, y4, g1, g2, g3, g4\}$
   c. $P(B \cap R) = 0,\ P(B \cup R) = 8/16,\ P(B \cap N) = 2/16,\ P(R \cup N) = 10/16,\ P(B^c) = 12/16,\ P((B \cup R)^c) = 8/16$
   d. Not mutually exclusive because they have an element in common.
6.
7. a. 0.36
   b. 0.78
   c. 0.64
   d. 0.27
   e. 0.87
8.
9. a. $P(A) = 0.38,\ P(B) = 0.62,\ P(A \cap B) = 0$
   b. $P(U) = 0.37,\ P(W) = 0.33,\ P(U \cap W) = 0$
   c. 0.7
   d. 0.7
   e. A and U are not mutually exclusive because $P(A \cap U)$ is the nonzero number 0.15. A and V are mutually exclusive because $P(A \cap V) = 0$.
10.
11. a. “four or less”
    b. “an odd number”
    c. “no heads” or “all tails”
    d. “a freshman”
12.
13. a. “All the children are boys.” Event: $\{bbg, bgb, bgg, gbb, gbg, ggb, ggg\}$, Complement: $\{bbb\}$
    b. “At least two of the children are girls” or “There are two or three girls.” Event: $\{bbb, bbg, bgb, gbb\}$, Complement: $\{bgg, gbg, ggb, ggg\}$
    c. “At least one child is a boy.” Event: $\{ggg\}$, Complement: $\{bbb, bbg, bgb, bgg, gbb, gbg, ggb\}$
    d. “There are either no girls, exactly one girl, or three girls.” Event: $\{bgg, gbg, ggb\}$, Complement: $\{bbb, bbg, bgb, gbb, ggg\}$
    e. “The first born is a boy.” Event: $\{gbb, gbg, ggb, ggg\}$, Complement: $\{bbb, bbg, bgb, bgg\}$

14.
15. **0.47**
16.
17. a. **0.0023**
    b. **0.9977**
    c. **0.0009**
    d. **0.3014**
18.
19. a. **920/1671**
    b. **668/1671**
    c. **368/1671**
    d. **1220/1671**
    e. **1003/1671**
20.
21. a. $\{hhh\}$
    b. $\{hht, hth, htt, thh, tht, tth, ttt\}$
    c. $\{ttt\}$

3.3: Conditional Probability and Independent Events

Basic

1. Q3.3.1For two events A and B, $P(A) = 0.73,\ P(B) = 0.48$ and $P(A \cap B) = 0.29$.
   a. Find $P(A \mid B)$.
   b. Find $P(B \mid A)$.
   c. Determine whether or not A and B are independent.
2. Q3.3.1For two events A and B, $P(A) = 0.26,\ P(B) = 0.37$ and $P(A \cap B) = 0.11$.
   a. Find $P(A \mid B)$.
   b. Find $P(B \mid A)$.
   c. Determine whether or not A and B are independent.
3. Q3.3.1For independent events A and B, $P(A) = 0.81$ and $P(B) = 0.27$.
   a. $P(A \cap B)$.
   b. Find $P(A \mid B)$.
   c. Find $P(B \mid A)$.
4. Q3.3.1For independent events A and B, $P(A) = 0.68$ and $P(B) = 0.37$.
   a. $P(A \cap B)$.
   b. Find $P(A \mid B)$.
   c. Find $P(B \mid A)$.
5. Q3.3.1For mutually exclusive events A and B, $P(A) = 0.17$ and $P(B) = 0.32$.
   a. Find $P(A \mid B)$.
   b. Find $P(B \mid A)$.
6. Q3.3.1For mutually exclusive events A and B, $P(A) = 0.45$ and $P(B) = 0.09$.
   a. Find $P(A \mid B)$.
   b. Find $P(B \mid A)$.
7. Q3.3.1Compute the following probabilities in connection with the roll of a single fair die.
   a. The probability that the roll is even.
   b. The probability that the roll is even, given that it is not a two.
   c. The probability that the roll is even, given that it is not a one.

8. Q3.3.1Compute the following probabilities in connection with two tosses of a fair coin.
    a. The probability that the second toss is heads.
    b. The probability that the second toss is heads, given that the first toss is heads.
    c. The probability that the second toss is heads, given that at least one of the two tosses is heads.
9. Q3.3.1A special deck of **16** cards has **4** that are blue, **4** yellow, **4** green, and **4** red. The four cards of each color are numbered from one to four. A single card is drawn at random. Find the following probabilities.
    a. The probability that the card drawn is red.
    b. The probability that the card is red, given that it is not green.
    c. The probability that the card is red, given that it is neither red nor yellow.
    d. The probability that the card is red, given that it is not a four.
10. Q3.3.1A special deck of **16** cards has **4** that are blue, **4** yellow, **4** green, and **4** red. The four cards of each color are numbered from one to four. A single card is drawn at random. Find the following probabilities.
    a. The probability that the card drawn is a two or a four.
    b. The probability that the card is a two or a four, given that it is not a one.
    c. The probability that the card is a two or a four, given that it is either a two or a three.
    d. The probability that the card is a two or a four, given that it is red or green.
11. Q3.3.1A random experiment gave rise to the two-way contingency table shown. Use it to compute the probabilities indicated.

	R	S
A	0.12	0.18
B	0.28	0.42

    a. $P(A),\ P(R),\ P(A \cap B)$.
    b. Based on the answer to (a), determine whether or not the events A and R are independent.
    c. Based on the answer to (b), determine whether or not $P(A \mid R)$ can be predicted without any computation. If so, make the prediction. In any case, compute $P(A \mid R)$ using the Rule for Conditional Probability.
12. Q3.3.1A random experiment gave rise to the two-way contingency table shown. Use it to compute the probabilities indicated.

	R	S
A	0.13	0.07
B	0.61	0.19

    a. $P(A),\ P(R),\ P(A \cap B)$.
    b. Based on the answer to (a), determine whether or not the events A and R are independent.
    c. Based on the answer to (b), determine whether or not $P(A \mid R)$ can be predicted without any computation. If so, make the prediction. In any case, compute $P(A \mid R)$ using the Rule for Conditional Probability.
13. Q3.3.1Suppose for events A and B in a random experiment $P(A) = 0.70$ and $P(B) = 0.30$.Compute the indicated probability, or explain why there is not enough information to do so.
    a. $P(A \cap B)$.
    b. $P(A \cap B)$, with the extra information that A and B are independent.
    c. $P(A \cap B)$, with the extra information that A and B are mutually exclusive.
14. Q3.3.1Suppose for events A and B in a random experiment $P(A) = 0.50$ and $P(B) = 0.50$. Compute the indicated probability, or explain why there is not enough information to do so.
    a. $P(A \cap B)$.
    b. $P(A \cap B)$, with the extra information that A and B are independent.
    c. $P(A \cap B)$, with the extra information that A and B are mutually exclusive.

15. Q3.3.1Suppose for events A, B, *and* C connected to some random experiment, A, B, *and* C are independent and $P(A) = 0.50$, $P(B) = 0.50$ **and** $P(C) = 0.44$. Compute the indicated probability, or explain why there is not enough information to do so.
    a. $P(A \cap B \cap C)$.
    b. $P(A^c \cap B^c \cap C^c)$.
16. Q3.3.1Suppose for events A, B, *and* C connected to some random experiment, A, B, *and* C are independent and $P(A) = 0.95$, $P(B) = 0.73$ **and** $P(C) = 0.62$. Compute the indicated probability, or explain why there is not enough information to do so.
    a. $P(A \cap B \cap C)$.
    b. $P(A^c \cap B^c \cap C^c)$.

Applications

Q3.3.17

The sample space that describes all three-child families according to the genders of the children with respect to birth order is

$$S = \{bbb, bbg, bgb, bgg, gbb, gbg, ggb, ggg\} \tag{4.E.7}$$

In the experiment of selecting a three-child family at random, compute each of the following probabilities, assuming all outcomes are equally likely.

a. The probability that the family has at least two boys.
b. The probability that the family has at least two boys, given that not all of the children are girls.
c. The probability that at least one child is a boy.
d. The probability that at least one child is a boy, given that the first born is a girl.

Q3.3.18

The following two-way contingency table gives the breakdown of the population in a particular locale according to age and number of vehicular moving violations in the past three years:

Age	Violations		
	0	**1**	**2+**
Under **21**	**0.04**	**0.06**	**0.02**
21 – 40	**0.25**	**0.16**	**0.01**
41 – 60	**0.23**	**0.10**	**0.02**
60+	**0.08**	**0.03**	**0.00**

A person is selected at random. Find the following probabilities.

a. The person is under **21**.
b. The person has had at least two violations in the past three years.
c. The person has had at least two violations in the past three years, given that he is under **21**.
d. The person is under **21**, given that he has had at least two violations in the past three years.
e. Determine whether the events "the person is under **21**" and "the person has had at least two violations in the past three years" are independent or not.

Q3.3.19

The following two-way contingency table gives the breakdown of the population in a particular locale according to party affiliation ($A, B, C,$ **or None**) and opinion on a bond issue:

Affiliation	Opinion		
	Favors	Opposes	Undecided
A	0.12	0.09	0.07
B	0.16	0.12	0.14
C	0.04	0.03	0.06
None	0.08	0.06	0.03

A person is selected at random. Find each of the following probabilities.

a. The person is in favor of the bond issue.
b. The person is in favor of the bond issue, given that he is affiliated with party A.
c. The person is in favor of the bond issue, given that he is affiliated with party B.

Q3.3.20

The following two-way contingency table gives the breakdown of the population of patrons at a grocery store according to the number of items purchased and whether or not the patron made an impulse purchase at the checkout counter:

Number of Items	Impulse Purchase	
	Made	Not Made
Few	0.01	0.19
Many	0.04	0.76

A patron is selected at random. Find each of the following probabilities.

a. The patron made an impulse purchase.
b. The patron made an impulse purchase, given that the total number of items purchased was many.
c. Determine whether or not the events “few purchases” and “made an impulse purchase at the checkout counter” are independent.

Q3.3.21

The following two-way contingency table gives the breakdown of the population of adults in a particular locale according to employment type and level of life insurance:

Employment Type	Level of Insurance		
	Low	Medium	High
Unskilled	0.07	0.19	0.00
Semi-skilled	0.04	0.28	0.08
Skilled	0.03	0.18	0.05
Professional	0.01	0.05	0.02

An adult is selected at random. Find each of the following probabilities.

a. The person has a high level of life insurance.
b. The person has a high level of life insurance, given that he does not have a professional position.
c. The person has a high level of life insurance, given that he has a professional position.
d. Determine whether or not the events “has a high level of life insurance” and “has a professional position” are independent.

Q3.3.22

The sample space of equally likely outcomes for the experiment of rolling two fair dice is

$$\begin{array}{cccccc} 11 & 12 & 13 & 14 & 15 & 16 \\ 21 & 22 & 23 & 24 & 25 & 26 \\ 31 & 32 & 33 & 34 & 35 & 36 \\ 41 & 42 & 43 & 44 & 45 & 46 \\ 51 & 52 & 53 & 54 & 55 & 56 \\ 61 & 62 & 63 & 64 & 65 & 66 \end{array} \tag{4.E.8}$$

Identify the events **N: the sum is at least nine, T: at least one of the dice is a two, and F: at least one of the dice is a five.**

a. Find $P(N)$.
b. Find $P(N \mid F)$.
c. Find $P(N \mid T)$.
d. Determine from the previous answers whether or not the events N and F are independent; whether or not N and T are.

Q3.3.23

The *sensitivity* of a drug test is the probability that the test will be positive when administered to a person who has actually taken the drug. Suppose that there are two independent tests to detect the presence of a certain type of banned drugs in athletes. One has sensitivity 0.75; the other has sensitivity 0.85. If both are applied to an athlete who has taken this type of drug, what is the chance that his usage will go undetected?

Q3.3.24

A man has two lights in his well house to keep the pipes from freezing in winter. He checks the lights daily. Each light has probability 0.002 of burning out before it is checked the next day (independently of the other light).

a. If the lights are wired in parallel one will continue to shine even if the other burns out. In this situation, compute the probability that at least one light will continue to shine for the full 24 hours. Note the greatly increased reliability of the system of two bulbs over that of a single bulb.
b. If the lights are wired in series neither one will continue to shine even if only one of them burns out. In this situation, compute the probability that at least one light will continue to shine for the full 24 hours. Note the slightly decreased reliability of the system of two bulbs over that of a single bulb.

Q3.3.25

An accountant has observed that 5% of all copies of a particular two-part form have an error in Part I, and 2% have an error in Part II. If the errors occur independently, find the probability that a randomly selected form will be error-free.

Q3.3.26

A box contains 20 screws which are identical in size, but 12 of which are zinc coated and 8 of which are not. Two screws are selected at random, without replacement.

a. Find the probability that both are zinc coated.
b. Find the probability that at least one is zinc coated.

Additional Exercises

Q3.3.27

Events A and B are mutually exclusive. Find $P(A \mid B)$.

Q3.3.28

The city council of a particular city is composed of five members of party A, four members of party B, and three independents. Two council members are randomly selected to form an investigative committee.

a. Find the probability that both are from party A.

b. Find the probability that at least one is an independent.
c. Find the probability that the two have different party affiliations (that is, not both A, not both B, and not both independent).

Q3.3.29

A basketball player makes 60% of the free throws that he attempts, except that if he has just tried and missed a free throw then his chances of making a second one go down to only 30%. Suppose he has just been awarded two free throws.

a. Find the probability that he makes both.
b. Find the probability that he makes at least one. (A tree diagram could help.)

Q3.3.30

An economist wishes to ascertain the proportion p of the population of individual taxpayers who have purposely submitted fraudulent information on an income tax return. To truly guarantee anonymity of the taxpayers in a random survey, taxpayers questioned are given the following instructions.

a. Flip a coin.
b. If the coin lands heads, answer "Yes" to the question "Have you ever submitted fraudulent information on a tax return?" even if you have not.
c. If the coin lands tails, give a truthful "Yes" or "No" answer to the question "Have you ever submitted fraudulent information on a tax return?"

The questioner is not told how the coin landed, so he does not know if a "Yes" answer is the truth or is given only because of the coin toss.

a. Using the Probability Rule for Complements and the independence of the coin toss and the taxpayers' status fill in the empty cells in the two-way contingency table shown. Assume that the coin is fair. Each cell except the two in the bottom row will contain the unknown proportion (or probability) p.

Status	Coin		Probability
	H	T	
Fraud			p
No fraud			
Probability			1

b. The only information that the economist sees are the entries in the following table:

$Response$	$"Yes"$	$"No"$
$Proportion$	r	s

(4.E.9)

Equate the entry in the one cell in the table in (a) that corresponds to the answer "No" to the number s to obtain the formula that expresses the unknown number p in terms of the known number s.
c. Equate the sum of the entries in the three cells in the table in (a) that together correspond to the answer "Yes" to the number r to obtain the formula that expresses the unknown number p in terms of the known number r.
d. Use the fact that $r + s = 1$(since they are the probabilities of complementary events) to verify that the formulas in (b) and (c) give the same value for p. (For example, insert $s = 1 - r$into the formula in (b) to obtain the formula in (c)).
e. Suppose a survey of $1,200$ taxpayers is conducted and 690 respond "Yes" (truthfully or not) to the question "Have you ever submitted fraudulent information on a tax return?" Use the answer to either (b) or (c) to estimate the true proportion p of all individual taxpayers who have purposely submitted fraudulent information on an income tax return.

Answers

1. a. 0.6
   b. 0.4
   c. not independent

2.

3. a. **0.22**
   b. **0.81**
   c. **0.27**

4.

5. a. **0**
   b. **0**

6.

7. a. **0.5**
   b. **0.4**
   c. **0.6**

8.

9. a. **0.25**
   b. **0.33**
   c. **0**
   d. **0.25**

10.

11. a. $P(A) = 0.3,\ P(R) = 0.4,\ P(A \cap R) = 0.12$
    b. independent
    c. without computation **0.3**

12.

13. a. Insufficient information. The events *A* and *B* are not known to be either independent or mutually exclusive.
    b. **0.21**
    c. **0**

14.

15. a. **0.25**
    b. **0.02**

16.

17. a. **0.5**
    b. **0.57**
    c. **0.875**
    d. **0.75**

18.

19. a. **0.4**
    b. **0.43**
    c. **0.38**

20.

21. a. **0.15**
    b. **0.14**
    c. **0.25**
    d. not independent

22.

23. **0.0375**

24.

25. **0.931**

26.

https://stats.libretexts.org/@go/page/14472

27. **0**
28.
29. a. **0.36**
    b. **0.72**

Contributor

- Anonymous

4.E: Basic Concepts of Probability (Exercises) is shared under a CC BY-SA license and was authored, remixed, and/or curated by LibreTexts.

- **3.E: Basic Concepts of Probability (Exercises)** has no license indicated.

CHAPTER OVERVIEW

5: Z-scores and the Standard Normal Distribution

5.1: Normal Distributions

5.2: Z-scores

5.3: Z-scores and the Area under the Curve

5.E: Z-scores and the Standard Normal Distribution (Exercises)

This page titled 5: Z-scores and the Standard Normal Distribution is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

5.1: Normal Distributions

The normal distribution is the most important and most widely used distribution in statistics. It is sometimes called the “bell curve,” although the tonal qualities of such a bell would be less than pleasing. It is also called the “Gaussian curve” of Gaussian distribution after the mathematician Karl Friedrich Gauss.

Strictly speaking, it is not correct to talk about “the normal distribution” since there are many normal distributions. Normal distributions can differ in their means and in their standard deviations. Figure 1 shows three normal distributions. The green (left-most) distribution has a mean of -3 and a standard deviation of 0.5, the distribution in red (the middle distribution) has a mean of 0 and a standard deviation of 1, and the distribution in black (right-most) has a mean of 2 and a standard deviation of 3. These as well as all other normal distributions are symmetric with relatively more values at the center of the distribution and relatively few in the tails. What is consistent about all normal distribution is the shape and the proportion of scores within a given distance along the x-axis. We will focus on the Standard Normal Distribution (also known as the Unit Normal Distribution), which has a mean of 0 and a standard deviation of 1 (i.e. the red distribution in Figure **5.1.1**).

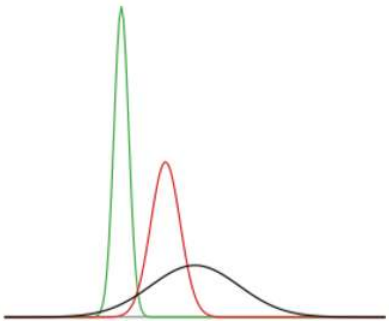


Figure **5.1.1**: Normal distributions differing in mean and standard deviation.

Seven features of normal distributions are listed below.

1. Normal distributions are symmetric around their mean.
2. The mean, median, and mode of a normal distribution are equal.
3. The area under the normal curve is equal to 1.0.
4. Normal distributions are denser in the center and less dense in the tails.
5. Normal distributions are defined by two parameters, the mean (μ) and the standard deviation (σ).
6. 68% of the area of a normal distribution is within one standard deviation of the mean.
7. Approximately 95% of the area of a normal distribution is within two standard deviations of the mean.

These properties enable us to use the normal distribution to understand how scores relate to one another within and across a distribution. But first, we need to learn how to calculate the standardized score than make up a standard normal distribution.

This page titled 5.1: Normal Distributions is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri’s Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **4.1: Normal Distributions** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

5.2: Z-scores

A z-score is a standardized version of a raw score (x) that gives information about the relative location of that score within its distribution. The formula for converting a raw score into a z-score is:

$$z = \frac{x - \mu}{\sigma} \tag{5.2.1}$$

for values from a population and for values from a sample:

$$z = \frac{x - M}{s} \tag{5.2.2}$$

As you can see, z-scores combine information about where the distribution is located (the mean/center) with how wide the distribution is (the standard deviation/spread) to interpret a raw score (x). Specifically, z-scores will tell us how far the score is away from the mean in units of standard deviations and in what direction.

The value of a z-score has two parts: the sign (positive or negative) and the magnitude (the actual number). The sign of the z-score tells you in which half of the distribution the z-score falls: a positive sign (or no sign) indicates that the score is above the mean and on the right hand-side or upper end of the distribution, and a negative sign tells you the score is below the mean and on the left-hand side or lower end of the distribution. The magnitude of the number tells you, in units of standard deviations, how far away the score is from the center or mean. The magnitude can take on any value between negative and positive infinity, but for reasons we will see soon, they generally fall between -3 and 3.

Let's look at some examples. A z-score value of -1.0 tells us that this z-score is 1 standard deviation (because of the magnitude 1.0) below (because of the negative sign) the mean. Similarly, a z-score value of 1.0 tells us that this z-score is 1 standard deviation above the mean. Thus, these two scores are the same distance away from the mean but in opposite directions. A z-score of -2.5 is two-and-a-half standard deviations below the mean and is therefore farther from the center than both of the previous scores, and a z -score of 0.25 is closer than all of the ones before. In Unit 2, we will learn to formalize the distinction between what we consider "close" to the center or "far" from the center. For now, we will use a rough cut-off of 1.5 standard deviations in either direction as the difference between close scores (those within 1.5 standard deviations or between z = -1.5 and z = 1.5) and extreme scores (those farther than 1.5 standard deviations – below z = -1.5 or above z = 1.5).

We can also convert raw scores into z-scores to get a better idea of where in the distribution those scores fall. Let's say we get a score of 68 on an exam. We may be disappointed to have scored so low, but perhaps it was just a very hard exam. Having information about the distribution of all scores in the class would be helpful to put some perspective on ours. We find out that the class got an average score of 54 with a standard deviation of 8. To find out our relative location within this distribution, we simply convert our test score into a z-score.

$$z = \frac{X - \mu}{\sigma} = \frac{68 - 54}{8} = 1.75$$

We find that we are 1.75 standard deviations above the average, above our rough cut off for close and far. Suddenly our 68 is looking pretty good!

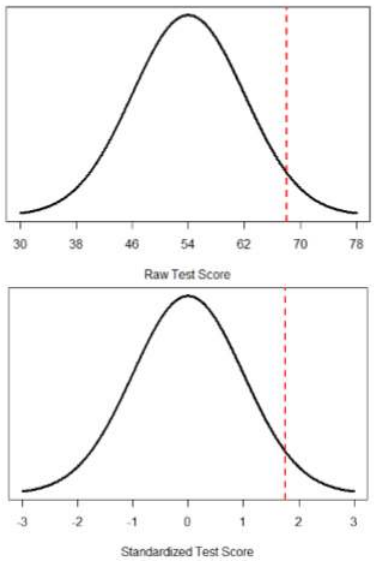


Figure 5.2.1: Raw and standardized versions of a single score

Figure 5.2.1 shows both the raw score and the z-score on their respective distributions. Notice that the red line indicating where each score lies is in the same relative spot for both. This is because transforming a raw score into a z-score does not change its relative location, it only makes it easier to know precisely where it is.

Z-scores are also useful for comparing scores from different distributions. Let's say we take the SAT and score 501 on both the math and critical reading sections. Does that mean we did equally well on both? Scores on the math portion are distributed normally with a mean of 511 and standard deviation of 120, so our z-score on the math section is

$$z_{math} = \frac{501 - 511}{120} = -0.08$$

which is just slightly below average (note that use of "math" as a subscript; subscripts are used when presenting multiple versions of the same statistic in order to know which one is which and have no bearing on the actual calculation). The critical reading section has a mean of 495 and standard deviation of 116, so

$$z_{CR} = \frac{501 - 495}{116} = 0.05$$

So even though we were almost exactly average on both tests, we did a little bit better on the critical reading portion relative to other people.

Finally, z-scores are incredibly useful if we need to combine information from different measures that are on different scales. Let's say we give a set of employees a series of tests on things like job knowledge, personality, and leadership. We may want to combine these into a single score we can use to rate employees for development or promotion, but look what happens when we take the average of raw scores from different scales, as shown in Table 5.2.1:

Table 5.2.1: Raw test scores on different scales (ranges in parentheses).

Raw Scores	Job Knowledge (0 – 100)	Personality (1 –5)	Leadership (1 – 5)	Average
Employee 1	98	4.2	1.1	34.43
Employee 2	96	3.1	4.5	34.53
Employee 3	97	2.9	3.6	34.50

Because the job knowledge scores were so big and the scores were so similar, they overpowered the other scores and removed almost all variability in the average. However, if we standardize these scores into z-scores, our averages retain more variability and it is easier to assess differences between employees, as shown in Table 5.2.2.

Table 5.2.2: Standardized scores.

z-Scores	Job Knowledge (0 – 100)	Personality (1 –5)	Leadership (1 – 5)	Average
Employee 1	1.00	1.14	-1.12	0.34
Employee 2	-1.00	-0.43	0.81	-0.20
Employee 3	0.00	-0.71	0.30	-0.14

Setting the scale of a distribution

Another convenient characteristic of z-scores is that they can be converted into any "scale" that we would like. Here, the term scale means how far apart the scores are (their spread) and where they are located (their central tendency). This can be very useful if we don't want to work with negative numbers or if we have a specific range we would like to present. The formulas for transforming z to x are:

$$x = z\sigma + \mu \tag{5.2.3}$$

for a population and

$$x = zs + M \tag{5.2.4}$$

for a sample. Notice that these are just simple rearrangements of the original formulas for calculating z from raw scores.

Let's say we create a new measure of intelligence, and initial calibration finds that our scores have a mean of 40 and standard deviation of 7. Three people who have scores of 52, 43, and 34 want to know how well they did on the measure. We can convert their raw scores into z-scores:

$$z = \frac{52 - 40}{7} = 1.71$$
$$z = \frac{43 - 40}{7} = 0.43$$
$$z = \frac{34 - 40}{7} = -0.80$$

A problem is that these new z-scores aren't exactly intuitive for many people. We can give people information about their relative location in the distribution (for instance, the first person scored well above average), or we can translate these z scores into the more familiar metric of IQ scores, which have a mean of 100 and standard deviation of 16:

$$\text{IQ} = 1.71 * 16 + 100 = 127.36$$
$$\text{IQ} = 0.43 * 16 + 100 = 106.88$$
$$\text{IQ} = -0.80 * 16 + 100 = 87.20$$

We would also likely round these values to 127, 107, and 87, respectively, for convenience.

Contributors

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 5.2: Z-scores is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

5.3: Z-scores and the Area under the Curve

Z-scores and the standard normal distribution go hand-in-hand. A *z*-score will tell you exactly where in the standard normal distribution a value is located, and any normal distribution can be converted into a standard normal distribution by converting all of the scores in the distribution into *z*-scores, a process known as standardization.

We saw in the previous chapter that standard deviations can be used to divide the normal distribution: 68% of the distribution falls within 1 standard deviation of the mean, 95% within (roughly) 2 standard deviations, and 99.7% within 3 standard deviations. Because *z*-scores are in units of standard deviations, this means that 68% of scores fall between z = -1.0 and z = 1.0 and so on. We call this 68% (or any percentage we have based on our *z*-scores) the proportion of the area under the curve. Any area under the curve is bounded by (defined by, delineated by, etc.) a single *z*-score or pair of *z*-scores.

An important property to point out here is that, by virtue of the fact that the total area under the curve of a distribution is always equal to 1.0 (see section on Normal Distributions at the beginning of this chapter), these areas under the curve can be added together or subtracted from 1 to find the proportion in other areas. For example, we know that the area between z = -1.0 and z = 1.0 (i.e. within one standard deviation of the mean) contains 68% of the area under the curve, which can be represented in decimal form at 0.6800 (to change a percentage to a decimal, simply move the decimal point 2 places to the left). Because the total area under the curve is equal to 1.0, that means that the proportion of the area outside z= -1.0 and z = 1.0 is equal to 1.0 – 0.6800 = 0.3200 or 32% (see Figure **5.3.1** below). This area is called the area in the tails of the distribution. Because this area is split between two tails and because the normal distribution is symmetrical, each tail has exactly one-half, or 16%, of the area under the curve.

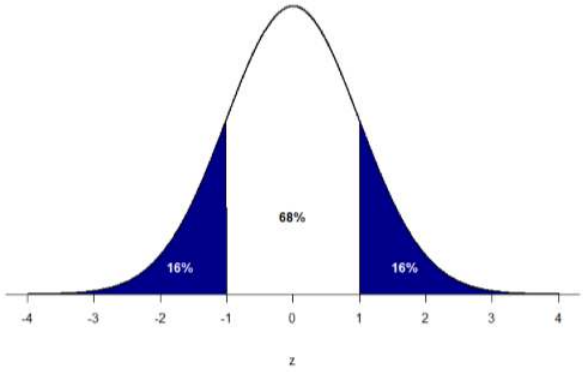


Figure **5.3.1**: Shaded areas represent the area under the curve in the tails

We will have much more to say about this concept in the coming chapters. As it turns out, this is a quite powerful idea that enables us to make statements about how likely an outcome is and what that means for research questions we would like to answer and hypotheses we would like to test. But first, we need to make a brief foray into some ideas about probability.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 5.3: Z-scores and the Area under the Curve is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

5.E: Z-scores and the Standard Normal Distribution (Exercises)

1. What are the two pieces of information contained in a z-score?

Answer:

The location above or below the mean (from the sign of the number) and the distance in standard deviations away from the mean (from the magnitude of the number).

2. A z-score takes a raw score and standardizes it into units of _____.
3. Assume the following 5 scores represent a sample: 2, 3, 5, 5, 6. Transform these scores into z-scores.

Answer:

$\overline{X}$= 4.2, s = 1.64; z = -1.34, -0.73, 0.49, 0.49, 1.10

4. True or false:
   a. All normal distributions are symmetrical
   b. All normal distributions have a mean of 1.0
   c. All normal distributions have a standard deviation of 1.0
   d. The total area under the curve of all normal distributions is equal to 1
5. Interpret the location, direction, and distance (near or far) of the following z-scores:
   a. -2.00
   b. 1.25
   c. 3.50
   d. -0.34

Answer:

   a. 2 standard deviations below the mean, far
   b. 1.25 standard deviations above the mean, near
   c. 3.5 standard deviations above the mean, far
   d. 0.34 standard deviations below the mean, near

6. Transform the following z-scores into a distribution with a mean of 10 and standard deviation of 2: -1.75, 2.20, 1.65, -0.95
7. Calculate z-scores for the following raw scores taken from a population with a mean of 100 and standard deviation of 16: 112, 109, 56, 88, 135, 99

Answer:

z = 0.75, 0.56, -2.75, -0.75, 2.19, -0.06

8. What does a z-score of 0.00 represent?
9. For a distribution with a standard deviation of 20, find z-scores that correspond to:
   a. One-half of a standard deviation below the mean
   b. 5 points above the mean
   c. Three standard deviations above the mean
   d. 22 points below the mean

Answer:

   a. -0.50
   b. 0.25
   c. 3.00
   d. 1.10

10. Calculate the raw score for the following z-scores from a distribution with a mean of 15 and standard deviation of 3:

10. Calculate the raw score for the following z-scores from a distribution with a mean of 15 and standard deviation of 3:

a. 4.0

b. 2.2

c. -1.3

d. 0.46

This page titled 5.E: Z-scores and the Standard Normal Distribution (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **4.E: Z-scores and the Standard Normal Distribution (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

6: Probability

NOTE: This is a second probability chapter. This one comes from this original text and gives a good overview of how probability will connect with z scores. There are no calculations in this chapter, unlike in the earlier probability chapter.

6.1: What is Probability

6.2: Probability in Graphs and Distributions

6.3: The Bigger Picture

6.E: Probability (Exercises)

This page titled 6: Probability is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

6.1: What is Probability

When we speak of the probability of something happening, we are talking how likely it is that "thing" will happen based on the conditions present. For instance, what is the probability that it will rain? That is, how likely do we think it is that it will rain today under the circumstances or conditions today? To define or understand the conditions that might affect how likely it is to rain, we might look out the window and say, "it's sunny outside, so it's not very likely that it will rain today." Stated using probability language: given that it is sunny outside, the probability of rain is low. "Given" is the word we use to state what the conditions are. As the conditions change, so does the probability. Thus, if it were cloudy and windy outside, we might say, "given the current weather conditions, there is a high probability that it is going to rain."

In these examples, we spoke about whether or not it is going to rain. Raining is an example of an event, which is the catch-all term we use to talk about any specific thing happening; it is a generic term that we specified to mean "rain" in exactly the same way that "conditions" is a generic term that we specified to mean "sunny" or "cloudy and windy."

It should also be noted that the terms "low" and "high" are relative and vague, and they will likely be interpreted different by different people (in other words: given how vague the terminology was, the probability of different interpretations is high). Most of the time we try to use more precise language or, even better, numbers to represent the probability of our event. Regardless, the basic structure and logic of our statements are consistent with how we speak about probability using numbers and formulas.

Let's look at a slightly deeper example. Say we have a regular, six-sided die (note that "die" is singular and "dice" is plural, a distinction that Dr. Foster has yet to get correct on his first try) and want to know how likely it is that we will roll a 1. That is, what is the probability of rolling a 1, given that the die is not weighted (which would introduce what we call a bias, though that is beyond the scope of this chapter). We could roll the die and see if it is a 1 or not, but that won't tell us about the probability, it will only tell us a single result. We could also roll the die hundreds or thousands of times, recording each outcome and seeing what the final list looks like, but this is time consuming, and rolling a die that many times may lead down a dark path to gambling or, worse, playing Dungeons & Dragons. What we need is a simple equation that represents what we are looking for and what is possible.

To calculate the probability of an event, which here is defined as rolling a 1 on an unbiased die, we need to know two things: how many outcomes satisfy the criteria of our event (stated different, how many outcomes would count as what we are looking for) and the total number of outcomes possible. In our example, only a single outcome, rolling a 1, will satisfy our criteria, and there are a total of six possible outcomes (rolling a 1, rolling a 2, rolling a 3, rolling a 4, rolling a 5, and rolling a 6). Thus, the probability of rolling a 1 on an unbiased die is 1 in 6 or 1/6. Put into an equation using generic terms, we get:

$$\text{Probability of an event } = \frac{\text{number of outcomes that satisfy our criteria}}{\text{total number of possible outcomes}} \tag{6.1.1}$$

We can also using P() as shorthand for probability and A as shorthand for an event:

$$P(A) = \frac{\text{number of outcomes that count a } A}{\text{total number of possible outcomes}} \tag{6.1.2}$$

Using this equation, let's now calculate the probability of rolling an even number on this die:

$$P(\text{Even Number}) = \frac{2, 4, \text{or } 6}{1, 2, 3, 4, 5, \text{or } 6} = \frac{3}{6} = \frac{1}{2}$$

So we have a 50% chance of rolling an even number of this die. The principles laid out here operate under a certain set of conditions and can be elaborated into ideas that are complex yet powerful and elegant. However, such extensions are not necessary for a basic understanding of statistics, so we will end our discussion on the math of probability here. Now, let's turn back to more familiar topics.

This page titled 6.1: What is Probability is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **5.1: What is Probability** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

6.2: Probability in Graphs and Distributions

We will see shortly that the normal distribution is the key to how probability works for our purposes. To understand exactly how, let's first look at a simple, intuitive example using pie charts.

Probability in Pie

Recall that a pie chart represents how frequently a category was observed and that all slices of the pie chart add up to 100%, or 1. This means that if we randomly select an observation from the data used to create the pie chart, the probability of it taking on a specific value is exactly equal to the size of that category's slice in the pie chart.

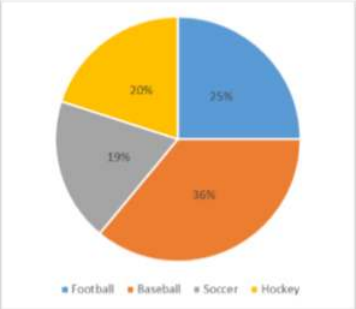


Figure **6.2.1**: Favorite sports

Take, for example, the pie chart in Figure **6.2.1** representing the favorite sports of 100 people. If you put this pie chart on a dart board and aimed blindly (assuming you are guaranteed to hit the board), the likelihood of hitting the slice for any given sport would be equal to the size of that slice. So, the probability of hitting the baseball slice is the highest at 36%. The probability is equal to the proportion of the chart taken up by that section.

