Lakewood City Schools
Course of Study for Statistics

Scope and Sequence – Is an elective course wherein students develop major theories and techniques for collecting, analyzing, and drawing conclusions from data. Students will be exposed to three or four broad conceptual themes: exploring data, planning a study, anticipating patterns in advance, and possibly statistical inference. Solving real-life problems that require the use of statistical inference and a combination of statistical techniques will be emphasized. Technology will be used to develop understanding, and work on large datasets. Students will present, defend, and critique statistical arguments. Students taking this course are strongly encouraged to have a TI-83+ or higher graphing calculator.

Prerequisite: One year of Algebra 2 ("C–" average or better).

Course Overview: The General Organizational for this Course

Course Skills

Skill 1 — Selecting Statistical Methods: *Select methods for collecting and/or analyzing data for statistical inference.* (20-30%) Students will identify key and relevant information to answer a question or solve a problem (including appropriate inference methods), describe appropriate methods for gathering and representing data, and identify null and alternative hypotheses.

- Identify the question to be answered or problem to be solved *(not assessed)*.
- Identify key and relevant information to answer a question or solve a problem.
- Describe an appropriate method for gathering and representing data.
- Identify an appropriate inference method for confidence intervals.
- Identify an appropriate inference method for significance tests.
- Identify null and alternative hypotheses.

Scope and Sequence – Is an elective course wherein students develop major theories and techniques for collecting, analyzing, and drawing conclusions from data. Students will be exposed to three or four broad conceptual themes: exploring data, planning a study, anticipating patterns in advance, and possibly statistical inference. Solving real-life problems that require the use of statistical inference and a combination of statistical techniques will be emphasized. Technology will be used to develop understanding, and work on large datasets. Students will present, defend, and critique statistical arguments. Students taking this course are strongly encouraged to have a TI-83+ or higher graphing calculator.

Prerequisite: One year of Algebra 2 (“C–” average or better).

Course Overview: The General Organizational for this Course

Skill 2 — Data Analysis: *Describe patterns, trends, associations, and relationships in data.* (20-30%) Students will describe data presented numerically or graphically, perform statistical calculations, and compare distributions.

- Describe data presented numerically and graphically.
- Construct numerical and graphical representations of distributions.
- Calculate summary statistics, relative positions of points within a distribution, correlation, and predicted response.
- Compare distributions and relative positions of points within a distribution.
- Identify an appropriate inference method for significance tests.
- Identify null and alternative hypotheses.

Scope and Sequence – Is an elective course wherein students develop major theories and techniques for collecting, analyzing, and drawing conclusions from data. Students will be exposed to three or four broad conceptual themes: exploring data, planning a study, anticipating patterns in advance, and possibly statistical inference. Solving real-life problems that require the use of statistical inference and a combination of statistical techniques will be emphasized. Technology will be used to develop understanding, and work on large datasets. Students will present, defend, and critique statistical arguments. Students taking this course are strongly encouraged to have a TI-83+ or higher graphing calculator.

Prerequisite: One year of Algebra 2 ("C–" average or better).

Course Overview: The General Organizational for this Course

Skill 3 — Using Probability and Simulation: *Explore random phenomena.* (30-40%) Students will determine relative frequencies, proportions, or probabilities using simulation or calculations. Additionally, students will describe and determine parameters for probability distributions. Students will also need to construct confidence intervals and calculate test statistics.

- Determine relative frequencies, proportions, and probabilities using simulations and calculations.
- Determine parameters for probability distributions.
- Describe probability distributions.
- Construct a confidence interval, provided conditions for inference are met.
- Calculate a test statistic and find a p-value, provided conditions for inference are met.

Scope and Sequence – Is an elective course wherein students develop major theories and techniques for collecting, analyzing, and drawing conclusions from data. Students will be exposed to three or four broad conceptual themes: exploring data, planning a study, anticipating patterns in advance, and possibly statistical inference. Solving real-life problems that require the use of statistical inference and a combination of statistical techniques will be emphasized. Technology will be used to develop understanding, and work on large datasets. Students will present, defend, and critique statistical arguments. Students taking this course are strongly encouraged to have a TI-83+ or higher graphing calculator.

Prerequisite: One year of Algebra 2 (“C–” average or better).

Course Overview: The General Organizational for this Course

★ Skill 4 — Statistical Argumentation: *Develop an explanation or justify a conclusion using evidence from data, definitions, or statistical inference.* (0-30%) Students will make appropriate claims, interpret statistical calculations, verify application of inference procedures, and justify statistical claims.

- Make an appropriate claim, and draw an appropriate conclusion.
- Interpret statistical calculations and findings to assign meaning and assess a claim.
- Verify that inference procedures apply in a given situation.
- Justify a claim based on a confidence interval.
- Justify a claim using a decision based on significance tests.

Scope and Sequence – Is an elective course wherein students develop major theories and techniques for collecting, analyzing, and drawing conclusions from data. Students will be exposed to three or four broad conceptual themes: exploring data, planning a study, anticipating patterns in advance, and possibly statistical inference. Solving real-life problems that require the use of statistical inference and a combination of statistical techniques will be emphasized. Technology will be used to develop understanding, and work on large datasets. Students will present, defend, and critique statistical arguments. Students taking this course are strongly encouraged to have a TI-83+ or higher graphing calculator.

Prerequisite: One year of Algebra 2 (“C–” average or better).

Course Overview: The General Organizational for this Course

This course is generally designed to cover a substantial portion of the material described in CollegeBoard’s AP Statistics Course and Exam Description, effective Fall 2019. It also follows and is guided by the Guidelines for Assessment and Instruction in Statistics Education (GAISE) REPORT.

Skills and sections marked with ★ are for enrichment and extension, for selected students as time permits.

★ skills and topics allow the results obtained through the use of the other skills and topics to be applied more broadly; these are the basis of inference and generalization.

Course Resources

Starnes, Daren S, and Josh Tabor. *The Practice of Statistics: For the Ap Exam*. , 2018. Print.

Tintle, Nathan, Beth L. Chance, George W. Cobb, Allan J. Rossman, Soma Roy, Todd Swanson, and Jill VanderStoep. *Introduction to Statistical Investigations*. , 2015. Print.

§ 1: Exploring One-Variable Data

§ 1 introduces students to data and the vocabulary of statistics. Students also learn to talk about data in real-world contexts. Variability in data may seem to suggest certain conclusions about the data distribution, but not all variation is meaningful. Statistics allows us to develop shared understandings of uncertainty and variation. In this unit, students will define and represent categorical and quantitative variables, describe and compare distributions of one-variable data, and interpret statistical calculations to assess claims about individual data points or samples. Students will also begin to apply the normal distribution model as an introduction to how theoretical models for populations can be used to describe some distributions of sample data. Later units will more fully develop probabilistic modeling and inference.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: What Can We Learn from Data (1.1)	Identify questions to be answered, based on variation in one-variable data.	Numbers may convey meaningful information, when placed in context.
The Language of Variation: Variables (1.2)	Identify variables in a set of data.	A variable is a characteristic that changes from one individual to another.
	Classify types of variables.	A categorical variable takes on values that are category names or group labels. A quantitative variable is one that takes on numerical values for a measured or counted quantity.
Representing a Categorical Variable with Tables (1.3)	Represent categorical data using frequency or relative frequency tables.	A frequency table gives the number of cases falling into each category. A relative frequency table gives the proportion of cases falling into each category.
	Describe categorical data represented in frequency or relative tables.	Percentages, relative frequencies, and rates all provide the same information as proportions. Counts and relative frequencies of categorical data reveal information that can be used to justify claims about the data in context.
Representing a Categorical Variable with Graphs (1.4)	Represent categorical data graphically.	Bar charts (or bar graphs) are used to display frequencies (counts) or relative frequencies (proportions) for categorical data. The height or length of each bar in a bar graph corresponds to either the number or proportion of observations falling within each category. There are many additional ways to represent frequencies (counts) or relative frequencies (proportions) for categorical data.
	Describe categorical data represented graphically.	Graphical representations of a categorical variable reveal information that can be used to justify claims about the data in context.
	Compare multiple sets of categorical data.	Frequency tables, bar graphs, or other representations can be used to compare two or more data sets in terms of the same categorical variable.

§ 1: Exploring One-Variable Data

§ 1 introduces students to data and the vocabulary of statistics. Students also learn to talk about data in real-world contexts. Variability in data may seem to suggest certain conclusions about the data distribution, but not all variation is meaningful. Statistics allows us to develop shared understandings of uncertainty and variation. In this unit, students will define and represent categorical and quantitative variables, describe and compare distributions of one-variable data, and interpret statistical calculations to assess claims about individual data points or samples. Students will also begin to apply the normal distribution model as an introduction to how theoretical models for populations can be used to describe some distributions of sample data. Later units will more fully develop probabilistic modeling and inference.

Representing a Quantitative Variable with Graphs (1.5)	Classify types of quantitative variables.	A discrete variable can take on a countable number of values. The number of values may be finite or countably infinite, as with the counting numbers. A continuous variable can take on infinitely many values, but those values cannot be counted. No matter how small the interval between two values of a continuous variable, it is always possible to determine another value between them.
	Represent quantitative data graphically.	In a histogram, the height of each bar shows the number or proportion of observations that fall within the interval corresponding to that bar. Altering the interval widths can change the appearance of the histogram. In a stem and leaf plot, each data value is split into a "stem" (the first digit or digits) and a "leaf" (usually the last digit). A dotplot represents each observation by a dot, with the position on the horizontal axis corresponding to the data value of that observation, with nearly identical values stacked on top of each other. A cumulative graph represents the number or proportion of a data set less than or equal to a given number. There are many additional ways to graphically represent distributions of quantitative data.
Describing the Distribution of a Quantitative Variable (1.6)	Describe the characteristics of quantitative data distributions.	Descriptions of the distribution of quantitative data include shape, center, and variability (spread), as well as any unusual features such as outliers, gaps, clusters, or multiple peaks. Outliers for one-variable data are data points that are unusually small or large relative to the rest of the data. A distribution is skewed to the right (positive skew) if the right tail is longer than the left.

§ 1: Exploring One-Variable Data

§ 1 introduces students to data and the vocabulary of statistics. Students also learn to talk about data in real-world contexts. Variability in data may seem to suggest certain conclusions about the data distribution, but not all variation is meaningful. Statistics allows us to develop shared understandings of uncertainty and variation. In this unit, students will define and represent categorical and quantitative variables, describe and compare distributions of one-variable data, and interpret statistical calculations to assess claims about individual data points or samples. Students will also begin to apply the normal distribution model as an introduction to how theoretical models for populations can be used to describe some distributions of sample data. Later units will more fully develop probabilistic modeling and inference.