We can also add slices together. For instance, maybe we want to know the probability to finding someone whose favorite sport is usually played on grass. The outcomes that satisfy this criteria are baseball, football, and soccer. To get the probability, we simply add their slices together to see what proportion of the area of the pie chart is in that region: 36% + 25% + 20% = 81%. We can also add sections together even if they do not touch. If we want to know the likelihood that someone's favorite sport is not called football somewhere in the world (i.e. baseball and hockey), we can add those slices even though they aren't adjacent or continuous in the chart itself: 36% + 20% = 56%. We are able to do all of this because 1) the size of the slice corresponds to the area of the chart taken up by that slice, 2) the percentage for a specific category can be represented as a decimal (this step was skipped for ease of explanation above), and 3) the total area of the chart is equal to 100% or 1.0, which makes the size of the slices interpretable.

Probability in Normal Distributions

If the language at the end of the last section sounded familiar, that's because its exactly the language used in the last chapter to describe the normal distribution. Recall that the normal distribution has an area under its curve that is equal to 1 and that it can be split into sections by drawing a line through it that corresponds to a given z-score. Because of this, we can interpret areas under the normal curve as probabilities that correspond to z-scores.

First, let's look back at the area between z = -1.00 and z = 1.00 presented in Figure **6.2.2**. We were told earlier that this region contains 68% of the area under the curve. Thus, if we randomly chose a z-score from all possible z-scores, there is a 68% chance that it will be between z = -1.00 and z = 1.00 because those are the z-scores that satisfy our criteria.

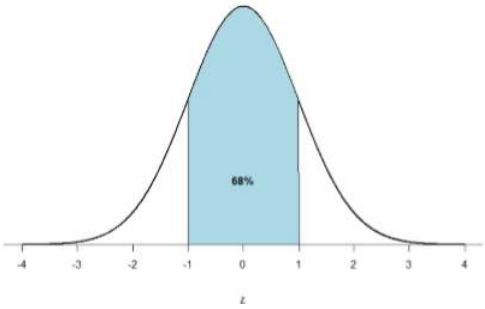


Figure **6.2.2**: There is a 68% chance of selection a z-score from the blue-shaded region

Just like a pie chart is broken up into slices by drawing lines through it, we can also draw a line through the normal distribution to split it into sections. Take a look at the normal distribution in Figure **6.2.3** which has a line drawn through it as $z = 1.25$. This line creates two sections of the distribution: the smaller section called the tail and the larger section called the body. Differentiating between the body and the tail does not depend on which side of the distribution the line is drawn. All that matters is the relative size of the pieces: bigger is always body.

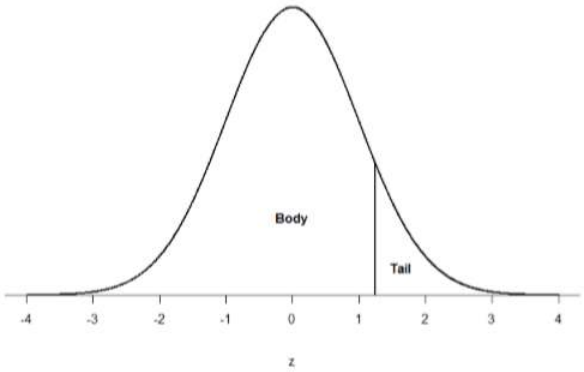
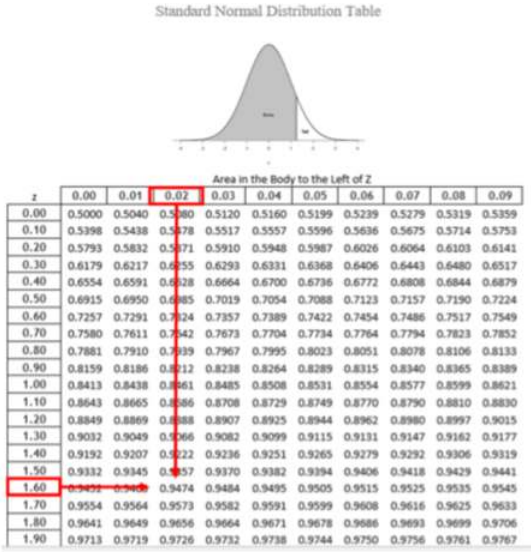


Figure **6.2.3**: Body and tail of the normal distribution

As you can see, we can break up the normal distribution into 3 pieces (lower tail, body, and upper tail) as in Figure **6.2.2** or into 2 pieces (body and tail) as in Figure **6.2.3**. We can then find the proportion of the area in the body and tail based on where the line was drawn (i.e. at what z-score). Mathematically this is done using calculus. Fortunately, the exact values are given to you in the Standard Normal Distribution Table, also known at the z-table. Using the values in this table, we can find the area under the normal curve in any body, tail, or combination of tails no matter which z-scores are used to define them.

The z-table presents the values for the area under the curve to the left of the positive z-scores from 0.00-3.00 (technically 3.09), as indicated by the shaded region of the distribution at the top of the table. To find the appropriate value, we first find the row corresponding to our z-score then follow it over until we get to the column that corresponds to the number in the hundredths place of our z-score. For example, suppose we want to find the area in the body for a z-score of 1.62. We would first find the row for 1.60 then follow it across to the column labeled 0.02 (1.60 + 0.02 = 1.62) and find 0.9474 (see Figure **6.2.4**). Thus, the odds of randomly selecting someone with a z-score less than (to the left of) $z = 1.62$ is 94.74% because that is the proportion of the area taken up by values that satisfy our criteria.



z	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0.00	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
0.10	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
0.20	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
0.30	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
0.40	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
0.50	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
0.60	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
0.70	0.7580	0.7611	0.7642	0.7673	0.7704	0.7734	0.7764	0.7794	0.7823	0.7852
0.80	0.7881	0.7910	0.7939	0.7967	0.7995	0.8023	0.8051	0.8078	0.8106	0.8133
0.90	0.8159	0.8186	0.8212	0.8238	0.8264	0.8289	0.8315	0.8340	0.8365	0.8389
1.00	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554	0.8577	0.8599	0.8621
1.10	0.8643	0.8665	0.8686	0.8708	0.8729	0.8749	0.8770	0.8790	0.8810	0.8830
1.20	0.8849	0.8869	0.8888	0.8907	0.8925	0.8944	0.8962	0.8980	0.8997	0.9015
1.30	0.9032	0.9049	0.9066	0.9082	0.9099	0.9115	0.9131	0.9147	0.9162	0.9177
1.40	0.9192	0.9207	0.9222	0.9236	0.9251	0.9265	0.9279	0.9292	0.9306	0.9319
1.50	0.9332	0.9345	0.9357	0.9370	0.9382	0.9394	0.9406	0.9418	0.9429	0.9441
1.60	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515	0.9525	0.9535	0.9545
1.70	0.9554	0.9564	0.9573	0.9582	0.9591	0.9599	0.9608	0.9616	0.9625	0.9633
1.80	0.9641	0.9649	0.9656	0.9664	0.9671	0.9678	0.9686	0.9693	0.9699	0.9706
1.90	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750	0.9756	0.9761	0.9767

Figure **6.2.4**: Using the z-table to find the area in the body to the left of z = 1.62

The z-table only presents the area in the body for positive z-scores because the normal distribution is symmetrical. Thus, the area in the body of z = 1.62 is equal to the area in the body for z = -1.62, though now the body will be the shaded area to the right of z (because the body is always larger). When in doubt, drawing out your distribution and shading the area you need to find will always help. The table also only presents the area in the body because the total area under the normal curve is always equal to 1.00, so if we need to find the area in the tail for z = 1.62, we simply find the area in the body and subtract it from 1.00 (1.00 – 0.9474 = 0.0526).

Let's look at another example. This time, let's find the area corresponding to z-scores more extreme than z = -1.96 and z = 1.96. That is, let's find the area in the tails of the distribution for values less than z = -1.96 (farther negative and therefore more extreme) and greater than z = 1.96 (farther positive and therefore more extreme). This region is illustrated in Figure **6.2.5**.

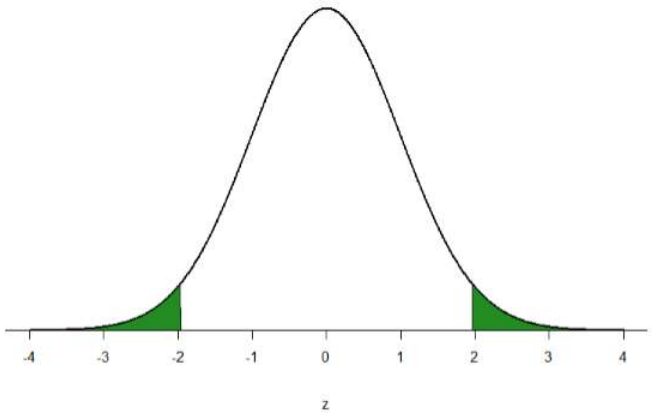


Figure **6.2.5**: Area in the tails beyond z = -1.96 and z = 1.96

Let's start with the tail for z = 1.96. If we go to the z-table we will find that the body to the left of z = 1.96 is equal to 0.9750. To find the area in the tail, we subtract that from 1.00 to get 0.0250. Because the normal distribution is symmetrical, the area in the tail for z = -1.96 is the exact same value, 0.0250. Finally, to get the total area in the shaded region, we simply add the areas together to get 0.0500. Thus, there is a 5% chance of randomly getting a value more extreme than z = -1.96 or z = 1.96 (this particular value and region will become incredibly important in Unit 2).

Finally, we can find the area between two z-scores by shading and subtracting. Figure **6.2.6** shows the area between z = 0.50 and z = 1.50. Because this is a subsection of a body (rather than just a body or a tail), we must first find the larger of the two bodies, in this case the body for z = 1.50, and subtract the smaller of the two bodies, or the body for z = 0.50. Aligning the distributions vertically,

as in Figure 6, makes this clearer. From the z-table, the area in the body for z = 1.50 is 0.9332 and the area in the body for z = 0.50 is 0.6915. Subtracting these gives us 0.9332 – 0.6915 = 0.2417.

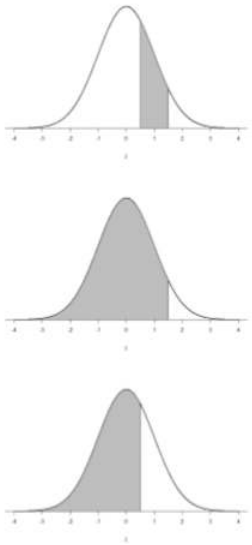


Figure 6.2.6: Area between z = 0.50 and 1.50, along with the corresponding areas in the body

Contributors

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 6.2: Probability in Graphs and Distributions is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **5.2: Probability in Graphs and Distributions** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

6.3: The Bigger Picture

The concepts and ideas presented in this chapter are likely not intuitive at first. Probability is a tough topic for everyone, but the tools it gives us are incredibly powerful and enable us to do amazing things with data analysis. They are the heart of how inferential statistics work.

To summarize, the probability that an event happens is the number of outcomes that qualify as that event (i.e. the number of ways the event could happen) compared to the total number of outcomes (i.e. how many things are possible). This extends to graphs like a pie chart, where the biggest slices take up more of the area and are therefore more likely to be chosen at random. This idea then brings us back around to our normal distribution, which can also be broken up into regions or areas, each of which are bounded by one or two z-scores and correspond to all z-scores in that region. The probability of randomly getting one of those z-scores in the specified region can then be found on the Standard Normal Distribution Table. Thus, the larger the region, the more likely an event is, and vice versa. Because the tails of the distribution are, by definition, smaller and we go farther out into the tail, the likelihood or probability of finding a result out in the extremes becomes small.

This page titled 6.3: The Bigger Picture is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **5.3: The Bigger Picture** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

6.E: Probability (Exercises)

1. In your own words, what is probability?

Answer:

Your answer should include information about an event happening under certain conditions given certain criteria. You could also discuss the relation between probability and the area under the curve or the proportion of the area in a chart.

2. There is a bag with 5 red blocks, 2 yellow blocks, and 4 blue blocks. If you reach in and grab one block without looking, what is the probability it is red?
3. Under a normal distribution, which of the following is more likely? (Note: this question can be answered without any calculations if you draw out the distributions and shade properly)
   - Getting a z-score greater than $z = 2.75$
   - Getting a z-score less than $z = -1.50$

Answer:

Getting a z-score less than $z = -1.50$ is more likely. $z = 2.75$ is farther out into the right tail than $z = -1.50$ is into the left tail, therefore there are fewer more extreme scores beyond 2.75 than -1.50, regardless of the direction

4. The heights of women in the United States are normally distributed with a mean of 63.7 inches and a standard deviation of 2.7 inches. If you randomly select a woman in the United States, what is the probability that she will be between 65 and 67 inches tall?
5. The heights of men in the United States are normally distributed with a mean of 69.1 inches and a standard deviation of 2.9 inches. What proportion of men are taller than 6 feet (72 inches)?

Answer:

15.87% or 0.1587

6. You know you need to score at least 82 points on the final exam to pass your class. After the final, you find out that the average score on the exam was 78 with a standard deviation of 7. How likely is it that you pass the class?
7. What proportion of the area under the normal curve is greater than $z = 1.65$?

Answer:

4.95% or 0.0495

8. Find the z-score that bounds 25% of the lower tail of the distribution.
9. Find the z-score that bounds the top 9% of the distribution.

Answer:

$z = 1.34$ (the top 9% means 9% of the area is in the upper tail and 91% is in the body to the left; finding the value in the normal table closest to .9100 is .9099, which corresponds to $z = 1.34$)

10. In a distribution with a mean of 70 and standard deviation of 12, what proportion of scores are lower than 55?

This page titled 6.E: Probability (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **5.E: Probability (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

7: Sampling Distributions

7.1: People, Samples, and Populations

7.2: The Sampling Distribution of Sample Means

7.3: Using Standard Error for Probability

7.4: Sampling Distribution, Probability and Inference

7.E: Sampling Distributions (Exercises)

This page titled 7: Sampling Distributions is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

7.1: People, Samples, and Populations

Most of what we have dealt with so far has concerned individual scores grouped into samples, with those samples being drawn from and, hopefully, representative of a population. We saw how we can understand the location of individual scores within a sample's distribution via z-scores, and how we can extend that to understand how likely it is to observe scores higher or lower than an individual score via probability.

Inherent in this work is the notion that an individual score will differ from the mean, which we quantify as a z-score. All of the individual scores will differ from the mean in different amounts and different directions, which is natural and expected. We quantify these differences as variance and standard deviation. Measures of spread and the idea of variability in observations is a key principle in inferential statistics. We know that any observation, whether it is a single score, a set of scores, or a particular descriptive statistic will differ from the center of whatever distribution it belongs in.

This is equally true of things outside of statistics and formal data collection and analysis. Some days you hear your alarm and wake up easily, other days you need to hit snooze a few [dozen] times. Some days traffic is light, other days it is very heavy. Some classes you are able to focus, pay attention, and take good notes, but other days you find yourself zoning out the entire time. Each individual observation is an insight but is not, by itself, the entire story, and it takes an extreme deviation from what we expect for us to think that something strange is going on. Being a little sleepy is normal, but being completely unable to get out of bed might indicate that we are sick. Light traffic is a good thing, but almost no cars on the road might make us think we forgot it is Saturday. Zoning out occasionally is fine, but if we cannot focus at all, we might be in a stats class rather than a fun one.

All of these principles carry forward from scores within samples to samples within populations. Just like an individual score will differ from its mean, an individual sample mean will differ from the true population mean. We encountered this principle in earlier chapters: sampling error. As mentioned way back in chapter 1, sampling error is an incredibly important principle. We know ahead of time that if we collect data and compute a sample, the observed value of that sample will be at least slightly off from what we expect it to be based on our supposed population mean; this is natural and expected. However, if our sample mean is extremely different from what we expect based on the population mean, there may be something going on.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 7.1: People, Samples, and Populations is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **6.1: People, Samples, and Populations** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

7.2: The Sampling Distribution of Sample Means

To see how we use sampling error, we will learn about a new, theoretical distribution known as the sampling distribution. In the same way that we can gather a lot of individual scores and put them together to form a distribution with a center and spread, if we were to take many samples, all of the same size, and calculate the mean of each of those, we could put those means together to form a distribution. This new distribution is, intuitively, known as the distribution of sample means. It is one example of what we call a sampling distribution, which can be formed from a set of any statistic, such as a mean, a test statistic, or a correlation coefficient (more on the latter two in Units 2 and 3). For our purposes, understanding the distribution of sample means will be enough to see how all other sampling distributions work to enable and inform our inferential analyses, so these two terms will be used interchangeably from here on out. Let's take a deeper look at some of its characteristics.

The sampling distribution of sample means can be described by its shape, center, and spread, just like any of the other distributions we have worked with. The shape of our sampling distribution is normal: a bell-shaped curve with a single peak and two tails extending symmetrically in either direction, just like what we saw in previous chapters. The center of the sampling distribution of sample means – which is, itself, the mean or average of the means – is the true population mean, μ. This will sometimes be written as μ_M to denote it as the mean of the sample means. The spread of the sampling distribution is called the standard error, the quantification of sampling error, denoted μ_M. The formula for standard error is:

$$\sigma_M = \frac{\sigma}{\sqrt{n}} \tag{7.2.1}$$

Notice that the sample size is in this equation. As stated above, the sampling distribution refers to samples of a specific size. That is, all sample means must be calculated from samples of the same size n, such n = 10, n = 30, or n = 100. This sample size refers to how many people or observations are in each individual sample, not how many samples are used to form the sampling distribution. This is because the sampling distribution is a theoretical distribution, not one we will ever actually calculate or observe. Figure **7.2.1** displays the principles stated here in graphical form.

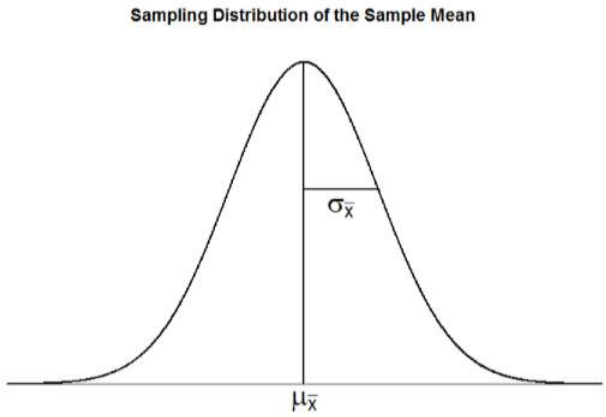


Figure **7.2.1**: The sampling distribution of sample means

Two Important Axioms

We just learned that the sampling distribution is theoretical: we never actually see it. If that is true, then how can we know it works? How can we use something that we don't see? The answer lies in two very important mathematical facts: the central limit theorem and the law of large numbers. We will not go into the math behind how these statements were derived, but knowing what they are and what they mean is important to understanding why inferential statistics work and how we can draw conclusions about a population based on information gained from a single sample.

Central Limit Theorem

The central limit theorem states:

Theorem 7.2.1

For samples of a single size n, drawn from a population with a given mean μ and variance σ^2, the sampling distribution of sample means will have a mean $\mu_M = \mu$ and variance $\sigma_M^2 = \frac{\sigma^2}{n}$. This distribution will approach normality as n increases.

From this, we are able to find the standard deviation of our sampling distribution, the standard error. As you can see, just like any other standard deviation, the standard error is simply the square root of the variance of the distribution.

The last sentence of the central limit theorem states that the sampling distribution will be normal as the sample size of the samples used to create it increases. What this means is that bigger samples will create a more normal distribution, so we are better able to use the techniques we developed for normal distributions and probabilities. So how large is large enough? In general, a sampling distribution will be normal if either of two characteristics is true:

1. the population from which the samples are drawn is normally distributed or
2. the sample size is equal to or greater than 30.

This second criteria is very important because it enables us to use methods developed for normal distributions even if the true population distribution is skewed.

Law of Large Numbers

The law of large numbers simply states that as our sample size increases, the probability that our sample mean is an accurate representation of the true population mean also increases. It is the formal mathematical way to state that larger samples are more accurate.

The law of large numbers is related to the central limit theorem, specifically the formulas for variance and standard error. Notice that the sample size appears in the denominators of those formulas. A larger denominator in any fraction means that the overall value of the fraction gets smaller (i.e 1/2 = 0.50, 1/3 – 0.33, 1/4 = 0.25, and so on). Thus, larger sample sizes will create smaller standard errors. We already know that standard error is the spread of the sampling distribution and that a smaller spread creates a narrower distribution. Therefore, larger sample sizes create narrower sampling distributions, which increases the probability that a sample mean will be close to the center and decreases the probability that it will be in the tails. This is illustrated in Figures **7.2.2** and **7.2.3**.

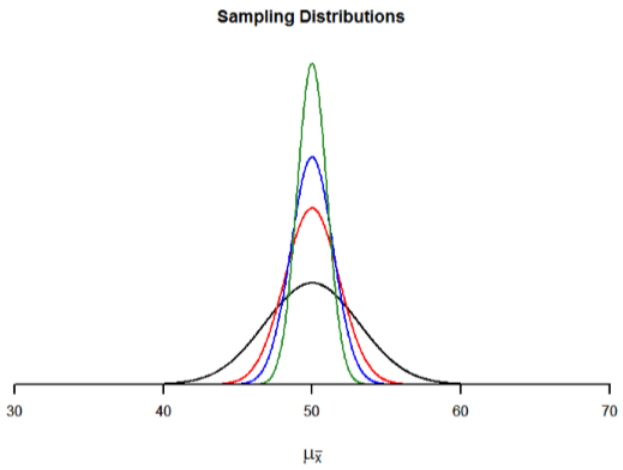


Figure **7.2.2**: Sampling distributions from the same population with μ = 50 and σ = 10 but different sample sizes (N = 10, N = 30, N = 50, N = 100)

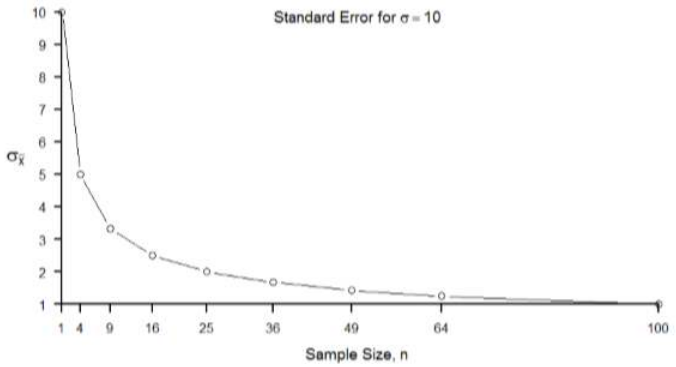


Figure 7.2.3: Relation between sample size and standard error for a constant $\sigma = 10$

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 7.2: The Sampling Distribution of Sample Means is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **6.2: The Sampling Distribution of Sample Means** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

7.3: Using Standard Error for Probability

We saw in chapter 6 that we can use z-scores to split up a normal distribution and calculate the proportion of the area under the curve in one of the new regions, giving us the probability of randomly selecting a z-score in that range. We can follow the exact sample process for sample means, converting them into z-scores and calculating probabilities. The only difference is that instead of dividing a raw score by the standard deviation, we divide the sample mean by the standard error.

$$z = \frac{M - \mu}{\sigma_M} = \frac{M - \mu}{\frac{\sigma}{\sqrt{n}}} \tag{7.3.1}$$

Let's say we are drawing samples from a population with a mean of 50 and standard deviation of 10 (the same values used in Figure 2). What is the probability that we get a random sample of size 10 with a mean greater than or equal to 55? That is, for n = 10, what is the probability that $M \geq 55$? First, we need to convert this sample mean score into a z-score:

$$z = \frac{55 - 50}{\frac{10}{\sqrt{10}}} = \frac{5}{3.16} = 1.58$$

Now we need to shade the area under the normal curve corresponding to scores greater than z = 1.58 as in Figure 7.3.1:

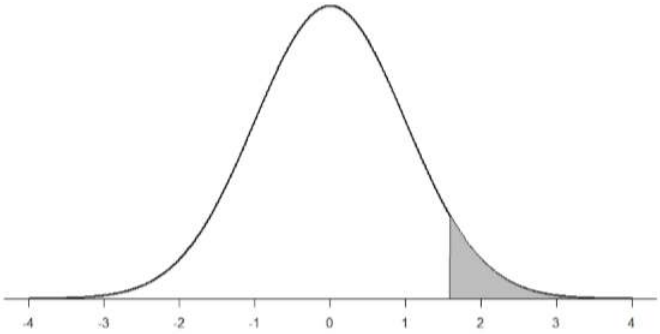


Figure 7.3.1: Area under the curve greater than z = 1.58

Now we go to our z-table and find that the area to the left of z = 1.58 is 0.9429. Finally, because we need the area to the right (per our shaded diagram), we simply subtract this from 1 to get 1.00 – 0.9429 = 0.0571. So, the probability of randomly drawing a sample of 10 people from a population with a mean of 50 and standard deviation of 10 whose sample mean is 55 or more is p = .0571, or 5.71%. Notice that we are talking about means that are 55 or more. That is because, strictly speaking, it's impossible to calculate the probability of a score taking on exactly 1 value since the "shaded region" would just be a line with no area to calculate.

Now let's do the same thing, but assume that instead of only having a sample of 10 people we took a sample of 50 people. First, we find z:

$$z = \frac{55 - 50}{\frac{10}{\sqrt{50}}} = \frac{5}{1.41} = 3.55 \tag{7.3.2}$$

Then we shade the appropriate region of the normal distribution:

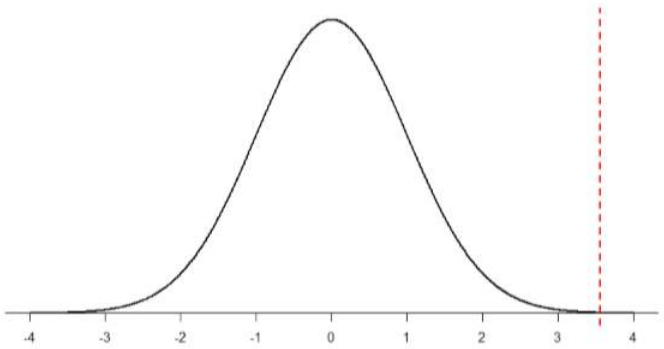


Figure **7.3.2**: Area under the curve greater than $z = 3.55$

Notice that no region of Figure **7.3.2** appears to be shaded. That is because the area under the curve that far out into the tail is so small that it can't even be seen (the red line has been added to show exactly where the region starts). Thus, we already know that the probability must be smaller for $N = 50$ than $N = 10$ because the size of the area (the proportion) is much smaller.

We run into a similar issue when we try to find $z = 3.55$ on our Standard Normal Distribution Table. The table only goes up to 3.09 because everything beyond that is almost 0 and changes so little that it's not worth printing values. The closest we can get is subtracting the largest value, 0.9990, from 1 to get 0.001. We know that, technically, the actual probability is smaller than this (since 3.55 is farther into the tail than 3.09), so we say that the probability is $p < 0.001$, or less than 0.1%.

This example shows what an impact sample size can have. From the same population, looking for exactly the same thing, changing only the sample size took us from roughly a 5% chance (or about 1/20 odds) to a less than 0.1% chance (or less than 1 in 1000). As the sample size n increased, the standard error decreased, which in turn caused the value of z to increase, which finally caused the p -value (a term for probability we will use a lot in Unit 2) to decrease. You can think of this relation like gears: turning the first gear (sample size) clockwise causes the next gear (standard error) to turn counterclockwise, which causes the third gear (z) to turn clockwise, which finally causes the last gear (probability) to turn counterclockwise. All of these pieces fit together, and the relations will always be the same:

$$\mathrm{n} \uparrow \sigma_M \downarrow \mathrm{z} \uparrow \mathrm{p} \downarrow \tag{7.3.3}$$

Let's look at this one more way. For the same population of sample size 50 and standard deviation 10, what proportion of sample means fall between 47 and 53 if they are of sample size 10 and sample size 50?

We'll start again with $n = 10$. Converting 47 and 53 into z-scores, we get $z = -0.95$ and $z = 0.95$, respectively. From our z-table, we find that the proportion between these two scores is 0.6578 (the process here is left off for the student to practice converting M to z and z to proportions). So, 65.78% of sample means of sample size 10 will fall between 47 and 53. For $n = 50$, our z-scores for 47 and 53 are ±2.13, which gives us a proportion of the area as 0.9668, almost 97%! Shaded regions for each of these sampling distributions is displayed in Figure **7.3.3**. The sampling distributions are shown on the original scale, rather than as z-scores, so you can see the effect of the shading and how much of the body falls into the range, which is marked off with dotted line.

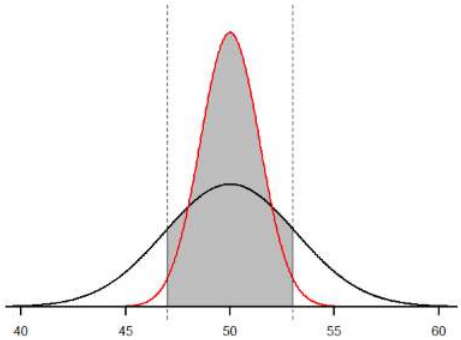


Figure **7.3.3**: Areas between 47 and 53 for sampling distributions of $n = 10$ and $n = 50$

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 7.3: Using Standard Error for Probability is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **6.3: Using Standard Error for Probability** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

7.4: Sampling Distribution, Probability and Inference

We've seen how we can use the standard error to determine probability based on our normal curve. We can think of the standard error as how much we would naturally expect our statistic (be it a mean or some other statistic) to vary. In our formula for z based on a sample mean, the numerator ($M - \mu$) is what we call an observed effect. That is, it is what we observe in our sample mean versus what we expected based on the population from which that sample mean was calculated. Because the sample mean will naturally move around due to sampling error, our observed effect will also change naturally. In the context of our formula for z, then, our standard error is how much we would naturally expect the observed effect to change. Changing by a little is completely normal, but changing by a lot might indicate something is going on. This is the basis of inferential statistics and the logic behind hypothesis testing, the subject of Unit 2.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 7.4: Sampling Distribution, Probability and Inference is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **6.4: Sampling Distribution, Probability and Inference** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

7.E: Sampling Distributions (Exercises)

1. What is a sampling distribution?

Answer:

The sampling distribution (or sampling distribution of the sample means) is the distribution formed by combining many sample means taken from the same population and of a single, consistent sample size.

2. What are the two mathematical facts that describe how sampling distributions work?
3. What is the difference between a sampling distribution and a regular distribution?

Answer:

A sampling distribution is made of statistics (e.g. the mean) whereas a regular distribution is made of individual scores.

4. What effect does sample size have on the shape of a sampling distribution?
5. What is standard error?

Answer:

Standard error is the spread of the sampling distribution and is the quantification of sampling error. It is how much we expect the sample mean to naturally change based on random chance.

6. For a population with a mean of 75 and a standard deviation of 12, what proportion of sample means of size n = 16 fall above 82?
7. For a population with a mean of 100 and standard deviation of 16, what is the probability that a random sample of size 4 will have a mean between 110 and 130?

Answer:

10.46% or 0.1046

8. Find the z-score for the following means taken from a population with mean 10 and standard deviation 2:
   a. M= 8, n = 12
   b. M= 8, n = 30
   c. M = 20, n = 4
   d. M = 20, n = 16
9. As the sample size increases, what happens to the p-value associated with a given sample mean?

Answer:

As sample size increases, the p-value will decrease

10. For a population with a mean of 35 and standard deviation of 7, find the sample mean of size n = 20 that cuts off the top 5% of the sampling distribution.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 7.E: Sampling Distributions (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **6.E: Sampling Distributions (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

8: Introduction to Hypothesis Testing

8.1: Logic and Purpose of Hypothesis Testing
8.2: The Probability Value
8.3: The Null Hypothesis
8.4: The Alternative Hypothesis
8.5: Critical values, p-values, and significance level
8.6: Steps of the Hypothesis Testing Process
8.7: Movie Popcorn
8.8: Effect Size
8.9: Office Temperature
8.10: Different Significance Level
8.11: Other Considerations in Hypothesis Testing
8.E: Introduction to Hypothesis Testing (Exercises)

This page titled 8: Introduction to Hypothesis Testing is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

8.1: Logic and Purpose of Hypothesis Testing

The statistician R. Fisher explained the concept of hypothesis testing with a story of a lady tasting tea. Here we will present an example based on James Bond who insisted that martinis should be shaken rather than stirred. Let's consider a hypothetical experiment to determine whether Mr. Bond can tell the difference between a shaken and a stirred martini. Suppose we gave Mr. Bond a series of 16 taste tests. In each test, we flipped a fair coin to determine whether to stir or shake the martini. Then we presented the martini to Mr. Bond and asked him to decide whether it was shaken or stirred. Let's say Mr. Bond was correct on 13 of the 16 taste tests. Does this prove that Mr. Bond has at least some ability to tell whether the martini was shaken or stirred?

This result does not prove that he does; it could be he was just lucky and guessed right 13 out of 16 times. But how plausible is the explanation that he was just lucky? To assess its plausibility, we determine the probability that someone who was just guessing would be correct 13/16 times or more. This probability can be computed to be 0.0106. This is a pretty low probability, and therefore someone would have to be very lucky to be correct 13 or more times out of 16 if they were just guessing. So either Mr. Bond was very lucky, or he can tell whether the drink was shaken or stirred. The hypothesis that he was guessing is not proven false, but considerable doubt is cast on it. Therefore, there is strong evidence that Mr. Bond can tell whether a drink was shaken or stirred.

Let's consider another example. The case study Physicians' Reactions sought to determine whether physicians spend less time with obese patients. Physicians were sampled randomly and each was shown a chart of a patient complaining of a migraine headache. They were then asked to estimate how long they would spend with the patient. The charts were identical except that for half the charts, the patient was obese and for the other half, the patient was of average weight. The chart a particular physician viewed was determined randomly. Thirty-three physicians viewed charts of average-weight patients and 38 physicians viewed charts of obese patients.

The mean time physicians reported that they would spend with obese patients was 24.7 minutes as compared to a mean of 31.4 minutes for normal-weight patients. How might this difference between means have occurred? One possibility is that physicians were influenced by the weight of the patients. On the other hand, perhaps by chance, the physicians who viewed charts of the obese patients tend to see patients for less time than the other physicians. Random assignment of charts does not ensure that the groups will be equal in all respects other than the chart they viewed. In fact, it is certain the groups differed in many ways by chance. The two groups could not have exactly the same mean age (if measured precisely enough such as in days). Perhaps a physician's age affects how long physicians see patients. There are innumerable differences between the groups that could affect how long they view patients. With this in mind, is it plausible that these chance differences are responsible for the difference in times?

To assess the plausibility of the hypothesis that the difference in mean times is due to chance, we compute the probability of getting a difference as large or larger than the observed difference (31.4 - 24.7 = 6.7 minutes) if the difference were, in fact, due solely to chance. Using methods presented in later chapters, this probability can be computed to be 0.0057. Since this is such a low probability, we have confidence that the difference in times is due to the patient's weight and is not due to chance.

This page titled 8.1: Logic and Purpose of Hypothesis Testing is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.1: Logic and Purpose of Hypothesis Testing** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.2: The Probability Value

It is very important to understand precisely what the probability values mean. In the James Bond example, the computed probability of 0.0106 is the probability he would be correct on 13 or more taste tests (out of 16) if he were just guessing.

It is easy to mistake this probability of 0.0106 as the probability he cannot tell the difference. This is not at all what it means.

The probability of 0.0106 is the probability of a certain outcome (13 or more out of 16) assuming a certain state of the world (James Bond was only guessing). It is not the probability that a state of the world is true. Although this might seem like a distinction without a difference, consider the following example. An animal trainer claims that a trained bird can determine whether or not numbers are evenly divisible by 7. In an experiment assessing this claim, the bird is given a series of 16 test trials. On each trial, a number is displayed on a screen and the bird pecks at one of two keys to indicate its choice. The numbers are chosen in such a way that the probability of any number being evenly divisible by 7 is 0.50. The bird is correct on 9/16 choices. We can compute that the probability of being correct nine or more times out of 16 if one is only guessing is 0.40. Since a bird who is only guessing would do this well 40% of the time, these data do not provide convincing evidence that the bird can tell the difference between the two types of numbers. As a scientist, you would be very skeptical that the bird had this ability. Would you conclude that there is a 0.40 probability that the bird can tell the difference? Certainly not! You would think the probability is much lower than 0.0001.