		A distribution is skewed to the left (negative skew) if the left tail is longer than the right. A distribution is symmetric if the left half is the mirror image of the right half. Univariate graphs with one main peak are known as unimodal. Graphs with two prominent peaks are bimodal. A graph where each bar height is approximately the same (no prominent peaks) is approximately uniform. A gap is a region of a distribution between two data values where there are no observed data. Clusters are concentrations of data usually separated by gaps. Descriptive statistics does not attribute properties of a data set to a larger population, but may provide the basis for conjectures for subsequent testing.
Summary Statistics for a Quantitative Variable (1.7)	Calculate measures of center and position for quantitative data.	A statistic is a numerical summary of sample data. The mean is the sum of all the data values divided by the number of values. For a sample, the mean is denoted by x-bar: $\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$ where x_i represents the ith data point in the sample and n represents the number of data values in the sample. The median of a data set is the middle value when data are ordered. When the number of data points is even, the median can take on any value between the two middle values. In Statistics, the most commonly used value for the median of a data set with an even number of values is the average of the two middle values. The first quartile, Q1, is the median of the half of the ordered data set from the minimum to the position of the median. The third quartile, Q3, is the median of the half of the ordered data set from the position of the median to the maximum.

§ 1: Exploring One-Variable Data

§ 1 introduces students to data and the vocabulary of statistics. Students also learn to talk about data in real-world contexts. Variability in data may seem to suggest certain conclusions about the data distribution, but not all variation is meaningful. Statistics allows us to develop shared understandings of uncertainty and variation. In this unit, students will define and represent categorical and quantitative variables, describe and compare distributions of one-variable data, and interpret statistical calculations to assess claims about individual data points or samples. Students will also begin to apply the normal distribution model as an introduction to how theoretical models for populations can be used to describe some distributions of sample data. Later units will more fully develop probabilistic modeling and inference.

		Q1 and Q3 form the boundaries for the middle 50% of values in an ordered data set. — — — The p_{th} percentile is interpreted as the value that has *p*% of the data less than or equal to it.
	Calculate measures of variability for quantitative data.	Three commonly used measures of variability (or spread) in a distribution are the range, interquartile range, and standard deviation. — — — The range is defined as the difference between the maximum data value and the minimum data value. The interquartile range (IQR) is defined as the difference between the third and first quartiles: Q3 − Q1. Both the range and the interquartile range are possible ways of measuring variability of the distribution of a quantitative variable. — — — Standard deviation is a way to measure variability of the distribution of a quantitative variable. For a sample, the standard deviation is denoted by *s*: $s_x = \sqrt{\frac{1}{n-1}\sum(x_i - \bar{x})^2}$ The square of the sample standard deviation, s^2, is called the sample variance. — — — Changing units of measurement affects the values of the calculated statistics.
	Explain the selection of a particular measure of center and/or variability for describing a set of quantitative data.	There are many methods for determining outliers. Two methods frequently used in this course are: i. An outlier is a value greater than 1.5 × IQR above the third quartile or more than 1.5 × IQR below the first quartile. ii. An outlier is a value located 2 or more standard deviations above, or below, the mean. — — — The mean, standard deviation, and range are considered nonresistant (or non-robust) because they are influenced by outliers. The median and IQR are

§ 1: Exploring One-Variable Data

§ 1 introduces students to data and the vocabulary of statistics. Students also learn to talk about data in real-world contexts. Variability in data may seem to suggest certain conclusions about the data distribution, but not all variation is meaningful. Statistics allows us to develop shared understandings of uncertainty and variation. In this unit, students will define and represent categorical and quantitative variables, describe and compare distributions of one-variable data, and interpret statistical calculations to assess claims about individual data points or samples. Students will also begin to apply the normal distribution model as an introduction to how theoretical models for populations can be used to describe some distributions of sample data. Later units will more fully develop probabilistic modeling and inference.

		considered resistant (or robust), because outliers do not greatly (if at all) affect their value.
Graphical Representations of Summary Statistics (1.8)	Represent summary statistics for quantitative data graphically.	Taken together, the minimum data value, the first quartile (Q1), the median, the third quartile (Q3), and the maximum data value make up the five-number summary. A boxplot is a graphical representation of the five-number summary (minimum, first quartile, median, third quartile, maximum). The box represents the middle 50% of data, with a line at the median and the ends of the box corresponding to the quartiles. Lines ("whiskers") extend from the quartiles to the most extreme point that is not an outlier, and outliers are indicated by their own symbol beyond this.
	Describe summary statistics of quantitative data represented graphically.	Summary statistics of quantitative data, or of sets of quantitative data, can be used to justify claims about the data in context. If a distribution is relatively symmetric, then the mean and median are relatively close to one another. If a distribution is skewed right, then the mean is usually to the right of the median. If the distribution is skewed left, then the mean is usually to the left of the median.
Comparing Distributions of a Quantitative Variable (1.9)	Compare graphical representations for multiple sets of quantitative data.	Any of the graphical representations, e.g., histograms, side-by-side boxplots, etc., can be used to compare two or more independent samples on center, variability, clusters, gaps, outliers, and other features.
	Compare summary statistics for multiple sets of quantitative data.	Any of the numerical summaries (e.g., mean, standard deviation, relative frequency, etc.) can be used to compare two or more independent samples.
The Normal Distribution (1.10)	Compare a data distribution to the normal distribution model.	A parameter is a numerical summary of a population. Some sets of data may be described as approximately normally distributed. A normal curve is mound-shaped and symmetric.

§ 1: Exploring One-Variable Data

§ 1 introduces students to data and the vocabulary of statistics. Students also learn to talk about data in real-world contexts. Variability in data may seem to suggest certain conclusions about the data distribution, but not all variation is meaningful. Statistics allows us to develop shared understandings of uncertainty and variation. In this unit, students will define and represent categorical and quantitative variables, describe and compare distributions of one-variable data, and interpret statistical calculations to assess claims about individual data points or samples. Students will also begin to apply the normal distribution model as an introduction to how theoretical models for populations can be used to describe some distributions of sample data. Later units will more fully develop probabilistic modeling and inference.

		The parameters of a normal distribution are the population mean, μ, and the population standard deviation, σ. – – – For a normal distribution, approximately 68% of the observations are within 1 standard deviation of the mean, approximately 95% of observations are within 2 standard deviations of the mean, and approximately 99.7% of observations are within 3 standard deviations of the mean. This is called the empirical rule. – – – Many variables can be modeled by a normal distribution.
	Determine proportions and percentiles from a normal distribution.	A standardized score for a particular data value is calculated as (data value – mean)/(standard deviation), and measures the number of standard deviations a data value falls above or below the mean. – – – One example of a standardized score is a *z*-score, which is calculated as $z \text{ score} = \left(\frac{x_i - \mu}{\sigma}\right)$. A z-score measures how many standard deviations a data value is from the mean. – – – Technology, such as a calculator, a standard normal table, or computer-generated output, can be used to find the proportion of data values located on a given interval of a normally distributed random variable. – – – Given the area of a region under the graph of the normal distribution curve, it is possible to use technology, such as a calculator, a standard normal table, or computer-generated output, to estimate parameters for some populations.
	Compare measures of relative position in data sets.	Percentiles and *z*-scores may be used to compare relative positions of points within a data set or between data sets.

§ 2: Exploring Two-Variable Data

Building on § 1, students will explore relationships in two-variable categorical or quantitative data sets. They will use graphical and numerical methods to investigate an association between two categorical variables. Skills learned while working with two-way tables will transfer to calculating probabilities in § 4.
Students will describe form, direction, strength, and unusual features for an association between two quantitative variables. They will assess correlation and, if appropriate, use a linear model to predict values of the response variable from values of the explanatory variable. Students will interpret the least-squares regression line in context, analyze prediction errors (residuals), and explore departures from a linear pattern.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Are Variables Related? (2.1)	Identify questions to be answered about possible relationships in data.	Apparent patterns and associations in data may be random or not.
Representing Two Categorical Variables (2.2)	Compare numerical and graphical representations for two categorical variables.	Side-by-side bar graphs, segmented bar graphs, and mosaic plots are examples of bar graphs for one categorical variable, broken down by categories of another categorical variable. – – – Graphical representations of two categorical variables can be used to compare distributions and/or determine if variables are associated. – – – A two-way table, also called a contingency table, is used to summarize two categorical variables. The entries in the cells can be frequency counts or relative frequencies. – – – A joint relative frequency is a cell frequency divided by the total for the entire table.
Statistics for Two Categorical Variables (2.3)	Calculate statistics for two categorical variables.	The marginal relative frequencies are the row and column totals in a two-way table divided by the total for the entire table. – – – A conditional relative frequency is a relative frequency for a specific part of the contingency table (e.g., cell frequencies in a row divided by the total for that row).
	Compare statistics for two categorical variables.	Summary statistics for two categorical variables can be used to compare distributions and/or determine if variables are associated.
Representing the Relationship Between Two Quantitative Variables (2.4)	Represent bivariate quantitative data using scatterplots.	A bivariate quantitative data set consists of observations of two different quantitative variables made on individuals in a sample or population. – – – A scatterplot shows two numeric values for each observation, one corresponding to the value on the *x*-axis and one corresponding to the value on the *y*-axis. – – –

§ 2: Exploring Two-Variable Data

Building on § 1, students will explore relationships in two-variable categorical or quantitative data sets. They will use graphical and numerical methods to investigate an association between two categorical variables. Skills learned while working with two-way tables will transfer to calculating probabilities in § 4.
Students will describe form, direction, strength, and unusual features for an association between two quantitative variables. They will assess correlation and, if appropriate, use a linear model to predict values of the response variable from values of the explanatory variable. Students will interpret the least-squares regression line in context, analyze prediction errors (residuals), and explore departures from a linear pattern.

		An explanatory variable is a variable whose values are used to explain or predict corresponding values for the response variable.
	Describe the characteristics of a scatter plot.	A description of a scatter plot includes form, direction, strength, and unusual features. — — — The direction of the association shown in a scatterplot, if any, can be described as positive or negative. — — — A positive association means that as values of one variable increase, the values of the other variable tend to increase. A negative association means that as values of one variable increase, values of the other variable tend to decrease. — — — The form of the association shown in a scatterplot, if any, can be described as linear or non-linear to varying degrees. — — — The strength of the association is how closely the individual points follow a specific pattern, e.g., linear, and can be shown in a scatterplot. Strength can be described as strong, moderate, or weak. — — — Unusual features of a scatter plot include clusters of points or points with relatively large discrepancies between the value of the response variable and a predicted value for the response variable.
Correlation (2.5)	Determine the correlation for a linear relationship.	The correlation, r, gives the direction and quantifies the strength of the linear association between two quantitative variables. — — — The correlation coefficient can be calculated by: $r = \frac{1}{n-1}\sum\left(\frac{x_i-\bar{x}}{s_x}\right)\left(\frac{y_i-\bar{y}}{s_y}\right)$. However, the most common way to determine r is by using technology. — — — A correlation coefficient close to 1 or -1 does not necessarily mean that a linear model is appropriate.