To reiterate, the probability value is the probability of an outcome (9/16 or better) and not the probability of a particular state of the world (the bird was only guessing). In statistics, it is conventional to refer to possible states of the world as hypotheses since they are hypothesized states of the world. Using this terminology, the probability value is the probability of an outcome given the hypothesis. It is not the probability of the hypothesis given the outcome.

This is not to say that we ignore the probability of the hypothesis. If the probability of the outcome given the hypothesis is sufficiently low, we have evidence that the hypothesis is false. However, we do not compute the probability that the hypothesis is false. In the James Bond example, the hypothesis is that he cannot tell the difference between shaken and stirred martinis. The probability value is low (0.0106), thus providing evidence that he can tell the difference. However, we have not computed the probability that he can tell the difference.

This page titled 8.2: The Probability Value is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.2: The Probability Value** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.3: The Null Hypothesis

The hypothesis that an apparent effect is due to chance is called the null hypothesis, written H_0 ("H-naught"). In the Physicians' Reactions example, the null hypothesis is that in the population of physicians, the mean time expected to be spent with obese patients is equal to the mean time expected to be spent with average-weight patients. This null hypothesis can be written as:

$$H_0 : \mu_{obese} - \mu_{average} = 0 \tag{8.3.1}$$

In words, the null hypothesis should be written in the format: There is no difference between the time spent with obese patients and the time spent with average-weight patients.

The null hypothesis in a correlational study of the relationship between high school grades and college grades would typically be that the population correlation is 0. This can be written as

$$H_0 : \rho = 0 \tag{8.3.2}$$

where ρ is the population correlation, which we will cover in chapter 12.

Although the null hypothesis is usually that the value of a parameter is 0, there are occasions in which the null hypothesis is a value other than 0. For example, if we are working with mothers in the U.S. whose children are at risk of low birth weight, we can use 7.47 pounds, the average birthweight in the US, as our null value and test for differences against that.

For now, we will focus on testing a value of a single mean against what we expect from the population. Using birthweight as an example, our null hypothesis takes the form:

$$H_0 : \mu = 7.47$$

The number on the right hand side is our null hypothesis value that is informed by our research question. Notice that we are testing the value for μ, the population parameter, NOT the sample statistic M. This is for two reasons: 1) once we collect data, we know what the value of M is – it's not a mystery or a question, it is observed and used for the second reason, which is 2) we are interested in understanding the population, not just our sample.

Keep in mind that the null hypothesis is typically the opposite of the researcher's hypothesis. In the Physicians' Reactions study, the researchers hypothesized that physicians would expect to spend less time with obese patients. The null hypothesis that the two types of patients are treated identically is put forward with the hope that it can be discredited and therefore rejected. If the null hypothesis were true, a difference as large or larger than the sample difference of 6.7 minutes would be very unlikely to occur. Therefore, the researchers rejected the null hypothesis of no difference and concluded that in the population, physicians intend to spend less time with obese patients.

In general, the null hypothesis is the idea that nothing is going on: there is no effect of our treatment, no relation between our variables, and no difference in our sample mean from what we expected about the population mean. This is always our baseline starting assumption, and it is what we seek to reject. If we are trying to treat depression, we want to find a difference in average symptoms between our treatment and control groups. If we are trying to predict job performance, we want to find a relation between conscientiousness and evaluation scores. However, until we have evidence against it, we must use the null hypothesis as our starting point.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.3: The Null Hypothesis is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

8.4: The Alternative Hypothesis

If the null hypothesis is rejected, then we will need some other explanation, which we call the alternative hypothesis, H_A or H_1. The alternative hypothesis is simply the reverse of the null hypothesis, and there are three options, depending on where we expect the difference to lie. Thus, our alternative hypothesis is the mathematical way of stating our research question. If we expect our obtained sample mean to be above or below the null hypothesis value, which we call a directional hypothesis, then our alternative hypothesis takes the form:

$$\mathbf{H_A : \mu > 7.47} \quad \text{or} \quad \mathbf{H_A : \mu < 7.47}$$

based on the research question itself. We should only use a directional hypothesis if we have good reason, based on prior observations or research, to suspect a particular direction. When we do not know the direction, such as when we are entering a new area of research, we use a non-directional alternative:

$$\mathbf{H_A : \mu \neq 7.47}$$

We will set different criteria for rejecting the null hypothesis based on the directionality (greater than, less than, or not equal to) of the alternative. To understand why, we need to see where our criteria come from and how they relate to z-scores and distributions.

Writing hypotheses in words

As we alluded to in the null hypothesis section, we can write our hypotheses in word statements (in addition to the statements with symbols). These statements should be specific enough to the particular experiment or situation being referred to. That is, don't make them generic enough so that they would apply to any hypothesis test that you would conduct.

Examples for how to write null and alternate hypotheses in words for directional and non-directional situations are given throughout the chapters.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.4: The Alternative Hypothesis is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.4: The Alternative Hypothesis** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.5: Critical values, p-values, and significance level

A low probability value casts doubt on the null hypothesis. How low must the probability value be in order to conclude that the null hypothesis is false? Although there is clearly no right or wrong answer to this question, it is conventional to conclude the null hypothesis is false if the probability value is less than 0.05. More conservative researchers conclude the null hypothesis is false only if the probability value is less than 0.01. When a researcher concludes that the null hypothesis is false, the researcher is said to have rejected the null hypothesis. The probability value below which the null hypothesis is rejected is called the α level or simply α ("alpha"). It is also called the significance level. If α is not explicitly specified, assume that $\alpha = 0.05$.

The significance level is a threshold we set before collecting data in order to determine whether or not we should reject the null hypothesis. We set this value beforehand to avoid biasing ourselves by viewing our results and then determining what criteria we should use. If our data produce values that meet or exceed this threshold, then we have sufficient evidence to reject the null hypothesis; if not, we fail to reject the null (we never "accept" the null).

There are two criteria we use to assess whether our data meet the thresholds established by our chosen significance level, and they both have to do with our discussions of probability and distributions. Recall that probability refers to the likelihood of an event, given some situation or set of conditions. In hypothesis testing, that situation is the assumption that the null hypothesis value is the correct value, or that there is no effect. The value laid out in H_0 is our condition under which we interpret our results. To reject this assumption, and thereby reject the null hypothesis, we need results that would be very unlikely if the null was true. Now recall that values of z which fall in the tails of the standard normal distribution represent unlikely values. That is, the proportion of the area under the curve as or more extreme than z is very small as we get into the tails of the distribution. Our significance level corresponds to the area under the tail that is exactly equal to α: if we use our normal criterion of $\alpha = .05$, then 5% of the area under the curve becomes what we call the rejection region (also called the critical region) of the distribution. This is illustrated in Figure 8.5.1.

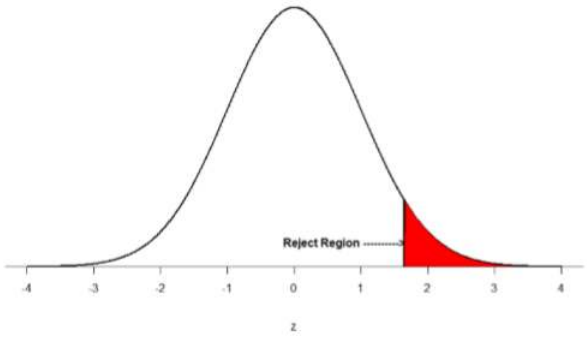


Figure 8.5.1: The rejection region for a one-tailed test.

The shaded rejection region takes us 5% of the area under the curve. Any result which falls in that region is sufficient evidence to reject the null hypothesis.

The rejection region is bounded by a specific z-value, as is any area under the curve. In hypothesis testing, the value corresponding to a specific rejection region is called the critical value, z_{crit} ("z-crit") or $z*$ (hence the other name "critical region"). Finding the critical value works exactly the same as finding the z-score corresponding to any area under the curve like we did in Unit 1. If we go to the normal table, we will find that the z-score corresponding to 5% of the area under the curve is equal to 1.645 ($z = 1.64$ corresponds to 0.0505 and $z = 1.65$ corresponds to 0.0495, so .05 is exactly in between them) if we go to the right and -1.645 if we go to the left. The direction must be determined by your alternative hypothesis, and drawing then shading the distribution is helpful for keeping directionality straight.

Suppose, however, that we want to do a non-directional test. We need to put the critical region in both tails, but we don't want to increase the overall size of the rejection region (for reasons we will see later). To do this, we simply split it in half so that an equal proportion of the area under the curve falls in each tail's rejection region. For $\alpha = .05$, this means 2.5% of the area is in each tail, which, based on the z-table, corresponds to critical values of $z* = \pm 1.96$. This is shown in Figure 8.5.2.

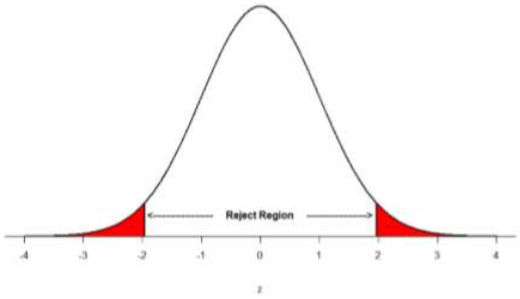


Figure 8.5.2: Two-tailed rejection region.

Thus, any z-score falling outside ±1.96 (greater than 1.96 in absolute value) falls in the rejection region. When we use z-scores in this way, the obtained value of z (sometimes called z-obtained) is something known as a test statistic, which is simply an inferential statistic used to test a null hypothesis. The formula for our z-statistic has not changed:

$$z = \frac{M - \mu}{\sigma/\sqrt{n}} \tag{8.5.1}$$

To formally test our hypothesis, we compare our obtained z-statistic to our critical z-value. If $Z_{obt} > Z_{crit}$, that means it falls in the rejection region (to see why, draw a line for z = 2.5 on Figure 8.5.1 or Figure 8.5.2) and so we reject H_0. If $Z_{obt} < Z_{crit}$, we fail to reject. Remember that as z gets larger, the corresponding area under the curve beyond z gets smaller. Thus, the proportion, or p-value, will be smaller than the area for α, and if the area is smaller, the probability gets smaller. Specifically, the probability of obtaining that result, or a more extreme result, under the condition that the null hypothesis is true gets smaller.

The z-statistic is very useful when we are doing our calculations by hand. However, when we use computer software, it will report to us a p-value, which is simply the proportion of the area under the curve in the tails beyond our obtained z-statistic. We can directly compare this p-value to α to test our null hypothesis: if $p < \alpha$, we reject H_0, but if $p > \alpha$, we fail to reject. Note also that the reverse is always true: if we use critical values to test our hypothesis, we will always know if p is greater than or less than α. If we reject, we know that $p < \alpha$ because the obtained z-statistic falls farther out into the tail than the critical z-value that corresponds to α, so the proportion (p-value) for that z-statistic will be smaller. Conversely, if we fail to reject, we know that the proportion will be larger than α because the z-statistic will not be as far into the tail. This is illustrated for a one-tailed test in Figure 8.5.3.

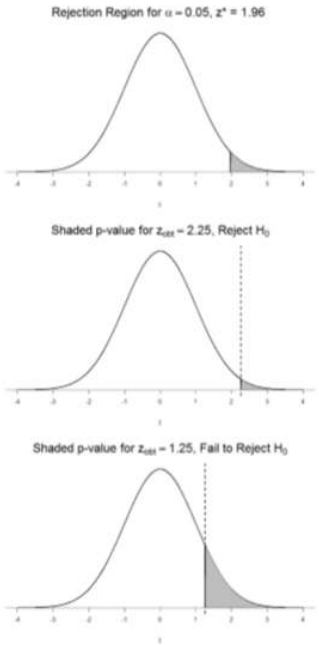


Figure 8.5.3: Relation between α, z_{obt}, and p

When the null hypothesis is rejected, the effect is said to be statistically significant. For example, in the Physicians Reactions case study, the probability value is 0.0057. Therefore, the effect of obesity is statistically significant and the null hypothesis that obesity

makes no difference is rejected. It is very important to keep in mind that statistical significance means only that the null hypothesis of exactly no effect is rejected; it does not mean that the effect is important, which is what “significant” usually means. When an effect is significant, you can have confidence the effect is not exactly zero. Finding that an effect is significant does not tell you about how large or important the effect is. Do not confuse statistical significance with practical significance. A small effect can be highly significant if the sample size is large enough. Why does the word “significant” in the phrase “statistically significant” mean something so different from other uses of the word? Interestingly, this is because the meaning of “significant” in everyday language has changed. It turns out that when the procedures for hypothesis testing were developed, something was “significant” if it signified something. Thus, finding that an effect is statistically significant signifies that the effect is real and not due to chance. Over the years, the meaning of “significant” changed, leading to the potential misinterpretation.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.5: Critical values, p-values, and significance level is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri’s Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.5: Critical values, p-values, and significance level** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.6: Steps of the Hypothesis Testing Process

The process of testing hypotheses follows a simple four-step procedure. This process will be what we use for the remainder of the textbook and course, and though the hypothesis and statistics we use will change, this process will not.

Step 1: State the Hypotheses

Your hypotheses are the first thing you need to lay out. Otherwise, there is nothing to test! You have to state the null hypothesis (which is what we test) and the alternative hypothesis (which is what we expect). These should be stated mathematically as they were presented above AND in words, explaining in normal English what each one means in terms of the research question.

Step 2: Find the Critical Values

Next, we formally lay out the criteria we will use to test our hypotheses. There are two pieces of information that inform our critical values: α, which determines how much of the area under the curve composes our rejection region, and the directionality of the test, which determines where the region will be.

Step 3: Compute the Test Statistic

Once we have our hypotheses and the standards we use to test them, we can collect data and calculate our test statistic, in this case z. This step is where the vast majority of differences in future chapters will arise: different tests used for different data are calculated in different ways, but the way we use and interpret them remains the same.

Step 4: Make the Decision

Finally, once we have our obtained test statistic, we can compare it to our critical value and decide whether we should reject or fail to reject the null hypothesis. When we do this, we must interpret the decision in relation to our research question, stating what we concluded, what we based our conclusion on, and the specific statistics we obtained.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.6: Steps of the Hypothesis Testing Process is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.6: Steps of the Hypothesis Testing Process** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.7: Movie Popcorn

Let's see how hypothesis testing works in action by working through an example. Say that a movie theater owner likes to keep a very close eye on how much popcorn goes into each bag sold, so he knows that the average bag has 8 cups of popcorn and that this varies a little bit, about half a cup. That is, the known population mean is μ = 8.00 and the known population standard deviation is σ = 0.50. The owner wants to make sure that the newest employee is filling bags correctly, so over the course of a week he randomly assesses 25 bags filled by the employee to test for a difference (N = 25). He doesn't want bags overfilled or under filled, so he looks for differences in both directions. This scenario has all of the information we need to begin our hypothesis testing procedure.

Step 1: State the Hypotheses

Our manager is looking for a difference in the mean weight of popcorn bags compared to the population mean of 8. We will need both a null and an alternative hypothesis written both mathematically and in words. We'll always start with the null hypothesis:

H_0: There is no difference in the weight of popcorn bags from this employee compared to normal.

H_0: $\mu = 8.00$

Notice that we phrase the hypothesis in terms of the population parameter μ, which in this case would be the true average weight of bags filled by the new employee. Our assumption of no difference, the null hypothesis, is that this mean is exactly the same as the known population mean value we want it to match, 8.00. Now let's do the alternative:

H_A: There is a difference in the weight of popcorn bags from this employee compared to normal.

H_A: $\mu \neq 8.00$

In this case, we don't know if the bags will be too full or not full enough, so we do a two-tailed alternative hypothesis that there is a difference.

Step 2: Find the Critical Values

Our critical values are based on two things: the directionality of the test and the level of significance. We decided in step 1 that a two-tailed test is the appropriate directionality. We were given no information about the level of significance, so we assume that α = 0.05 is what we will use. As stated earlier in the chapter, the critical values for a two-tailed z-test at α = 0.05 are $z*$ = ±1.96. This will be the criteria we use to test our hypothesis. We can now draw out our distribution so we can visualize the rejection region and make sure it makes sense.

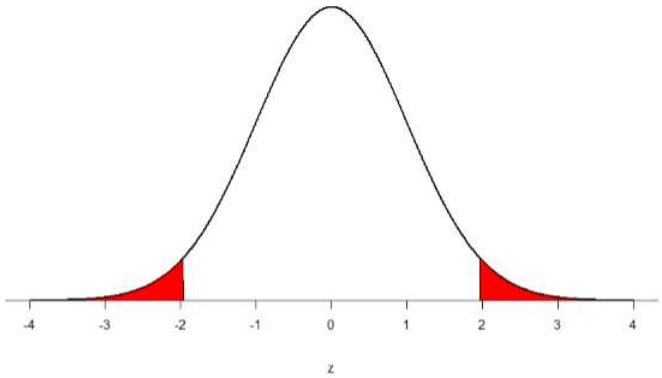


Figure 8.7.1: Rejection region for $z*$ = ±1.96

Step 3: Calculate the Test Statistic

Now we come to our formal calculations. Let's say that the manager collects data and finds that the average weight of this employee's popcorn bags is M= 7.75 cups. We can now plug this value, along with the values presented in the original problem, into our equation for z:

$$z = \frac{7.75 - 8.00}{0.50/\sqrt{25}} = \frac{-0.25}{0.10} = -2.50$$

So our test statistic is z = -2.50, which we can draw onto our rejection region distribution:

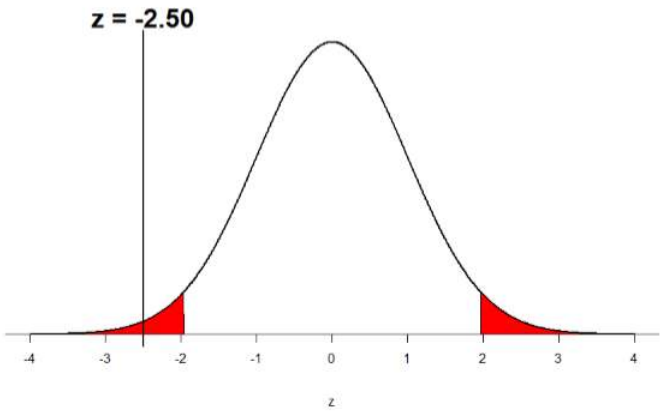


Figure **8.7.2**: Test statistic location

Step 4: Make the Decision

Looking at Figure **8.7.2**, we can see that our obtained z-statistic falls in the rejection region. We can also directly compare it to our critical value: in terms of absolute value, -2.50 > -1.96, so we reject the null hypothesis. We can now write our conclusion:

> Reject H_0. Based on the sample of 25 bags, we can conclude that the average popcorn bag from this employee is different (M = 7.75 cups) than the average weight of popcorn bags at this movie theater, z = 2.50, p < 0.05.

When we write our conclusion, we write out the words to communicate what it actually means, but we also include the average sample size we calculated (the exact location doesn't matter, just somewhere that flows naturally and makes sense) and the z-statistic and p-value. We don't know the exact p-value, but we do know that because we rejected the null, it must be less than α.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.7: Movie Popcorn is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.7: Movie Popcorn** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.8: Effect Size

When we reject the null hypothesis, we are stating that the difference we found was statistically significant, but we have mentioned several times that this tells us nothing about practical significance. To get an idea of the actual size of what we found, we can compute a new statistic called an effect size. Effect sizes give us an idea of how large, important, or meaningful a statistically significant effect is. For mean differences like we calculated here, our effect size is Cohen's ***d***:

$$d = \frac{M - \mu}{\sigma} \tag{8.8.1}$$

This is very similar to our formula for ***z***, but we no longer take into account the sample size (since overly large samples can make it too easy to reject the null). Cohen's ***d*** is interpreted in units of standard deviations, just like ***z***. For our example:

$$d = \frac{7.75 - 8.00}{0.50} = \frac{-0.25}{0.50} = 0.50$$

Cohen's ***d*** is interpreted as small, moderate, or large. Specifically, ***d*** = 0.20 is small, ***d*** = 0.50 is moderate, and ***d*** = 0.80 is large. Obviously values can fall in between these guidelines, so we should use our best judgment and the context of the problem to make our final interpretation of size. Our effect size happened to be exactly equal to one of these, so we say that there was a moderate effect.

Effect sizes are incredibly useful and provide important information and clarification that overcomes some of the weakness of hypothesis testing. Whenever you find a significant result, you should always calculate an effect size.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.8: Effect Size is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.8: Effect Size** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.9: Office Temperature

Let's do another example to solidify our understanding. Let's say that the office building you work in is supposed to be kept at 74 degree Fahrenheit but is allowed to vary by 1 degree in either direction. You suspect that, as a cost saving measure, the temperature was secretly set higher. You set up a formal way to test your hypothesis.

Step 1: State the Hypotheses

You start by laying out the null hypothesis:

H_0: The average building temperature is not higher than claimed

$H_0 : \mu \leq 74$

Next you state the alternative hypothesis. You have reason to suspect a specific direction of change, so you make a one-tailed test:

H_A: The average building temperature is higher than claimed

$H_A : \mu > 74$

Step 2: Find the Critical Values

You know that the most common level of significance is α = 0.05, so you keep that the same and know that the critical value for a one-tailed z-test is $z*$ = 1.645. To keep track of the directionality of the test and rejection region, you draw out your distribution:

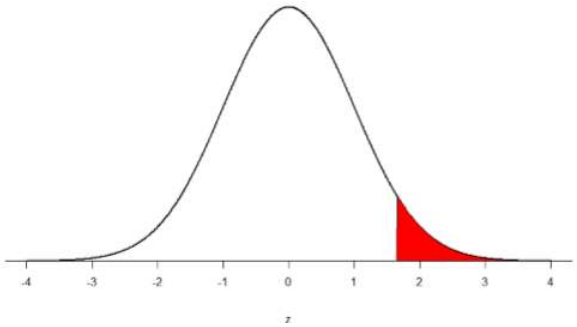


Figure 8.9.1: Rejection region

Step 3: Calculate the Test Statistic

Now that you have everything set up, you spend one week collecting temperature data:

Table 8.9.1: Temperature data for a week

Day	Temperature
Monday	77
Tuesday	76
Wednesday	74
Thursday	78
Friday	78

You calculate the average of these scores to be M= 76.6 degrees. You use this to calculate the test statistic, using μ = 74 (the supposed average temperature), σ = 1.00 (how much the temperature should vary), and n = 5 (how many data points you collected):

$$z = \frac{76.60 - 74.00}{1.00/\sqrt{5}} = \frac{2.60}{0.45} = 5.78$$

This value falls so far into the tail that it cannot even be plotted on the distribution!

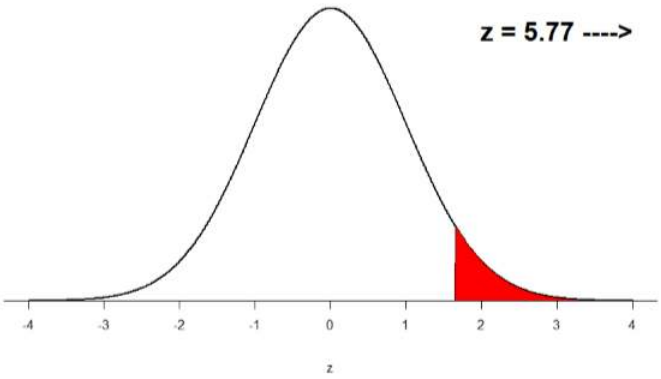


Figure 8.9.2: Obtained z-statistic

Step 4: Make the Decision

You compare your obtained z-statistic, z = 5.77, to the critical value, $z*$ = 1.645, and find that $z > z*$. Therefore you reject the null hypothesis, concluding:

Based on 5 observations, the average temperature (M= 76.6 degrees) is statistically significantly higher than it is supposed to be, z = 5.77, p < .05.

Because the result is significant, you also calculate an effect size:

$$d = \frac{76.60 - 74.00}{1.00} = \frac{2.60}{1.00} = 2.60$$

The effect size you calculate is definitely large, meaning someone has some explaining to do!

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.9: Office Temperature is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.9: Office Temperature** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.10: Different Significance Level

Finally, let's take a look at an example phrased in generic terms, rather than in the context of a specific research question, to see the individual pieces one more time. This time, however, we will use a stricter significance level, α = 0.01, to test the hypothesis.

Step 1: State the Hypotheses

We will use 60 as an arbitrary null hypothesis value:

H_0: The average score does not differ from the population

$H_0 : \mu = 60$

We will assume a two-tailed test:

H_A: The average score does differ

$H_A : \mu \neq 60$

Step 2: Find the Critical Values

We have seen the critical values for z-tests at α = 0.05 levels of significance several times. To find the values for α = 0.01, we will go to the standard normal table and find the z-score cutting of 0.005 (0.01 divided by 2 for a two-tailed test) of the area in the tail, which is $z*$ = ±2.575. Notice that this cutoff is much higher than it was for α = 0.05. This is because we need much less of the area in the tail, so we need to go very far out to find the cutoff. As a result, this will require a much larger effect or much larger sample size in order to reject the null hypothesis.

Step 3: Calculate the Test Statistic

We can now calculate our test statistic. We will use σ = 10 as our known population standard deviation and the following data to calculate our sample mean:

61	62
65	61
58	59
54	61
60	63

The average of these scores is M= 60.40. From this we calculate our z-statistic as:

$$z = \frac{60.40 - 60.00}{10.00/\sqrt{10}} = \frac{0.40}{3.16} = 0.13$$

Step 4: Make the Decision

Our obtained z-statistic, z = 0.13, is very small. It is much less than our critical value of 2.575. Thus, this time, we fail to reject the null hypothesis. Our conclusion would look something like:

Based on the sample of 10 scores, we cannot conclude that there is no effect causing the mean (M= 60.40) to be statistically significantly different from 60.00, z = 0.13, p > 0.01.

Notice two things about the end of the conclusion. First, we wrote that p is greater than instead of p is less than, like we did in the previous two examples. This is because we failed to reject the null hypothesis. We don't know exactly what the p-value is, but we know it must be larger than the α level we used to test our hypothesis. Second, we used 0.01 instead of the usual 0.05, because this time we tested at a different level. The number you compare to the p-value should always be the significance level you test at.

Finally, because we did not detect a statistically significant effect, we do not need to calculate an effect size.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 8.10: Different Significance Level is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.10: Different Significance Level** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.11: Other Considerations in Hypothesis Testing

There are several other considerations we need to keep in mind when performing hypothesis testing.

Errors in Hypothesis Testing

In the Physicians' Reactions case study, the probability value associated with the significance test is 0.0057. Therefore, the null hypothesis was rejected, and it was concluded that physicians intend to spend less time with obese patients. Despite the low probability value, it is possible that the null hypothesis of no true difference between obese and average-weight patients is true and that the large difference between sample means occurred by chance. If this is the case, then the conclusion that physicians intend to spend less time with obese patients is in error. This type of error is called a Type I error. More generally, a Type I error occurs when a significance test results in the rejection of a true null hypothesis.

By one common convention, if the probability value is below 0.05 then the null hypothesis is rejected. Another convention, although slightly less common, is to reject the null hypothesis if the probability value is below 0.01. The threshold for rejecting the null hypothesis is called the α level or simply α. It is also called the significance level. As discussed in the introduction to hypothesis testing, it is better to interpret the probability value as an indication of the weight of evidence against the null hypothesis than as part of a decision rule for making a reject or do-not-reject decision. Therefore, keep in mind that rejecting the null hypothesis is not an all-or-nothing decision.

The Type I error rate is affected by the α level: the lower the α level the lower the Type I error rate. It might seem that α is the probability of a Type I error. However, this is not correct. Instead, α is the probability of a Type I error given that the null hypothesis is true. If the null hypothesis is false, then it is impossible to make a Type I error.

The second type of error that can be made in significance testing is failing to reject a false null hypothesis. This kind of error is called a Type II error. Unlike a Type I error, a Type II error is not really an error. When a statistical test is not significant, it means that the data do not provide strong evidence that the null hypothesis is false. Lack of significance does not support the conclusion that the null hypothesis is true. Therefore, a researcher should not make the mistake of incorrectly concluding that the null hypothesis is true when a statistical test was not significant. Instead, the researcher should consider the test inconclusive. Contrast this with a Type I error in which the researcher erroneously concludes that the null hypothesis is false when, in fact, it is true.

A Type II error can only occur if the null hypothesis is false. If the null hypothesis is false, then the probability of a Type II error is called β (beta). The probability of correctly rejecting a false null hypothesis equals 1- β and is called power. Power is simply our ability to correctly detect an effect that exists. It is influenced by the size of the effect (larger effects are easier to detect), the significance level we set (making it easier to reject the null makes it easier to detect an effect, but increases the likelihood of a Type I Error), and the sample size used (larger samples make it easier to reject the null).

Misconceptions in Hypothesis Testing

Misconceptions about significance testing are common. This section lists three important ones.

1. Misconception: The probability value is the probability that the null hypothesis is false.
    - Proper interpretation: The probability value is the probability of a result as extreme or more extreme given that the null hypothesis is true. It is the probability of the data given the null hypothesis. It is not the probability that the null hypothesis is false.
2. Misconception: A low probability value indicates a large effect.
    - Proper interpretation: A low probability value indicates that the sample outcome (or one more extreme) would be very unlikely if the null hypothesis were true. A low probability value can occur with small effect sizes, particularly if the sample size is large.
3. Misconception: A non-significant outcome means that the null hypothesis is probably true.
    - Proper interpretation: A non-significant outcome means that the data do not conclusively demonstrate that the null hypothesis is false.

This page titled 8.11: Other Considerations in Hypothesis Testing is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.11: Other Considerations in Hypothesis Testing** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

8.E: Introduction to Hypothesis Testing (Exercises)

1. In your own words, explain what the null hypothesis is.

Answer:

Your answer should include mention of the baseline assumption of no difference between the sample and the population.

2. What are Type I and Type II Errors?
3. What is α?

Answer:

Alpha is the significance level. It is the criteria we use when decided to reject or fail to reject the null hypothesis, corresponding to a given proportion of the area under the normal distribution and a probability of finding extreme scores assuming the null hypothesis is true.

4. Why do we phrase null and alternative hypotheses with population parameters and not sample means?
5. If our null hypothesis is "$H_0 : \mu = 40$", what are the three possible alternative hypotheses?

Answer:

$H_A : \mu \neq 40$, $H_A : \mu > 40$, $H_A : \mu < 40$

6. Why do we state our hypotheses and decision criteria before we collect our data?
7. When and why do you calculate an effect size?

Answer:

We calculate an effect size when we find a statistically significant result to see if our result is practically meaningful or important

8. Determine whether you would reject or fail to reject the null hypothesis in the following situations:
   a. z= 1.99, two-tailed test at α = 0.05
   b. z = 0.34, $z*$ = 1.645
   c. p = 0.03, α = 0.05
   d. p = 0.015, α = 0.01
9. You are part of a trivia team and have tracked your team's performance since you started playing, so you know that your scores are normally distributed with μ = 78 and σ = 12. Recently, a new person joined the team, and you think the scores have gotten better. Use hypothesis testing to see if the average score has improved based on the following 8 weeks' worth of score data: 82, 74, 62, 68, 79, 94, 90, 81, 80.

Answer:

Step 1: $H_0 : \mu = 78$ "The average score is not different after the new person joined", $H_A : \mu > 78$ "The average score has gone up since the new person joined."

Step 2: One-tailed test to the right, assuming α = 0.05, $z*$ = 1.645.

Step 3: $\overline{X}$= 88.75, $\sigma_{\overline{X}}$ = 4.24, z = 2.54.

Step 4: $z > z*$, Reject H_0. Based on 8 weeks of games, we can conclude that our average score ($\overline{X}$ = 88.75) is higher now that the new person is on the team, z = 2.54, p < .05. Since the result is significant, we need an effect size: Cohen's d = 0.90, which is a large effect.

10. You get hired as a server at a local restaurant, and the manager tells you that servers' tips are $42 on average but vary about $12 ($\mu$ = 42, σ = 12). You decide to track your tips to see if you make a different amount, but because this is your first job as a

server, you don’t know if you will make more or less in tips. After working 16 shifts, you find that your average nightly amount is $44.50 from tips. Test for a difference between this value and the population mean at the $\alpha = 0.05$ level of significance.

This page titled 8.E: Introduction to Hypothesis Testing (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri’s Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **7.E: Introduction to Hypothesis Testing (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

9: Introduction to t-tests

9.1: The t-statistic

9.2: Hypothesis Testing with t

9.3: Confidence Intervals

9.E: Introduction to t-tests (Exercises)

This page titled 9: Introduction to t-tests is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

9.1: The t-statistic

Last chapter, we were introduced to hypothesis testing using the z-statistic for sample means that we learned in Unit 1. This was a useful way to link the material and ease us into the new way to looking at data, but it isn't a very common test because it relies on knowing the populations standard deviation, σ, which is rarely going to be the case. Instead, we will estimate that parameter σ using the sample statistic s in the same way that we estimate μ using M (μ will still appear in our formulas because we suspect something about its value and that is what we are testing). Our new statistic is called t, and for testing one population mean using a single sample (called a 1-sample t-test) it takes the form:

$$t = \frac{M - \mu}{s_M} = \frac{M - \mu}{s/\sqrt{n}} \tag{9.1.1}$$

Notice that t looks almost identical to z; this is because they test the exact same thing: the value of a sample mean compared to what we expect of the population. The only difference is that the standard error is now denoted s_M to indicate that we use the sample statistic for standard deviation, s, instead of the population parameter σ. The process of using and interpreting the standard error and the full test statistic remain exactly the same.

In chapter 3 we learned that the formulae for sample standard deviation and population standard deviation differ by one key factor: the denominator for the parameter is N but the denominator for the statistic is N–1, also known as degrees of freedom, df. Because we are using a new measure of spread, we can no longer use the standard normal distribution and the z-table to find our critical values. For t-tests, we will use the t-distribution and t-table to find these values.

The t-distribution, like the standard normal distribution, is symmetric and normally distributed with a mean of 0 and standard error (as the measure of standard deviation for sampling distributions) of 1. However, because the calculation of standard error uses degrees of freedom, there will be a different t-distribution for every degree of freedom. Luckily, they all work exactly the same, so in practice this difference is minor.

Figure **9.1.1** shows four curves: a normal distribution curve labeled z, and three tdistribution curves for 2, 10, and 30 degrees of freedom. Two things should stand out: First, for lower degrees of freedom (e.g. 2), the tails of the distribution are much fatter, meaning the a larger proportion of the area under the curve falls in the tail. This means that we will have to go farther out into the tail to cut off the portion corresponding to 5% or α = 0.05, which will in turn lead to higher critical values. Second, as the degrees of freedom increase, we get closer and closer to the z curve. Even the distribution with df = 30, corresponding to a sample size of just 31 people, is nearly indistinguishable from z. In fact, a t-distribution with infinite degrees of freedom (theoretically, of course) is exactly the standard normal distribution. Because of this, the bottom row of the t-table also includes the critical values for z-tests at the specific significance levels. Even though these curves are very close, it is still important to use the correct table and critical values, because small differences can add up quickly.