§ 2: Exploring Two-Variable Data

Building on § 1, students will explore relationships in two-variable categorical or quantitative data sets. They will use graphical and numerical methods to investigate an association between two categorical variables. Skills learned while working with two-way tables will transfer to calculating probabilities in § 4.
Students will describe form, direction, strength, and unusual features for an association between two quantitative variables. They will assess correlation and, if appropriate, use a linear model to predict values of the response variable from values of the explanatory variable. Students will interpret the least-squares regression line in context, analyze prediction errors (residuals), and explore departures from a linear pattern.

	Interpret the correlation for a linear relationship.	The correlation, r, is unit-free, and always between -1 and 1, inclusive. A value of $r = 0$ indicates that there is no linear association. A value of $r = 1$ or $r = -1$ indicates that there is a perfect linear association. A perceived or real relationship between two variables does not mean that changes in one variable cause changes in the other. That is, correlation does not necessarily imply causation.
Linear Regression Models (2.6)	Calculate a predicted response value using a linear regression model.	A simple linear regression model is an equation that uses an explanatory variable, x, to predict the response variable, y. The predicted response value, denoted by $\hat{y}$, is calculated as $\hat{y} = a + bx$, where a is the y-intercept and b is the slope of the regression line, and x is the value of the explanatory variable. Extrapolation is predicting a response value using a value for the explanatory variable that is beyond the interval of x-values used to determine the regression line. The predicted value is less reliable as an estimate the further we extrapolate.
Residuals (2.7)	Represent differences between measured and predicted responses using residual plots.	The residual is the difference between the actual value and the predicted value: residual $= y - \hat{y}$. A residual plot is a plot of residuals versus explanatory variable values or predicted response values.
	Describe the form of association of bivariate data using residual plots.	Apparent randomness in a residual plot for a linear model is evidence of a linear form to the association between the variables. Residual plots can be used to investigate the appropriateness of a selected model.
Least Squares Regression (2.8)	Estimate parameters for the least-squares regression line model.	The least-squares regression model minimizes the sum of the squares of the residuals and contains the point $(\bar{x}, \bar{y})$.

§ 2: Exploring Two-Variable Data

Building on § 1, students will explore relationships in two-variable categorical or quantitative data sets. They will use graphical and numerical methods to investigate an association between two categorical variables. Skills learned while working with two-way tables will transfer to calculating probabilities in § 4.
Students will describe form, direction, strength, and unusual features for an association between two quantitative variables. They will assess correlation and, if appropriate, use a linear model to predict values of the response variable from values of the explanatory variable. Students will interpret the least-squares regression line in context, analyze prediction errors (residuals), and explore departures from a linear pattern.

		The slope, b, of the regression line can be calculated as $b = r\left(\frac{s_y}{s_x}\right)$ where r is the correlation between x and y, s_y is the sample standard deviation of the response variable, y, and s_x is the sample standard deviation of the explanatory variable, x. – – – Sometimes, the y-intercept of the line does not have a logical interpretation in context. – – – In simple linear regression, r^2 is the square of the correlation, r. It is also called the coefficient of determination. r^2 is the proportion of variation in the response variable that is explained by the explanatory variable in the model.
	Interpret the coefficients for the least-squares regression line model.	The coefficients of the least-squares regression model are the estimated slope and y-intercept. – – – The slope is the amount that the predicted y-value changes for every unit increase in x. – – – The y-intercept value is the predicted value of the response variable when the explanatory variable is equal to 0. The formula for the y-intercept, a, is $a = \bar{y} - b\bar{x}$.
Analyzing Departures from Linearity (2.9)	Identify influential points in regression.	An outlier in regression is a point that does not follow the general trend shown in the rest of the data and has a large residual when the Least Squares Regression Line (LSRL) is calculated. – – – A high-leverage point in regression has a substantially larger or smaller x-value than the other observations have. – – – An influential point in regression is any point that, if removed, changes the relationship substantially. Examples include much different slope, y-intercept, and/or correlation. Outliers and high leverage points are often influential.

§ 2: Exploring Two-Variable Data

Building on § 1, students will explore relationships in two-variable categorical or quantitative data sets. They will use graphical and numerical methods to investigate an association between two categorical variables. Skills learned while working with two-way tables will transfer to calculating probabilities in § 4.
Students will describe form, direction, strength, and unusual features for an association between two quantitative variables. They will assess correlation and, if appropriate, use a linear model to predict values of the response variable from values of the explanatory variable. Students will interpret the least-squares regression line in context, analyze prediction errors (residuals), and explore departures from a linear pattern.

	Calculate a predicted response using a least- squares regression line for a transformed data set.	Transformations of variables, such as evaluating the natural logarithm of each value of the response variable or squaring each value of the explanatory variable, can be used to create transformed data sets, which may be more linear in form than the untransformed data. Increased randomness in residual plots after transformation of data and/or movement of r^2 to a value closer to 1 offers evidence that the least-squares regression line for the transformed data is a more appropriate model to use to predict responses to the explanatory variable than the regression line for the untransformed data.

§ 3: Collecting Data

Depending on how data are collected, we may or may not be able to generalize findings or establish evidence of causal relationships. For example, if random selection is not used to obtain a sample from a population, bias may result and statistics from the sample cannot be assumed to generalize to the population. For data collected using well-designed experiments, statistically significant differences between or among experimental treatment groups are evidence that the treatments caused the effect. Students learn important principles of sampling and experimental design in this unit; they will learn about statistical inference in §s 6–9.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Do the Data We Collected Tell the Truth? (3.1)	Identify questions to be answered about data collection methods.	Methods for data collection that do not rely on chance result in untrustworthy conclusions.
Introduction to Planning a Study (3.2)	Identify the type of a study.	A population consists of all items or subjects of interest. A sample selected for study is a subset of the population. In an observational study, treatments are not imposed. Investigators examine data for a sample of individuals (retrospective) or follow a sample of individuals into the future collecting data (prospective) in order to investigate a topic of interest about the population. A sample survey is a type of observational study that collects data from a sample in an attempt to learn about the population from which the sample was taken. In an experiment, different conditions (treatments) are assigned to experimental units (participants or subjects).
	Identify appropriate generalizations and determinations based on observational studies.	It is only appropriate to make generalizations about a population based on samples that are randomly selected or otherwise representative of that population. A sample is only generalizable to the population from which the sample was selected. It is not possible to determine causal relationships between variables using data collected in an observational study.
Random Sampling and Data Collection (3.3)	Identify a sampling method, given a description of a study.	When an item from a population can be selected only once, this is called sampling without replacement. When an item from the population can be selected more than once, this is called sampling with replacement. A simple random sample (SRS) is a sample in which every group of a given size has an equal chance of being chosen. This method is the basis for many

§ 3: Collecting Data

Depending on how data are collected, we may or may not be able to generalize findings or establish evidence of causal relationships. For example, if random selection is not used to obtain a sample from a population, bias may result and statistics from the sample cannot be assumed to generalize to the population. For data collected using well-designed experiments, statistically significant differences between or among experimental treatment groups are evidence that the treatments caused the effect. Students learn important principles of sampling and experimental design in this unit; they will learn about statistical inference in §s 6–9.

		types of sampling mechanisms. A few examples of mechanisms used to obtain SRSs include numbering individuals and using a random number generator to select which ones to include in the sample, ignoring repeats, using a table of random numbers, or drawing a card from a deck without replacement. – – – A stratified random sample involves the division of a population into separate groups, called strata, based on shared attributes or characteristics (homogeneous grouping). Within each stratum a simple random sample is selected, and the selected units are combined to form the sample. – – – A cluster sample involves the division of a population into smaller groups, called clusters. Ideally, there is heterogeneity within each cluster, and clusters are similar to one another in their composition. A simple random sample of clusters is selected from the population to form the sample of clusters. Data are collected from all observations in the selected clusters. – – – A systematic random sample is a method in which sample members from a population are selected according to a random starting point and a fixed, periodic interval. – – – A census selects all items/subjects in a population.
	Explain why a particular sampling method is or is not appropriate for a given situation.	There are advantages and disadvantages for each sampling method depending upon the question that is to be answered and the population from which the sample will be drawn.
Potential Problems with Sampling (3.4)	Identify potential sources of bias in sampling methods.	Bias occurs when certain responses are systematically favored over others. – – – When a sample is comprised entirely of volunteers or people who choose to participate, the sample will typically not be representative of the population (voluntary response bias). – – –

§ 3: Collecting Data

Depending on how data are collected, we may or may not be able to generalize findings or establish evidence of causal relationships. For example, if random selection is not used to obtain a sample from a population, bias may result and statistics from the sample cannot be assumed to generalize to the population. For data collected using well-designed experiments, statistically significant differences between or among experimental treatment groups are evidence that the treatments caused the effect. Students learn important principles of sampling and experimental design in this unit; they will learn about statistical inference in §s 6–9.

		When part of the population has a reduced chance of being included in the sample, the sample will typically not be representative of the population (undercoverage bias). – – – Individuals chosen for the sample for whom data cannot be obtained (or who refuse to respond) may differ from those for whom data can be obtained (nonresponse bias). – – – Problems in the data gathering instrument or process result in response bias. Examples include questions that are confusing or leading (question wording bias) and self-reported responses. – – – Non-random sampling methods (for example, samples chosen by convenience or voluntary response) introduce potential for bias because they do not use chance to select the individuals.
Introduction to Experimental Design (3.5)	Identify the components of an experiment.	The experimental units are the individuals (which may be people or other objects of study) that are assigned treatments. When experimental units consist of people, they are sometimes referred to as participants or subjects. – – – An explanatory variable (or factor) in an experiment is a variable whose levels are manipulated intentionally. The levels or combination of levels of the explanatory variable(s) are called treatments. – – – A response variable in an experiment is an outcome from the experimental units that is measured after the treatments have been administered. – – – A confounding variable in an experiment is a variable that is related to the explanatory variable and influences the response variable and may create a false perception of association between the two.
	Describe elements of a well-designed experiment.	A well-designed experiment should include the following:

§ 3: Collecting Data

Depending on how data are collected, we may or may not be able to generalize findings or establish evidence of causal relationships. For example, if random selection is not used to obtain a sample from a population, bias may result and statistics from the sample cannot be assumed to generalize to the population. For data collected using well-designed experiments, statistically significant differences between or among experimental treatment groups are evidence that the treatments caused the effect. Students learn important principles of sampling and experimental design in this unit; they will learn about statistical inference in §s 6–9.

		a. Comparisons of at least two treatment groups, one of which could be a control group. b. Random assignment/allocation of treatments to experimental units. c. Replication (more than one experimental unit in each treatment group). d. Control of potential confounding variables where appropriate.
	Compare experimental designs and methods.	In a completely randomized design, treatments are assigned to experimental units completely at random. Random assignment tends to balance the effects of uncontrolled (confounding) variables so that differences in responses can be attributed to the treatments. – – – Methods for randomly assigning treatments to experimental units in a completely randomized design include using a random number generator, a table of random values, drawing chips without replacement, etc. – – – In a single-blind experiment, subjects do not know which treatment they are receiving, but members of the research team do, or vice versa. – – – In a double-blind experiment neither the subjects nor the members of the research team who interact with them know which treatment a subject is receiving. – – – A control group is a collection of experimental units either not given a treatment of interest or given a treatment with an inactive substance (placebo) in order to determine if the treatment of interest has an effect. – – – The placebo effect occurs when experimental units have a response to a placebo. – – – For randomized complete block designs, treatments are assigned completely at random within each block. – – –

§ 3: Collecting Data

Depending on how data are collected, we may or may not be able to generalize findings or establish evidence of causal relationships. For example, if random selection is not used to obtain a sample from a population, bias may result and statistics from the sample cannot be assumed to generalize to the population. For data collected using well-designed experiments, statistically significant differences between or among experimental treatment groups are evidence that the treatments caused the effect. Students learn important principles of sampling and experimental design in this unit; they will learn about statistical inference in §s 6–9.

		Blocking ensures that at the beginning of the experiment the units within each block are similar to each other with respect to at least one blocking variable. A randomized block design helps to separate natural variability from differences due to the blocking variable. – – – A matched pairs design is a special case of a randomized block design. Using a blocking variable, subjects (whether they are people or not) are arranged in pairs matched on relevant factors. Matched pairs may be formed naturally or by the experimenter. Every pair receives both treatments by randomly assigning one treatment to one member of the pair and subsequently assigning the remaining treatment to the second member of the pair. Alternately, each subject may get both treatments.
Selecting an Experimental Design (3.6)	Explain why a particular experimental design is appropriate.	There are advantages and disadvantages for each experimental design depending on the question of interest, the resources available, and the nature of the experimental units.
Inference and Experiments (3.7)	Interpret the results of a well-designed experiment.	Statistical inference attributes conclusions based on data to the distribution from which the data were collected. – – – Random assignment of treatments to experimental units allows researchers to conclude that some observed changes are so large as to be unlikely to have occurred by chance. Such changes are said to be statistically significant. – – – Statistically significant differences between or among experimental treatment groups are evidence that the treatments caused the effect. – – – If the experimental units used in an experiment are representative of some larger group of units, the results of an experiment can be generalized to the larger group. Random selection of experimental units gives a better chance that the units will be representative.

§ 4: Probability, Random Variables, and Probability Distributions

Probabilistic reasoning allows statisticians to quantify the likelihood of random events over the long run and to make statistical inferences. Simulations and concrete examples can help students to understand the abstract definitions and calculations of probability. This unit builds on understandings of simulated or empirical data distributions and fundamental principles of probability to represent, interpret, and calculate parameters for theoretical probability distributions for discrete random variables. Interpretations of probabilities and parameters associated with a probability distribution should use appropriate units and relate to the context of the situation.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Random and Non-Random Patterns? (4.1)	Identify questions suggested by patterns in data.	Patterns in data do not necessarily mean that variation is not random.
Estimating Probabilities Using Simulation (4.2)	Estimate probabilities using simulation.	A random process generates results that are determined by chance. – – – An outcome is the result of a trial of a random process. – – – An event is a collection of outcomes. – – – Simulation is a way to model random events, such that simulated outcomes closely match real-world outcomes. All possible outcomes are associated with a value to be determined by chance. Record the counts of simulated outcomes and the count total. – – – The relative frequency of an outcome or event in simulated or empirical data can be used to estimate the probability of that outcome or event. – – – The law of large numbers states that simulated (empirical) probabilities tend to get closer to the true probability as the number of trials increases.
Introduction to Probability (4.3)	Calculate probabilities for events and their complements.	The sample space of a random process is the set of all possible non-overlapping outcomes. – – – If all outcomes in the sample space are equally likely, then the probability an event E will occur is defined as the fraction: $\frac{\text{number of outcomes in event E}}{\text{total \# of outcomes in sample space}}$ – – – The probability of an event is a number between 0 and 1, inclusive. – – – The probability of the complement of an event E, E' or E^C, (i.e., not E) is equal to $1 - P(E)$.

§ 4: Probability, Random Variables, and Probability Distributions

Probabilistic reasoning allows statisticians to quantify the likelihood of random events over the long run and to make statistical inferences. Simulations and concrete examples can help students to understand the abstract definitions and calculations of probability. This unit builds on understandings of simulated or empirical data distributions and fundamental principles of probability to represent, interpret, and calculate parameters for theoretical probability distributions for discrete random variables. Interpretations of probabilities and parameters associated with a probability distribution should use appropriate units and relate to the context of the situation.

	Interpret probabilities for events.	Probabilities of events in repeatable situations can be interpreted as the relative frequency with which the event will occur in the long run.
Mutually Exclusive Events (4.4)	Explain why two events are (or are not) mutually exclusive.	The probability that events A and B both will occur, sometimes called the joint probability, is the probability of the intersection of A and B, denoted P(A ∩ B). Two events are mutually exclusive or disjoint if they cannot occur at the same time. So P(A ∩ B) = 0.
Conditional Probability (4.5)	Calculate conditional probabilities.	The probability that event *A* will occur given that event *B* has occurred is called a conditional probability and denoted $P(A\|B) = \frac{P(A \cap B)}{P(B)}$. The multiplication rule states that the probability that events *A* and *B* both will occur is equal to the probability that event A will occur multiplied by the probability that event B will occur, given that A has occurred. This is denoted $P(A \cap B) = P(A) \cdot P(B\|A)$.
Independent Events and Unions of Events (4.6)	Calculate probabilities for independent events and for the union of two events.	Events *A* and *B* are independent if, and only if, knowing whether event *A* has occurred (or will occur) does not change the probability that event *B* will occur. If, and only if, events *A* and *B* are independent, then $P(A\|B) = P(A)$, $P(B\|A) = P(B)$ and $P(A \cap B) = P(A) \cdot P(B)$. The probability that event *A* or event *B* (or both) will occur is the probability of the union of *A* and *B*, denoted $P(A \cup B)$. The addition rule states that the probability that event *A* or event *B* or both will occur is equal to the probability that event *A* will occur plus the probability that event *B* will occur minus the probability that both events *A* and *B* will occur. This is denoted

§ 4: Probability, Random Variables, and Probability Distributions

Probabilistic reasoning allows statisticians to quantify the likelihood of random events over the long run and to make statistical inferences. Simulations and concrete examples can help students to understand the abstract definitions and calculations of probability. This unit builds on understandings of simulated or empirical data distributions and fundamental principles of probability to represent, interpret, and calculate parameters for theoretical probability distributions for discrete random variables. Interpretations of probabilities and parameters associated with a probability distribution should use appropriate units and relate to the context of the situation.

		$P(A \cup B) = P(A) + P(B) - P(A \cap B)$.
Introduction to Random Variables and Probability Distributions (4.7)	Represent the probability distribution for a discrete random variable.	The values of a random variable are the numerical outcomes of random behavior. A discrete random variable is a variable that can only take a countable number of values. Each value has a probability associated with it. The sum of the probabilities over all of the possible values must be 1. A probability distribution can be represented as a graph, table, or function showing the probabilities associated with values of a random variable. A cumulative probability distribution can be represented as a table or function showing the probability of being less than or equal to each value of the random variable.
	Interpret a probability distribution.	An interpretation of a probability distribution provides information about the shape, center, and spread of a population and allows one to make conclusions about the population of interest.
Mean and Standard Deviation of Random Variables (4.8)	Calculate parameters for a discrete random variable.	A numerical value measuring a characteristic of a population or the distribution of a random variable is known as a parameter, which is a single, fixed value. The mean, or expected value, for a discrete random variable *X* is $\mu_X = \sum x_i \cdot P(x_i)$. The standard deviation for a discrete random variable *X* is $\sigma_X = \sqrt{\sum (x_i - \mu_X)^2 \cdot P(x_i)}$.
	Interpret parameters for a discrete random variable.	Parameters for a discrete random variable should be interpreted using appropriate units and within the context of a specific population.
Combining Random Variables (4.9)	Calculate parameters for linear combinations of random variables.	For random variables *X* and *Y* and real numbers *a* and *b*, the mean of *aX* + *bY* is $a\mu_x + b\mu_y$.

§ 4: Probability, Random Variables, and Probability Distributions

Probabilistic reasoning allows statisticians to quantify the likelihood of random events over the long run and to make statistical inferences. Simulations and concrete examples can help students to understand the abstract definitions and calculations of probability. This unit builds on understandings of simulated or empirical data distributions and fundamental principles of probability to represent, interpret, and calculate parameters for theoretical probability distributions for discrete random variables. Interpretations of probabilities and parameters associated with a probability distribution should use appropriate units and relate to the context of the situation.

		Two random variables are independent if knowing information about one of them does not change the probability distribution of the other. For independent random variables X and Y and real numbers a and b, the mean of $aX + bY$ is $a\mu_x + b\mu_y$, and the variance of $aX + bY$ is $a^2\sigma_X^2 + b^2\sigma_Y^2$.
	Describe the effects of linear transformations of parameters of random variables.	For $Y = a + bX$, the probability distribution of the transformed random variable, Y, has the same shape as the probability distribution for X, so long as $a > 0$ and $b > 0$. The mean of Y is $\mu_Y = a + b\mu_X$. The standard deviation of Y is $\sigma_Y = \lvert b \rvert \sigma_X$.
Introduction to the Binomial Distribution (4.10)	Estimate probabilities of binomial random variables using data from a simulation.	A probability distribution can be constructed using the rules of probability or estimated with a simulation using random number generators. A binomial random variable, X, counts the number of successes in n repeated independent trials, each trial having two possible outcomes (success or failure), with the probability of success p and the probability of failure $1 – p$.
	Calculate probabilities for a binomial distribution.	The probability that a binomial random variable, X, has exactly x successes for n independent trials, when the probability of success is p, is calculated as: $P(X = x) = \binom{n}{x} p^x (1-p)^{n-x}$, $x = 0, 1, 2, 3, \ldots n$ This is the binomial probability function.
Parameters for a Binomial Distribution (4.11)	Calculate parameters for a binomial distribution.	If a random variable is binomial, its mean, μ_x, is np and its standard deviation, σ_x, is $np(1-p)$.
	Interpret probabilities and parameters for a binomial distribution.	Probabilities and parameters for a binomial distribution should be interpreted using appropriate units and within the context of a specific population or situation.
The Geometric Distribution (4.12)	Calculate probabilities for geometric random variables.	For a sequence of independent trials, a geometric random variable, X, gives the number of the trial on which the first success occurs. Each trial has two possible

§ 4: Probability, Random Variables, and Probability Distributions

Probabilistic reasoning allows statisticians to quantify the likelihood of random events over the long run and to make statistical inferences. Simulations and concrete examples can help students to understand the abstract definitions and calculations of probability. This unit builds on understandings of simulated or empirical data distributions and fundamental principles of probability to represent, interpret, and calculate parameters for theoretical probability distributions for discrete random variables. Interpretations of probabilities and parameters associated with a probability distribution should use appropriate units and relate to the context of the situation.

		outcomes (success or failure) with the probability of success p and the probability of failure $1 - p$. The probability that the first success for repeated independent trials with probability of success p occurs on trial x is calculated as: $P(X = x) = (1 - p)^{x-1} \cdot p,$ $x = 1, 2, 3, \ldots$. This is the geometric probability function.
	Calculate parameters of a geometric distribution.	If a random variable is geometric, its mean, μ_x is $\frac{1}{p}$ and its standard deviation, σ_x is $\frac{\sqrt{(1-p)}}{p}$.
	Interpret probabilities and parameters for a geometric distribution.	Probabilities and parameters for a geometric distribution should be interpreted using appropriate units and within the context of a specific population or situation.