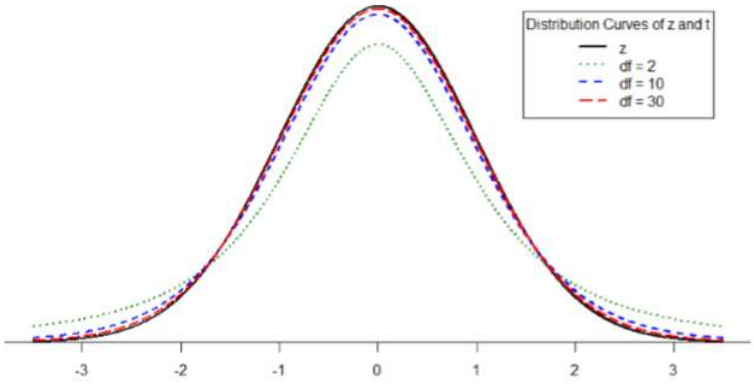


Figure **9.1.1**: Distributions comparing effects of degrees of freedom

The t-distribution table lists critical values for one- and two-tailed tests at several levels of significance arranged into columns. The rows of the t-table list degrees of freedom up to df = 100 in order to use the appropriate distribution curve. It does not, however, list all possible degrees of freedom in this range, because that would take too many rows. Above df = 40, the rows jump in

increments of 10. If a problem requires you to find critical values and the exact degrees of freedom is not listed, you always round down to the next smallest number. For example, if you have 48 people in your sample, the degrees of freedom are $N - 1 = 48 - 1 = 47$; however, 47 doesn't appear on our table, so we round down and use the critical values for $df = 40$, even though 50 is closer. We do this because it avoids inflating Type I Error (false positives, see chapter 7) by using criteria that are too lax.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 9.1: The t-statistic is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **8.1: The t-statistic** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

9.2: Hypothesis Testing with t

Hypothesis testing with the t-statistic works exactly the same way as z-tests did, following the four-step process of

1. Stating the Hypothesis
2. Finding the Critical Values
3. Computing the Test Statistic
4. Making the Decision.

We will work though an example: let's say that you move to a new city and find an auto shop to change your oil. Your old mechanic did the job in about 30 minutes (though you never paid close enough attention to know how much that varied), and you suspect that your new shop takes much longer. After 4 oil changes, you think you have enough evidence to demonstrate this.

Step 1 State the Hypotheses:

Our hypotheses for 1-sample t-tests are identical to those we used for z-tests. We still state the null and alternative hypotheses mathematically in terms of the population parameter and written out in readable English. For our example:

H_0: This shop does not take longer to change oil than your old mechanic

$$H_0 : \mu \leq 30$$

H_A: This shop takes longer to change oil than your old mechanic

$$H_A : \mu > 30$$

Step 2 Find the Critical Values:

As noted above, our critical values still delineate the area in the tails under the curve corresponding to our chosen level of significance. Because we have no reason to change significance levels, we will use α = 0.05, and because we suspect a direction of effect, we have a one-tailed test. To find our critical values for t, we need to add one more piece of information: the degrees of freedom. For this example:

$$df = N–1 = 4–1 = 3$$

Going to our t-table, we find the column corresponding to our one-tailed significance level and find where it intersects with the row for 3 degrees of freedom. As shown in Figure **9.2.1**: our critical value is $t*$ = 2.353

t-distribution Table

df	0.05 0.10	0.025 0.05	0.01 0.02	0.005 0.01	1-tailed α 2-tailed α
1	6.314	12.706	31.821	63.657	
2	2.920	4.303	6.965	9.925	
3	2.353	3.182	4.541	5.841	
4	2.132	2.776	3.747	4.604	
5	2.015	2.571	3.365	4.032	
6	1.943	2.447	3.143	3.707	

Figure **9.2.1**: t-table

We can then shade this region on our t-distribution to visualize our rejection region

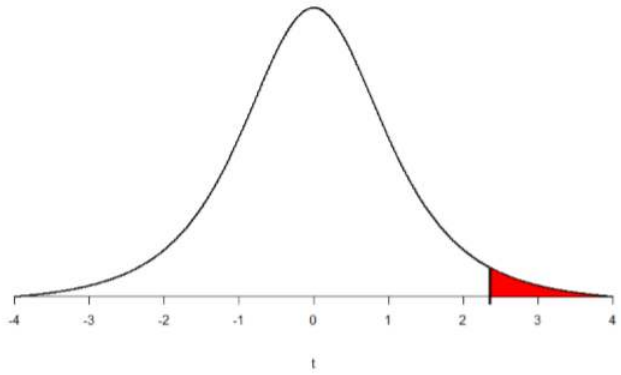


Figure 9.2.2: Rejection Region

Step 3 Compute the Test Statistic:

The four wait times you experienced for your oil changes at the new shop were 46 minutes, 58 minutes, 40 minutes, and 71 minutes. We will use these to calculate M and s by first filling in the sum of squares table in Table 9.2.1:

Table 9.2.1: Sum of Squares Table

X	$X - M$	$(X - M)^2$
46	-7.75	60.06
58	4.25	18.06
40	-13.75	189.06
71	17.25	297.56
Σ=215	Σ=0	Σ=564.74

After filling in the first row to get Σ=215, we find that the mean is M = 53.75 (215 divided by sample size 4), which allows us to fill in the rest of the table to get our sum of squares SS = 564.74, which we then plug in to the formula for standard deviation from chapter 3:

$$s = \sqrt{\frac{\sum (X - M)^2}{N - 1}} = \sqrt{\frac{SS}{df}} = \sqrt{\frac{564.74}{3}} = 13.72$$

Next, we take this value and plug it in to the formula for standard error:

$$s_M = \frac{s}{\sqrt{n}} = \frac{13.72}{2} = 6.86$$

And, finally, we put the standard error, sample mean, and null hypothesis value into the formula for our test statistic t:

$$t = \frac{M - \mu}{s_M} = \frac{53.75 - 30}{6.86} = \frac{23.75}{6.68} = 3.46$$

This may seem like a lot of steps, but it is really just taking our raw data to calculate one value at a time and carrying that value forward into the next equation: data, sample size/degrees of freedom, mean, sum of squares, standard deviation, standard error, and finally the test statistic. At each step, we simply match the symbols of what we just calculated to where they appear in the next formula to make sure we are plugging everything in correctly.

Step 4 Make the Decision:

Now that we have our critical value and test statistic, we can make our decision using the same criteria we used for a z-test. Our obtained t-statistic was t = 3.46 and our critical value was $t* = 2.353 : t > t*$, so we reject the null hypothesis and conclude:

Based on our four oil changes, the new mechanic takes longer on average (M = 53.75) to change oil than our old mechanic, $t(3)$ = 3.46, $p < .05$.

Notice that we also include the degrees of freedom in parentheses next to t. And because we found a significant result, we need to calculate an effect size, which is still Cohen's d, but now we use s in place of σ:

$$d = \frac{M - \mu}{s} = \frac{53.75 - 30.00}{13.72} = 1.73$$

This is a large effect. It should also be noted that for some things, like the minutes in our current example, we can also interpret the magnitude of the difference we observed (23 minutes and 45 seconds) as an indicator of importance since time is a familiar metric.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 9.2: Hypothesis Testing with t is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **8.2: Hypothesis Testing with t** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

9.3: Confidence Intervals

Up to this point, we have learned how to estimate the population parameter for the mean using sample data and a sample statistic. From one point of view, this makes sense: we have one value for our parameter so we use a single value (called a point estimate) to estimate it. However, we have seen that all statistics have sampling error and that the value we find for the sample mean will bounce around based on the people in our sample, simply due to random chance. Thinking about estimation from this perspective, it would make more sense to take that error into account rather than relying just on our point estimate. To do this, we calculate what is known as a confidence interval.

A confidence interval starts with our point estimate then creates a range of scores considered plausible based on our standard deviation, our sample size, and the level of confidence with which we would like to estimate the parameter. This range, which extends equally in both directions away from the point estimate, is called the margin of error. We calculate the margin of error by multiplying our two-tailed critical value by our standard error:

$$\text{Margin of Error } = t^*(s/\sqrt{n}) \tag{9.3.1}$$

One important consideration when calculating the margin of error is that it can only be calculated using the critical value for a two-tailed test. This is because the margin of error moves away from the point estimate in both directions, so a one-tailed value does not make sense.

The critical value we use will be based on a chosen level of confidence, which is equal to 1 – α. Thus, a 95% level of confidence corresponds to α = 0.05. Thus, at the 0.05 level of significance, we create a 95% Confidence Interval. How to interpret that is discussed further on.

Once we have our margin of error calculated, we add it to our point estimate for the mean to get an upper bound to the confidence interval and subtract it from the point estimate for the mean to get a lower bound for the confidence interval:

$$\begin{aligned} \text{Upper Bound} &= M + \text{Margin of Error} \\ \text{Lower Bound } &= M - \text{Margin of Error} \end{aligned} \tag{9.3.2}$$

Or simply:

$$\text{Confidence Interval } = M \pm t^*(s/\sqrt{n}) \tag{9.3.3}$$

To write out a confidence interval, we always use soft brackets and put the lower bound, a comma, and the upper bound:

$$\text{Confidence Interval } = \text{ (Lower Bound, Upper Bound)} \tag{9.3.4}$$

Let's see what this looks like with some actual numbers by taking our oil change data and using it to create a 95% confidence interval estimating the average length of time it takes at the new mechanic. We already found that our average was M= 53.75 and our standard error was s_M = 6.86. We also found a critical value to test our hypothesis, but remember that we were testing a one-tailed hypothesis, so that critical value won't work. To see why that is, look at the column headers on the t-table. The column for one-tailed α = 0.05 is the same as a two-tailed α = 0.10. If we used the old critical value, we'd actually be creating a 90% confidence interval (1.00-0.10 = 0.90, or 90%). To find the correct value, we use the column for two-tailed α = 0.05 and, again, the row for 3 degrees of freedom, to find $t*$ = 3.182.

Now we have all the pieces we need to construct our confidence interval:

$$95\%CI = 53.75 \pm 3.182(6.86)$$

$$\begin{aligned} \text{Upper Bound} &= 53.75 + 3.182(6.86) \\ UB &= 53.75 + 21.83 \\ UB &= 75.58 \end{aligned}$$

$$\begin{aligned} \text{Lower Bound} &= 53.75 - 3.182(6.86) \\ LB &= 53.75 - 21.83 \\ LB &= 31.92 \end{aligned}$$

$$95\%CI = (31.92, 75.58)$$

So we find that our 95% confidence interval runs from 31.92 minutes to 75.58 minutes, but what does that actually mean? The range (31.92, 75.58) represents values of the mean that we consider reasonable or plausible based on our observed data. It includes our point estimate of the mean, M= 53.75, in the center, but it also has a range of values that could also have been the case based on what we know about how much these scores vary (i.e. our standard error).

It is very tempting to also interpret this interval by saying that we are 95% confident that the true population mean falls within the range (31.92, 75.58), but this is not true. The reason it is not true is that phrasing our interpretation this way suggests that we have firmly established an interval and the population mean does or does not fall into it, suggesting that our interval is firm and the population mean will move around. However, the population mean is an absolute that does not change; it is our interval that will vary from data collection to data collection, even taking into account our standard error. The correct interpretation, then, is that we are 95% confident that the range (31.92, 75.58) brackets the true population mean. This is a very subtle difference, but it is an important one.

Hypothesis Testing with Confidence Intervals

As a function of how they are constructed, we can also use confidence intervals to test hypotheses. However, we are limited to testing two-tailed hypotheses only, because of how the intervals work, as discussed above.

Once a confidence interval has been constructed, using it to test a hypothesis is simple. The range of the confidence interval brackets (or contains, or is around) the null hypothesis value, we fail to reject the null hypothesis. If it does not bracket the null hypothesis value (i.e. if the entire range is above the null hypothesis value or below it), we reject the null hypothesis. The reason for this is clear if we think about what a confidence interval represents. Remember: a confidence interval is a range of values that we consider reasonable or plausible based on our data. Thus, if the null hypothesis value is in that range, then it is a value that is plausible based on our observations. If the null hypothesis is plausible, then we have no reason to reject it. Thus, if our confidence interval brackets the null hypothesis value, thereby making it a reasonable or plausible value based on our observed data, then we have no evidence against the null hypothesis and fail to reject it. However, if we build a confidence interval of reasonable values based on our observations and it does not contain the null hypothesis value, then we have no empirical (observed) reason to believe the null hypothesis value and therefore reject the null hypothesis.

Let's see an example. You hear that the national average on a measure of friendliness is 38 points. You want to know if people in your community are more or less friendly than people nationwide, so you collect data from 30 random people in town to look for a difference. We'll follow the same four step hypothesis testing procedure as before.

Step 1: State the Hypotheses:

We will start by laying out our null and alternative hypotheses:

H_0: There is no difference in how friendly the local community is compared to the national average

$$H_0 : \mu = 38$$

H_A: There is a difference in how friendly the local community is compared to the national average

$$H_A : \mu \neq 38$$

Step 2: Find the Critical Values:

We need our critical values in order to determine the width of our margin of error. We will assume a significance level of α = 0.05 (which will give us a 95% CI). From the t-table, a two-tailed critical value at α = 0.05 with 29 degrees of freedom (N – 1 = 30 – 1 = 29) is $t*$ = 2.045.

Step 3: Calculations:

Now we can construct our confidence interval. After we collect our data, we find that the average person in our community scored 39.85, or M= 39.85, and our standard deviation was s = 5.61. First, we need to use this standard deviation, plus our sample size of N = 30, to calculate our standard error:

$$s_M = \frac{s}{\sqrt{n}} = \frac{5.61}{5.48} = 1.02$$

Now we can put that value, our point estimate for the sample mean, and our critical value from step 2 into the formula for a confidence interval:

$$95\%CI = 39.85 \pm 2.045(1.02)$$

$$\text{Upper Bound} = 39.85 + 2.045(1.02)$$
$$UB = 39.85 + 2.09$$
$$UB = 41.94$$

$$\text{Lower Bound} = 39.85 - 2.045(1.02)$$
$$LB = 39.85 - 2.09$$
$$LB = 37.76$$

$$95\%CI = (37.76, 41.94)$$

Step 4: Make the Decision:

Finally, we can compare our confidence interval to our null hypothesis value. The null value of 38 is higher than our lower bound of 37.76 and lower than our upper bound of 41.94. Thus, the confidence interval brackets our null hypothesis value, and we fail to reject the null hypothesis:

> Fail to Reject H_0. Based on our sample of 30 people, our community not different in average friendliness (M= 39.85) than the nation as a whole, 95% CI = (37.76, 41.94).

Note that we don't report a test statistic or p-value because that is not how we tested the hypothesis, but we do report the value we found for our confidence interval.

An important characteristic of hypothesis testing is that both methods will always give you the same result. That is because both are based on the standard error and critical values in their calculations. To check this, we can calculate a t-statistic for the example above and find it to be t = 1.81, which is smaller than our critical value of 2.045 and fails to reject the null hypothesis.

Confidence Intervals using z Confidence intervals can also be constructed using z-score criteria, if one knows the population standard deviation. The format, calculations, and interpretation are all exactly the same, only replacing $t*$ with $z*$ and s_M with σ_M.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 9.3: Confidence Intervals is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **8.3: Confidence Intervals** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

9.E: Introduction to t-tests (Exercises)

1. What is the difference between a z-test and a 1-sample t-test?

Answer:

A z-test uses population standard deviation for calculating standard error and gets critical values based on the standard normal distribution. A t-test uses sample standard deviation as an estimate when calculating standard error and gets critical values from the t-distribution based on degrees of freedom.

2. What does a confidence interval represent?
3. What is the relationship between a chosen level of confidence for a confidence interval and how wide that interval is? For instance, if you move from a 95% CI to a 90% CI, what happens? Hint: look at the t-table to see how critical values change when you change levels of significance.

Answer:

As the level of confidence gets higher, the interval gets wider. In order to speak with more confidence about having found the population mean, you need to cast a wider net. This happens because critical values for higher confidence levels are larger, which creates a wider margin of error.

4. Construct a confidence interval around the sample mean M= 25 for the following conditions:
   a. N = 25, s = 15, 95% confidence level
   b. N = 25, s = 15, 90% confidence level
   c. s_M = 4.5, α = 0.05, df = 20
   d. s = 12, df = 16 (yes, that is all the information you need)
5. True or False: a confidence interval represents the most likely location of the true population mean.

Answer:

False: a confidence interval is a range of plausible scores that may or may not bracket the true population mean.

6. You hear that college campuses may differ from the general population in terms of political affiliation, and you want to use hypothesis testing to see if this is true and, if so, how big the difference is. You know that the average political affiliation in the nation is μ = 4.00 on a scale of 1.00 to 7.00, so you gather data from 150 college students across the nation to see if there is a difference. You find that the average score is 3.76 with a standard deviation of 1.52. Use a 1-sample t-test to see if there is a difference at the α = 0.05 level.
7. You hear a lot of talk about increasing global temperature, so you decide to see for yourself if there has been an actual change in recent years. You know that the average land temperature from 1951-1980 was 8.79 degrees Celsius. You find annual average temperature data from 1981-2017 and decide to construct a 99% confidence interval (because you want to be as sure as possible and look for differences in both directions, not just one) using this data to test for a difference from the previous average.

Year	Temp	Year	Temp	Year	Temp	Year	Temp
1981	9.301	1991	9.336	2001	9.542	2011	9.65
1982	8.788	1992	8.974	2002	9.695	2012	9.635
1983	9.173	1993	9.008	2003	9.649	2013	9.753
1984	8.824	1994	9.175	2004	9.451	2014	9.714
1985	8.799	1995	9.484	2005	9.829	2015	9.962
1986	8.985	1996	9.168	2006	9.662	2016	10.16
1987	9.141	1997	9.326	2007	9.876	2017	10.049
1988	9.345	1998	9.66	2008	9.581		
1989	9.076	1999	9.406	2009	9.657		
1990	9.378	2000	9.332	2010	9.828		

Answer:

M= 9.44, s = 0.35, s_M= 0.06, df = 36, $t*$ = 2.719, 99% CI = (9.28, 9.60); CI does not bracket μ, reject null hypothesis. d = 1.83

8. Determine whether you would reject or fail to reject the null hypothesis in the following situations:
   a. t = 2.58, N = 21, two-tailed test at α = 0.05

b. t = 1.99, N = 49, one-tailed test at α = 0.01
c. μ = 47.82, 99% CI = (48.71, 49.28)
d. μ = 0, 95% CI = (-0.15, 0.20)

9. You are curious about how people feel about craft beer, so you gather data from 55 people in the city on whether or not they like it. You code your data so that 0 is neutral, positive scores indicate liking craft beer, and negative scores indicate disliking craft beer. You find that the average opinion was M= 1.10 and the spread was s = 0.40, and you test for a difference from 0 at the α = 0.05 level.

Answer:

Step 1: $H_0 \neq 0$ "The average person has a neutral opinion towards craft beer", $H_A \neq 0$ "Overall people will have an opinion about craft beer, either good or bad."

Step 2: Two-tailed test, df = 54, $t*$ = 2.009.

Step 3: $\overline{X}$= 1.10, $s_{\overline{X}}$= 0.05, t = 22.00.

Step 4: $t > t*$, Reject H_0.

Based on opinions from 55 people, we can conclude that the average opinion of craft beer (M= 1.10) is positive, $t(54)$ = 22.00, p < .05. Since the result is significant, we need an effect size: Cohen's d = 2.75, which is a large effect.

10. You want to know if college students have more stress in their daily lives than the general population (μ = 12), so you gather data from 25 people to test your hypothesis. Your sample has an average stress score of M= 13.11 and a standard deviation of s = 3.89. Use a 1-sample t-test to see if there is a difference.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 9.E: Introduction to t-tests (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **8.E: Introduction to t-tests (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

10: Repeated Measures

10.1: Change and Differences

10.2: Hypotheses of Change and Differences

10.3: Increasing Satisfaction at Work

10.4: Bad Press

10.E: Repeated Measures (Exercises)

This page titled 10: Repeated Measures is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

10.1: Change and Differences

Researchers are often interested in change over time. Sometimes we want to see if change occurs naturally, and other times we are hoping for change in response to some manipulation. In each of these cases, we measure a single variable at different times, and what we are looking for is whether or not we get the same score at time 2 as we did at time 1. The absolute value of our measurements does not matter – all that matters is the change. Let's look at an example:

Table **10.1.1**: Raw and difference scores before and after training.

Before	After	Improvement
6	9	3
7	7	0
4	10	6
1	3	2
8	10	2

Table **10.1.1** shows scores on a quiz that five employees received before they took a training course and after they took the course. The difference between these scores (i.e. the score after minus the score before) represents improvement in the employees' ability. This third column is what we look at when assessing whether or not our training was effective. We want to see positive scores, which indicate that the employees' performance went up. What we are not interested in is how good they were before they took the training or after the training. Notice that the lowest scoring employee before the training (with a score of 1) improved just as much as the highest scoring employee before the training (with a score of 8), regardless of how far apart they were to begin with. There's also one improvement score of 0, meaning that the training did not help this employee. An important factor in this is that the participants received the same assessment at both time points. To calculate improvement or any other difference score, we must measure only a single variable.

When looking at change scores like the ones in Table **10.1.1**, we calculate our difference scores by taking the time 2 score and subtracting the time 1 score. That is:

$$\mathbf{X_d = X_{T2} - X_{T1}} \tag{10.1.1}$$

Where $\mathbf{X_d}$ is the difference score, $\mathbf{X_{T1}}$ is the score on the variable at time 1, and $\mathbf{X_{T2}}$ is the score on the variable at time 2. The difference score, $\mathbf{X_d}$, will be the data we use to test for improvement or change. We subtract time 2 minus time 1 for ease of interpretation; if scores get better, then the difference score will be positive. Similarly, if we're measuring something like reaction time or depression symptoms that we are trying to reduce, then better outcomes (lower scores) will yield negative difference scores.

We can also test to see if people who are matched or paired in some way agree on a specific topic. For example, we can see if a parent and a child agree on the quality of home life, or we can see if two romantic partners agree on how serious and committed their relationship is. In these situations, we also subtract one score from the other to get a difference score. This time, however, it doesn't matter which score we subtract from the other because what we are concerned with is the agreement.

In both of these types of data, what we have are multiple scores on a single variable. That is, a single observation or data point is comprised of two measurements that are put together into one difference score. This is what makes the analysis of change unique – our ability to link these measurements in a meaningful way. This type of analysis would not work if we had two separate samples of people that weren't related at the individual level, such as samples of people from different states that we gathered independently. Such datasets and analyses are the subject of the following chapter.

A rose by any other name…

It is important to point out that this form of t-test has been called many different things by many different people over the years: "matched pairs", "paired samples", "repeated measures", "dependent measures", "dependent samples", and many others. What all of these names have in common is that they describe the analysis of two scores that are related in a systematic way within people or within pairs, which is what each of the datasets usable in this analysis have in common. As such, all of these names are equally

appropriate, and the choice of which one to use comes down to preference. In this text, we will refer to paired samples, though the appearance of any of the other names throughout this chapter should not be taken to refer to a different analysis: they are all the same thing.

Now that we have an understanding of what difference scores are and know how to calculate them, we can use them to test hypotheses. As we will see, this works exactly the same way as testing hypotheses about one sample mean with a tstatistic. The only difference is in the format of the null and alternative hypotheses.

This page titled 10.1: Change and Differences is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **9.1: Change and Differences** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

10.2: Hypotheses of Change and Differences

When we work with difference scores, our research questions have to do with change. Did scores improve? Did symptoms get better? Did prevalence go up or down? Our hypotheses will reflect this. Remember that the null hypothesis is the idea that there is nothing interesting, notable, or impactful represented in our dataset. In a paired samples t-test, that takes the form of 'no change'. There is no change in scores or in symptoms. Thus, our null hypothesis is:

H_0: There is no change or difference

$$H_0 : \mu_D = 0$$

As with our other null hypotheses, we express the null hypothesis for paired samples t-tests in both words and mathematical notation. The exact wording of the written-out version should be changed to match whatever research question we are addressing (e.g. " There is no change in ability scores after training"). However, the mathematical version of the null hypothesis is always exactly the same (for two tailed tests): the average change score is equal to zero. Our population parameter for the average is still μ, but it now has a subscript D to denote the fact that it is the average change score and not the average raw observation before or after our manipulation. Obviously individual difference scores can go up or down, but the null hypothesis states that these positive or negative change values are just random chance and that the true average change score across all people is 0.

Our alternative hypotheses will also follow the same format that they did before. They can be directional if we suspect a change or difference in a specific direction or we can use an inequality sign to test for any change:

H_A: There is a change or difference

$$H_A : \mu_D \neq 0$$

H_A: The average score increases

$$H_A : \mu_D > 0$$

H_A: The average score decreases

$$H_A : \mu_D < 0$$

Remember that if you do directional alternative hypotheses, your null should match it.

For "The average score increases:"

H_0: The average score does not increase

$$H_A : \mu_D \leq 0$$

For "The average score decreases:"

H_0: The average score does not decrease

$$H_A : \mu_D \geq 0$$

As before, your choice of which alternative hypothesis to use should be specified before you collect data based on your research question and any evidence you might have that would indicate a specific directional (or non-directional) change.

Critical Values and Decision Criteria

As with before, once we have our hypotheses laid out, we need to find our critical values that will serve as our decision criteria. This step has not changed at all from the last chapter. Our critical values are based on our level of significance (still usually α = 0.05), the directionality of our test (one-tailed or two-tailed), and the degrees of freedom, which are still calculated as $df = n–1$. Because this is a t-test like the last chapter, we will find our critical values on the same t-table using the same process of identifying the correct column based on our significance level and directionality and the correct row based on our degrees of freedom or the next lowest value if our exact degrees of freedom are not presented. After we calculate our test statistic, our decision criteria are the same as well: $p < \alpha$ or $t_{obt} > t*$.

Test Statistic

Our test statistic for our change scores follows exactly the same format as it did for our 1-sample t-test. In fact, the only difference is in the data that we use. For our change test, we first calculate a difference score as shown above. Then, we use those scores as the raw data in the same mean calculation, standard error formula, and t-statistic. Let's look at each of these.

The mean difference score is calculated in the same way as any other mean: sum each of the individual difference scores and divide by the sample size.

$$M_D = \frac{\Sigma X_D}{n} \tag{10.2.1}$$

Here we are using the subscript D to keep track of that fact that these are difference scores instead of raw scores; it has no actual effect on our calculation. Using this, we calculate the standard deviation of the difference scores the same way as well:

$$s_D = \sqrt{\frac{\sum (X_D - M_D)^2}{n-1}} = \sqrt{\frac{SS}{df}} \tag{10.2.2}$$

We will find the numerator, the Sum of Squares, using the same table format that we learned in chapter 3. Once we have our standard deviation, we can find the standard error:

$$s_{M_D} = {}^{s_D}/\sqrt{n} \tag{10.2.3}$$

Finally, our test statistic t has the same structure as well:

$$t = \frac{M_D - \mu_D}{s_{M_D}} \tag{10.2.4}$$

As we can see, once we calculate our difference scores from our raw measurements, everything else is exactly the same. Let's see an example.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 10.2: Hypotheses of Change and Differences is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **9.2: Hypotheses of Change and Differences** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

10.3: Increasing Satisfaction at Work

Workers at a local company have been complaining that working conditions have gotten very poor, hours are too long, and they don't feel supported by the management. The company hires a consultant to come in and help fix the situation before it gets so bad that the employees start to quit. The consultant first assesses 40 of the employee's level of job satisfaction as part of focus groups used to identify specific changes that might help. The company institutes some of these changes, and six months later the consultant returns to measure job satisfaction again. Knowing that some interventions miss the mark and can actually make things worse, the consultant tests for a difference in either direction (i.e. and increase or a decreased in average job satisfaction) at the α = 0.05 level of significance.

Step 1: State the Hypotheses:

First, we state our null and alternative hypotheses:

H_0: There is no increase in average job satisfaction after the changes

$$H_0 : \mu_D \leq 0$$

H_A: There is an increase in average job satisfaction after the changes

$$H_A : \mu_D > 0$$

In this case, we are hoping that the changes we made will improve employee satisfaction, and, because we based the changes on employee recommendations, we have good reason to believe that they will. Thus, we will use a one-directional alternative hypothesis.

Step 2: Find the Critical Values:

Our critical values will once again be based on our level of significance, which we know is α = 0.05, the directionality of our test, which is one-tailed to the right, and our degrees of freedom. For our dependent-samples t-test, the degrees of freedom are still given as $df = n–1$. For this problem, we have 40 people, so our degrees of freedom are 39. Going to our t-table, we find that the critical value is $t*$ = 1.685 as shown in Figure **10.3.1**.

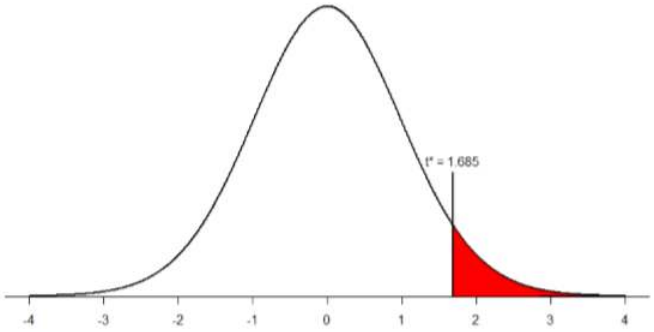


Figure **10.3.1**: Critical region for one-tailed t-test at α = 0.05

Step 3: Calculate the Test Statistic:

Now that the criteria are set, it is time to calculate the test statistic. The data obtained by the consultant found that the difference scores from time 1 to time 2 had a mean of M_D = 2.96 and a standard deviation of s_D = 2.85. Using this information, plus the size of the sample (N = 40), we first calculate the standard error:

$$s_{M_D} = s_{D/\sqrt{n}} = 2.85/\sqrt{40} = 2.85/6.32 = 0.46$$

Now, we can put that value, along with our sample mean and null hypothesis value, into the formula for t and calculate the test statistic:

$$t = \frac{M_D - \mu_D}{s_{M_D}} = \frac{2.96 - 0}{0.46} = 6.43$$

Notice that, because the μ_D of a paired samples t-test is 0, we can simply divide our obtained sample mean by the standard error.

Step 4: Make the Decision:

We have obtained a test statistic of t = 6.43 that we can compare to our previously established critical value of $t*$ = 1.685. 6.43 is larger than 1.685, so $t > t*$ and we reject the null hypothesis:

> Reject H_0. Based on the sample data from 40 workers, we can say that the intervention statistically significantly improved job satisfaction (M_D= 2.96) among the workers, $t(39) = 6.43$, $p < 0.05$.

Because this result was statistically significant, we will want to calculate Cohen's d as an effect size using the same format as we did for the last t-test:

$$t = \frac{M_D - \mu_D}{s_D} = \frac{2.96}{2.85} = 1.04$$

This is a large effect size. Notice again that we can omit the μ_D here because it is equal to 0.

Hopefully the above example made it clear that running a paired samples t-test to look for differences before and after some treatment works exactly the same way as a regular 1-sample t-test does, which was just a small change in how z-tests were performed in chapter 7. At this point, this process should feel familiar, and we will continue to make small adjustments to this familiar process as we encounter new types of data to test new types of research questions.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 10.3: Increasing Satisfaction at Work is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **9.3: Increasing Satisfaction at Work** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

10.4: Bad Press

Let's say that a bank wants to make sure that their new commercial will make them look good to the public, so they recruit 7 people to view the commercial as a focus group. The focus group members fill out a short questionnaire about how they view the company, then watch the commercial and fill out the same questionnaire a second time. The bank really wants to find significant results, so they test for a change at α = 0.10. However, they use a 2-tailed test since they know that past commercials have not gone over well with the public, and they want to make sure the new one does not backfire. They decide to test their hypothesis using a confidence interval to see just how spread out the opinions are. As we will see, confidence intervals work the same way as they did before, just like with the test statistic.

Step 1: State the Hypotheses:

As always, we start with hypotheses:

H_0: There is no difference in how people view the bank after the commercial compared to before.

$$H_0 : \mu_D = 0$$

H_A: There is a difference in how people view the bank after the commercial compared to before.

$$H_A : \mu_D \neq 0$$

Step 2: Find the Critical Values:

Just like with our regular hypothesis testing procedure, we will need critical values from the appropriate level of significance and degrees of freedom in order to form our confidence interval. Because we have 7 participants, our degrees of freedom are df = 6. Remember that our α = 0.10. From our t-table, we find that the critical value corresponding to this df at this level of significance is $t*$ = +/-1.943.

Step 3: Calculate the Confidence Interval:

The data collected before (time 1) and after (time 2) the participants viewed the commercial is presented in Table 10.4.1. In order to build our confidence interval, we will first have to calculate the mean and standard deviation of the difference scores, which are also in Table 10.4.1. As a reminder, the difference scores are calculated as Time 2 – Time 1.

Table 10.4.1: Opinions of the bank

Time 1	Time 2	X_D
3	2	-1
3	6	3
5	3	-2
8	4	-4
3	9	6
1	2	1
4	5	1

The mean of the difference scores is:

$$M_D = \frac{\sum X_D}{n} = \frac{4}{7} = 0.57$$

The standard deviation will be solved by first using the Sum of Squares Table:

Table 10.4.2: Sum of Squares

X_D	$X_D - M_D$	$(X_D - M_D)^2$
-1	-1.57	2.46
3	2.43	5.90
-2	-2.57	6.60
-4	-4.57	20.88
6	5.43	29.48
1	0.43	0.18
1	0.43	0.18
$\Sigma = 4$	$\Sigma = 0$	$\Sigma = 65.68$

$$s_D = \sqrt{\frac{SS}{df}} = \sqrt{\frac{65.68}{6}} = \sqrt{10.95} = 3.31$$

Finally, we find the standard error:

$$s_{M_D} = {}^{S_D}/\sqrt{n} = 3.31/\sqrt{7} = 1.25$$

We now have all the pieces needed to compute our confidence interval:

$$95\%CI = M_D \pm t^*(s_{M_D})$$
$$95\%CI = 0.57 \pm 1.943(1.25)$$

$$\text{Upper Bound} = 0.57 + 1.943(1.25)$$
$$UB = 0.57 + 2.43$$
$$UB = 3.00$$

$$\text{Lower Bound} = 0.57 - 1.943(1.25)$$
$$LB = 0.57 - 2.43$$
$$LB = -1.86$$

$$95\%CI = (-1.86, 3.00)$$

Step 4: Make the Decision:

Remember that the confidence interval represents a range of values that seem plausible or reasonable based on our observed data. The interval spans -1.86 to 3.00, which includes 0, our null hypothesis value. Because the null hypothesis value is in the interval, it is considered a reasonable value, and because it is a reasonable value, we have no evidence against it. We fail to reject the null hypothesis.

Fail to Reject H_0. Based on our focus group of 7 people, we cannot say that the average change in opinion (M_D = 0.57) was any better or worse after viewing the commercial, CI: (-1.86, 3.00).

As with before, we only report the confidence interval to indicate how we performed the test.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 10.4: Bad Press is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **9.4: Bad Press** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

10.E: Repeated Measures (Exercises)

1. What is the difference between a 1-sample t-test and a paired samples t-test? How are they alike?

Answer:

A 1-sample t-test uses raw scores to compare an average to a specific value. A paired samples t-test uses two raw scores from each person to calculate difference scores and test for an average difference score that is equal to zero. The calculations, steps, and interpretation is exactly the same for each.

2. Name 3 research questions that could be addressed using a paired samples t-test.
3. What are difference scores and why do we calculate them?

Answer:

Difference scores indicate change or discrepancy relative to a single person or pair of people. We calculate them to eliminate individual differences in our study of change or agreement.

4. Why is the null hypothesis for a paired samples t-test always $\mu_D = 0$?
5. A researcher is interested in testing whether explaining the processes of statistics helps increase trust in computer algorithms. He wants to test for a difference at the α = 0.05 level and knows that some people may trust the algorithms less after the training, so he uses a two-tailed test. He gathers pre-post data from 35 people and finds that the average difference score is M_D = 12.10 with a standard deviation of s_D= 17.39. Conduct a hypothesis test to answer the research question.

Answer:

Step 1: $H_0 : \mu = 0$ “The average change in trust of algorithms is 0”, $H_A : \mu \neq 0$ “People’s opinions of how much they trust algorithms changes.”

Step 2: Two-tailed test, df = 34, $t*$ = 2.032.

Step 3: M_D = 12.10, s_{M_D}= 2.94, t = 4.12.

Step 4: $t > t*$, Reject H_0. Based on opinions from 35 people, we can conclude that people trust algorithms more (M_D= 12.10) after learning statistics, $t(34) = 4.12, p < .05$. Since the result is significant, we need an effect size: Cohen’s d = 0.70, which is a moderate to large effect.