§ 5: Sampling Distribution ★

This unit applies probabilistic reasoning to sampling, introducing students to sampling distributions of statistics they will use when performing inference in §s 6 and 7. Students should understand that sample statistics can be used to estimate corresponding population parameters and that measures of center (mean) and variability (standard deviation) for these sampling distributions can be determined directly from the population parameters when certain sampling criteria are met. For large enough samples from any population, these sampling distributions can be approximated by a normal distribution. Simulating sampling distributions helps students to understand how the values of statistics vary in repeated random sampling from populations with known parameters.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Why Is My Sample Not Like Yours? (5.1)	Identify questions suggested by variation in statistics for samples collected from the same population.	Variation in statistics for samples taken from the same population may be random or not.
The Normal Distribution, Revisited (5.2)	Calculate the probability that a particular value lies in a given interval of a normal distribution.	A continuous random variable is a variable that can take on any value within a specified domain. Every interval within the domain has a probability associated with it. – – – A continuous random variable with a normal distribution is commonly used to describe populations. The distribution of a normal random variable can be described by a normal, or "bell-shaped," curve. – – – The area under a normal curve over a given interval represents the probability that a particular value lies in that interval.
	Determine the interval associated with a given area in a normal distribution.	The boundaries of an interval associated with a given area in a normal distribution can be determined using z-scores or technology, such as a calculator, a standard normal table, or computer-generated output. – – – Intervals associated with a given area in a normal distribution can be determined by assigning appropriate inequalities to the boundaries of the intervals: a. $P(X < x_a) = \frac{p}{100}$ means that the lowest p% of values lie to the left of x_a. b. $P(x_a < X < x_b) = \frac{p}{100}$ means that p% of values lie between x_a and x_b. c. $P(X > x_b) = \frac{p}{100}$ means that the highest p% of values lie to the right of x_a. d. To determine the most extreme p% of values requires dividing the area associated with p% into two equal areas on either extreme of the distribution:

§ 5: Sampling Distribution ★

This unit applies probabilistic reasoning to sampling, introducing students to sampling distributions of statistics they will use when performing inference in §s 6 and 7. Students should understand that sample statistics can be used to estimate corresponding population parameters and that measures of center (mean) and variability (standard deviation) for these sampling distributions can be determined directly from the population parameters when certain sampling criteria are met. For large enough samples from any population, these sampling distributions can be approximated by a normal distribution. Simulating sampling distributions helps students to understand how the values of statistics vary in repeated random sampling from populations with known parameters.

		$P(X < x_a) = \frac{1}{2} \cdot \frac{p}{100}$ and $P(X > x_b) = \frac{1}{2} \cdot \frac{p}{100}$ means that half of the p% most extreme values lie to the left of x_a and half of the p% most extreme values lie to the right of x_b. 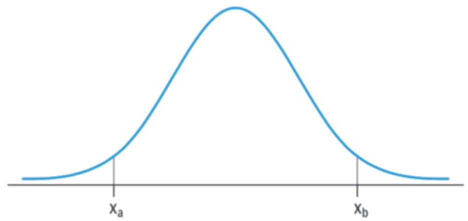
	Determine the appropriateness of using the normal distribution to approximate probabilities for unknown distributions.	Normal distributions are symmetrical and "bell-shaped." As a result, normal distributions can be used to approximate distributions with similar characteristics.
The Central Limit Theorem (5.3)	Estimate sampling distributions using simulation.	A sampling distribution of a statistic is the distribution of values for the statistic for all possible samples of a given size from a given population. – – – The central limit theorem (CLT) states that when the sample size is sufficiently large, a sampling distribution of the mean of a random variable will be approximately normally distributed. – – – The central limit theorem requires that the sample values are independent of each other and that n is sufficiently large. – – – A randomization distribution is a collection of statistics generated by simulation assuming known values for the parameters. For a randomized experiment, this

§ 5: Sampling Distribution ★

This unit applies probabilistic reasoning to sampling, introducing students to sampling distributions of statistics they will use when performing inference in §s 6 and 7. Students should understand that sample statistics can be used to estimate corresponding population parameters and that measures of center (mean) and variability (standard deviation) for these sampling distributions can be determined directly from the population parameters when certain sampling criteria are met. For large enough samples from any population, these sampling distributions can be approximated by a normal distribution. Simulating sampling distributions helps students to understand how the values of statistics vary in repeated random sampling from populations with known parameters.

		means repeatedly randomly reallocating/reassigning the response values to treatment groups. – – – The sampling distribution of a statistic can be simulated by generating repeated random samples from a population.
Biased and Unbiased Point Estimates (5.4)	Explain why an estimator is or is not unbiased.	When estimating a population parameter, an estimator is unbiased if, on average, the value of the estimator is equal to the population parameter.
	Calculate estimates for a population parameter.	When estimating a population parameter, an estimator exhibits variability that can be modeled using probability. – – – A sample statistic is a point estimator of the corresponding population parameter.
Sampling Distributions for Sample Proportions (5.5)	Determine parameters of a sampling distribution for sample proportions.	For independent samples (sampling with replacement) of a categorical variable from a population with population proportion, p, the sampling distribution of the sample proportion, $\hat{p}$, has a mean, $\mu_{\hat{p}} = p$ and a standard deviation, $$\sigma_{\hat{p}} = \sqrt{\frac{p(1-p)}{n}}$$ – – – If sampling without replacement, the standard deviation of the sample proportion is smaller than what is given by the formula above. If the sample size is less than 10% of the population size, the difference is negligible.
	Determine whether a sampling distribution for a sample proportion can be described as approximately normal.	For a categorical variable, the sampling distribution of the sample proportion, $\hat{p}$, will have an approximate normal distribution, provided the sample size is large enough: $np \geq 10$ and $n(1-p) \geq 10$.

§ 5: Sampling Distribution ★

This unit applies probabilistic reasoning to sampling, introducing students to sampling distributions of statistics they will use when performing inference in §s 6 and 7. Students should understand that sample statistics can be used to estimate corresponding population parameters and that measures of center (mean) and variability (standard deviation) for these sampling distributions can be determined directly from the population parameters when certain sampling criteria are met. For large enough samples from any population, these sampling distributions can be approximated by a normal distribution. Simulating sampling distributions helps students to understand how the values of statistics vary in repeated random sampling from populations with known parameters.

	Interpret probabilities and parameters for a sampling distribution for a sample proportion.	Probabilities and parameters for a sampling distribution for a sample proportion should be interpreted using appropriate units and within the context of a specific population.
Sampling Distributions for Differences in Sample Proportions (5.6)	Determine parameters of a sampling distribution for a difference in sample proportions.	For a categorical variable, when randomly sampling with replacement from two independent populations with population proportions p_1 and p_2, the sampling distribution of the difference in sample proportions $\hat{p}_1 - \hat{p}_2$ has mean, $\mu_{\hat{p}_1-\hat{p}_2} = p_1 - p_2$ and standard deviation, $\sigma_{\hat{p}_1-\hat{p}_2} = \sqrt{\frac{p_1(1-p_1)}{n_1} + \frac{p_2(1-p_2)}{n_2}}$ If sampling without replacement, the standard deviation of the difference in sample proportions is smaller than what is given by the formula above. If the sample sizes are less than 10% of the population sizes, the difference is negligible.
	Determine whether a sampling distribution for a difference of sample proportions can be described as approximately normal.	The sampling distribution of the difference in sample proportions $p_1 - p_2$ will have an approximate normal distribution provided the sample sizes are large enough: $n_1p_1 \geq 10, n_1(1-p_1) \geq 10)$, $n_2p_2 \geq 10, n_2(1-p_2) \geq 10)$
	Interpret probabilities and parameters for a sampling distribution for a difference in proportions.	Parameters for a sampling distribution for a difference of proportions should be interpreted using appropriate units and within the context of a specific populations.
Sampling Distributions for Sample Means (5.7	Determine parameters for a sampling distribution for sample means.	For a numerical variable, when random sampling with replacement from a population with mean μ and standard deviation, σ , the sampling distribution of the sample mean has mean, $\mu_{\bar{x}} = \mu$ and standard deviation $\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$.

§ 5: Sampling Distribution ★

This unit applies probabilistic reasoning to sampling, introducing students to sampling distributions of statistics they will use when performing inference in §s 6 and 7. Students should understand that sample statistics can be used to estimate corresponding population parameters and that measures of center (mean) and variability (standard deviation) for these sampling distributions can be determined directly from the population parameters when certain sampling criteria are met. For large enough samples from any population, these sampling distributions can be approximated by a normal distribution. Simulating sampling distributions helps students to understand how the values of statistics vary in repeated random sampling from populations with known parameters.

		If sampling without replacement, the standard deviation of the sample mean is smaller than what is given by the formula above. If the sample size is less than 10% of the population size, the difference is negligible.
	Determine whether a sampling distribution of a sample mean can be described as approximately normal.	For a numerical variable, if the population distribution can be modeled with a normal distribution, the sampling distribution of the sample mean, $\bar{x}$, can be modeled with a normal distribution. For a numerical variable, if the population distribution cannot be modeled with a normal distribution, the sampling distribution of the sample mean, $\bar{x}$, can be modeled approximately by a normal distribution, provided the sample size is large enough, e.g., greater than or equal to 30.
	Interpret probabilities and parameters for a sampling distribution for a sample mean.	Probabilities and parameters for a sampling distribution for a sample mean should be interpreted using appropriate units and within the context of a specific population.
Sampling Distributions for Differences in Sample Means (5.8)	Determine parameters of a sampling distribution for a difference in sample means.	For a numerical variable, when randomly sampling with replacement from two independent populations with population means μ_1 and μ_2 and population standard deviations σ_1 and σ_2, the sampling distribution of the difference in sample means $\bar{x}_1 - \bar{x}_2$ has mean $\mu_{(\bar{x}_1 - \bar{x}_2)} = \mu_1 - \mu_2$ and standard deviation, $$\sigma_{(\bar{x}_1 - \bar{x}_2)} = \sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}$$ If sampling without replacement, the standard deviation of the difference in sample means is smaller than what is given by the formula above. If the sample sizes are less than 10% of the population sizes, the difference is negligible.
	Determine whether a sampling distribution of a difference in sample	The sampling distribution of the difference in sample means $\bar{x}_1 - \bar{x}_2$ can be modeled with a Normal distribution if the two population distributions can be modeled with a normal distribution.

§ 5: Sampling Distribution ★

This unit applies probabilistic reasoning to sampling, introducing students to sampling distributions of statistics they will use when performing inference in §s 6 and 7. Students should understand that sample statistics can be used to estimate corresponding population parameters and that measures of center (mean) and variability (standard deviation) for these sampling distributions can be determined directly from the population parameters when certain sampling criteria are met. For large enough samples from any population, these sampling distributions can be approximated by a normal distribution. Simulating sampling distributions helps students to understand how the values of statistics vary in repeated random sampling from populations with known parameters.

	means can be described as approximately normal.	The sampling distribution of the difference in sample means $\bar{x}_1 - \bar{x}_2$ can be modeled approximately by a normal distribution if the two population distributions cannot be modeled with a normal distribution but both sample sizes are greater than or equal to 30.
	Interpret probabilities and parameters for a sampling distribution for a difference in sample means.	Probabilities and parameters for a sampling distribution for a difference of sample means should be interpreted using appropriate units and within the context of a specific populations.

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Why Be Normal? (6.1)	Identify questions suggested by variation in the shapes of distributions of samples taken from the same population.	Variation in shapes of data distributions may be random or not.
Constructing a Confidence Interval for a Population Proportion (6.2)	Identify an appropriate confidence interval procedure for a population proportion.	The appropriate confidence interval procedure for a one-sample proportion for one categorical variable is a one sample z-interval for a proportion.
	Verify the conditions for calculating confidence intervals for a population proportion.	In order to make assumptions necessary for inference on population proportions, means, and slopes, we must check for independence in data collection methods and for selection of the appropriate sampling distribution. In order to calculate a confidence interval to estimate a population proportion, p, we must check for independence and that the sampling distribution is approximately normal. a. To check for independence: i. Data should be collected using a random sample or a randomized experiment. ii. When sampling without replacement, check that $n \leq 10\% \ N$, where N is the size of the population. b. To check that the sampling distribution of p is approximately normal (shape): i. For categorical variables, check that both the number of successes, $n\hat{p}$, and the number of failures, $n(1-\hat{p})$ are at least 10 so that the sample size is large enough to support an assumption of normality.
	Determine the margin of error for a given sample size and an estimate for the sample size that will result in a given margin of error for a population proportion.	Based on sample data, the standard error of a statistic is an estimate for the standard deviation for the statistic. The standard error of $\hat{p}$ is: $$SE_{\hat{p}} = \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

		A margin of error gives how much a value of a sample statistic is likely to vary from the value of the corresponding population parameter. For categorical variables, the margin of error is the critical value (z*) times the standard error (SE) of the relevant statistic, which equals $z^* \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$ for a one-sample proportion. The formula for margin of error can be rearranged to solve for *n*, the minimum sample size needed to achieve a given margin of error. For this purpose, use a guess for *p* or use *p* = 0.5 in order to find an upper bound for the sample size that will result in a given margin of error.
	Calculate an appropriate confidence interval for a population proportion.	In general, an interval estimate can be constructed as point estimate ± (margin of error). For a one-sample proportion, the interval estimate is: $$\hat{p} \pm z^* \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$ Critical values represent the boundaries encompassing the middle C% of the standard normal distribution, where C% is an approximate confidence level for a proportion.
	Calculate an interval estimate based on a confidence interval for a population proportion.	Confidence intervals for population proportions can be used to calculate interval estimates with specified units.
Justifying a Claim Based on a Confidence Interval for a Population Proportion (6.3)	Interpret a confidence interval for a population proportion.	A confidence interval for a population proportion either contains the population proportion or it does not, because each interval is based on random sample data, which varies from sample to sample. We are C% confident that the confidence interval for a population proportion captures the population proportion.

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

		In repeated random sampling with the same sample size, approximately C% of confidence intervals created will capture the population proportion. Interpreting a confidence interval for a one- sample proportion should include a reference to the sample taken and details about the population it represents.
	Justify a claim based on a confidence interval for a population proportion.	A confidence interval for a population proportion provides an interval of values that may provide sufficient evidence to support a particular claim in context.
	Identify the relationships between sample size, width of a confidence interval, confidence level, and margin of error for a population proportion.	When all other things remain the same, the width of the confidence interval for a population proportion tends to decrease as the sample size increases. For a population proportion, the width of the interval is proportional to $\frac{1}{\sqrt{n}}$. For a given sample, the width of the confidence interval for a population proportion increases as the confidence level increases. The width of a confidence interval for a population proportion is exactly twice the margin of error.
Setting Up a Test for a Population Proportion (6.4)	Identify the null and alternative hypotheses for a population proportion.	The null hypothesis is the situation that is assumed to be correct unless evidence suggests otherwise, and the alternative hypothesis is the situation for which evidence is being collected. For hypotheses about parameters, the null hypothesis contains an equality reference ($=, \leq,$ or $\geq$), while the alternative hypothesis contains a strict inequality ($<, >, or \neq$). The type of inequality in the alternative hypothesis is based on the question of interest. Alternative hypotheses with $<$ or $>$ are called one-sided, and alternative hypotheses with $\neq$ are called two- sided. Although the null hypothesis for a one-sided test may include an inequality symbol, it is still tested at the boundary of equality.

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

		The null hypothesis for a population proportion is: $H_0: p = p_0$, where p_0 is the null hypothesized value for the population proportion. A one-sided alternative hypothesis for a proportion is either $H_a: p < p_0$ or $H_a: p > p_0$. A two-sided alternate hypothesis is $H_a: p \neq p_0$. For a one-sample *z*-test for a population proportion, the null hypothesis specifies a value for the population proportion, usually one indicating no difference or effect.
	Identify an appropriate testing method for a population proportion.	For a single categorical variable, the appropriate testing method for a population proportion is a one-sample *z*-test for a population proportion.
	Verify the conditions for making statistical inferences when testing a population proportion.	In order to make statistical inferences when testing a population proportion, we must check for independence and that the sampling distribution is approximately normal: a. To check for independence: i. Data should be collected using a random sample or a randomized experiment. ii. When sampling without replacement, check that $n \leq 10\%\ N$. b. To check that the sampling distribution of *p* is approximately normal (shape): i. Assuming that H_0 is true ($p = p_0$), verify that both the number of successes, np_0, and the number of failures, $n(1 - p_0)$ are at least 10 so that that the sample size is large enough to support an assumption of normality.
Interpreting *p*-Values (6.5)	Calculate an appropriate test statistic and p-value for a population proportion.	The distribution of the test statistic assuming the null hypothesis is true (null distribution) can be either a randomization distribution or when a probability model is assumed to be true, a theoretical distribution (*z*). When using a z-test, the standardized test statistic can be written: $\text{test statistic} = \frac{\text{sample statistic null value of the parameter}}{\text{standard deviation of the statistic}}$

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

		This is called a *z*-statistic for proportions. – – – The test statistic for a population proportion is: $$z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}$$ – – – A *p*-value is the probability of obtaining a test statistic as extreme or more extreme than the observed test statistic when the null hypothesis and probability model are assumed to be true. The significance level may be given or determined by the researcher.
	Interpret the *p*-value of a significance test for a population proportion.	The p-value is the proportion of values for the null distribution that are as extreme or more extreme than the observed value of the test statistic. This is: a. The proportion at or above the observed value of the test statistic, if the alternative is >. b. The proportion at or below the observed value of the test statistic, if the alternative is <. c. The proportion less than or equal to the negative of the absolute value of the test statistic plus the proportion greater than or equal to the absolute value of the test statistic, if the alternative is ≠. – – – An interpretation of the *p*-value of a significance test for a one-sample proportion should recognize that the *p*-value is computed by assuming that the probability model and null hypothesis are true, i.e., by assuming that the true population proportion is equal to the particular value stated in the null hypothesis.
Concluding a Test for a Population Proportion (6.6)	Justify a claim about the population based on the results of a significance test for a population proportion.	The significance level, α, is the predetermined probability of rejecting the null hypothesis given that it is true. – – –

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

		A formal decision explicitly compares the p-value to the significance level, α. If the p-value $\leq \alpha$, reject the null hypothesis. If the p-value $> \alpha$, fail to reject the null hypothesis. — — — Rejecting the null hypothesis means there is sufficient statistical evidence to support the alternative hypothesis. Failing to reject the null means there is insufficient statistical evidence to support the alternative hypothesis. — — — The conclusion about the alternative hypothesis must be stated in context. — — — A significance test can lead to rejecting or not rejecting the null hypothesis, but can never lead to concluding or proving that the null hypothesis is true. Lack of statistical evidence for the alternative hypothesis is not the same as evidence for the null hypothesis. — — — Small p-values indicate that the observed value of the test statistic would be unusual if the null hypothesis and probability model were true, and so provide evidence for the alternative. The lower the p-value, the more convincing the statistical evidence for the alternative hypothesis. — — — p-values that are not small indicate that the observed value of the test statistic would not be unusual if the null hypothesis and probability model were true, so do not provide convincing statistical evidence for the alternative hypothesis nor do they provide evidence that the null hypothesis is true. — — — A formal decision explicitly compares the p-value to the significance α. If the p-value $\leq \alpha$, then reject the null hypothesis, $H_0: p = p_0$. If the p-value $> \alpha$, then fail to reject the null hypothesis. — — —

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

<table>
<tr><td></td><td></td><td>The results of a significance test for a population proportion can serve as the statistical reasoning to support the answer to a research question about the population that was sampled.</td></tr>
<tr><td>Potential Errors When Performing Tests (6.7)</td><td>Identify Type I and Type II errors.</td><td>A Type I error occurs when the null hypothesis is true and is rejected (false positive).

A Type II error occurs when the null hypothesis is false and is not rejected (false negative).

Table of Errors
<table>
<tr><th></th><th></th><th colspan="2">Actual Population Value</th></tr>
<tr><th></th><th></th><th>H_0 true</th><th>H_a true</th></tr>
<tr><td rowspan="2">Decision</td><td>Reject H_0</td><td>Type I Error</td><td>Correct Decision</td></tr>
<tr><td>Fail to Reject H_0</td><td>Correct Decision</td><td>Type II Error</td></tr>
</table></td></tr>
<tr><td></td><td>Calculate the probability of a Type I and Type II errors.</td><td>The significance level, α, is the probability of making a Type I error, if the null hypothesis is true.

The power of a test is the probability that a test will correctly reject a false null hypothesis.