6. Decide whether you would reject or fail to reject the null hypothesis in the following situations:
   a. M_D = 3.50, s_D = 1.10, n = 12, α = 0.05, two-tailed test
   b. 95% CI= (0.20,1.85)
   c. t = 2.98, $t*$ = -2.36, one-tailed test to the left
   d. 90% CI = (-1.12, 4.36)
7. Calculate difference scores for the following data:

Time 1	Time 2	X_D
61	83	
75	89	
91	98	
83	92	
74	80	
82	88	
98	98	
82	77	
69	88	
76	79	
91	91	
70	80	

Answer:

Time 1	Time 2	X_D
61	83	22
75	89	14
91	98	7
83	92	9
74	80	6
82	88	6
98	98	0
82	77	-5
69	88	19
76	79	3
91	91	0
70	80	10

8. You want to know if an employee's opinion about an organization is the same as the opinion of that employee's boss. You collect data from 18 employee-supervisor pairs and code the difference scores so that positive scores indicate that the employee has a higher opinion and negative scores indicate that the boss has a higher opinion (meaning that difference scores of 0 indicate no difference and complete agreement). You find that the mean difference score is M_D= -3.15 with a standard deviation of s_D = 1.97. Test this hypothesis at the α = 0.01 level.
9. Construct confidence intervals from a mean of M_D = 1.25, standard error of s_{M_D}= 0.45, and df = 10 at the 90%, 95%, and 99% confidence level. Describe what happens as confidence changes and whether to reject H_0.

Answer:

At the 90% confidence level, $t* = 1.812$ and CI = (0.43, 2.07) so we reject H_0. At the 95% confidence level, $t* = 2.228$ and CI = (0.25, 2.25) so we reject H_0. At the 99% confidence level, $t* = 3.169$ and CI = (-0.18, 2.68) so we fail to reject H_0. As the confidence level goes up, our interval gets wider (which is why we have higher confidence), and eventually we do not reject the null hypothesis because the interval is so wide that it contains 0.

10. A professor wants to see how much students learn over the course of a semester. A pre-test is given before the class begins to see what students know ahead of time, and the same test is given at the end of the semester to see what students know at the end. The data are below. Test for an improvement at the $\alpha = 0.05$ level. Did scores increase? How much did scores increase?

Pretest	Posttest	X_D
90	89	
60	66	
95	99	
93	91	
95	100	
67	64	
89	91	
90	95	
94	95	
83	89	
75	82	
87	92	
82	83	
82	85	
88	93	
66	69	
90	90	
93	100	
86	95	
91	96	

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 10.E: Repeated Measures (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **9.E: Repeated Measures (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

11: Independent Samples

11.1: Difference of Means
11.2: Research Questions about Independent Means
11.3: Hypotheses and Decision Criteria
11.4: Independent Samples t-statistic
11.5: Standard Error and Pooled Variance
11.6: Movies and Mood
11.7: Effect Sizes and Confidence Intervals
11.8: Homogeneity of Variance
11.E: Independent Samples (Exercises)

This page titled 11: Independent Samples is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

11.1: Difference of Means

Last chapter, we learned about mean differences, that is, the average value of difference scores. Those difference scores came from ONE group and TWO time points (or two perspectives). Now, we will deal with the difference of the means, that is, the average values of separate groups that are represented by separate descriptive statistics. This analysis involves TWO groups and ONE time point. As with all of our other tests as well, both of these analyses are concerned with a single variable.

It is very important to keep these two tests separate and understand the distinctions between them because they assess very different questions and require different approaches to the data. When in doubt, think about how the data were collected and where they came from. If they came from two time points with the same people (sometimes referred to as "longitudinal" data), you know you are working with repeated measures data (the measurement literally was repeated) and will use a repeated measures/paired samples t-test. If it came from a single time point that used separate groups, you need to look at the nature of those groups and if they are related. Can individuals in one group being meaningfully matched up with one and only one individual from the other group? For example, are they a romantic couple? If so, we call those data matched and we use a matched pairs/paired samples t-test. However, if there's no logical or meaningful way to link individuals across groups, or if there is no overlap between the groups, then we say the groups are independent and use the independent samples t-test, the subject of this chapter.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.1: Difference of Means is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.1: Difference of Means** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.2: Research Questions about Independent Means

Many research ideas in the behavioral sciences and other areas of research are concerned with whether or not two means are the same or different. Logically, we therefore say that these research questions are concerned with group mean differences. That is, on average, do we expect a person from Group A to be higher or lower on some variable that a person from Group B. In any type of research design looking at group mean differences, there are some key criteria we must consider: the groups must be mutually exclusive (i.e. you can only be part of one group at any given time) and the groups have to be measured on the same variable (i.e. you can't compare personality in one group to reaction time in another group since those values would not be the same anyway).

Let's look at one of the most common and logical examples: testing a new medication. When a new medication is developed, the researchers who created it need to demonstrate that it effectively treats the symptoms they are trying to alleviate. The simplest design that will answer this question involves two groups: one group that receives the new medication (the "treatment" group) and one group that receives a placebo (the "control" group). Participants are randomly assigned to one of the two groups (remember that random assignment is the hallmark of a true experiment), and the researchers test the symptoms in each person in each group after they received either the medication or the placebo. They then calculate the average symptoms in each group and compare them to see if the treatment group did better (i.e. had fewer or less severe symptoms) than the control group.

In this example, we had two groups: treatment and control. Membership in these two groups was mutually exclusive: each individual participant received either the experimental medication or the placebo. No one in the experiment received both, so there was no overlap between the two groups. Additionally, each group could be measured on the same variable: symptoms related to the disease or ailment being treated. Because each group was measured on the same variable, the average scores in each group could be meaningfully compared. If the treatment was ineffective, we would expect that the average symptoms of someone receiving the treatment would be the same as the average symptoms of someone receiving the placebo (i.e. there is no difference between the groups). However, if the treatment WAS effective, we would expect fewer symptoms from the treatment group, leading to a lower group average.

Now let's look at an example using groups that already exist. A common, and perhaps salient, question is how students feel about their job prospects after graduation. Suppose that we have narrowed our potential choice of college down to two universities and, in the course of trying to decide between the two, we come across a survey that has data from each university on how students at those universities feel about their future job prospects. As with our last example, we have two groups: University A and University B, and each participant is in only one of the two groups (assuming there are no transfer students who were somehow able to rate both universities). Because students at each university completed the same survey, they are measuring the same thing, so we can use a ***t***-test to compare the average perceptions of students at each university to see if they are the same. If they are the same, then we should continue looking for other things about each university to help us decide on where to go. But, if they are different, we can use that information in favor of the university with higher job prospects.

As we can see, the grouping variable we use for an independent samples ***t***-test can be a set of groups we create (as in the experimental medication example) or groups that already exist naturally (as in the university example). There are countless other examples of research questions relating to two group means, making the independent samples ***t***-test one of the most widely used analyses around.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.2: Research Questions about Independent Means is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.2: Research Questions about Independent Means** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.3: Hypotheses and Decision Criteria

The process of testing hypotheses using an independent samples *t*-test is the same as it was in the last three chapters, and it starts with stating our hypotheses and laying out the criteria we will use to test them.

Our null hypothesis for an independent samples *t*-test follows a similar format as all others. For a two tailed test, the means of the two groups are the same under the null hypothesis, no matter how those groups were formed. Mathematically, this takes on two equivalent forms:

$$H_0 : \mu_1 = \mu_2$$

or

$$H_0 : \mu_1 - \mu_2 = 0$$

Both of these formulations of the null hypothesis tell us exactly the same thing: that the numerical value of the means is the same in both groups. This is more clear in the first formulation, but the second formulation also makes sense (any number minus itself is always zero) and helps us out a little when we get to the math of the test statistic. Either one is acceptable and you only need to report one. The English interpretation of both of them is also the same:

$$H_0 : \text{ There is no difference between the means of the two groups}$$

Note that you should make sure that you make the hypothesis specific to your experiment/situation when you write it out.

Our alternative hypotheses also follow a similar format.

For two tailed tests:

$$H_A : \mu_1 \neq \mu_2$$

Or

$$H_A : \mu_1 - \mu_2 \neq 0$$

In words:

$$H_A : \text{There is a difference between the means of the two groups}$$

For one tailed tests, the alternative hypothesis will include predictions about one group being higher or lower than the other group:

$$H_A : \mu_1 > \mu_2 \; H_A : \mu_1 < \mu_2$$

And the appropriate null hypothesis:

$$H_0 : \mu_1 \leq \mu_2 \; H_0 : \mu_1 \geq \mu_2$$

In words, these would be:

$$H_A : \text{Group 1 is higher than group 2}$$

OR

$$H_A : \text{Group 1 is lower than group 2}$$

$$H_0 : \text{Group 1 is not higher than group 2}$$

OR

$$H_0 : \text{Group 1 is not lower than group 2}$$

Notice that we are now dealing with two means instead of just one, so it will be very important to keep track of which mean goes with which population and, by extension, which dataset and sample data. We use subscripts to differentiate between the populations, so make sure to keep track of which is which. If it is helpful, you can also use more descriptive subscripts. To use the experimental medication example:

H_0: There is no difference between the scores of the those who received the treatment and the control groups

$$H_0 : \mu_{\text{treatment}} = \mu_{\text{control}}$$

H_A: There is a difference between the scores of those who received the treatment and the control groups

$$H_A : \mu_{\text{treatment}} \neq \mu_{\text{control}}$$

Once we have our hypotheses laid out, we can set our criteria to test them using the same three pieces of information as before: significance level (α), directionality (left, right, or two-tailed), and degrees of freedom, which for an independent samples t-test are:

$$df = n_1 + n_2 - 2$$

This looks different than before, but it is just adding the individual degrees of freedom from each group (n–1) together. Notice that the sample sizes, n, also get subscripts so we can tell them apart.

For an independent samples t-test, it is often the case that our two groups will have slightly different sample sizes, either due to chance or some characteristic of the groups themselves. Generally, this is not as issue, so long as one group is not massively larger than the other group. What is of greater concern is keeping track of which is which using the subscripts.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.3: Hypotheses and Decision Criteria is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.3: Hypotheses and Decision Criteria** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.4: Independent Samples t-statistic

The test statistic for our independent samples t-test takes on the same logical structure and format as our other t-tests: our observed effect minus our null hypothesis value, all divided by the standard error:

$$t = \frac{(M_1 - M_2) - (\mu_1 - \mu_2)}{s_{M_1 - M_2}} \tag{11.4.1}$$

This looks like more work to calculate, but remember that, generally, our null hypothesis states that the quantity $\mu_1 - \mu_2 = 0$, so we can drop that out of the equation and are left with:

$$t = \frac{(M_1 - M_2)}{s_{M_1 - M_2}} \tag{11.4.2}$$

Our standard error in the denomination is still standard deviation (s) with a subscript denoting what it is the standard error of. Because we are dealing with the difference between two separate means, rather than a single mean or single mean of difference scores, we put both means in the subscript. Calculating our standard error, as we will see next, is where the biggest differences between this t-test and other t-tests appears. However, once we do calculate it and use it in our test statistic, everything else goes back to normal. Our decision criteria is still comparing our obtained test statistic to our critical value, and our interpretation based on whether or not we reject the null hypothesis is unchanged as well.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.4: Independent Samples t-statistic is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.4: Independent Samples t-statistic** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.5: Standard Error and Pooled Variance

Recall that the standard error is the average distance between any given sample mean and the center of its corresponding sampling distribution, and it is a function of the standard deviation of the population (either given or estimated) and the sample size. This definition and interpretation hold true for our independent samples t-test as well, but because we are working with two samples drawn from two populations, we have to first combine their estimates of standard deviation – or, more accurately, their estimates of variance – into a single value that we can then use to calculate our standard error.

The combined estimate of variance using the information from each sample is called the pooled variance and is denoted s_p^2; the subscript p serves as a reminder indicating that it is the pooled variance. The term "pooled variance" is a literal name because we are simply pooling or combining the information on variance – the Sum of Squares and Degrees of Freedom – from both of our samples into a single number. The result is a weighted average of the observed sample variances, the weight for each being determined by the sample size, and will always fall between the two observed variances. The computational formula for the pooled variance is:

$$s_p^2 = \frac{(n_1 - 1)\, s_1^2 + (n_2 - 1)\, s_2^2}{n_1 + n_2 - 2} \tag{11.5.1}$$

This formula can look daunting at first, but it is in fact just a weighted average. Even more conveniently, some simple algebra can be employed to greatly reduce the complexity of the calculation. The simpler and more appropriate formula to use when calculating pooled variance is:

$$s_p^2 = \frac{SS_1 + SS_2}{df_1 + df_2} \tag{11.5.2}$$

Using this formula, it's very simple to see that we are just adding together the same pieces of information we have been calculating since chapter 3. Thus, when we use this formula, the pooled variance is not nearly as intimidating as it might have originally seemed.

Once we have our pooled variance calculated, we can drop it into the equation for our standard error:

$$S_{M_1 - M_2} = \sqrt{\frac{S_p^2}{n_1} + \frac{S_p^2}{n_2}} \tag{11.5.3}$$

Once again, although this formula may seem different than it was before, in reality it is just a different way of writing the same thing. An alternative but mathematically equivalent way of writing our old standard error is:

$$s_M = \frac{s}{\sqrt{n}} = \sqrt{\frac{s^2}{n}} \tag{11.5.4}$$

Looking at that, we can now see that, once again, we are simply adding together two pieces of information: no new logic or interpretation required. Once the standard error is calculated, it goes in the denominator of our test statistic, as shown above and as was the case in all previous chapters. Thus, the only additional step to calculating an independent samples t-statistic is computing the pooled variance. Let's see an example in action.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.5: Standard Error and Pooled Variance is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.5: Standard Error and Pooled Variance** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.6: Movies and Mood

We are interested in whether the type of movie someone sees at the theater affects their mood when they leave. We decide to ask people about their mood as they leave one of two movies: a comedy (group 1, n = 35) or a horror film (group 2, n = 29). Our data are coded so that higher scores indicate a more positive mood. We have good reason to believe that people leaving the comedy will be in a better mood, so we use a one-tailed test at α = 0.05 to test our hypothesis.

Step 1: State the Hypotheses:

As always, we start with hypotheses:

H_0: The comedy film will not give a better average mood than the horror film

$$H_0 : \mu_1 - \mu_2 \leq 0$$

or

$$H_0 : \mu_1 \leq \mu_2$$

H_A: The comedy film will give a better average mood than the horror film

$$H_A : \mu_1 - \mu_2 > 0$$

or

$$H_A : \mu_1 > \mu_2$$

Notice that in the first formulation of the alternative hypothesis we say that the first mean minus the second mean will be greater than zero. This is based on how we code the data (higher is better), so we suspect that the mean of the first group will be higher. Thus, we will have a larger number minus a smaller number, which will be greater than zero. Be sure to pay attention to which group is which and how your data are coded (higher is almost always used as better outcomes) to make sure your hypothesis makes sense!

Step 2: Find the Critical Values:

Just like before, we will need critical values, which come from our t-table. In this example, we have a one-tailed test at α = 0.05 and expect a positive answer (because we expect the difference between the means to be greater than zero). Our degrees of freedom for our independent samples t-test is just the degrees of freedom from each group added together: 35 + 29 – 2 = 62. From our t-table, we find that our critical value is $t*$ = 1.671. Note that because 62 does not appear on the table, we use the next lowest value, which in this case is 60.

Step 3: Compute the Test Statistic:

The data from our two groups are presented in the tables below. Table **11.6.1** shows the values for the Comedy group, and Table **11.6.2** shows the values for the Horror group. Values for both have already been placed in the Sum of Squares tables since we will need to use them for our further calculations. As always, the column on the left is our raw data.

Table **11.6.1**: Raw scores and Sum of Squares for Group 1

X	(X − M)	(X − M)²
Group 1: Comedy Film		
39.10	15.10	228.01
38.00	14.00	196.00
14.90	-9.10	82.81
20.70	-3.30	10.89
19.50	-4.50	20.25
32.20	8.20	67.24
11.00	-13.00	169.00
20.70	-3.30	10.89
26.40	2.40	5.76
35.70	11.70	136.89
26.40	2.40	5.76
28.80	4.80	23.04
33.40	9.40	88.36
13.70	-10.30	106.09
46.10	22.10	488.41
13.70	-10.30	106.09
23.00	-1.00	1.00
20.70	-3.30	10.89
19.50	-4.50	20.25
11.40	-12.60	158.76
24.10	0.10	0.01
17.20	-6.80	46.24
38.00	14.00	196.00
10.30	-13.70	187.69
35.70	11.70	136.89
41.50	17.50	306.25
18.40	-5.60	31.36
36.80	12.80	163.84
54.10	30.10	906.01
11.40	-12.60	158.76
8.70	-15.30	234.09
23.00	-1.00	1.00

X	$(X - M)$	$(X - M)^2$
14.30	-9.70	94.09
5.30	-18.70	349.69
6.30	-17.70	313.29
$\Sigma = 840$	$\Sigma = 0$	$\Sigma = 5061.60$

Table 11.6.2: Raw scores and Sum of Squares for Group 2

X	$(X - M)$	$(X - M)^2$
Group 2: Horror Film		
24.00	7.50	56.25
17.00	0.50	0.25
35.80	19.30	372.49
18.00	1.50	2.25
-1.70	-18.20	331.24
11.10	-5.40	29.16
10.10	-6.40	40.96
16.10	-0.40	0.16
-0.70	-17.20	295.84
14.10	-2.40	5.76
25.90	9.40	88.36
23.00	6.50	42.25
20.00	3.50	12.25
14.10	-2.40	5.76
-1.70	-18.20	331.24
19.00	2.50	6.25
20.00	3.50	12.25
30.90	14.40	207.36
30.90	14.40	207.36
22.00	5.50	30.25
6.20	-10.30	106.09
27.90	11.40	129.96
14.10	-2.40	5.76
33.80	17.30	299.29
26.90	10.40	108.16
5.20	-11.30	127.69
13.10	-3.40	11.56
19.00	2.50	6.25
-15.50	-32.00	1024.00
$\Sigma = 478.6$	$\Sigma = 0.10$	$\Sigma = 3896.45$

Using the sum of the first column for each table, we can calculate the mean for each group:

$$M_1 = \frac{840}{35} = 24.00$$

And

$$M_2 = \frac{478.60}{29} = 16.50$$

These values were used to calculate the middle rows of each table, which sum to zero as they should (the middle column for group 2 sums to a very small value instead of zero due to rounding error – the exact mean is 16.50344827586207, but that's far more than we need for our purposes). Squaring each of the deviation scores in the middle columns gives us the values in the third columns, which sum to our next important value: the Sum of Squares for each group: SS_1 = 5061.60 and SS_2 = 3896.45. These values have all been calculated and take on the same interpretation as they have since chapter 3 – no new computations yet. Before we move on to the pooled variance that will allow us to calculate standard error, let's compute our standard deviation for each group; even though we will not use them in our calculation of the test statistic, they are still important descriptors of our data:

$$s_1 = \sqrt{\frac{5061.60}{34}} = 12.20$$

And

$$s_2 = \sqrt{\frac{3896.45}{28}} = 11.80$$

Now we can move on to our new calculation, the pooled variance, which is just the Sums of Squares that we calculated from our table and the degrees of freedom, which is just n–1 for each group:

$$s_p^2 = \frac{SS_1 + SS_2}{df_1 + df_2} = \frac{5061.60 + 3896.45}{34 + 28} = \frac{8958.05}{62} = 144.48$$

As you can see, if you follow the regular process of calculating standard deviation using the Sum of Squares table, finding the pooled variance is very easy. Now we can use that value to calculate our standard error, the last step before we can find our test statistic:

$$s_{M_1 - M_2} = \sqrt{\frac{s_p^2}{n_1} + \frac{s_p^2}{n_2}} = \sqrt{\frac{144.48}{35} + \frac{144.48}{29}} = \sqrt{4.13 + 4.98} = \sqrt{9.11} = 3.02$$

Finally, we can use our standard error and the means we calculated earlier to compute our test statistic. Because the null hypothesis value of $\mu_1 - \mu_2$ is 0.00, we will leave that portion out of the equation for simplicity:

$$t = \frac{(M_1 - M_2)}{s_{M_1 - M_2}} = \frac{24.00 - 16.50}{3.02} = 2.48$$

The process of calculating our obtained test statistic t = 2.48 followed the same sequence of steps as before: use raw data to compute the mean and sum of squares (this time for two groups instead of one), use the sum of squares and degrees of freedom to calculate standard error (this time using pooled variance instead of standard deviation), and use that standard error and the observed means to get t. Now we can move on to the final step of the hypothesis testing procedure.

Step 4: Make the Decision:

Our test statistic has a value of t = 2.48, and in step 2 we found that the critical value is $t*$ = 1.671. 2.48 > 1.671, so we reject the null hypothesis:

Reject H_0. Based on our sample data from people who watched different kinds of movies, we can say that the average mood after a comedy movie ($M_1 = 24.00$) is better than the average mood after a horror movie ($M_2 = 16.50$), $t(62) = 2.48, p < .05$.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.6: Movies and Mood is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.6: Movies and Mood** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.7: Effect Sizes and Confidence Intervals

We have seen in previous chapters that even a statistically significant effect needs to be interpreted along with an effect size to see if it is practically meaningful. We have also seen that our sample means, as a point estimate, are not perfect and would be better represented by a range of values that we call a confidence interval. As with all other topics, this is also true of our independent samples *t*-tests.

Our effect size for the independent samples *t*-test is still Cohen's *d*, and it is still just our observed effect divided by the standard deviation. Remember that standard deviation is just the square root of the variance, and because we work with pooled variance in our test statistic, we will use the square root of the pooled variance as our denominator in the formula for Cohen's *d*. This gives us:

$$d = \frac{M_1 - M_2}{\sqrt{s_p^2}} \tag{11.7.1}$$

For our example above, we can calculate the effect size to be:

$$d = \frac{24.00 - 16.50}{\sqrt{144.48}} = \frac{7.50}{12.02} = 0.62$$

We interpret this using the same guidelines as before, so we would consider this a moderate or moderately large effect.

Our confidence intervals also take on the same form and interpretation as they have in the past. The value we are interested in is the difference between the two means, so our point estimate is the value of one mean minus the other, or M_1 minus M_2. Just like before, this is our observed effect and is the same value as the one we place in the numerator of our test statistic. We calculate this value then place the margin of error – still our critical value times our standard error – above and below it. That is:

$$\text{Confidence Interval } = (M_1 - M_2) \pm t^* \left(s_{M_1 - M_2}\right) \tag{11.7.2}$$

Because our hypothesis testing example used a one-tailed test, it would be inappropriate to calculate a confidence interval on those data (remember that we can only calculate a confidence interval for a two-tailed test because the interval extends in both directions). Let's say we find summary statistics on the average life satisfaction of people from two different towns and want to create a confidence interval to see if the difference between the two might actually be zero.

Our sample data are $M_1 = 28.65\ s_1 = 12.40\ n_1 = 40$ and $M_2 = 25.40 s_2 = 15.68 n_2 = 42$. At face value, it looks like the people from the first town have higher life satisfaction (28.65 vs. 25.40), but it will take a confidence interval (or complete hypothesis testing process) to see if that is true or just due to random chance. First, we want to calculate the difference between our sample means, which is 28.65 – 25.40 = 3.25. Next, we need a critical value from our *t*-table. If we want to test at the normal 95% level of confidence, then our sample sizes will yield degrees of freedom equal to 40 + 42 – 2 = 80. From our table, that gives us a critical value of $t* = 1.990$. Finally, we need our standard error. Recall that our standard error for an independent samples *t*-test uses pooled variance, which requires the Sum of Squares and degrees of freedom. Up to this point, we have calculated the Sum of Squares using raw data, but in this situation, we do not have access to it. So, what are we to do?

If we have summary data like standard deviation and sample size, it is very easy to calculate the pooled variance, and the key lies in rearranging the formulas to work backwards through them. We need the Sum of Squares and degrees of freedom to calculate our pooled variance. Degrees of freedom is very simple: we just take the sample size minus 1.00 for each group. Getting the Sum of Squares is also easy: remember that variance is standard deviation squared and is the Sum of Squares divided by the degrees of freedom. That is:

$$s^2 = (s)^2 = \frac{SS}{df} \tag{11.7.3}$$

To get the Sum of Squares, we just multiply both sides of the above equation to get:

$$s^2 * df = SS \tag{11.7.4}$$

Which is the squared standard deviation multiplied by the degrees of freedom $(n - 1)$ equals the Sum of Squares.

Using our example data:

$$(s_1)^2 * df_1 = SS_1$$
$$(12.40)^2 * (40 - 1) = 5996.64$$

$$(s_2)^2 * df_2 = SS_2$$
$$(15.68)^2 * (42 - 1) = 10080.36$$

And thus our pooled variance equals:

$$s_p^2 = \frac{SS_1 + SS_2}{df_1 + df_2} = \frac{5996.64 + 10080.36}{39 + 41} = \frac{16077}{80} = 200.96$$

And our standard error equals:

$$s_{M_1 - M_2} = \sqrt{\frac{s_p^2}{n_1} + \frac{s_p^2}{n_2}} = \sqrt{\frac{200.96}{40} + \frac{200.96}{42}} = \sqrt{5.02 + 4.78} = \sqrt{9.89} = 3.14$$

All of these steps are just slightly different ways of using the same formulae, numbers, and ideas we have worked with up to this point. Once we get out standard error, it's time to build our confidence interval.

$$95\%CI = 3.25 \pm 1.990(3.14)$$

$$\text{Upper Bound} = 3.25 + 1.990(3.14)$$
$$UB = 3.25 + 6.25$$
$$UB = 9.50$$

$$\text{Lower Bound } = 3.25 - 1.990(3.14)$$
$$LB = 3.25 - 6.25$$
$$LB = -3.00$$

$$95\%CI = (-3.00, 9.50)$$

Our confidence interval, as always, represents a range of values that would be considered reasonable or plausible based on our observed data. In this instance, our interval (-3.00, 9.50) does contain zero. Thus, even though the means look a little bit different, it may very well be the case that the life satisfaction in both of these towns is the same. Proving otherwise would require more data.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.7: Effect Sizes and Confidence Intervals is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.7: Effect Sizes and Confidence Intervals** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.8: Homogeneity of Variance

Before wrapping up the coverage of independent samples t-tests, there is one other important topic to cover. Using the pooled variance to calculate the test statistic relies on an assumption known as homogeneity of variance. In statistics, an assumption is some characteristic that we assume is true about our data, and our ability to use our inferential statistics accurately and correctly relies on these assumptions being true. If these assumptions are not true, then our analyses are at best ineffective (e.g. low power to detect effects) and at worst inappropriate (e.g. too many Type I errors). A detailed coverage of assumptions is beyond the scope of this course, but it is important to know that they exist for all analyses.

For the current analysis, one important assumption is homogeneity of variance. This is fancy statistical talk for the idea that the true population variance for each group is the same and any difference in the observed sample variances is due to random chance (if this sounds eerily similar to the idea of testing the null hypothesis that the true population means are equal, that's because it is exactly the same!) This notion allows us to compute a single pooled variance that uses our easily calculated degrees of freedom. If the assumption is shown to not be true, then we have to use a very complicated formula to estimate the proper degrees of freedom. There are formal tests to assess whether or not this assumption is met, but we will not discuss them here.

Many statistical programs incorporate the test of homogeneity of variance automatically and can report the results of the analysis assuming it is true or assuming it has been violated. You can easily tell which is which by the degrees of freedom: the corrected degrees of freedom (which is used when the assumption of homogeneity of variance is violated) will have decimal places. Fortunately, the independent samples t-test is very robust to violations of this assumption (an analysis is "robust" if it works well even when its assumptions are not met), which is why we do not bother going through the tedious work of testing and estimating new degrees of freedom by hand.

This page titled 11.8: Homogeneity of Variance is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.8: Homogeneity of Variance** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

11.E: Independent Samples (Exercises)

1. What is meant by “the difference of the means” when talking about an independent samples t-test? How does it differ from the “mean of the differences” in a repeated measures t-test?

Answer:

The difference of the means is one mean, calculated from a set of scores, compared to another mean which is calculated from a different set of scores; the independent samples t-test looks for whether the two separate values are different from one another. This is different than the “mean of the differences” because the latter is a single mean computed on a single set of difference scores that come from one data collection of matched pairs. So, the difference of the means deals with two numbers but the mean of the differences is only one number.

2. Describe three research questions that could be tested using an independent samples t-test.
3. Calculate pooled variance from the following raw data:

Group 1	Group 2
16	4
11	10
9	15
7	13
5	12
4	9
12	8

Answer:

$SS_1 = 106.86, SS_2 = 78.86, s_p^2 = 15.48$

4. Calculate the standard error from the following descriptive statistics
   a. s_1 = 24, s_2 = 21, n_1 = 36, n_2 = 49
   b. s_1 = 15.40, s_2 = 14.80, n_1 = 20, n_2 = 23
   c. s_1 = 12, s_2 = 10, n_1 = 25, n_2 = 25
5. Determine whether to reject or fail to reject the null hypothesis in the following situations:
   a. $t(40) = 2.49$, $\alpha = 0.01$, one-tailed test to the right
   b. $M_1 = 64, M_2 = 54, n_1 = 14, n_2 = 12, s_{M_1 - M_2} = 9.75, \alpha = 0.05$, two-tailed test
   c. 95% Confidence Interval: (0.50, 2.10)

Answer:

a. Reject
b. Fail to Reject
c. Reject

6. A professor is interest in whether or not the type of software program used in a statistics lab affects how well students learn the material. The professor teaches the same lecture material to two classes but has one class use a point-and-click software program in lab and has the other class use a basic programming language. The professor tests for a difference between the two classes on their final exam scores.

Point-and-Click	Programming
83	86
83	79
63	100
77	74
86	70
84	67
78	83
61	85
65	74
75	86
100	87
60	61
90	76
66	100
54	

7. A researcher wants to know if there is a difference in how busy someone is based on whether that person identifies as an early bird or a night owl. The researcher gathers data from people in each group, coding the data so that higher scores represent higher levels of being busy, and tests for a difference between the two at the .05 level of significance.

Early Bird	Night Owl
23	26
28	10
27	20
33	19
26	26
30	18
22	12
25	25
26	

Answer:

Step 1: $H_0 : \mu_1 - \mu_2 = 0$ “There is not difference in the average business of early birds versus night owls”, $H_A : \mu_1 - \mu_2 \neq 0$ “There is a difference in the average business of early birds versus night owls.”

Step 2: Two-tailed test, df = 15, $t*$ = 2.131.

Step 3: $M_1 = 26.67, M_2 = 19.50, s_p^2 = 27.73, s_{X_1} - M_2 = 2.37$

Step 4: $t > t*$, Reject H_0. Based on our data of early birds and night owls, we can conclude that early birds are busier ($M_1 = 26.67$) than night owls ($M_2 = 19.50$), $t(15) = 3.03$, $p < .05$. Since the result is significant, we need an effect size: Cohen's $d = 1.47$, which is a large effect.

8. Lots of people claim that having a pet helps lower their stress level. Use the following summary data to test the claim that there is a lower average stress level among pet owners (group 1) than among non-owners (group 2) at the .05 level of significance.

$$M_1 = 16.25, M_2 = 20.95, s_1 = 4.00, s_2 = 5.10, n_1 = 29, n_2 = 25$$

9. Administrators at a university want to know if students in different majors are more or less extroverted than others. They provide you with descriptive statistics they have for English majors (coded as 1) and History majors (coded as 2) and ask you to create a confidence interval of the difference between them. Does this confidence interval suggest that the students from the majors differ?

$$M_1 = 3.78, M_2 = 2.23, s_1 = 2.60, s_2 = 1.15, n_1 = 45, n_2 = 40$$

Answer:

$M_1 - M_2 = 1.55, \mathrm{t}^* = 1.990, s_{M_1 - M_2} = 0.45$, CI = (0.66, 2.44). This confidence interval does not contain zero, so it does suggest that there is a difference between the extroversion of English majors and History majors.

10. Researchers want to know if people's awareness of environmental issues varies as a function of where they live. The researchers have the following summary data from two states, Alaska and Hawaii, that they want to use to test for a difference.

$$M_H = 47.50, M_A = 45.70, \mathrm{s_H} = 14.65, \mathrm{s_A} = 13.20, \mathrm{n_H} = 139, \mathrm{n_A} = 150$$

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 11.E: Independent Samples (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **10.E: Independent Samples (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

12: Analysis of Variance

12.1: Observing and Interpreting Variability
12.2: Sources of Variance
12.3: ANOVA Table
12.4: ANOVA and Type I Error
12.5: Hypotheses in ANOVA
12.6: Scores on Job Application Tests
12.7: Variance Explained
12.8: Post Hoc Tests
12.9: Other ANOVA Designs
12.10: Analysis of Variance (Exercises)

This page titled 12: Analysis of Variance is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

12.1: Observing and Interpreting Variability

We have seen time and again that scores, be they individual data or group means, will differ naturally. Sometimes this is due to random chance, and other times it is due to actual differences. Our job as scientists, researchers, and data analysts is to determine if the observed differences are systematic and meaningful (via a hypothesis test) and, if so, what is causing those differences. Through this, it becomes clear that, although we are usually interested in the mean or average score, it is the variability in the scores that is key.

Take a look at Figure **12.1.1**, which shows scores for many people on a test of skill used as part of a job application. The x-axis has each individual person, in no particular order, and the y-axis contains the score each person received on the test. As we can see, the job applicants differed quite a bit in their performance, and understanding why that is the case would be extremely useful information. However, there's no interpretable pattern in the data, especially because we only have information on the test, not on any other variable (remember that the x-axis here only shows individual people and is not ordered or interpretable).

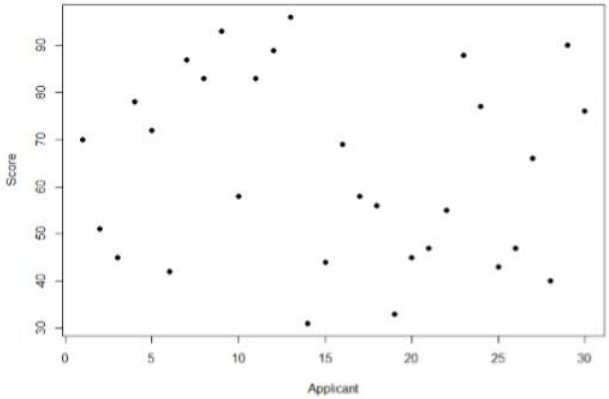


Figure **12.1.1**: Scores on a job test

Our goal is to explain this variability that we are seeing in the dataset. Let's assume that as part of the job application procedure we also collected data on the highest degree each applicant earned. With knowledge of what the job requires, we could sort our applicants into three groups: those applicants who have a college degree related to the job, those applicants who have a college degree that is not related to the job, and those applicants who did not earn a college degree. This is a common way that job applicants are sorted, and we can use ANOVA to test if these groups are actually different. Figure **12.1.2** presents the same job applicant scores, but now they are color coded by group membership (i.e. which group they belong in). Now that we can differentiate between applicants this way, a pattern starts to emerge: those applicants with a relevant degree (coded red) tend to be near the top, those applicants with no college degree (coded black) tend to be near the bottom, and the applicants with an unrelated degree (coded green) tend to fall into the middle. However, even within these groups, there is still some variability, as shown in Figure **12.1.2**.

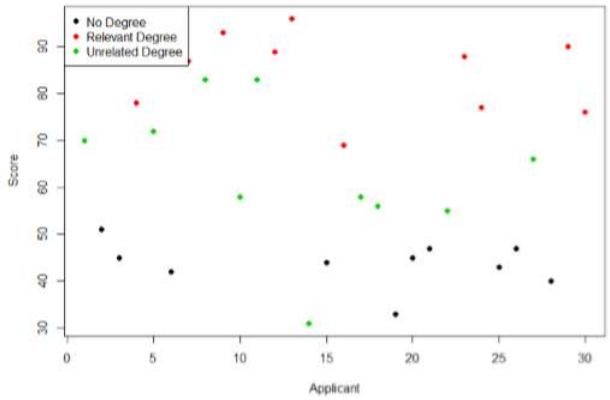


Figure **12.1.2**: Applicant scores coded by degree earned

This pattern is even easier to see when the applicants are sorted and organized into their respective groups, as shown in Figure **12.1.3**.

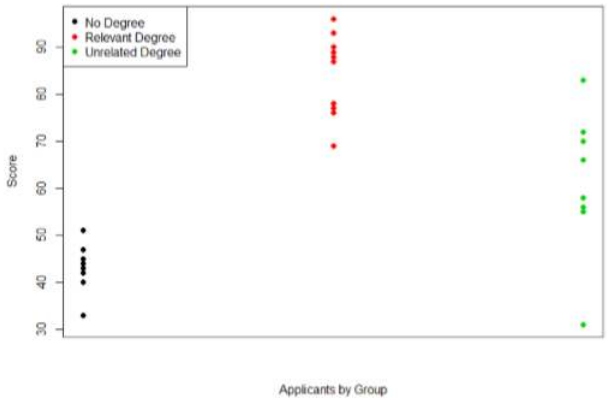


Figure **12.1.3**: Applicant scores by group

Now that we have our data visualized into an easily interpretable format, we can clearly see that our applicants' scores differ largely along group lines. Those applicants who do not have a college degree received the lowest scores, those who had a degree relevant to the job received the highest scores, and those who did have a degree but one that is not related to the job tended to fall somewhere in the middle. Thus, we have systematic variance between our groups.

We can also clearly see that within each group, our applicants' scores differed from one another. Those applicants without a degree tended to score very similarly, since the scores are clustered close together. Our group of applicants with relevant degrees varied a little but more than that, and our group of applicants with unrelated degrees varied quite a bit. It may be that there are other factors that cause the observed score differences within each group, or they could just be due to random chance. Because we do not have any other explanatory data in our dataset, the variability we observe within our groups is considered random error, with any deviations between a person and that person's group mean caused only by chance. Thus, we have unsystematic (random) variance within our groups.

The process and analyses used in ANOVA will take these two sources of variance (systematic variance between groups and random error within groups, or how much groups differ from each other and how much people differ within each group) and compare them to one another to determine if the groups have any explanatory value in our outcome variable. By doing this, we will test for statistically significant differences between the group means, just like we did for ***t***-tests. We will go step by step to break down the math to see how ANOVA actually works.

This page titled 12.1: Observing and Interpreting Variability is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.1: Observing and Interpreting Variability** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.2: Sources of Variance

ANOVA is all about looking at the different sources of variance (i.e. the reasons that scores differ from one another) in a dataset. Fortunately, the way we calculate these sources of variance takes a very familiar form: the Sum of Squares. Before we get into the calculations themselves, we must first lay out some important terminology and notation.

In ANOVA, we are working with two variables: a grouping or explanatory variable and a continuous outcome variable. The grouping variable is our predictor (it predicts or explains the values in the outcome variable) or, in experimental terms, our independent variable, and it is made up of k groups, with k being any whole number 2 or greater. That is, ANOVA requires two or more groups to work, and it is usually conducted with three or more groups. In ANOVA, we refer to groups as "levels", so the number of levels is just the number of groups, which again is k. In the above example, our grouping variable was education, which had 3 levels, so $k = 3$. When we report any descriptive value (e.g. mean, sample size, standard deviation) for a specific group, we will use a subscript $1 \ldots k$ to denote which group it refers to. For example, if we have three groups and want to report the standard deviation s for each group, we would report them as s_1, s_2, and s_3.

Our second variable is our outcome variable. This is the variable on which people differ, and we are trying to explain or account for those differences based on group membership. In the example above, our outcome was the score each person earned on the test. Our outcome variable will still use X for scores as before. When describing the outcome variable using means, we will use subscripts to refer to specific group means. So if we have $k = 3$ groups, our means will be M_1, M_2, and M_3. We will also have a single mean representing the average of all participants across all groups. This is known as the grand mean, and we use the symbol GM. These different means – the individual group means and the overall grand mean – will be how we calculate our sums of squares.

Finally, we now have to differentiate between several different sample sizes. Our data will now have sample sizes for each group, and we will denote these with a lower case "n" and a subscript, just like with our other descriptive statistics: n_1, n_2, and n_3. We also have the overall sample size in our dataset, and we will denote this with a capital N. The total sample size is just the group sample sizes added together.

Between Groups Sum of Squares

One source of variability we can identified in 11.1.3 of the above example was differences or variability between the groups. That is, the groups clearly had different average levels. The variability arising from these differences is known as the between groups variability, and it is quantified using Between Groups Sum of Squares.

Our calculations for sums of squares in ANOVA will take on the same form as it did for regular calculations of variance. Each observation, in this case the group means, is compared to the overall mean, in this case the grand mean, to calculate a deviation score. These deviation scores are squared so that they do not cancel each other out and sum to zero. The squared deviations are then added up, or summed. There is, however, one small difference. Because each group mean represents a group composed of multiple people, before we sum the deviation scores we must multiply them by the number of people within that group. Incorporating this, we find our equation for Between Groups Sum of Squares to be:

$$SS_B = \sum n_j (M_J - GM)^2 \tag{12.2.1}$$

The subscript j refers to the "j^{th}" group where $j = 1 \ldots k$ to keep track of which group mean and sample size we are working with. As you can see, the only difference between this equation and the familiar sum of squares for variance is that we are adding in the sample size. Everything else logically fits together in the same way.

Within Groups Sum of Squares

The other source of variability in the figures comes from differences that occur within each group. That is, each individual deviates a little bit from their respective group mean, just like the group means differed from the grand mean. We therefore label this source the Within Groups Sum of Squares. Because we are trying to account for variance based on group-level means, any deviation from the group means indicates an inaccuracy or error. Thus, our within groups variability represents our error in ANOVA.

The formula for this sum of squares is again going to take on the same form and logic. What we are looking for is the distance between each individual person and the mean of the group to which they belong. We calculate this deviation score, square it so that

they can be added together, then sum all of them into one overall value:

$$SS_W = \sum (X_{ij} - M_j)^2 \tag{12.2.2}$$

In this instance, because we are calculating this deviation score for each individual person, there is no need to multiply by how many people we have. The subscript j again represents a group and the subscript i refers to a specific person. So, X_{ij} is read as "the ith person of the jth group." It is important to remember that the deviation score for each person is only calculated relative to their group mean: do not calculate these scores relative to the other group means.

Total Sum of Squares

The Between Groups and Within Groups Sums of Squares represent all variability in our dataset. We also refer to the total variability as the Total Sum of Squares, representing the overall variability with a single number. The calculation for this score is exactly the same as it would be if we were calculating the overall variance in the dataset (because that's what we are interested in explaining) without worrying about or even knowing about the groups into which our scores fall:

$$SS_T = \sum (X_i - GM)^2 \tag{12.2.3}$$

We can see that our Total Sum of Squares is just each individual score minus the grand mean. As with our Within Groups Sum of Squares, we are calculating a deviation score for each individual person, so we do not need to multiply anything by the sample size; that is only done for Between Groups Sum of Squares.

An important feature of the sums of squares in ANOVA is that they all fit together. We could work through the algebra to demonstrate that if we added together the formulas for SS_B and SS_W, we would end up with the formula for SS_T. That is:

$$SS_T = SS_B + SS_W \tag{12.2.4}$$

This will prove to be very convenient, because if we know the values of any two of our sums of squares, it is very quick and easy to find the value of the third. It is also a good way to check calculations: if you calculate each SS by hand, you can make sure that they all fit together as shown above, and if not, you know that you made a math mistake somewhere.

We can see from the above formulas that calculating an ANOVA by hand from raw data can take a very, very long time. For this reason, you will not be required to calculate the SS values by hand, but you should still take the time to understand how they fit together and what each one represents to ensure you understand the analysis itself.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 12.2: Sources of Variance is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.2: Sources of Variance** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.3: ANOVA Table

All of our sources of variability fit together in meaningful, interpretable ways as we saw above, and the easiest way to do this is to organize them into a table. The ANOVA table, shown in Table **12.3.1**, is how we calculate our test statistic.

Table **12.3.1**: ANOVA Table

Source	SS	df	MS	F
Between	SS_B	$k-1$	$\frac{SS_B}{df_B}$	$\frac{MS_B}{MS_W}$
Within	SS_W	$N-k$	$\frac{SS_W}{df_W}$	
Total	SS_T	$N-1$		

The first column of the ANOVA table, labeled "Source", indicates which of our sources of variability we are using: between groups, within groups, or total. The second column, labeled "SS", contains our values for the sums of squares that we learned to calculate above. As noted previously, calculating these by hand takes too long, and so the formulas are not presented in Table **12.3.1**. However, remember that the Total is the sum of the other two, in case you are only given two SS values and need to calculate the third.

The next column, labeled "df", is our degrees of freedom. As with the sums of squares, there is a different df for each group, and the formulas are presented in the table. Notice that the total degrees of freedom, N–1, is the same as it was for our regular variance. This matches the SS_T formulation to again indicate that we are simply taking our familiar variance term and breaking it up into difference sources. Also remember that the capital N in the df calculations refers to the overall sample size, not a specific group sample size. Notice that the total row for degrees of freedom, just like for sums of squares, is just the Between and Within rows added together. If you take $N–k+k–1$, then the "$–k$" and "$+k$" portions will cancel out, and you are left with N–1. This is a convenient way to quickly check your calculations.

The third column, labeled "MS", is our Mean Squares for each source of variance. A "mean square" is just another way to say variability. Each mean square is calculated by dividing the sum of squares by its corresponding degrees of freedom. Notice that we do this for the Between row and the Within row, but not for the Total row. There are two reasons for this. First, our Total Mean Square would just be the variance in the full dataset (put together the formulas to see this for yourself), so it would not be new information. Second, the Mean Square values for Between and Within would not add up to equal the Mean Square Total because they are divided by different denominators. This is in contrast to the first two columns, where the Total row was both the conceptual total (i.e. the overall variance and degrees of freedom) and the literal total of the other two rows.

The final column in the ANOVA table, labeled "F", is our test statistic for ANOVA. The F statistic, just like a t- or z-statistic, is compared to a critical value to see whether we can reject for fail to reject a null hypothesis. Thus, although the calculations look different for ANOVA, we are still doing the same thing that we did in all of Unit 2. We are simply using a new type of data to test our hypotheses. We will see what these hypotheses look like shortly, but first, we must take a moment to address why we are doing our calculations this way.

This page titled 12.3: ANOVA Table is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.3: ANOVA Table** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.4: ANOVA and Type I Error

You may be wondering why we do not just use another t-test to test our hypotheses about three or more groups the way we did in Unit 2. After all, we are still just looking at group mean differences. The reason is that our t-statistic formula can only handle up to two groups, one minus the other. With only two groups, we can move our population parameters for the group means around in our null hypothesis and still get the same interpretation: the means are equal, which can also be concluded if one mean minus the other mean is equal to zero. However, if we tried adding a third mean, we would no longer be able to do this. So, in order to use t-tests to compare three or more means, we would have to run a series of individual group comparisons.

For only three groups, we would have three t-tests: group 1 vs group 2, group 1 vs group 3, and group 2 vs group 3. This may not sound like a lot, especially with the advances in technology that have made running an analysis very fast, but it quickly scales up. With just one additional group, bringing our total to four, we would have six comparisons: group 1 vs group 2, group 1 vs group 3, group 1 vs group 4, group 2 vs group 3, group 2 vs group 4, and group 3 vs group 4. This makes for a logistical and computation nightmare for five or more groups.

A bigger issue, however, is our probability of committing a Type I Error. Remember that a Type I error is a false positive, and the chance of committing a Type I error is equal to our significance level, α. This is true if we are only running a single analysis (such as a t-test with only two groups) on a single dataset. However, when we start running multiple analyses on the same dataset, our Type I error rate increases, raising the probability that we are capitalizing on random chance and rejecting a null hypothesis when we should not. ANOVA, by comparing all groups simultaneously with a single analysis, averts this issue and keeps our error rate at the α we set.

This page titled 12.4: ANOVA and Type I Error is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.4: ANOVA and Type I Error** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.5: Hypotheses in ANOVA

So far we have seen what ANOVA is used for, why we use it, and how we use it. Now we can turn to the formal hypotheses we will be testing. As with before, we have a null and an alternative hypothesis to lay out. Our null hypothesis is still the idea of "no difference" in our data. Because we have multiple group means, we simply list them out as equal to each other:

$$\mathbf{H_0: \text{ There is no difference in the group means}}$$
$$\mathbf{H_0: \mu_1 = \mu_2 = \mu_3}$$

We list as many μ parameters as groups we have. In the example above, we have three groups to test, so we have three parameters in our null hypothesis. If we had more groups, say, four, we would simply add another μ to the list and give it the appropriate subscript, giving us:

$$\mathbf{H_0: \text{ There is no difference in the group means}}$$
$$\mathbf{H_0: \mu_1 = \mu_2 = \mu_3 = \mu_4}$$

Notice that we do not say that the means are all equal to zero, we only say that they are equal to one another; it does not matter what the actual value is, so long as it holds for all groups equally.

Our alternative hypothesis for ANOVA is a little bit different. Let's take a look at it and then dive deeper into what it means:

$$\mathbf{H_A: \text{ At least one mean is different}}$$

The first difference in obvious: there is no mathematical statement of the alternative hypothesis in ANOVA. This is due to the second difference: we are not saying which group is going to be different, only that at least one will be. Because we do not hypothesize about which mean will be different, there is no way to write it mathematically. Related to this, we do not have directional hypotheses (greater than or less than) like we did in Unit 2. Due to this, our alternative hypothesis is always exactly the same: at least one mean is different.

In Unit 2, we saw that, if we reject the null hypothesis, we can adopt the alternative, and this made it easy to understand what the differences looked like. In ANOVA, we will still adopt the alternative hypothesis as the best explanation of our data if we reject the null hypothesis. However, when we look at the alternative hypothesis, we can see that it does not give us much information. We will know that a difference exists somewhere, but we will not know where that difference is. Is only group 1 different but groups 2 and 3 the same? Is it only group 2? Are all three of them different? Based on just our alternative hypothesis, there is no way to be sure. We will come back to this issue later and see how to find out specific differences. For now, just remember that we are testing for any difference in group means, and it does not matter where that difference occurs.

Now that we have our hypotheses for ANOVA, let's work through an example. We will continue to use the data from Figures 11.1.1 through 11.1.3 for continuity.

This page titled 12.5: Hypotheses in ANOVA is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.5: Hypotheses in ANOVA** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.6: Scores on Job Application Tests

Our data come from three groups of 10 people each, all of whom applied for a single job opening: those with no college degree, those with a college degree that is not related to the job opening, and those with a college degree from a relevant field. We want to know if we can use this group membership to account for our observed variability and, by doing so, test if there is a difference between our three group means. We will start, as always, with our hypotheses.

Step 1: State the Hypotheses

Our hypotheses are concerned with the test scores of groups based on education level, so:

$$\mathbf{H_0:}\ \textbf{There is no difference between the test scores of the education groups}$$
$$\mathbf{H_0}: \mu_1 = \mu_2 = \mu_3$$

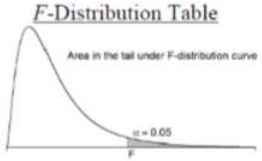
$$\mathbf{H_A:}\ \textbf{At least one group of test scores is different from the rest}$$

Again, we phrase our null hypothesis in terms of what we are actually looking for, and we use a number of population parameters equal to our number of groups.

Step 2: Find the Critical Values

Our test statistic for ANOVA, as we saw above, is F. Because we are using a new test statistic, we will get a new table: the F distribution table, the top of which is shown in Figure **12.6.1**:

F-Distribution Table



df denom.	Degrees of Freedom: Numerator 1	2	3	4	5	6	7	8	9	10
1	161.45	199.50	215.71	224.58	230.16	233.99	236.77	238.88	240.54	241.88
2	18.51	19.00	19.16	19.25	19.30	19.33	19.35	19.37	19.39	19.40
3	10.13	9.55	9.28	9.12	9.01	8.94	8.89	8.85	8.81	8.79
4	7.71	6.94	6.59	6.39	6.26	6.16	6.09	6.04	6.00	5.96
5	6.61	5.79	5.41	5.19	5.05	4.95	4.88	4.82	4.77	4.74
6	5.99	5.14	4.76	4.53	4.39	4.28	4.21	4.15	4.10	4.06
7	5.59	4.74	4.35	4.12	3.97	3.87	3.79	3.73	3.68	3.64
8	5.32	4.46	4.07	3.84	3.69	3.58	3.50	3.44	3.39	3.35
9	5.12	4.26	3.86	3.63	3.48	3.37	3.29	3.23	3.18	3.14
10	4.97	4.10	3.71	3.48	3.33	3.22	3.14	3.07	3.02	2.98
11	4.84	3.98	3.59	3.36	3.20	3.10	3.01	2.95	2.90	2.85
12	4.75	3.89	3.49	3.26	3.11	3.00	2.91	2.85	2.80	2.75
13	4.67	3.81	3.41	3.18	3.03	2.92	2.83	2.77	2.71	2.67
14	4.60	3.74	3.34	3.11	2.96	2.85	2.76	2.70	2.65	2.60
15	4.54	3.68	3.29	3.06	2.90	2.79	2.71	2.64	2.59	2.54
16	4.49	3.63	3.24	3.01	2.85	2.74	2.66	2.59	2.54	2.49
17	4.45	3.59	3.20	2.97	2.81	2.70	2.61	2.55	2.49	2.45
18	4.41	3.56	3.16	2.93	2.77	2.66	2.58	2.51	2.46	2.41
19	4.38	3.52	3.13	2.90	2.74	2.63	2.54	2.48	2.42	2.38
20	4.35	3.49	3.10	2.87	2.71	2.60	2.51	2.45	2.39	2.35
21	4.33	3.47	3.07	2.84	2.69	2.57	2.49	2.42	2.37	2.32
22	4.30	3.44	3.05	2.82	2.66	2.55	2.46	2.40	2.34	2.30
23	4.28	3.42	3.03	2.80	2.64	2.53	2.44	2.38	2.32	2.28
24	4.26	3.40	3.01	2.78	2.62	2.51	2.42	2.36	2.30	2.26
25	4.24	3.39	2.99	2.76	2.60	2.49	2.41	2.34	2.28	2.24
26	4.23	3.37	2.98	2.74	2.59	2.47	2.39	2.32	2.27	2.22
27	4.21	3.35	2.96	2.73	2.57	2.46	2.37	2.31	2.25	2.20

Figure **12.6.1**: F distribution table.

The F table only displays critical values for α = 0.05. This is because other significance levels are uncommon and so it is not worth it to use up the space to present them. There are now two degrees of freedom we must use to find our critical value: Numerator and Denominator. These correspond to the numerator and denominator of our test statistic, which, if you look at the ANOVA table presented earlier, are our Between Groups and Within Groups rows, respectively. The df_B is the "Degrees of Freedom: Numerator" because it is the degrees of freedom value used to calculate the Mean Square Between, which in turn was the numerator of our F statistic. Likewise, the df_W is the "df denom." (short for denominator) because it is the degrees of freedom value used to calculate the Mean Square Within, which was our denominator for F.

The formula for df_B is k–1, and remember that k is the number of groups we are assessing. In this example, $k = 3$ so our df_B = 2. This tells us that we will use the second column, the one labeled 2, to find our critical value. To find the proper row, we simply

calculate the df_W, which was $N–k$. The original prompt told us that we have "three groups of 10 people each," so our total sample size is 30. This makes our value for df_W = 27. If we follow the second column down to the row for 27, we find that our critical value is 3.35. We use this critical value the same way as we did before: it is our criterion against which we will compare our obtained test statistic to determine statistical significance.

Step 3: Calculate the Test Statistic

Now that we have our hypotheses and the criterion we will use to test them, we can calculate our test statistic. To do this, we will fill in the ANOVA table. When we do so, we will work our way from left to right, filling in each cell to get our final answer. We will assume that we are given the SS values as shown below:

Table 12.6.1: ANOVA Table

Source	SS	df	MS	F
Between	8246			
Within	3020			
Total				

These may seem like random numbers, but remember that they are based on the distances between the groups themselves and within each group. Figure 12.6.2 shows the plot of the data with the group means and grand mean included. If we wanted to, we could use this information, combined with our earlier information that each group has 10 people, to calculate the Between Groups Sum of Squares by hand. However, doing so would take some time, and without the specific values of the data points, we would not be able to calculate our Within Groups Sum of Squares, so we will trust that these values are the correct ones.

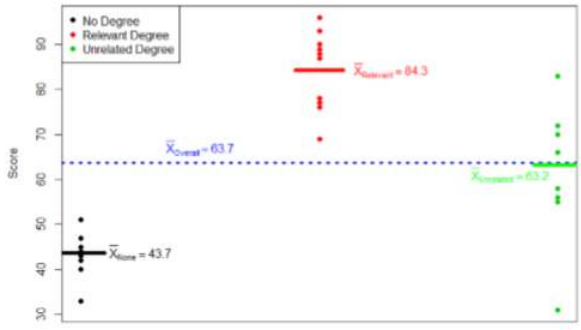


Figure 12.6.2: Means

We were given the sums of squares values for our first two rows, so we can use those to calculate the Total Sum of Squares.

Table 12.6.2: Total Sum of Squares

Source	SS	df	MS	F
Between	8246			
Within	3020			
Total	11266			

We also calculated our degrees of freedom earlier, so we can fill in those values. Additionally, we know that the total degrees of freedom is $N–1$, which is 29. This value of 29 is also the sum of the other two degrees of freedom, so everything checks out.

Table 12.6.3: Total Sum of Squares

Source	SS	df	MS	F
Between	8246	2		
Within	3020	27		
Total	11266	29		

Now we have everything we need to calculate our mean squares. Our MS values for each row are just the SS divided by the df for that row, giving us:

Table **12.6.4**: Total Sum of Squares

Source	SS	df	MS	F
Between	8246	2	4123	
Within	3020	27	111.85	
Total	11266	29		

Remember that we do not calculate a Total Mean Square, so we leave that cell blank. Finally, we have the information we need to calculate our test statistic. F is our MS_B divided by MS_W.

Table **12.6.5**: Total Sum of Squares

Source	SS	df	MS	F
Between	8246	2	4123	36.86
Within	3020	27	111.85	
Total	11266	29		

So, working our way through the table given only two SS values and the sample size and group size given before, we calculate our test statistic to be $F_{obt} = 36.86$, which we will compare to the critical value in step 4.

Step 4: Make the Decision

Our obtained test statistic was calculated to be $F_{obt} = 36.86$ and our critical value was found to be $F^* = 3.35$. Our obtained statistic is larger than our critical value, so we can reject the null hypothesis.

Reject H_0. Based on our 3 groups of 10 people, we can conclude that job test scores are statistically significantly different based on education level, $F(2, 27) = 36.86, p < .05$.

Notice that when we report F, we include both degrees of freedom. We always report the numerator then the denominator, separated by a comma. We must also note that, because we were only testing for any difference, we cannot yet conclude which groups are different from the others. We will do so shortly, but first, because we found a statistically significant result, we need to calculate an effect size to see how big of an effect we found.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 12.6: Scores on Job Application Tests is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.6: Scores on Job Application Tests** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.7: Variance Explained

Recall that the purpose of ANOVA is to take observed variability and see if we can explain those differences based on group membership. To that end, our effect size will be just that: the variance explained. You can think of variance explained as the proportion or percent of the differences we are able to account for based on our groups. We know that the overall observed differences are quantified as the Total Sum of Squares, and that our observed effect of group membership is the Between Groups Sum of Squares. Our effect size, therefore, is the ratio of these to sums of squares. Specifically:

$$\eta^2 = \frac{SS_B}{SS_T} \tag{12.7.1}$$

The effect size η^2 is called "eta-squared" and represents variance explained. For our example, our values give an effect size of:

$$\eta^2 = \frac{8246}{11266} = 0.73$$

So, we are able to explain 73% of the variance in job test scores based on education. This is, in fact, a huge effect size, and most of the time we will not explain nearly that much variance. Our guidelines for the size of our effects are:

Table **12.7.1**: Guidelines for the size of our effects

η^2	Size
0.01	Small
0.09	Medium
0.25	Large

So, we found that not only do we have a statistically significant result, but that our observed effect was very large! However, we still do not know specifically which groups are different from each other. It could be that they are all different, or that only those who have a relevant degree are different from the others, or that only those who have no degree are different from the others. To find out which is true, we need to do a special analysis called a post hoc test.

This page titled 12.7: Variance Explained is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.7: Variance Explained** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.8: Post Hoc Tests

A post hoc test is used only after we find a statistically significant result and need to determine where our differences truly came from. The term “post hoc” comes from the Latin for “after the event”. There are many different post hoc tests that have been developed, and most of them will give us similar answers. We will only focus here on the most commonly used ones. We will also only discuss the concepts behind each and will not worry about calculations.

Bonferroni Test

A Bonferroni test is perhaps the simplest post hoc analysis. A Bonferroni test is a series of t-tests performed on each pair of groups. As we discussed earlier, the number of groups quickly grows the number of comparisons, which inflates Type I error rates. To avoid this, a Bonferroni test divides our significance level α by the number of comparisons we are making so that when they are all run, they sum back up to our original Type I error rate. Once we have our new significance level, we simply run independent samples t-tests to look for differences between our pairs of groups. This adjustment is sometimes called a Bonferroni Correction, and it is easy to do by hand if we want to compare obtained p-values to our new corrected α level, but it is more difficult to do when using critical values like we do for our analyses so we will leave our discussion of it to that.

Tukey’s Honest Significant Difference

Tukey’s Honest Significant Difference (HSD) is a very popular post hoc analysis. This analysis, like Bonferroni’s, makes adjustments based on the number of comparisons, but it makes adjustments to the test statistic when running the comparisons of two groups. These comparisons give us an estimate of the difference between the groups and a confidence interval for the estimate. We use this confidence interval in the same way that we use a confidence interval for a regular independent samples t-test: if it contains 0.00, the groups are not different, but if it does not contain 0.00 then the groups are different.

Below are the differences between the group means and the Tukey’s HSD confidence intervals for the differences:

Table **12.8.1**: Differences between the group means and the Tukey’s HSD confidence intervals

Comparison	Difference	Tukey’s HSD CI
None vs Relevant	40.60	(28.87, 52.33)
None vs Unrelated	19.50	(7.77, 31.23)
Relevant vs Unrelated	21.10	(9.37, 32.83)

As we can see, none of these intervals contain 0.00, so we can conclude that all three groups are different from one another.

Scheffe’s Test

Another common post hoc test is Scheffe’s Test. Like Tukey’s HSD, Scheffe’s test adjusts the test statistic for how many comparisons are made, but it does so in a slightly different way. The result is a test that is “conservative,” which means that it is less likely to commit a Type I Error, but this comes at the cost of less power to detect effects. We can see this by looking at the confidence intervals that Scheffe’s test gives us:

Table **12.8.2**: Confidence intervals given by Scheffe’s test

Comparison	Difference	Tukey’s HSD CI
None vs Relevant	40.60	(28.35, 52.85)
None vs Unrelated	19.50	(7.25, 31.75)
Relevant vs Unrelated	21.10	(8.85, 33.35)

As we can see, these are slightly wider than the intervals we got from Tukey’s HSD. This means that, all other things being equal, they are more likely to contain zero. In our case, however, the results are the same, and we again conclude that all three groups

differ from one another.

There are many more post hoc tests than just these three, and they all approach the task in different ways, with some being more conservative and others being more powerful. In general, though, they will give highly similar answers. What is important here is to be able to interpret a post hoc analysis. If you are given post hoc analysis confidence intervals, like the ones seen above, read them the same way we read confidence intervals in chapter 10: if they contain zero, there is no difference; if they do not contain zero, there is a difference.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 12.8: Post Hoc Tests is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.8: Post Hoc Tests** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.9: Other ANOVA Designs

We have only just scratched the surface on ANOVA in this chapter. There are many other variations available for the one-way ANOVA presented here. There are also other types of ANOVAs that you are likely to encounter. The first is called a factorial ANOVA. Factorial ANOVAs use multiple grouping variables, not just one, to look for group mean differences. Just as there is no limit to the number of groups in a one-way ANOVA, there is no limit to the number of grouping variables in a Factorial ANOVA, but it becomes very difficult to find and interpret significant results with many factors, so usually they are limited to two or three grouping variables with only a small number of groups in each. Another ANOVA is called a Repeated Measures ANOVA. This is an extension of a repeated measures or matched pairs t-test, but in this case we are measuring each person three or more times to look for a change. We can even combine both of these advanced ANOVAs into mixed designs to test very specific and valuable questions. These topics are far beyond the scope of this text, but you should know about their existence. Our treatment of ANOVA here is a small first step into a much larger world!

This page titled 12.9: Other ANOVA Designs is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.9: Other ANOVA Designs** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

12.10: Analysis of Variance (Exercises)

1. What are the three pieces of variance analyzed in ANOVA?

Answer:

Variance between groups (SSB), variance within groups (SSW) and total variance (SST).

2. What does rejecting the null hypothesis in ANOVA tell us? What does it not tell us?
3. What is the purpose of post hoc tests?

Answer:

Post hoc tests are run if we reject the null hypothesis in ANOVA; they tell us which specific group differences are significant.

4. Based on the ANOVA table below, do you reject or fail to reject the null hypothesis? What is the effect size?

Source	SS	df	MS	F
Between	60.72	3	20.24	3.88
Within	213.61	41	5.21	
Total	274.33	44		

5. Finish filling out the following ANOVA tables:

a. $K = 4$

Source	SS	df	MS	F
Between	87.40			
Within				
Total	199.22	33		

b. $N = 14$

Source	SS	df	MS	F
Between		2	14.10	
Within				
Total	64.65			

c.

Source	SS	df	MS	F
Between		2		42.36
Within		54	2.48	
Total				

Answer:

a. $K = 4$

Source	SS	df	MS	F
Between	87.40	3	29.13	7.81
Within	111.82	30	3.73	
Total	199.22	33		

b. $N = 14$

Source	SS	df	MS	F
Between	28.20	2	14.10	4.26
Within	36.45	11	3.31	
Total	64.65	13		

c.

Source	SS	df	MS	F
Between	210.10	2	105.05	42.36
Within	133.92	54	2.48	
Total	344.02			

6. You know that stores tend to charge different prices for similar or identical products, and you want to test whether or not these differences are, on average, statistically significantly different. You go online and collect data from 3 different stores, gathering information on 15 products at each store. You find that the average prices at each store are: Store 1 xbar = $27.82, Store 2 xbar = $38.96, and Store 3 xbar = $24.53. Based on the overall variability in the products and the variability within each store, you find the following values for the Sums of Squares: SST = 683.22, SSW = 441.19. Complete the ANOVA table and use the 4 step hypothesis testing procedure to see if there are systematic price differences between the stores.
7. You and your friend are debating which type of candy is the best. You find data on the average rating for hard candy (e.g. jolly ranchers, M= 3.60), chewable candy (e.g. starburst, M = 4.20), and chocolate (e.g. snickers, M= 4.40); each type of candy was rated by 30 people. Test for differences in average candy rating using SSB = 16.18 and SSW = 28.74.

Answer:

Step 1: $H_0 : \mu_1 = \mu_2 = \mu_3$ "There is no difference in average rating of candy quality", H_A: "At least one mean is different."

Step 2: 3 groups and 90 total observations yields $df_{num} = 2$ and $df_{den} = 87$, $\alpha = 0.05$, $F^* = 3.11$.

Step 3: based on the given SSB and SSW and the computed df from step 2, is:

Source	SS	df	MS	F
Between	16.18	2	8.09	24.52
Within	28.74	87	0.33	
Total	44.92	89		

Step 4: $F > F^*$, reject H_0. Based on the data in our 3 groups, we can say that there is a statistically significant difference in the quality of different types of candy, $F(2, 87) = 24.52, p < .05$. Since the result is significant, we need an effect size: $\eta^2 = 16.18/44.92 = .36$, which is a large effect

8. Administrators at a university want to know if students in different majors are more or less extroverted than others. They provide you with data they have for English majors (M = 3.78, n = 45), History majors (M= 2.23, n = 40), Psychology majors (M= 4.41, n = 51), and Math majors (M = 1.15, n = 28). You find the $SSB = 75.80$ and $SSW = 47.40$ and test at $\alpha = 0.05$.
9. You are assigned to run a study comparing a new medication (M= 17.47, n = 19), an existing medication (M= 17.94, n = 18), and a placebo (M= 13.70, n = 20), with higher scores reflecting better outcomes. Use $SSB = 210.10$ and $SSW = 133.90$ to test for differences.

Answer:

Step 1: $H_0 : \mu_1 = \mu_2 = \mu_3$ "There is no difference in average outcome based on treatment", H_A: "At least one mean is different."

Step 2: 3 groups and 57 total participants yields $df_{num} = 2$ and $df_{den} = 54$, $\alpha = 0.05, F^* = 3.18$.

Step 3: based on the given SSB and SSW and the computed df from step 2, is:

Source	SS	df	MS	F
Between	210.10	2	105.02	42.36
Within	133.90	54	2.48	
Total	344.00	56		

Step 4: $F > F^*$, reject H_0. Based on the data in our 3 groups, we can say that there is a statistically significant difference in the effectiveness of the treatments, $F(2, 54) = 42.36, p < .05$. Since the result is significant, we need an effect size: $\eta^2 = 210.10/344.00 = .61$, which is a large effect.

10. You are in charge of assessing different training methods for effectiveness. You have data on 4 methods: Method 1 (M = 87, n = 12), Method 2 (M= 92, n = 14), Method 3 (M = 88, n = 15), and Method 4 (M= 75, n = 11). Test for differences among these means, assuming $SSB = 64.81$ and $SST = 399.45$.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 12.10: Analysis of Variance (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **11.1E: Analysis of Variance (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

13: Correlations

A common theme throughout statistics is the notion that individuals will differ on different characteristics and traits, which we call variance. In inferential statistics and hypothesis testing, our goal is to find systematic reasons for differences and rule out random chance as the cause. By doing this, we are using information on a different variable – which so far has been group membership like in ANOVA – to explain this variance. In correlations, we will instead use a continuous variable to account for the variance.

13.1: Variability and Covariance
13.2: Visualizing Relations
13.3: Three Characteristics
13.4: Pearson's r
13.5: Anxiety and Depression
13.6: Effect Size
13.7: Correlation versus Causation
13.8: Final Considerations
13.E: Correlations (Exercises)

Thumbnail: Correlation shown when the two variables' ranges are unrestricted, and when the range of is restricted to the interval (0,1). (CC BY 3.0 Unported; Skbkekas via Wikipedia)

This page titled 13: Correlations is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

13.1: Variability and Covariance

Because we have two continuous variables, we will have two characteristics or scores on which people will vary. What we want to know is do people vary on the scores together. That is, as one score changes, does the other score also change in a predictable or consistent way? This notion of variables differing together is called covariance (the prefix "co" meaning "together").

Let's look at our formula for variance on a single variable:

$$s^2 = \frac{\sum (X - M)^2}{N - 1} \tag{13.1.1}$$

We use X to represent a person's score on the variable at hand, and M to represent the mean of that variable. The numerator of this formula is the Sum of Squares, which we have seen several times for various uses. Recall that squaring a value is just multiplying that value by itself. Thus, we can write the same equation as:

$$s^2 = \frac{\sum((X - M)(X - M))}{N - 1} \tag{13.1.2}$$

This is the same formula and works the same way as before, where we multiply the deviation score by itself (we square it) and then sum across squared deviations.

Now, let's look at the formula for covariance. In this formula, we will still use X to represent the score on one variable, and we will now use Y to represent the score on the second variable. The formula for covariance (cov_{XY} with the subscript XY to indicate covariance across the X and Y variables) is:

$$\text{cov}_{XY} = \frac{\sum((X - M_X)(Y - M_Y))}{N - 1} \tag{13.1.3}$$

As we can see, this is the exact same structure as the previous formula. Now, instead of multiplying the deviation score by itself on one variable, we take the deviation scores from a single person on each variable and multiply them together. We do this for each person (exactly the same as we did for variance) and then sum them to get our numerator. The numerator in this is called the Sum of Products.

$$SP = \sum((X - M_X)(Y - M_Y)) \tag{13.1.4}$$

We will calculate the sum of products using the same table we used to calculate the sum of squares. In fact, the table for sum of products is simply a sum of squares table for X, plus a sum of squares table for Y, with a final column of products, as shown below.