The probability of making a Type II error = 1 − *power*.</td></tr>
<tr><td></td><td>Identify factors that affect the probability of errors in significance testing.</td><td>The probability of a Type II error decreases when any of the following occurs, provided the others do not change:
i. Sample size(s) increases.
ii. Significance level (α) of a test increases.
iii. Standard error decreases.
iv. True parameter value is farther from the null.</td></tr>
</table>

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

	Interpret Type I and Type II errors.	Whether a Type I or a Type II error is more consequential depends upon the situation. Since the significance level, α, is the probability of a Type I error, the consequences of a Type I error influence decisions about a significance level.
Confidence Intervals for the Difference of Two Proportions (6.8)	Identify an appropriate confidence interval procedure for a comparison of population proportions.	The appropriate confidence interval procedure for a two-sample comparison of proportions for one categorical variable is a two-sample *z*-interval for a difference between population proportions.
	Verify the conditions for calculating confidence intervals for a difference between population proportions.	In order to calculate confidence intervals to estimate a difference between proportions, we must check for independence and that the sampling distribution is approximately normal: a. To check for independence: i. Data should be collected using random samples or a randomized experiment. ii. When sampling without replacement, check that $n_1 \leq 10\%\ N_1$ and $n_2 \leq 10\%\ N_2$. b. To check that sampling distribution of $p_1 - p_2$ is approximately normal (shape). i. For categorical variables, check that $n_1\hat{p}_1, n_1(1-\hat{p}_1), n_2\hat{p}_2, \& \ n_2(1-\hat{p}_2)$ are all greater than or equal to some predetermined value, typically either 5 or 10.
	Calculate an appropriate confidence interval for a comparison of population proportions.	For a comparison of proportions, the interval estimate is $(\hat{p}_1 - \hat{p}_2) \pm z^* \sqrt{\frac{\hat{p}_1(1-\hat{p}_1)}{n_1} + \frac{\hat{p}_2(1-\hat{p}_2)}{n_2}}$
	Calculate an interval estimate based on a confidence interval for a difference of proportions.	Confidence intervals for a difference in proportions can be used to calculate interval estimates with specified units.

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

Justifying a Claim Based on a Confidence Interval for a Difference of Population Proportions (6.9)	Interpret a confidence interval for a difference of proportions.	In repeated random sampling with the same sample size, approximately C% of confidence intervals created will capture the difference in population proportions. – – – Interpreting a confidence interval for difference between population proportions should include a reference to the sample taken and details about the population it represents.
	Justify a claim based on a confidence interval for a difference of proportions.	A confidence interval for difference in population proportions provides an interval of values that may provide sufficient evidence to support a particular claim in context.
Setting Up a Test for the Difference of Two Population Proportion (6.10)	Identify the null and alternative hypotheses for a difference of two population proportions.	For a two-sample test for a difference of two proportions, the null hypothesis specifies a value of 0 for the difference in population proportions, indicating no difference or effect. – – – The null hypothesis for a difference in proportions is: $H_0: p_1 = p_2$ or $H_0: (p_1 - p_2) = 0$. – – – A one-sided alternative hypothesis for a difference in proportions is $H_a: p_1 < p_2$, or, $H_a: p_1 > p_2$. A two-sided alternative hypothesis for a difference of proportions is $H_a: p_1 \neq p_2$.
	Identify an appropriate testing method for the difference of two population proportions.	For a single categorical variable, the appropriate testing method for the difference of two population proportions is a two-sample *z*-test for a difference between two population proportions.
	Verify the conditions for making statistical inferences when testing a difference of two population proportions.	In order to make statistical inferences when testing a difference between population proportions, we must check for independence and that the sampling distribution is approximately normal: a. To check for independence: i. Data should be collected using random samples or a randomized experiment. ii. When sampling without replacement,

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

		check that $n_1 \leq 10\%\ N_1$ and $n_2 \leq 10\%\ N_2$. b. To check that sampling distribution of $\hat{p}_1 - \hat{p}_2$ is approximately normal (shape): i. For the combined sample, define the combined (or pooled) proportion, $\hat{p}_c = \frac{n_1\hat{p}_1 + n_2\hat{p}_2}{n_1 + n_2}$. Assuming that H_0 is true ($p_1 - p_2 = 0$ or $p_1 = p_2$), check that $n_1\hat{p}_c, n_1(1 - \hat{p}_c), n_2\hat{p}_c, n_2(1 - \hat{p}_c)$ are all greater than or equal to some predetermined value, typically either 5 or 10.

§ 6: Inference for Categorical Data: Proportions ★

This unit introduces statistical inference, which will continue through the end of the course. Students will analyze categorical data to make inferences about binomial population proportions. Provided conditions are met, students will use statistical inference to construct and interpret confidence intervals to estimate population proportions and perform significance tests to evaluate claims about population proportions. Students begin by learning inference procedures for one proportion and then examine inference methods for a difference between two proportions. They will also interpret the two types of errors that can be made in a significance test, their probabilities, and possible consequences in context.

Carrying Out a Test for the Difference of Two Population Proportions (6.11)	Calculate an appropriate test statistic for the difference of two population proportions.	The test statistic for a difference in proportions is: $$z = \frac{(\hat{p}_1 - \hat{p}_2) - 0}{\sqrt{\hat{p}_c(1-\hat{p}_c)}\sqrt{\frac{1}{n_1}+\frac{1}{n_2}}},$$ where $\hat{p}_c = \frac{n_1\hat{p}_1 + n_2\hat{p}_2}{n_1+n_2}$.
	Interpret the p-value of a significance test for a difference of population proportions.	An interpretation of the *p*-value of a significance test for a difference of two population proportions should recognize that the *p*-value is computed by assuming that the null hypothesis is true, i.e., by assuming that the true population proportions are equal to each other.
	Justify a claim about the population based on the results of a significance test for a difference of population proportions.	A formal decision explicitly compares the *p*-value to the significance α. If the p-value ≤ α, then reject the null hypothesis, $H_0: p_1 = p_2$, or $H_0: (p_1 - p_2) = 0$. If the p-value > α, then fail to reject the null hypothesis. The results of a significance test for a difference of two population proportions can serve as the statistical reasoning to support the answer to a research question about the two populations that were sampled.

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and t-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, t, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a t-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Why Should I Worry About Error? (7.1)	Identify questions suggested by probabilities of errors in statistical inference.	Random variation may result in errors in statistical inference.
Constructing a Confidence Interval for a Population Mean (7.2)	Describe t-distributions.	When s is used instead of σ to calculate a test statistic, the corresponding distribution, known as the t-distribution, varies from the normal distribution in shape, in that more of the area is allocated to the tails of the density curve than in a normal distribution. As the degrees of freedom increase, the area in the tails of a t-distribution decreases.
	Identify an appropriate confidence interval procedure for a population mean, including the mean difference between values in matched pairs.	Because σ is typically not known for distributions of quantitative variables, the appropriate confidence interval procedure for estimating the population mean of one quantitative variable for one sample is a one-sample t-interval for a mean. For one quantitative variable, X, that is normally distributed, the distribution of $t = \frac{(\bar{x}-\mu)}{\frac{s}{\sqrt{n}}}$ is a t-distribution with $n-1$ degrees of freedom. Matched pairs can be thought of as one sample of pairs. Once differences between pairs of values are found, inference for confidence intervals proceeds as for a population mean.
	Verify the conditions for calculating confidence intervals for a population mean, including the mean difference between values in matched pairs.	In order to calculate confidence intervals to estimate a population mean, we must check for independence and that the sampling distribution is approximately normal: a. To check for independence: i. Data should be collected using a random sample or a randomized experiment. ii. When sampling without replacement, check that $n \leq 10\%\ N$, where N is the size of the population.

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and *t*-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, *t*, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a *t*-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

		b. To check that the sampling distribution of x is approximately normal (shape): i. If the observed distribution is skewed, n should be greater than 30. ii. If the sample size is less than 30, the distribution of the sample data should be free from strong skewness and outliers.
	Determine the margin of error for a given sample size for a one-sample *t*-interval.	The critical value t^* with $n-1$ degrees of freedom can be found using a table or computer-generated output. — — — The standard error for a sample mean is given by $SE=\frac{s}{\sqrt{n}}$, where s is the sample standard deviation. — — — For a one-sample *t*-interval for a mean, the margin of error is the critical value (t^*) times the standard error (SE), which equals $t^*\left(\frac{s}{\sqrt{n}}\right)$.
	Calculate an appropriate confidence interval for a population mean, including the mean difference between values in matched pairs.	The point estimate for a population mean is the sample mean, $\bar{x}$. — — — For the population mean for one sample with unknown population standard deviation, the confidence interval is $\bar{x} \pm t^* \frac{s}{\sqrt{n}}$.
Justifying a Claim About a Population Mean Based on a Confidence Interval (7.3)	Interpret a confidence interval for a population mean, including the mean difference between values in matched pairs.	A confidence interval for a population mean either contains the population mean or it does not, because each interval is based on data from a random sample, which varies from sample to sample. — — — We are C% confident that the confidence interval for a population mean captures the population mean. — — — An interpretation of a confidence interval for a population mean includes a reference to the sample taken and details about the population it represents.
	Justify a claim based on a confidence interval for a population mean,	A confidence interval for a population mean provides an interval of values that may provide sufficient evidence to support a particular claim in context.

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and t-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, t, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a t-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

	including the mean difference between values in matched pairs.	
	Identify the relationships between sample size, width of a confidence interval, confidence level, and margin of error for a population mean.	When all other things remain the same, the width of a confidence interval for a population mean tends to decrease as the sample size increases. For a single mean, the width of the interval is proportional to $\frac{1}{\sqrt{n}}$. For a given sample, the width of the confidence interval for a population mean increases as the confidence level increases.
Setting Up a Test for a Population Mean (7.4)	Identify an appropriate testing method for a population mean with unknown σ, including the mean difference between values in matched pairs.	The appropriate test for a population mean with unknown σ is a one-sample t-test for a population mean. Matched pairs can be thought of as one sample of pairs. Once differences between pairs of values are found, inference for significance testing proceeds as for a population mean.
	Identify the null and alternative hypotheses for a population mean with unknown σ , including the mean difference between values in matched pairs.	The null hypothesis for a one-sample t-test for a population mean is $H_0: \mu = \mu_0$, where μ_0 is the hypothesized value. Depending upon the situation, the alternative hypothesis is $H_a: \mu < \mu_0$, or $H_a: \mu > \mu_0$, or $H_a: \mu \neq \mu_0$. When finding the mean difference, μ_d , between values in a matched pair, it is important to define the order of subtraction.
	Verify the conditions for the test for a population mean, including the mean difference between values in matched pairs.	In order to make statistical inferences when testing a population mean, we must check for independence and that the sampling distribution is approximately normal: a. To check for independence: i. Data should be collected using a random sample or a randomized experiment. ii. When sampling without replacement, check that $n \leq 10\% \ N$. b. To check that the sampling distribution of x is approximately normal (shape):