Table **13.1.1**: Sum of Products table

X	$(X - M_X)$	$(X - M_X)^2$	Y	$(Y - M_Y)$	$(Y - M_Y)^2$	$(X - M_X)(Y - M_Y)$

This table works the same way that it did before (remember that the column headers tell you exactly what to do in that column). We list our raw data for the X and Y variables in the X and Y columns, respectively, then add them up so we can calculate the mean of each variable. We then take those means and subtract them from the appropriate raw score to get our deviation scores for each person on each variable, and the columns of deviation scores will both add up to zero. We will square our deviation scores for each variable to get the sum of squares for X and Y so that we can compute the variance and standard deviation of each (we will use the standard deviation in our equation below). Finally, we take the deviation score from each variable and multiply them together to get our product score. Summing this column will give us our sum of products. It is very important that you multiply the raw deviation scores from each variable, NOT the squared deviation scores.

Our sum of products will go into the numerator of our formula for covariance, and then we only have to divide by $N-1$ to get our covariance. Unlike the sum of squares, both our sum of products and our covariance can be positive, negative, or zero, and they will always match (e.g. if our sum of products is positive, our covariance will always be positive). A positive sum of products and covariance indicates that the two variables are related and move in the same direction. That is, as one variable goes up, the other will also go up, and vice versa. A negative sum of products and covariance means that the variables are related but move in opposite directions when they change, which is called an inverse relation. In an inverse relation, as one variable goes up, the other variable goes down. If the sum of products and covariance are zero, then that means that the variables are not related. As one variable goes up or down, the other variable does not change in a consistent or predictable way.

The previous paragraph brings us to an important definition about relations between variables. What we are looking for in a relation is a consistent or predictable pattern. That is, the variables change together, either in the same direction or opposite directions, in the same way each time. It doesn't matter if this relation is positive or negative, only that it is not zero. If there is no consistency in how the variables change within a person, then the relation is zero and does not exist. We will revisit this notion of direction vs zero relation later on.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.1: Variability and Covariance is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.1: Variability and Covariance** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.2: Visualizing Relations

Chapter 2 covered many different forms of data visualization, and visualizing data remains an important first step in understanding and describing our data before we move into inferential statistics. Nowhere is this more important than in correlation. Correlations are visualized by a scatterplot, where our X variable values are plotted on the X-axis, the Y variable values are plotted on the Y-axis, and each point or marker in the plot represents a single person's score on X and Y. Figure **13.2.1** shows a scatterplot for hypothetical scores on job satisfaction (X) and worker well-being (Y). We can see from the axes that each of these variables is measured on a 10-point scale, with 10 being the highest on both variables (high satisfaction and good health and well-being) and 1 being the lowest (dissatisfaction and poor health). When we look at this plot, we can see that the variables do seem to be related. The higher scores on job satisfaction tend to also be the higher scores on well-being, and the same is true of the lower scores.

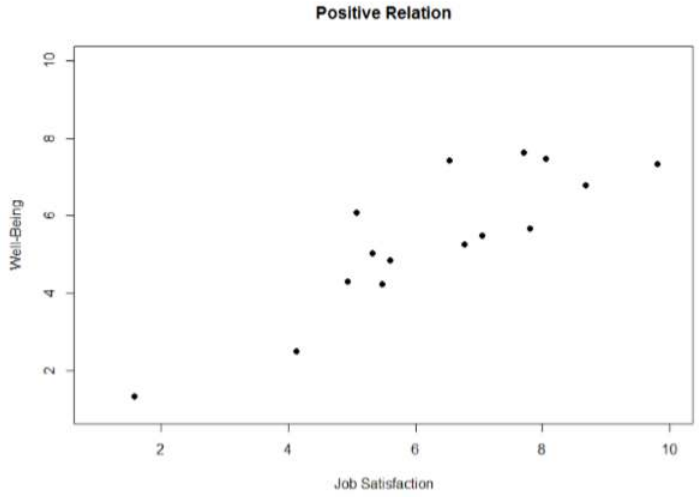


Figure **13.2.1**: Plotting satisfaction and well-being scores.

Figure **13.2.1** demonstrates a positive relation. As scores on X increase, scores on Y also tend to increase. Although this is not a perfect relation (if it were, the points would form a single straight line), it is nonetheless very clearly positive. This is one of the key benefits to scatterplots: they make it very easy to see the direction of the relation. As another example, Figure **13.2.2** shows a negative relation between job satisfaction (X) and burnout (Y). As we can see from this plot, higher scores on job satisfaction tend to correspond to lower scores on burnout, which is how stressed, un-energetic, and unhappy someone is at their job. As with Figure **13.2.1**, this is not a perfect relation, but it is still a clear one. As these figures show, points in a positive relation moves from the bottom left of the plot to the top right, and points in a negative relation move from the top left to the bottom right.

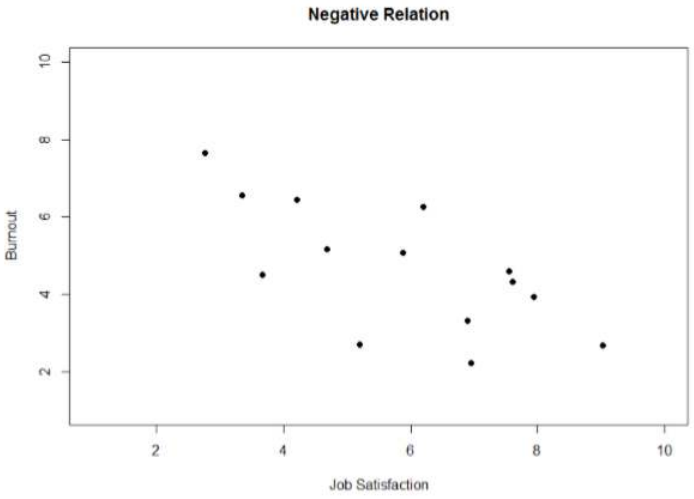


Figure **13.2.2**: Plotting satisfaction and burnout scores.

Scatterplots can also indicate that there is no relation between the two variables. In these scatterplots (an example is shown below in Figure **13.2.3** plotting job satisfaction and job performance) there is no interpretable shape or line in the scatterplot. The points appear randomly throughout the plot. If we tried to draw a straight line through these points, it would basically be flat. The low scores on job satisfaction have roughly the same scores on job performance as do the high scores on job satisfaction. Scores in the middle or average range of job satisfaction have some scores on job performance that are about equal to the high and low levels and some scores on job performance that are a little higher, but the overall picture is one of inconsistency.

As we can see, scatterplots are very useful for giving us an approximate idea of whether or not there is a relation between the two variables and, if there is, if that relation is positive or negative. They are also useful for another reason: they are the only way to determine one of the characteristics of correlations that are discussed next: form.

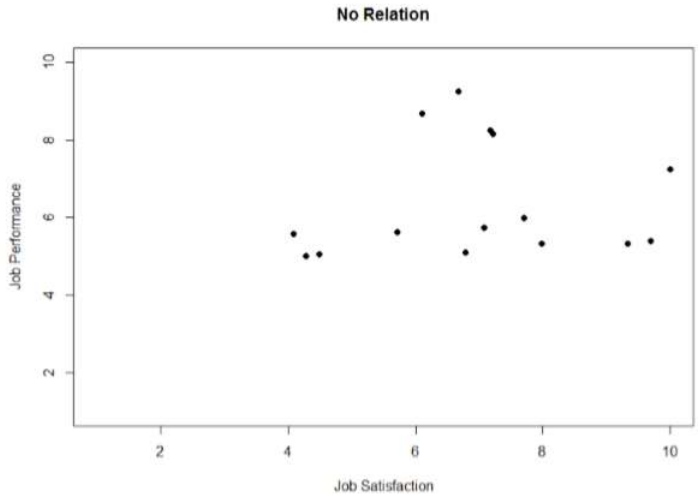


Figure **13.2.3**: Plotting no relation between satisfaction and job performance.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.2: Visualizing Relations is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.2: Visualizing Relations** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.3: Three Characteristics

When we talk about correlations, there are three traits that we need to know in order to truly understand the relation (or lack of relation) between X and Y: form, direction, and magnitude. We will discuss each of them in turn.

Form

The first characteristic of relations between variables is their form. The form of a relation is the shape it takes in a scatterplot, and a scatterplot is the only way it is possible to assess the form of a relation. There are three forms we look for: linear, curvilinear, or no relation. A linear relation is what we saw in Figures 12.2.1, 12.2.2, and 12.2.3. If we drew a line through the middle points in the any of the scatterplots, we would be best suited with a straight line. The term "linear" comes from the word "line". A linear relation is what we will always assume when we calculate correlations. All of the correlations presented here are only valid for linear relations. Thus, it is important to plot our data to make sure we meet this assumption.

The relation between two variables can also be curvilinear. As the name suggests, a curvilinear relation is one in which a line through the middle of the points in a scatterplot will be curved rather than straight. Two examples are presented in Figures **13.3.1** and **13.3.2**.

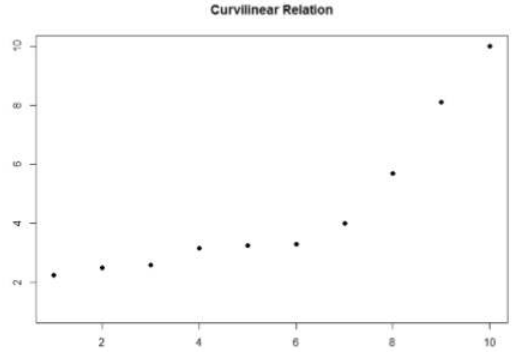


Figure **13.3.1**: Exponentially increasing curvilinear relation

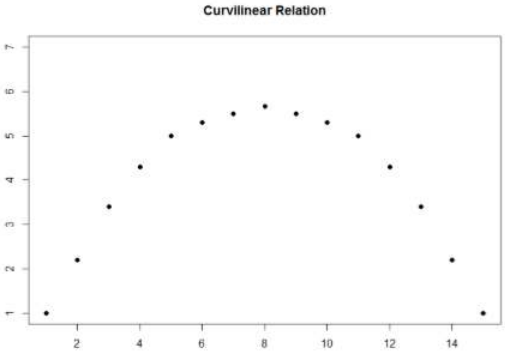


Figure **13.3.2**: Inverted-U curvilinear relation.

Curvilinear relations can take many shapes, and the two examples above are only a small sample of the possibilities. What they have in common is that they both have a very clear pattern but that pattern is not a straight line. If we try to draw a straight line through them, we would get a result similar to what is shown in Figure **13.3.3**.

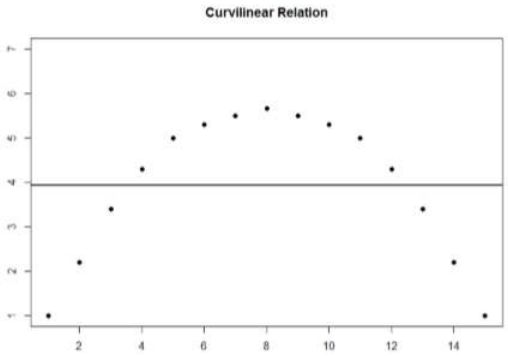


Figure 13.3.3: Overlaying a straight line on a curvilinear relation.

Although that line is the closest it can be to all points at the same time, it clearly does a very poor job of representing the relation we see. Additionally, the line itself is flat, suggesting there is no relation between the two variables even though the data show that there is one. This is important to keep in mind, because the math behind our calculations of correlation coefficients will only ever produce a straight line – we cannot create a curved line with the techniques discussed here.

Finally, sometimes when we create a scatterplot, we end up with no interpretable relation at all. An example of this is shown below in Figure 13.3.4. The points in this plot show no consistency in relation, and a line through the middle would once again be a straight, flat line.

Sometimes when we look at scatterplots, it is tempting to get biased by a few points that fall far away from the rest of the points and seem to imply that there may be some sort of relation. These points are called outliers, and we will discuss them in more detail later in the chapter. These can be common, so it is important to formally test for a relation between our variables, not just rely on visualization. This is the point of hypothesis testing with correlations, and we will go in depth on it soon. First, however, we need to describe the other two characteristics of relations: direction and magnitude.

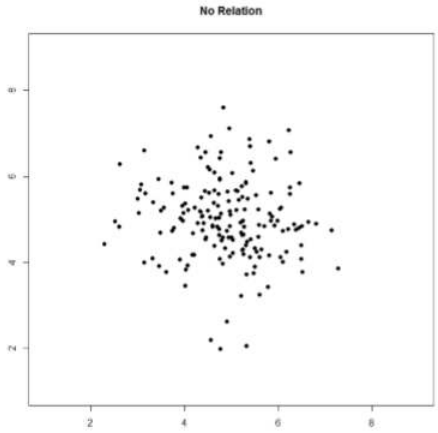


Figure 13.3.4: No relation

Direction

The direction of the relation between two variables tells us whether the variables change in the same way at the same time or in opposite ways at the same time. We saw this concept earlier when first discussing scatterplots, and we used the terms positive and negative. A positive relation is one in which X and Y change in the same direction: as X goes up, Y goes up, and as X goes down, Y also goes down. A negative relation is just the opposite: X and Y change together in opposite directions: as X goes up, Y goes down, and vice versa.

As we will see soon, when we calculate a correlation coefficient, we are quantifying the relation demonstrated in a scatterplot. That is, we are putting a number to it. That number will be either positive, negative, or zero, and we interpret the sign of the number as our direction. If the number is positive, it is a positive relation, and if it is negative, it is a negative relation. If it is zero, then there is no relation. The direction of the relation corresponds directly to the slope of the hypothetical line we draw through scatterplots

when assessing the form of the relation. If the line has a positive slope that moves from bottom left to top right, it is positive, and vice versa for negative. If the line it flat, that means it has no slope, and there is no relation, which will in turn yield a zero for our correlation coefficient.

Magnitude

The number we calculate for our correlation coefficient, which we will describe in detail below, corresponds to the magnitude of the relation between the two variables. The magnitude is how strong or how consistent the relation between the variables is. Higher numbers mean greater magnitude, which means a stronger relation.

Our correlation coefficients will take on any value between -1.00 and 1.00, with 0.00 in the middle, which again represents no relation. A correlation of -1.00 is a perfect negative relation; as X goes up by some amount, Y goes down by the same amount, consistently. Likewise, a correlation of 1.00 indicates a perfect positive relation; as X goes up by some amount, Y also goes up by the same amount. Finally, a correlation of 0.00, which indicates no relation, means that as X goes up by some amount, Y may or may not change by any amount, and it does so inconsistently.

The vast majority of correlations do not reach -1.00 or positive 1.00. Instead, they fall in between, and we use rough cut offs for how strong the relation is based on this number. Importantly, the sign of the number (the direction of the relation) has no bearing on how strong the relation is. The only thing that matters is the magnitude, or the absolute value of the correlation coefficient. A correlation of -1 is just as strong as a correlation of 1. We generally use values of 0.10, 0.30, and 0.50 as indicating weak, moderate, and strong relations, respectively.

The strength of a relation, just like the form and direction, can also be inferred from a scatterplot, though this is much more difficult to do. Some examples of weak and strong relations are shown in Figures **13.3.5** and **13.3.6**, respectively. Weak correlations still have an interpretable form and direction, but it is much harder to see. Strong correlations have a very clear pattern, and the points tend to form a line. The examples show two different directions, but remember that the direction does not matter for the strength, only the consistency of the relation and the size of the number, which we will see next.

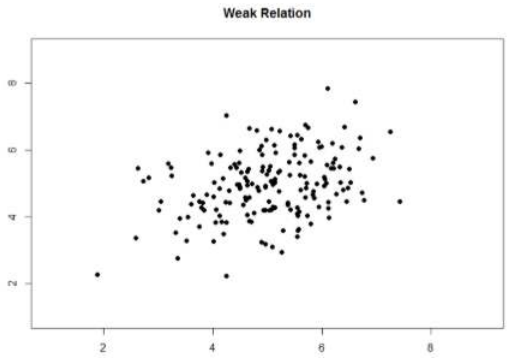


Figure **13.3.5**: Weak positive correlation.

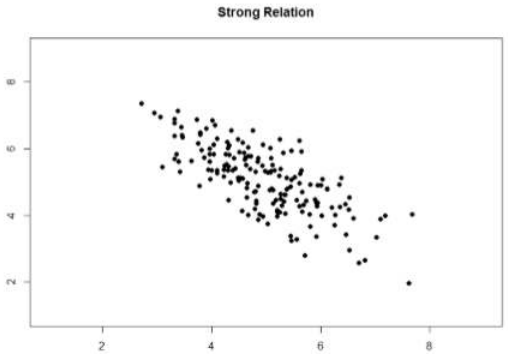


Figure **13.3.6**: Strong negative correlation.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.3: Three Characteristics is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.3: Three Characteristics** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.4: Pearson's r

There are several different types of correlation coefficients, but we will only focus on the most common: Pearson's r. r is a very popular correlation coefficient for assessing linear relations, and it serves as both a descriptive statistic (like M) and as a test statistic (like t). It is descriptive because it describes what is happening in the scatterplot; r will have both a sign (+/–) for the direction and a number (0 – 1 in absolute value) for the magnitude. As noted above, r assumes a linear relation, so nothing about r will suggest what the form is – it will only tell what the direction and magnitude would be if the form is linear (Remember: always make a scatterplot first!). r also works as a test statistic because the magnitude of r will correspond directly to a t value as the specific degrees of freedom, which can then be compared to a critical value. Luckily, we do not need to do this conversion by hand. Instead, we will have a table of r critical values that looks very similar to our t table, and we can compare our r directly to those.

The formula for r is very simple: it is just the covariance (defined above) divided by the standard deviations of X and Y:

$$r = \frac{\text{cov}_{XY}}{s_X s_Y} = \frac{SP}{\sqrt{SS_X * SS_Y}} \tag{13.4.1}$$

The first formula gives a direct sense of what a correlation is: a covariance standardized onto the scale of X and Y; the second formula is computationally simpler and faster. Both of these equations will give the same value, and as we saw at the beginning of the chapter, all of these values are easily computed by using the sum of products table. When we do this calculation, we will find that our answer is always between -1.00 and 1.00 (if it's not, check the math again), which gives us a standard, interpretable metric, similar to what z-scores did.

It was stated earlier that r is a descriptive statistic like M, and just like M, it corresponds to a population parameter. For correlations, the population parameter is the lowercase Greek letter ρ ("rho"); be careful not to confuse ρ with a p-value – they look quite similar. r is an estimate of ρ just like M is an estimate of μ. Thus, we will test our observed value of r that we calculate from the data and compare it to a value of ρ specified by our null hypothesis to see if the relation between our variables is significant, as we will see in our example next.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.4: Pearson's r is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.4: Pearson's r** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.5: Anxiety and Depression

Anxiety and depression are often reported to be highly linked (or "comorbid"). Our hypothesis testing procedure follows the same four-step process as before, starting with our null and alternative hypotheses. We will look for a positive relation between our variables among a group of 10 people because that is what we would expect based on them being comorbid.

Step 1: State the Hypotheses

Our hypotheses for correlations start with a baseline assumption of no relation, and our alternative will be directional if we expect to find a specific type of relation. For this example, we expect a positive relation:

$$\mathbf{H_0:}\ \textbf{There is no positive relation between anxiety and depression}$$
$$\mathbf{H_0}: \rho \leq 0$$

$$\mathbf{H_A}: \textbf{There is a positive relation between anxiety and depression}$$
$$\mathbf{H_0}: \rho > 0$$

Remember that ρ ("rho") is our population parameter for the correlation that we estimate with r, just like M and μ for means. Remember also that if there is no relation between variables, the magnitude will be 0, which is where we get the null and alternative hypothesis values.

Step 2: Find the Critical Values

The critical values for correlations come from the correlation table, which looks very similar to the t-table (see Figure 13.5.1). Just like our t-table, the column of critical values is based on our significance level (α) and the directionality of our test. The row is determined by our degrees of freedom. For correlations, we have N–2 degrees of freedom, rather than N–1 (why this is the case is not important). For our example, we have 10 people, so our degrees of freedom = 10 – 2 = 8.

Critical Values for Pearson's r

df	0.05 0.10	0.025 0.05	0.01 0.02	0.005 0.01	1-tailed α 2-tailed α
1	0.988	0.997	1.000	1.000	
2	0.900	0.950	0.980	0.990	
3	0.805	0.878	0.934	0.959	
4	0.729	0.811	0.882	0.917	
5	0.669	0.754	0.833	0.874	
6	0.622	0.707	0.789	0.834	
7	0.582	0.666	0.750	0.798	
8	0.549	0.632	0.716	0.765	
9	0.521	0.602	0.685	0.735	
10	0.497	0.576	0.658	0.708	
11	0.476	0.553	0.634	0.684	
12	0.458	0.532	0.612	0.661	
13	0.441	0.514	0.592	0.641	
14	0.426	0.497	0.574	0.623	
15	0.412	0.482	0.558	0.606	

Figure 13.5.1: Correlation table

We were not given any information about the level of significance at which we should test our hypothesis, so we will assume $\alpha = 0.05$ as always. From our table, we can see that a 1-tailed test (because we expect only a positive relation) at the $\alpha = 0.05$ level has a critical value of $r^* = 0.549$. Thus, if our observed correlation is greater than 0.549, it will be statistically significant. This is a rather high bar (remember, the guideline for a strong relation is r = 0.50); this is because we have so few people. Larger samples make it easier to find significant relations.

Step 3: Calculate the Test Statistic

We have laid out our hypotheses and the criteria we will use to assess them, so now we can move on to our test statistic. Before we do that, we must first create a scatterplot of the data to make sure that the most likely form of our relation is in fact linear. Figure 13.5.2 below shows our data plotted out, and it looks like they are, in fact, linearly related, so Pearson's r is appropriate.

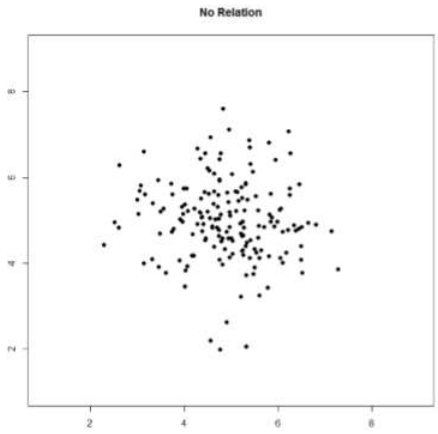


Figure **13.5.2**: Scatterplot of anxiety and depression

The data we gather from our participants is as follows:

Table **13.5.1**: Data from participants

Dep	2.81	1.96	3.43	3.40	4.71	1.80	4.27	3.68	2.44	3.13
Anx	3.54	3.05	3.81	3.43	4.03	3.59	4.17	3.46	3.19	4.12

We will need to put these values into our Sum of Products table to calculate the standard deviation and covariance of our variables. We will use X for depression and Y for anxiety to keep track of our data, but be aware that this choice is arbitrary and the math will work out the same if we decided to do the opposite. Our table is thus:

Table **13.5.2**: Sum of Products table

X	$(X - M_X)$	$(X - M_X)^2$	Y	$(Y - M_Y)$	$(Y - M_Y)^2$	$(X - M_X)(Y - M_Y)$
2.81	-0.35	0.12	3.54	-0.10	0.01	0.04
1.96	-1.20	1.44	3.05	-0.59	0.35	0.71
3.43	0.27	0.07	3.81	0.17	0.03	0.05
3.40	0.24	0.06	3.43	-0.21	0.04	-0.05
4.71	1.55	2.40	4.03	0.39	0.15	0.60
1.80	-1.36	1.85	3.59	-0.05	0.00	0.07
4.27	1.11	1.23	4.17	0.53	0.28	0.59
3.68	0.52	0.27	3.46	-0.18	0.03	-0.09
2.44	-0.72	0.52	3.19	-0.45	0.20	0.32
3.13	-0.03	0.00	4.12	0.48	0.23	-0.01
31.63	0.03	7.97	36.39	-0.01	1.33	2.22

The bottom row is the sum of each column. We can see from this that the sum of the X observations is 31.63, which makes the mean of the X variable $M_X = 3.16$. The deviation scores for X sum to 0.03, which is very close to 0, given rounding error, so everything looks right so far. The next column is the squared deviations for X, so we can see that the sum of squares for X is $SS_X = 7.97$. The same is true of the Y columns, with an average of $M_Y = 3.64$, deviations that sum to zero within rounding error, and a

sum of squares as $SS_Y = 1.33$. The final column is the product of our deviation scores (NOT of our squared deviations), which gives us a sum of products of $SP = 2.22$.

There are now three pieces of information we need to calculate before we compute our correlation coefficient: the covariance of X and Y and the standard deviation of each.

The covariance of two variable, remember, is the sum of products divided by $N–1$. For our data:

$$\text{cov}_{XY} = \frac{SP}{N-1} = \frac{2.22}{9} = 0.25$$

The formula for standard deviation are the same as before. Using subscripts X and Y to denote depression and anxiety:

$$s_X = \sqrt{\frac{\sum (X - M_X)^2}{N-1}} = \sqrt{\frac{7.97}{9}} = 0.94$$

$$s_Y = \sqrt{\frac{\sum (Y - M_Y)^2}{N-1}} = \sqrt{\frac{1.33}{9}} = 0.38$$

Now we have all of the information we need to calculate r:

$$r = \frac{\text{cov}_{XY}}{s_X s_Y} = \frac{0.25}{0.94 * 0.38} = 0.70$$

We can verify this using our other formula, which is computationally shorter:

$$r = \frac{SP}{\sqrt{SS_X * SS_Y}} = \frac{2.22}{\sqrt{7.97 * 1.33}} = .70$$

So our observed correlation between anxiety and depression is $r = 0.70$, which, based on sign and magnitude, is a strong, positive correlation. Now we need to compare it to our critical value to see if it is also statistically significant.

Step 4: Make a Decision

Our critical value was $r^* = 0.549$ and our obtained value was $r = 0.70$. Our obtained value was larger than our critical value, so we can reject the null hypothesis.

Reject H_0. Based on our sample of 10 people, there is a statistically significant, strong, positive relation between anxiety and depression, $r(8) = 0.70, p < .05$.

Notice in our interpretation that, because we already know the magnitude and direction of our correlation, we can interpret that. We also report the degrees of freedom, just like with t, and we know that $p < \alpha$ because we rejected the null hypothesis. As we can see, even though we are dealing with a very different type of data, our process of hypothesis testing has remained unchanged.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.5: Anxiety and Depression is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.5: Anxiety and Depression** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.6: Effect Size

Pearson's r is an incredibly flexible and useful statistic. Not only is it both descriptive and inferential, as we saw above, but because it is on a standardized metric (always between -1.00 and 1.00), it can also serve as its own effect size. In general, we use r = 0.10, r = 0.30, and r = 0.50 as our guidelines for small, medium, and large effects. Just like with Cohen's d, these guidelines are not absolutes, but they do serve as useful indicators in most situations. Notice as well that these are the same guidelines we used earlier to interpret the magnitude of the relation based on the correlation coefficient.

In addition to r being its own effect size, there is an additional effect size we can calculate for our results. This effect size is r^2, and it is exactly what it looks like – it is the squared value of our correlation coefficient. Just like η^2 in ANOVA, r^2 is interpreted as the amount of variance explained in the outcome variance, and the cut scores are the same as well: 0.01, 0.09, and 0.25 for small, medium, and large, respectively. Notice here that these are the same cutoffs we used for regular r effect sizes, but squared (0.102 = 0.01, 0.302 = 0.09, 0.502 = 0.25) because, again, the r^2 effect size is just the squared correlation, so its interpretation should be, and is, the same. The reason we use r^2 as an effect size is because our ability to explain variance is often important to us.

The similarities between η^2 and r^2 in interpretation and magnitude should clue you in to the fact that they are similar analyses, even if they look nothing alike. That is because, behind the scenes, they actually are! Those details, however, are beyond the score of this class.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.6: Effect Size is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.6: Effect Size** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.7: Correlation versus Causation

We cover a great deal of material in introductory statistics and, as mentioned chapter 1, many of the principles underlying what we do in statistics can be used in your day to day life to help you interpret information objectively and make better decisions. We now come to what may be the most important lesson in introductory statistics: the difference between correlation and causation.

It is very, very tempting to look at variables that are correlated and assume that this means they are causally related; that is, it gives the impression that X is causing Y.

However, in reality, correlation do not – and cannot – do this. Correlations DO NOT prove causation. No matter how logical or how obvious or how convenient it may seem, no correlational analysis can demonstrate causality. The ONLY way to demonstrate a causal relation is with a properly designed and controlled experiment.

Many times, we have good reason for assessing the correlation between two variables, and often that reason will be that we suspect that one causes the other. Thus, when we run our analyses and find strong, statistically significant results, it is very tempting to say that we found the causal relation that we are looking for. The reason we cannot do this is that, without an experimental design that includes random assignment and control variables, the relation we observe between the two variables may be caused by something else that we failed to measure. These “third variables” are lurking variables or confound variables, and they are impossible to detect and control for without an experiment.

Confound variables, which we will represent with Z, can cause two variables X and Y to appear related when in fact they are not. They do this by being the hidden – or lurking – cause of each variable independently. That is, if Z causes X and Z causes Y, the X and Y will appear to be related . However, if we control for the effect of Z (the method for doing this is beyond the scope of this text), then the relation between X and Y will disappear.

A popular example for this effect is the correlation between ice cream sales and deaths by drowning. These variables are known to correlate very strongly over time. However, this does not prove that one causes the other. The lurking variable in this case is the weather – people enjoy swimming and enjoy eating ice cream more during hot weather as a way to cool off. As another example, consider shoe size and spelling ability in elementary school children. Although there should clearly be no causal relation here, the variables are nonetheless consistently correlated. The confound in this case? Age. Older children spell better than younger children and are also bigger, so they have larger shoes.

When there is the possibility of confounding variables being the hidden cause of our observed correlation, we will often collect data on Z as well and control for it in our analysis. This is good practice and a wise thing for researchers to do. Thus, it would seem that it is easy to demonstrate causation with a correlation that controls for Z. However, the number of variables that could potentially cause a correlation between X and Y is functionally limitless, so it would be impossible to control for everything. That is why we use experimental designs; by randomly assigning people to groups and manipulating variables in those groups, we can balance out individual differences in any variable that may be our cause.

It is not always possible to do an experiment, however, so there are certain situations in which we will have to be satisfied with our observed relation and do the best we can to control for known confounds. However, in these situations, even if we do an excellent job of controlling for many extraneous (a statistical and research term for “outside”) variables, we must be very careful not to use causal language. That is because, even after controls, sometimes variables are related just by chance.

Sometimes, variables will end up being related simply due to random chance, and we call these correlation spurious. Spurious just means random, so what we are seeing is random correlations because, given enough time, enough variables, and enough data, sampling error will eventually cause some variables to be related when they should not. Sometimes, this even results in incredibly strong, but completely nonsensical, correlations. This becomes more and more of a problem as our ability to collect massive datasets and dig through them improves, so it is very important to think critically about any relation you encounter.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.7: Correlation versus Causation is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.7: Correlation versus Causation** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.8: Final Considerations

Correlations, although simple to calculate, can be very complex, and there are many additional issues we should consider. We will look at two of the most common issues that affect our correlations, as well as discuss some other correlations and reporting methods you may encounter.

Range Restriction

The strength of a correlation depends on how much variability is in each of the variables X and Y. This is evident in the formula for Pearson's r, which uses both covariance (based on the sum of products, which comes from deviation scores) and the standard deviation of both variables (which are based on the sums of squares, which also come from deviation scores). Thus, if we reduce the amount of variability in one or both variables, our correlation will go down. Failure to capture the full variability of a variable is called range restriction.

Take a look at Figures **13.8.1** and **13.8.2** below. The first shows a strong relation (r = 0.67) between two variables. An oval is overlain on top of it to make the relation even more distinct. The second shows the same data, but the bottom half of the X variable (all scores below 5) have been removed, which causes our relation (again represented by a red oval) to become much weaker (r = 0.38). Thus range restriction has truncated (made smaller) our observed correlation.

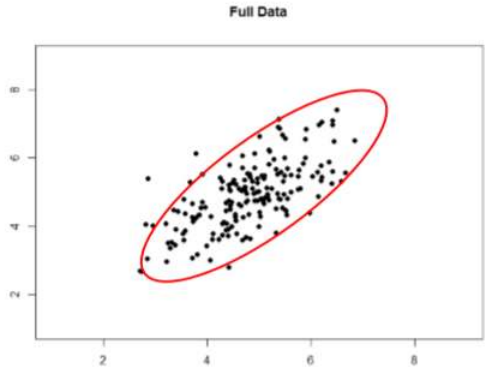


Figure **13.8.1**: Strong, positive correlation.

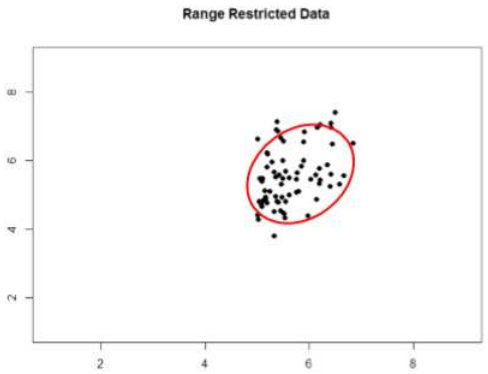


Figure **13.8.2**: Effect of range restriction.

Sometimes range restriction happens by design. For example, we rarely hire people who do poorly on job applications, so we would not have the lower range of those predictor variables. Other times, we inadvertently cause range restriction by not properly sampling our population. Although there are ways to correct for range restriction, they are complicated and require much information that may not be known, so it is best to be very careful during the data collection process to avoid it.

Outliers

Another issue that can cause the observed size of our correlation to be inappropriately large or small is the presence of outliers. An outlier is a data point that falls far away from the rest of the observations in the dataset. Sometimes outliers are the result of incorrect data entry, poor or intentionally misleading responses, or simple random chance. Other times, however, they represent

real people with meaningful values on our variables. The distinction between meaningful and accidental outliers is a difficult one that is based on the expert judgment of the researcher. Sometimes, we will remove the outlier (if we think it is an accident) or we may decide to keep it (if we find the scores to still be meaningful even though they are different).

The plots below in Figure **13.8.3** show the effects that an outlier can have on data. In the first, we have our raw dataset. You can see in the upper right corner that there is an outlier observation that is very far from the rest of our observations on both the X and Y variables. In the middle, we see the correlation computed when we include the outlier, along with a straight line representing the relation; here, it is a positive relation. In the third image, we see the correlation after removing the outlier, along with a line showing the direction once again. Not only did the correlation get stronger, it completely changed direction!

In general, there are three effects that an outlier can have on a correlation: it can change the magnitude (make it stronger or weaker), it can change the significance (make a non-significant correlation significant or vice versa), and/or it can change the direction (make a positive relation negative or vice versa). Outliers are a big issue in small datasets where a single observation can have a strong weight compared to the rest. However, as our samples sizes get very large (into the hundreds), the effects of outliers diminishes because they are outweighed by the rest of the data. Nevertheless, no matter how large a dataset you have, it is always a good idea to screen for outliers, both statistically (using analyses that we do not cover here) and/or visually (using scatterplots).

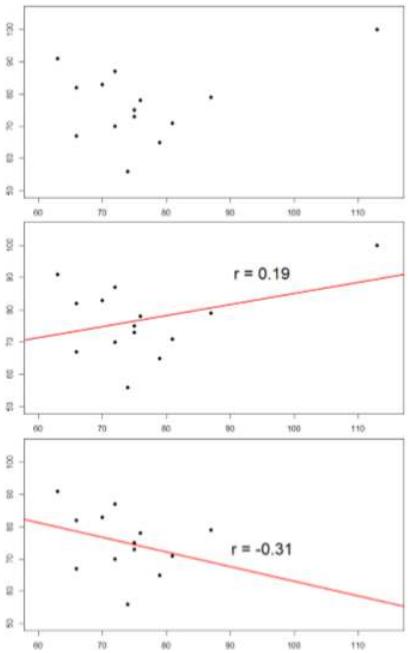


Figure **13.8.3**: Three plots showing correlations with and without outliers.

Other Correlation Coefficients

In this chapter we have focused on Pearson's r as our correlation coefficient because it very common and very useful. There are, however, many other correlations out there, each of which is designed for a different type of data. The most common of these is Spearman's rho (ρ), which is designed to be used on ordinal data rather than continuous data. This is a very useful analysis if we have ranked data or our data do not conform to the normal distribution. There are even more correlations for ordered categories, but they are much less common and beyond the scope of this chapter.

Additionally, the principles of correlations underlie many other advanced analyses. For example, a regression analysis is a formal way of running and analyzing a correlation that can be extended to more than two variables. Regression is a very powerful technique that serves as the basis for even our most advanced statistical models, so what we have learned in this chapter will open the door to an entire world of possibilities in data analysis.

Correlation Matrices

Many research studies look at the relation between more than two continuous variables. In such situations, we could simply list our all of our correlations, but that would take up a lot of space and make it difficult to quickly find the relation we are looking for.

Instead, we create correlation matrices so that we can quickly and simply display our results. A matrix is like a grid that contains our values. There is one row and one column for each of our variables, and the intersections of the rows and columns for different variables contain the correlation for those two variables.

At the beginning of the chapter, we saw scatterplots presenting data for correlations between job satisfaction, well-being, burnout, and job performance. We can create a correlation matrix to quickly display the numerical values of each. Such a matrix is shown below.

Table **13.8.1**: Correlation matrix to display the numerical values

	Satisfaction	Well-Being	Burnout	Performance
Satisfaction	1.00			
Well-Being	0.41	1.00		
Burnout	-0.54	-0.87	1.00	
Performance	0.08	0.21	-0.33	1.00

Notice that there are values of 1.00 where each row and column of the same variable intersect. This is because a variable correlates perfectly with itself, so the value is always exactly 1.00. Also notice that the upper cells are left blank and only the cells below the diagonal of 1s are filled in. This is because correlation matrices are symmetrical: they have the same values above the diagonal as below it. Filling in both sides would provide redundant information and make it a bit harder to read the matrix, so we leave the upper triangle blank.

Correlation matrices are a very condensed way of presenting many results quickly, so they appear in almost all research studies that use continuous variables. Many matrices also include columns that show the variable means and standard deviations, as well as asterisks showing whether or not each correlation is statistically significant.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.8: Final Considerations is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.8: Final Considerations** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

13.E: Correlations (Exercises)

1. What does a correlation assess?

Answer:

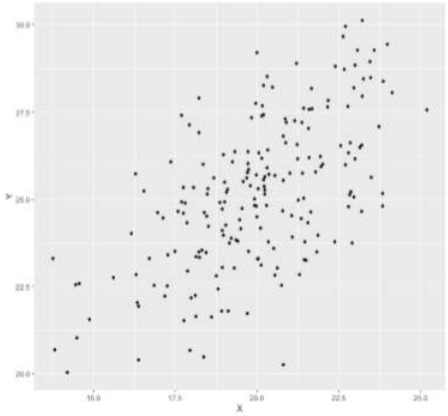
Correlations assess the linear relation between two continuous variables

2. What are the three characteristics of a correlation coefficient?
3. What is the difference between covariance and correlation?

Answer:

Covariance is an unstandardized measure of how related two continuous variables are. Correlations are standardized versions of covariance that fall between negative 1 and positive 1.

4. Why is it important to visualize correlational data in a scatterplot before performing analyses?
5. What sort of relation is displayed in the scatterplot below?



Answer:

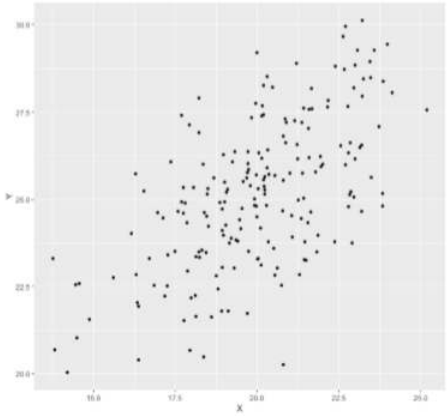
Strong, positive, linear relation

6. What is the direction and magnitude of the following correlation coefficients
   a. -0.81
   b. 0.40
   c. 0.15
   d. -0.08
   e. 0.29
7. Create a scatterplot from the following data:

Hours Studying	Overall Class Performance
0.62	2.02
1.50	4.62
0.34	2.60
0.97	1.59
3.54	4.67
0.69	2.52
1.53	2.28
0.32	1.68
1.94	2.50
1.25	4.04
1.42	2.63
3.07	3.53
3.99	3.90
1.73	2.75
1.9	2.95

Answer:

Your scatterplot should look similar to this:



8. In the following correlation matrix, what is the relation (number, direction, and magnitude) between…
   a. Pay and Satisfaction
   b. Stress and Health

Workplace	Pay	Satisfaction	Stress	Health
Pay	1.00			
Satisfaction	0.68	1.00		
Stress	0.02	-0.23	1.00	
Health	0.05	0.15	-0.48	1.00

9. Using the data from problem 7, test for a statistically significant relation between the variables.

Answer:

Step 1: $H_0 : \rho = 0$, "There is no relation between time spent studying and overall performance in class", $H_A : \rho > 0$, "There is a positive relation between time spent studying and overall performance in class."

Step 2: $df = 15–2 = 13, \alpha = 0.05$, 1-tailed test, $r^* = 0.441$.

Step 3: Using the Sum of Products table, you should find: $\overline{X} = 1.61, SS_X = 17.44, \overline{Y} = 2.95, SS_Y = 13.60, SP = 10.06, r = 0.65$.

Step 4: Obtained statistic is greater than critical value, reject H_0. There is a statistically significant, strong, positive relation between time spent studying and performance in class, $r(13) = 0.65, p < .05$.

10. A researcher collects data from 100 people to assess whether there is any relation between level of education and levels of civic engagement. The researcher finds the following descriptive values: $M_X = 4.02, s_x = 1.15, M_Y = 15.92, s_y = 5.01, SS_X = 130.93, SS_Y = 2484.91, SP = 159.39$. Test for a significant relation using the four step hypothesis testing procedure.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 13.E: Correlations (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **12.E: Correlations (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

CHAPTER OVERVIEW

14: Chi-square

14.1: Categories and Frequency Tables
14.2: Goodness-of-Fit
14.3: χ^2 Statistic
14.4: Pineapple on Pizza
14.5: Contingency Tables for Two Variables
14.6: Test for Independence
14.7: College Sports
14.E: Chi-square (Exercises)

This page titled 14: Chi-square is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

14.1: Categories and Frequency Tables

Our data for the χ^2 test are categorical, specifically nominal, variables. Recall from unit 1 that nominal variables have no specified order and can only be described by their names and the frequencies with which they occur in the dataset. Thus, unlike our other variables that we have tested, we cannot describe our data for the χ^2 test using means and standard deviations. Instead, we will use frequencies tables.

Table **14.1.1**: Pet Preferences

	Cat	Dog	Other	Total
Observed	14	17	5	36
Expected	12	12	12	36

Table **14.1.1** gives an example of a frequency table used for a χ^2 test. The columns represent the different categories within our single variable, which in this example is pet preference. The χ^2 test can assess as few as two categories, and there is no technical upper limit on how many categories can be included in our variable, although, as with ANOVA, having too many categories makes our computations long and our interpretation difficult. The final column in the table is the total number of observations, or N. The χ^2 test assumes that each observation comes from only one person and that each person will provide only one observation, so our total observations will always equal our sample size.

There are two rows in this table. The first row gives the observed frequencies of each category from our dataset; in this example, 14 people reported liking preferring cats as pets, 17 people reported preferring dogs, and 5 people reported a different animal. The second row gives expected values; expected values are what would be found if each category had equal representation. The calculation for an expected value is:

$$E = \frac{N}{k} \tag{14.1.1}$$

Where N is the total number of people in our sample and k is the number of categories in our variable (also the number of columns in our table). The expected values correspond to the null hypothesis for χ^2 tests: equal representation of categories. Our first of two χ^2 tests, the Goodness-of-Fit test, will assess how well our data lines up with, or deviates from, this assumption.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 14.1: Categories and Frequency Tables is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.1: Categories and Frequency Tables** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

14.2: Goodness-of-Fit

The first of our two χ^2 tests assesses one categorical variable against a null hypothesis of equally sized frequencies. Equal frequency distributions are what we would expect to get if categorization was completely random. We could, in theory, also test against a specific distribution of category sizes if we have a good reason to (e.g. we have a solid foundation of how the regular population is distributed), but this is less common, so we will not deal with it in this text.

Hypotheses

All χ^2 tests, including the goodness-of-fit test, are non-parametric. This means that there is no population parameter we are estimating or testing against; we are working only with our sample data. Because of this, there are no mathematical statements for χ^2 hypotheses. This should make sense because the mathematical hypothesis statements were always about population parameters (e.g. μ), so if we are non-parametric, we have no parameters and therefore no mathematical statements.

We do, however, still state our hypotheses verbally. For goodness-of-fit χ^2 tests, our null hypothesis is that there is an equal number of observations in each category. That is, there is no difference between the categories in how prevalent they are. Our alternative hypothesis says that the categories do differ in their frequency. We do not have specific directions or one-tailed tests for χ^2, matching our lack of mathematical statement.

Degrees of Freedom and the χ^2 table

Our degrees of freedom for the χ^2 test are based on the number of categories we have in our variable, not on the number of people or observations like it was for our other tests. Luckily, they are still as simple to calculate:

$$df = k - 1 \tag{14.2.1}$$

So for our pet preference example, we have 3 categories, so we have 2 degrees of freedom. Our degrees of freedom, along with our significance level (still defaulted to $\alpha = 0.05$) are used to find our critical values in the χ^2 table, which is shown in figure 1. Because we do not have directional hypotheses for χ^2 tests, we do not need to differentiate between critical values for 1- or 2-tailed tests. In fact, just like our F tests for regression and ANOVA, all χ^2 tests are 1-tailed tests.

	1-tailed α			
df	0.10	0.05	0.02	0.01
1	2.706	3.841	5.024	6.635
2	4.605	5.991	7.378	9.210
3	6.251	7.815	9.348	11.345
4	7.779	9.488	11.143	13.277
5	9.236	11.070	12.833	15.086
6	10.645	12.592	14.449	16.812
7	12.017	14.067	16.013	18.475
8	13.362	15.507	17.535	20.090
9	14.684	16.919	19.023	21.666
10	15.987	18.307	20.483	23.209

Figure **14.2.1**: First 10 rows of the χ^2 table

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 14.2: Goodness-of-Fit is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.2: Goodness-of-Fit** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

14.3: χ² Statistic

The calculations for our test statistic in χ^2 tests combine our information from our observed frequencies (O) and our expected frequencies (E) for each level of our categorical variable. For each cell (category) we find the difference between the observed and expected values, square them, and divide by the expected values. We then sum this value across cells for our test statistic. This is shown in the formula:

$$\chi^2 = \sum \frac{(O - E)^2}{E} \tag{14.3.1}$$

For our pet preference data, we would have:

$$\chi^2 = \frac{(14-12)^2}{12} + \frac{(17-12)^2}{12} + \frac{(5-12)^2}{12} = 0.33 + 2.08 + 4.08 = 6.49$$

Notice that, for each cell's calculation, the expected value in the numerator and the expected value in the denominator are the same value. Let's now take a look at an example from start to finish.

This page titled 14.3: χ² Statistic is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.3: χ² Statistic** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

14.4: Pineapple on Pizza

There is a very passionate and on-going debate on whether or not pineapple should go on pizza. Being the objective, rational data analysts that we are, we will collect empirical data to see if we can settle this debate once and for all. We gather data from a group of adults asking for a simple Yes/No answer.

Step 1: State the Hypotheses

We start, as always, with our hypotheses. Our null hypothesis of no difference will state that an equal number of people will say they do or do not like pineapple on pizza, and our alternative will be that one side wins out over the other:

$$\mathbf{H_0: \text{ There is no difference between the number of people who do and do not like pineapple on pizza}}$$

$$\mathbf{H_A: \text{ There is a difference between the number of people who do and do not like pineapple on pizza}}$$

Step 2: Find the Critical Value

To avoid any potential bias in this crucial analysis, we will leave α at its typical level. We have two options in our data (Yes or No), which will give us two categories. Based on this, we will have 1 degree of freedom. From our χ^2 table, we find a critical value of 3.84.

Step 3: Calculate the Test Statistic

The results of the data collection are presented in Table **14.4.1**. We had data from 45 people in all and 2 categories, so our expected values are $E = 45/2 = 22.50$.

Table **14.4.1**: Results of Data collection

	Yes	No	Total
Observed	26	19	45
Expected	22.50	22.50	45

We can use these to calculate our χ^2 statistic:

$$\chi^2 = \frac{(26-22.50)^2}{22.50} + \frac{(19-22.50)^2}{22.50} = 0.54 + 0.54 = 1.08$$

Step 4: Make the Decision

Our observed test statistic had a value of 1.08 and our critical value was 3.84. Our test statistic was smaller than our critical value, so we fail to reject the null hypothesis, and the debate rages on.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 14.4: Pineapple on Pizza is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.4: Pineapple on Pizza** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

14.5: Contingency Tables for Two Variables

The goodness-of-fit test is a useful tool for assessing a single categorical variable. However, what is more common is wanting to know if two categorical variables are related to one another. This type of analysis is similar to a correlation, the only difference being that we are working with nominal data, which violates the assumptions of traditional correlation coefficients. This is where the χ^2 test for independence comes in handy.

As noted above, our only description for nominal data is frequency, so we will again present our observations in a frequency table. When we have two categorical variables, our frequency table is crossed. That is, each combination of levels from each categorical variable are presented. This type of frequency table is called a contingency table because it shows the frequency of each category in one variable, contingent upon the specific level of the other variable.

An example contingency table is shown in Table **14.5.1**, which displays whether or not 168 college students watched college sports growing up (Yes/No) and whether the students' final choice of which college to attend was influenced by the college's sports teams (Yes – Primary, Yes – Somewhat, No):

Table **14.5.1**: Contingency table of college sports and decision making

College Sports		Affected Decision			
		Primary	Somewhat	No	Total
Watched	Yes	47	26	14	87
	No	21	23	37	81
	Total	68	49	51	168

In contrast to the frequency table for our goodness-of-fit test, our contingency table does not contain expected values, only observed data. Within our table, wherever our rows and columns cross, we have a cell. A cell contains the frequency of observing it's corresponding specific levels of each variable at the same time. The top left cell in Table **14.5.1** shows us that 47 people in our study watched college sports as a child AND had college sports as their primary deciding factor in which college to attend.

Cells are numbered based on which row they are in (rows are numbered top to bottom) and which column they are in (columns are numbered left to right). We always name the cell using (R,C), with the row first and the column second. A quick and easy way to remember the order is that R/C Cola exists but C/R Cola does not. Based on this convention, the top left cell containing our 47 participants who watched college sports as a child and had sports as a primary criteria is cell (1,1). Next to it, which has 26 people who watched college sports as a child but had sports only somewhat affect their decision, is cell (1,2), and so on. We only number the cells where our categories cross. We do not number our total cells, which have their own special name: marginal values.

Marginal values are the total values for a single category of one variable, added up across levels of the other variable. In table 3, these marginal values have been italicized for ease of explanation, though this is not normally the case. We can see that, in total, 87 of our participants (47+26+14) watched college sports growing up and 81 (21+23+37) did not. The total of these two marginal values is 168, the total number of people in our study. Likewise, 68 people used sports as a primary criteria for deciding which college to attend, 50 considered it somewhat, and 50 did not use it as criteria at all. The total of these marginal values is also 168, our total number of people. The marginal values for rows and columns will always both add up to the total number of participants, N, in the study. If they do not, then a calculation error was made and you must go back and check your work.

Expected Values of Contingency Tables

Our expected values for contingency tables are based on the same logic as they were for frequency tables, but now we must incorporate information about how frequently each row and column was observed (the marginal values) and how many people were in the sample overall (N) to find what random chance would have made the frequencies out to be. Specifically:

$$E_{ij} = \frac{R_i C_j}{N} \tag{14.5.1}$$

The subscripts i and j indicate which row and column, respectively, correspond to the cell we are calculating the expected frequency for, and the R_i and C_j are the row and column marginal values, respectively. N is still the total sample size. Using the data from Table **14.5.1**, we can calculate the expected frequency for cell (1,1), the college sport watchers who used sports at their primary criteria, to be:

$$E_{1,1} = \frac{87 * 68}{168} = 35.21$$

We can follow the same math to find all the expected values for this table:

Table **14.5.2**: Contingency table of college sports and decision making

College Sports		Affected Decision			
		Primary	Somewhat	No	Total
Watched	Yes	35.21	25.38	26.41	87
	No	32.79	23.62	24.59	81
	Total	68	49	51	

Notice that the marginal values still add up to the same totals as before. This is because the expected frequencies are just row and column averages simultaneously. Our total N will also add up to the same value.

The observed and expected frequencies can be used to calculate the same χ^2 statistic as we did for the goodness-of-fit test. Before we get there, though, we should look at the hypotheses and degrees of freedom used for contingency tables.

This page titled 14.5: Contingency Tables for Two Variables is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.5: Contingency Tables for Two Variables** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

14.6: Test for Independence

The χ^2 test performed on contingency tables is known as the **test for independence**. In this analysis, we are looking to see if the values of each categorical variable (that is, the frequency of their levels) is related to or independent of the values of the other categorical variable. Because we are still doing a χ^2 test, which is nonparametric, we still do not have mathematical versions of our hypotheses. The actual interpretations of the hypotheses are quite simple: the null hypothesis says that the variables are independent or not related, and alternative says that they are not independent or that they are related. Using this set up and the data provided in Table 14.5.2, let's formally test for whether or not watching college sports as a child is related to using sports as a criteria for selecting a college to attend.

This page titled 14.6: Test for Independence is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.6: Test for Independence** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

14.7: College Sports

We will follow the same 4 step procedure as we have since chapter 7.

Step 1: State the Hypotheses

Our null hypothesis of no difference will state that there is no relation between our variables, and our alternative will state that our variables are related:

$$\mathbf{H_0: \text{ College choice criteria is independent of college sports viewership as a child}}$$

$$\mathbf{H_A: \text{ College choice criteria is related to college sports viewership as a child}}$$

Step 2: Find the Critical Value

Our critical value will come from the same table that we used for the goodness-of-fit test, but our degrees of freedom will change. Because we now have rows and columns (instead of just columns) our new degrees of freedom use information on both:

$$df = (R-1)(C-1) \tag{14.7.1}$$

In our example:

$$df = (2-1)(3-1) = 1*2 = 2$$

Based on our 2 degrees of freedom, our critical value from our table is 5.991.

Step 3: Calculate the Test Statistic

The same formula for χ^2 is used once again:

$$\chi^2 = \sum \frac{(0-\mathrm{E})^2}{\mathrm{E}} \tag{14.7.2}$$

$$\chi^2 = \frac{(47-35.21)^2}{35.21} + \frac{(26-25.38)^2}{25.38} + \frac{(14-26.41)^2}{26.41} + \frac{(21-32.79)^2}{32.79} + \frac{(23-23.62)^2}{23.62} + \frac{(37-24.59)^2}{24.59}$$

Step 4: Make the Decision

The final decision for our test of independence is still based on our observed value (20.31) and our critical value (5.991). Because our observed value is greater than our critical value, we can reject the null hypothesis.

Reject H_0. Based on our data from 168 people, we can say that there is a statistically significant relation between whether or not someone watches college sports growing up and how much a college's sports team factor in to that person's decision on which college to attend, $\chi^2(2) = 20.31, p < 0.05$.

Effect Size for χ^2

Like all other significance tests, χ^2 tests – both goodness-of-fit and tests for independence – have effect sizes that can and should be calculated for statistically significant results. There are many options for which effect size to use, and the ultimate decision is based on the type of data, the structure of your frequency or contingency table, and the types of conclusions you would like to draw. For the purpose of our introductory course, we will focus only on a single effect size that is simple and flexible: Cramer's V.

Cramer's V is a type of correlation coefficient that can be computed on categorical data. Like any other correlation coefficient (e.g. Pearson's r), the cutoffs for small, medium, and large effect sizes of Cramer's V are 0.10, 0.30, and 0.50, respectively. The calculation of Cramer's V is very simple:

$$V = \sqrt{\frac{\chi^2}{N(k-1)}} \tag{14.7.3}$$

For this calculation, k is the smaller value of either R (the number of rows) or C (the number of columns). The numerator is simply the test statistic we calculate during step 3 of the hypothesis testing procedure. For our example, we had 2 rows and 3 columns, so $k = 2$:

$$V = \sqrt{\frac{\chi^2}{N(k-1)}} = \sqrt{\frac{20.38}{168(2-1)}} = \sqrt{0.12} = 0.35$$

So the statistically significant relation between our variables was moderately strong.

Contributors and Attributions

- Foster et al. (University of Missouri-St. Louis, Rice University, & University of Houston, Downtown Campus)

This page titled 14.7: College Sports is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.7: College Sports** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

14.E: Chi-square (Exercises)

1. What does a frequency table display? What does a contingency table display?

Answer:

Frequency tables display observed category frequencies and (sometimes) expected category frequencies for a single categorical variable. Contingency tables display the frequency of observing people in crossed category levels for two categorical variables, and (sometimes) the marginal totals of each variable level.

2. What does a goodness-of-fit test assess?
3. How do expected frequencies relate to the null hypothesis?

Answer:

Expected values are what we would observe if the proportion of categories was completely random (i.e. no consistent difference other than chance), which is the same was what the null hypothesis predicts to be true.

4. What does a test-for-independence assess?
5. Compute the expected frequencies for the following contingency table:

	Category A	Category B
Category C	22	38
Category D	16	14

Answer:

Observed	Category A	Category B	Total
Category C	22	38	60
Category D	16	14	30
Total	38	52	90

Expected	Category A	Category B	Total
Category C	$((60 * 38)/90) = 25.33$	$((60*52)/90) = 34.67$	60
Category D	$((30 * 38)/90) = 12.67$	$((30 * 52)/90) = 17.33$	30
Total	38	52	90

6. Test significance and find effect sizes (if significant) for the following tests:
   1. $N = 19, R = 3, C = 2, \chi 2(2) = 7.89, \alpha = .05$
   2. $N = 12, R = 2, C = 2, \chi 2(1) = 3.12, \alpha = .05$
   3. $N = 74, R = 3, C = 3, \chi 2(4) = 28.41, \alpha = .01$
7. You hear a lot of people claim that The Empire Strikes Back is the best movie in the original Star Wars trilogy, and you decide to collect some data to demonstrate this empirically (pun intended). You ask 48 people which of the original movies they liked best; 8 said A New Hope was their favorite, 23 said The Empire Strikes Back was their favorite, and 17 said Return of the Jedi was their favorite. Perform a chi-square test on these data at the .05 level of significance.

Answer:

Step 1: H_0: "There is no difference in preference for one movie", H_A: "There is a difference in how many people prefer one movie over the others."

Step 2: 3 categories (columns) gives $df = 2$, $\chi^2_{crit} = 5.991$.

Step 3: Based on the given frequencies:

	New Hope	Empire	Jedi	Total
Observed	8	23	17	48
Expected	16	16	16	

$\chi^2 = 7.13$.

Step 4: Our obtained statistic is greater than our critical value, reject H_0. Based on our sample of 48 people, there is a statistically significant difference in the proportion of people who prefer one Star Wars movie over the others, $\chi^2(2) = 7.13$, $p < .05$. Since this is a statistically significant result, we should calculate an effect size: Cramer's $V = \sqrt{\frac{7.13}{48(3-1)}} = 0.27$, which is a moderate effect size.

8. A pizza company wants to know if people order the same number of different toppings. They look at how many pepperoni, sausage, and cheese pizzas were ordered in the last week; fill out the rest of the frequency table and test for a difference.

	Pepperoni	Sausage	Cheese	Total
Observed	320	275	251	
Expected				

9. A university administrator wants to know if there is a difference in proportions of students who go on to grad school across different majors. Use the data below to test whether there is a relation between college major and going to grad school.

		Major		
		Psychology	Business	Math
Graduate School	Yes	32	8	36
	No	15	41	12

Answer:

Step 1: H_0: "There is no relation between college major and going to grad school", H_A: "Going to grad school is related to college major."

Step 2: $df = 2$, $\chi^2_{crit} = 5.991$.

Step 3: Based on the given frequencies:

Expected Values		Major		
		Psychology	Business	Math
Graduate School	Yes	24.81	25.86	25.33
	No	22.19	23.14	22.67

$\chi^2 = 2.09 + 12.34 + 4.49 + 2.33 + 13.79 + 5.02 = 40.05$.

Step 4: Obtained statistic is greater than the critical value, reject H_0. Based on our data, there is a statistically significant relation between college major and going to grad school, $\chi^2(2) = 40.05, p < .05$, Cramer's $V = 0.53$, which is a large effect

10. A company you work for wants to make sure that they are not discriminating against anyone in their promotion process. You have been asked to look across gender to see if there are differences in promotion rate (i.e. if gender and promotion rate are independent or not). The following data should be assessed at the normal level of significance:

		Promoted in last two years?	
		Yes	No
Gender	Women	8	5
	Men	9	7

This page titled 14.E: Chi-square (Exercises) is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Foster et al. (University of Missouri's Affordable and Open Access Educational Resources Initiative) via source content that was edited to the style and standards of the LibreTexts platform.

- **14.E: Chi-square (Exercises)** by Foster et al. is licensed CC BY-NC-SA 4.0. Original source: https://irl.umsl.edu/oer/4.

Index

A

Adding probabilities
4.2: Complements, Intersections, and Unions

B

Bonferroni Test
12.8: Post Hoc Tests

C

comorbid
13.5: Anxiety and Depression
complement
4.2: Complements, Intersections, and Unions
conditional probability
4.3: Conditional Probability and Independent Events
Covariance
13.1: Variability and Covariance

D

Direction (Correlation)
13.3: Three Characteristics

E

effect size
13.6: Effect Size
ELEMENT
4.1: Sample Spaces, Events, and Their Probabilities

F

Form (Correlation)
13.3: Three Characteristics

H

Hypothesis Testing
8: Introduction to Hypothesis Testing

I

independent events
4.3: Conditional Probability and Independent Events
INTERSECTIONS
4.2: Complements, Intersections, and Unions

M

Magnitude (Correlation)
13.3: Three Characteristics
mutually exclusive
4.2: Complements, Intersections, and Unions

N

normal distribution
5.1: Normal Distributions

O

OCCURRENCE
4.1: Sample Spaces, Events, and Their Probabilities

P

pearson's correlation
13.4: Pearson’s r
pie chart
6.2: Probability in Graphs and Distributions
post hoc test
12.8: Post Hoc Tests
Probability Rule for Complements
4.2: Complements, Intersections, and Unions

R

random experiment
4.1: Sample Spaces, Events, and Their Probabilities

S

sample space
4.1: Sample Spaces, Events, and Their Probabilities
SPECIFICITY OF A DIAGNOSTIC TEST
4.3: Conditional Probability and Independent Events
standard deviation
3.3: Spread and Variability

T

tree diagram
4.1: Sample Spaces, Events, and Their Probabilities
Tukey’s Honest Significant Difference (HSD)
12.8: Post Hoc Tests

U

unions
4.2: Complements, Intersections, and Unions

V

variability
3.3: Spread and Variability
variance
3.3: Spread and Variability
Venn diagram
4.1: Sample Spaces, Events, and Their Probabilities

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: PSY 190: Statistics for the Behavioral Sciences

Webpages: 114

All licenses found:

- CC BY-SA 4.0: 99.1% (113 pages)
- Undeclared: 0.9% (1 page)

By Page

- PSY 190: Statistics for the Behavioral Sciences - *CC BY-SA 4.0*
  - Front Matter - *CC BY-SA 4.0*
    - TitlePage - *CC BY-SA 4.0*
    - InfoPage - *CC BY-SA 4.0*
    - Table of Contents - *Undeclared*
    - Licensing - *CC BY-SA 4.0*
    - Note about this book/version - *CC BY-SA 4.0*
  - 1: Introduction - *CC BY-SA 4.0*
    - 1.1: What are statistics? - *CC BY-SA 4.0*
    - 1.2: Why do we study statistics? - *CC BY-SA 4.0*
    - 1.3: Types of Data and How to Collect Them - *CC BY-SA 4.0*
    - 1.4: Collecting Data - *CC BY-SA 4.0*
    - 1.5: Type of Research Designs - *CC BY-SA 4.0*
    - 1.6: Types of Statistical Analyses - *CC BY-SA 4.0*
    - 1.7: Mathematical Notation - *CC BY-SA 4.0*
    - 1.E: Introduction (Exercises) - *CC BY-SA 4.0*
  - 2: Describing Data using Distributions and Graphs - *CC BY-SA 4.0*
    - 2.1: Graphing Qualitative Variables - *CC BY-SA 4.0*
    - 2.2: Graphing Quantitative Variables - *CC BY-SA 4.0*
    - 2.E: Describing Data using Distributions and Graphs (Exercises) - *CC BY-SA 4.0*
  - 3: Measures of Central Tendency and Spread - *CC BY-SA 4.0*
    - 3.1: What is Central Tendency? - *CC BY-SA 4.0*
    - 3.2: Measures of Central Tendency - *CC BY-SA 4.0*
    - 3.3: Spread and Variability - *CC BY-SA 4.0*
    - 3.E: Measures of Central Tendency and Spread (Exercises) - *CC BY-SA 4.0*
  - 4: Basic Concepts of Probability: The math - *CC BY-SA 4.0*
    - 4.1: Sample Spaces, Events, and Their Probabilities - *CC BY-SA 4.0*
    - 4.2: Complements, Intersections, and Unions - *CC BY-SA 4.0*
    - 4.3: Conditional Probability and Independent Events - *CC BY-SA 4.0*
    - 4.E: Basic Concepts of Probability (Exercises) - *CC BY-SA 4.0*
  - 5: Z-scores and the Standard Normal Distribution - *CC BY-SA 4.0*
    - 5.1: Normal Distributions - *CC BY-SA 4.0*
    - 5.2: Z-scores - *CC BY-SA 4.0*
    - 5.3: Z-scores and the Area under the Curve - *CC BY-SA 4.0*
    - 5.E: Z-scores and the Standard Normal Distribution (Exercises) - *CC BY-SA 4.0*
  - 6: Probability - *CC BY-SA 4.0*
    - 6.1: What is Probability - *CC BY-SA 4.0*
    - 6.2: Probability in Graphs and Distributions - *CC BY-SA 4.0*
    - 6.3: The Bigger Picture - *CC BY-SA 4.0*
    - 6.E: Probability (Exercises) - *CC BY-SA 4.0*
  - 7: Sampling Distributions - *CC BY-SA 4.0*
    - 7.1: People, Samples, and Populations - *CC BY-SA 4.0*
    - 7.2: The Sampling Distribution of Sample Means - *CC BY-SA 4.0*
    - 7.3: Using Standard Error for Probability - *CC BY-SA 4.0*
    - 7.4: Sampling Distribution, Probability and Inference - *CC BY-SA 4.0*
    - 7.E: Sampling Distributions (Exercises) - *CC BY-SA 4.0*
  - 8: Introduction to Hypothesis Testing - *CC BY-SA 4.0*
    - 8.1: Logic and Purpose of Hypothesis Testing - *CC BY-SA 4.0*
    - 8.2: The Probability Value - *CC BY-SA 4.0*
    - 8.3: The Null Hypothesis - *CC BY-SA 4.0*

- 8.4: The Alternative Hypothesis - *CC BY-SA 4.0*
- 8.5: Critical values, p-values, and significance level - *CC BY-SA 4.0*
- 8.6: Steps of the Hypothesis Testing Process - *CC BY-SA 4.0*
- 8.7: Movie Popcorn - *CC BY-SA 4.0*
- 8.8: Effect Size - *CC BY-SA 4.0*
- 8.9: Office Temperature - *CC BY-SA 4.0*
- 8.10: Different Significance Level - *CC BY-SA 4.0*
- 8.11: Other Considerations in Hypothesis Testing - *CC BY-SA 4.0*
- 8.E: Introduction to Hypothesis Testing (Exercises) - *CC BY-SA 4.0*

9: Introduction to t-tests - *CC BY-SA 4.0*

- 9.1: The t-statistic - *CC BY-SA 4.0*
- 9.2: Hypothesis Testing with t - *CC BY-SA 4.0*
- 9.3: Confidence Intervals - *CC BY-SA 4.0*
- 9.E: Introduction to t-tests (Exercises) - *CC BY-SA 4.0*

10: Repeated Measures - *CC BY-SA 4.0*

- 10.1: Change and Differences - *CC BY-SA 4.0*
- 10.2: Hypotheses of Change and Differences - *CC BY-SA 4.0*
- 10.3: Increasing Satisfaction at Work - *CC BY-SA 4.0*
- 10.4: Bad Press - *CC BY-SA 4.0*
- 10.E: Repeated Measures (Exercises) - *CC BY-SA 4.0*

11: Independent Samples - *CC BY-SA 4.0*

- 11.1: Difference of Means - *CC BY-SA 4.0*
- 11.2: Research Questions about Independent Means - *CC BY-SA 4.0*
- 11.3: Hypotheses and Decision Criteria - *CC BY-SA 4.0*
- 11.4: Independent Samples t-statistic - *CC BY-SA 4.0*
- 11.5: Standard Error and Pooled Variance - *CC BY-SA 4.0*
- 11.6: Movies and Mood - *CC BY-SA 4.0*
- 11.7: Effect Sizes and Confidence Intervals - *CC BY-SA 4.0*
- 11.8: Homogeneity of Variance - *CC BY-SA 4.0*
- 11.E: Independent Samples (Exercises) - *CC BY-SA 4.0*

12: Analysis of Variance - *CC BY-SA 4.0*

- 12.1: Observing and Interpreting Variability - *CC BY-SA 4.0*
- 12.2: Sources of Variance - *CC BY-SA 4.0*
- 12.3: ANOVA Table - *CC BY-SA 4.0*
- 12.4: ANOVA and Type I Error - *CC BY-SA 4.0*
- 12.5: Hypotheses in ANOVA - *CC BY-SA 4.0*
- 12.6: Scores on Job Application Tests - *CC BY-SA 4.0*
- 12.7: Variance Explained - *CC BY-SA 4.0*
- 12.8: Post Hoc Tests - *CC BY-SA 4.0*
- 12.9: Other ANOVA Designs - *CC BY-SA 4.0*
- 12.10: Analysis of Variance (Exercises) - *CC BY-SA 4.0*

13: Correlations - *CC BY-SA 4.0*

- 13.1: Variability and Covariance - *CC BY-SA 4.0*
- 13.2: Visualizing Relations - *CC BY-SA 4.0*
- 13.3: Three Characteristics - *CC BY-SA 4.0*
- 13.4: Pearson's r - *CC BY-SA 4.0*
- 13.5: Anxiety and Depression - *CC BY-SA 4.0*
- 13.6: Effect Size - *CC BY-SA 4.0*
- 13.7: Correlation versus Causation - *CC BY-SA 4.0*
- 13.8: Final Considerations - *CC BY-SA 4.0*
- 13.E: Correlations (Exercises) - *CC BY-SA 4.0*

14: Chi-square - *CC BY-SA 4.0*

- 14.1: Categories and Frequency Tables - *CC BY-SA 4.0*
- 14.2: Goodness-of-Fit - *CC BY-SA 4.0*
- 14.3: χ^2 Statistic - *CC BY-SA 4.0*
- 14.4: Pineapple on Pizza - *CC BY-SA 4.0*
- 14.5: Contingency Tables for Two Variables - *CC BY-SA 4.0*
- 14.6: Test for Independence - *CC BY-SA 4.0*
- 14.7: College Sports - *CC BY-SA 4.0*
- 14.E: Chi-square (Exercises) - *CC BY-SA 4.0*

Back Matter - *CC BY-SA 4.0*

- Index - *CC BY-SA 4.0*
- Glossary - *CC BY-SA 4.0*
- Detailed Licensing - *CC BY-SA 4.0*