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and *t*-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, *t*, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a *t*-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

		i. If the observed distribution is skewed, *n* should be greater than 30. ii. If the sample size is less than 30, the distribution of the sample data should be free from strong skewness and outliers.
Carrying Out a Test for a Population Mean (7.5)	Calculate an appropriate test statistic for a population mean, including the mean difference between values in matched pairs.	For a single quantitative variable when random sampling with replacement from a population that can be modeled with a normal distribution with mean μ and standard deviation σ , the sampling distribution of $t = \frac{\bar{x}-\mu}{\frac{s}{\sqrt{n}}}$ has a *t*-distribution with $n - 1$ degrees of freedom.
	Interpret the *p*-value of a significance test for a population mean, including the mean difference between values in matched pairs.	An interpretation of the *p*-value of a significance test for a population mean should recognize that the *p*-value is computed by assuming that the null hypothesis is true, i.e., by assuming that the true population mean is equal to the particular value stated in the null hypothesis.
	Justify a claim about the population based on the results of a significance test for a population mean.	A formal decision explicitly compares the p-value to the significance α . If the p-value $\leq \alpha$, then reject the null hypothesis, $H_0: \mu = \mu_0$. If the p-value > α , then fail to reject the null hypothesis. The results of a significance test for a population mean can serve as the statistical reasoning to support the answer to a research question about the population that was sampled.
Confidence Intervals for the Difference of Two Means (7.6)	Identify an appropriate confidence interval procedure for a difference of two population means.	Consider a simple random sample from population 1 of size n_1 , mean μ_1 , and standard deviation σ_1 and a second simple random sample from population 2 of size n_2 , mean μ_2 , and standard deviation σ_2 . If the distributions of populations 1 and 2 are normal or if both n_1 and n_2 are greater than 30, then the sampling distribution of the difference of means, *x* is also normal. The mean for

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and *t*-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, *t*, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a *t*-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

		the sampling distribution of $\bar{x}_1 - \bar{x}_2$ is $\mu_1 - \mu_2$. The standard deviation of $\bar{x}_1 - \bar{x}_2$ is $\sqrt{\frac{(\sigma_1)^2}{n_1} + \frac{(\sigma_2)^2}{n_2}}$. The appropriate confidence interval procedure for one quantitative variable for two independent samples is a two-sample *t*-interval for a difference between population means.
	Verify the conditions to calculate confidence intervals for the difference of two population means.	In order to calculate confidence intervals to estimate a difference of population means, we must check for independence and that the sampling distribution is approximately normal: a. To check for independence: i. Data should be collected using two independent, random samples or a randomized experiment. ii. When sampling without replacement, check that $n_1 \leq 10\%\ N_1$ and $n_2 \leq 10\%\ N_2$. b. To check that the sampling distribution of $(\bar{x}_1 - \bar{x}_2)$ should be approximately normal (shape): i. If the observed distributions are skewed, both n_1 and n_2 should be greater than 30.
	Determine the margin of error for the difference of two population means.	For the difference of two sample means, the margin of error is the critical value (t*) times the standard error (*SE*) of the difference of two means. The standard error for the difference in two sample means with sample standard deviations, s_1 and s_2 , is $\sqrt{\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}}$.

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and t-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, t, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a t-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

	Calculate an appropriate confidence interval for a difference of two population means.	The point estimate for the difference of two population means is the difference in sample means, $\bar{x}_1 - \bar{x}_2$. For a difference of two population means where the population standard deviations are not known, the confidence interval is $(\bar{x}_1 - \bar{x}_2) \pm t^* \sqrt{\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}}$ where $\pm t^*$ are the critical values for the central C% of a t-distribution with appropriate degrees of freedom that can be found using technology.
Justifying a Claim About the Difference of Two Means Based on a Confidence Interval (7.7)	Interpret a confidence interval for a difference of population means.	In repeated random sampling with the same sample size, approximately C% of confidence intervals created will capture the difference of population means. An interpretation for a confidence interval for the difference of two population means should include a reference to the samples taken and details about the populations they represent.
	Justify a claim based on a confidence interval for a difference of population means.	A confidence interval for a difference of population means provides an interval of values that may provide sufficient evidence to support a particular claim in context.
	Identify the effects of sample size on the width of a confidence interval for the difference of two means.	When all other things remain the same, the width of the confidence interval for the difference of two means tends to decrease as the sample sizes increase.
Setting Up a Test for the Difference of Two Population Means (7.8)	Identify an appropriate selection of a testing method for a difference of two population means.	For a quantitative variable, the appropriate test for a difference of two population means is a two-sample t-test for a difference of two population means.
	Identify the null and alternative hypotheses for a difference of two population means.	The null hypothesis for a two-sample t-test for a difference of two population means, μ_1 and μ_2 is $H_0: \mu_1 - \mu_2 = 0$ or $H_0: \mu_1 = \mu_2$. The alternative hypothesis is $H_a: \mu_1 - \mu_2 < 0$, $H_a: \mu_1 - \mu_2 > 0$, or $H_a: \mu_1 - \mu_2 \neq 0$, or $H_a: \mu_1 < \mu_2$, $H_a: \mu_1 > \mu_2$, or $H_a: \mu_1 \neq \mu_2$.

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and *t*-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, *t*, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a *t*-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

	Verify the conditions for the significance test for the difference of two population means.	In order to make statistical inferences when testing a difference between population means, we must check for independence and that the sampling distribution is approximately normal: a. Individual observations should be independent: i. Data should be collected using simple random samples or a randomized experiment. ii. When sampling without replacement, check that $n_1 \leq 10\%\ N_1$ and $n_2 \leq 10\%\ N_2$. b. To check that the sampling distribution of $(\bar{x}_1 - \bar{x}_2)$ should be approximately normal (shape): i. If the observed distributions are skewed, both n_1 and n_2 should be greater than 30. ii. If the sample size is less than 30, the distribution of the sample data should be free from strong skewness and outliers. This should be checked for BOTH samples.
Carrying Out a Test for the Difference of Two Population Means (7.9)	Calculate an appropriate test statistic for a difference of two means.	For a single quantitative variable, data collected using independent random samples or a randomized experiment from two populations, each of which can be modeled with a normal distribution, the sampling distribution of $t = \frac{(\bar{x}_1 - \bar{x}_2) - (\mu_1 - \mu_2)}{\sqrt{\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}}}$ is an approximate *t*-distribution with degrees of freedom that can be found using technology. The degrees of freedom fall between the smaller of $n_1 - 1$ and $n_2 - 1$ and $n_1 + n_2 - 2$.

§ 7: Inference for Quantitative Data: Means ★

In this unit, students will analyze quantitative data to make inferences about population means. Students should understand that t^* and *t*-tests are used for inference with means when the population standard deviation, σ, is not known. Using s for σ in the formula for z gives a slightly different value, *t*, whose distribution, which depends on sample size, has more area in the tails than a normal distribution. The boundaries for rejecting a null hypothesis using a *t*-distribution tend to be further from the mean than for a normal distribution. Students should understand how and why conditions for inference with proportions and means are similar and different.

	Interpret the *p*-value of a significance test for a difference of population means.	An interpretation of the *p*-value of a significance test for a two-sample difference of population means should recognize that the *p*-value is computed by assuming that the null hypothesis is true, i.e., by assuming that the true population means are equal to each other.
	Justify a claim about the population based on the results of a significance test for a difference of the two population means in context.	A formal decision explicitly compares the p-value to the significance α . If the p-value $\leq \alpha$, then reject the null hypothesis, $H_0: \mu = \mu_0$. If the p-value > α , then fail to reject the null hypothesis. The results of a significance test for a two-sample test for a difference between two population means can serve as the statistical reasoning to support the answer to a research question about the populations that were sampled.

§ 8: Inference for Categorical Data: Chi-Square ★

§ 6 introduced inference for proportions of categorical data. § 8 introduces chi-square tests, which can be used when there are two or more categories. Students need to understand how to select from the following tests: the chi-square test for goodness of fit (for a distribution of proportions of one categorical variable in a population), the chi-square test for independence (for associations between categorical variables within a single population), or the chi-square test for homogeneity (for comparing distributions of a categorical variable across populations or treatments). To integrate conceptual understanding, teachers can make connections between frequency tables, conditional probability, and calculating expected counts. The chi-square statistic is introduced to measure the distance between observed and expected counts relative to expected counts.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Are My Results Unexpected? (8.1)	Identify questions suggested by variation between observed and expected counts in categorical data.	Variation between what we find and what we expect to find may be random or not.
Setting Up a Chi-Square Goodness of Fit Test (8.2)	Describe chi-square distributions.	Expected counts of categorical data are counts consistent with the null hypothesis. In general, an expected count is a sample size times a probability. The chi-square statistic measures the distance between observed and expected counts relative to expected counts. Chi-square distributions have positive values and are skewed right. Within a family of density curves, the skew becomes less pronounced with increasing degrees of freedom.
	Identify the null and alternative hypotheses in a test for a distribution of proportions in a set of categorical data.	For a chi-square goodness-of-fit test, the null hypothesis specifies null proportions for each category, and the alternative hypothesis is that at least one of these proportions is not as specified in the null hypothesis.
	Identify an appropriate testing method for a distribution of proportions in a set of categorical data.	When considering a distribution of proportions for one categorical variable, the appropriate test is the chi-square test for goodness of fit.
	Calculate expected counts for the chi-square test for goodness of fit.	Expected counts for a chi-square goodness-of-fit test are (sample size) × (null proportion).
	Verify the conditions for making statistical inferences when testing goodness of fit for a chi-square distribution.	In order to make statistical inferences for a chi-square test for goodness of fit we must check the following: a. To check for independence: i. Data should be collected using a random sample or randomized experiment. ii. When sampling without replacement,

§ 8: Inference for Categorical Data: Chi-Square ★

§ 6 introduced inference for proportions of categorical data. § 8 introduces chi-square tests, which can be used when there are two or more categories. Students need to understand how to select from the following tests: the chi-square test for goodness of fit (for a distribution of proportions of one categorical variable in a population), the chi-square test for independence (for associations between categorical variables within a single population), or the chi-square test for homogeneity (for comparing distributions of a categorical variable across populations or treatments). To integrate conceptual understanding, teachers can make connections between frequency tables, conditional probability, and calculating expected counts. The chi-square statistic is introduced to measure the distance between observed and expected counts relative to expected counts.

		check that $n \leq 10\% \ N$. b. The chi-square test for goodness of fit becomes more accurate with more observations, so large counts should be used (shape). i. A conservative check for large counts is that all expected counts should be greater than 5.

§ 8: Inference for Categorical Data: Chi-Square ★

§ 6 introduced inference for proportions of categorical data. § 8 introduces chi-square tests, which can be used when there are two or more categories. Students need to understand how to select from the following tests: the chi-square test for goodness of fit (for a distribution of proportions of one categorical variable in a population), the chi-square test for independence (for associations between categorical variables within a single population), or the chi-square test for homogeneity (for comparing distributions of a categorical variable across populations or treatments). To integrate conceptual understanding, teachers can make connections between frequency tables, conditional probability, and calculating expected counts. The chi-square statistic is introduced to measure the distance between observed and expected counts relative to expected counts.

Carrying Out a Chi-Square Test for Goodness of Fit (8.3)	Calculate the appropriate statistic for the chi-square test for goodness of fit.	The test statistic for the chi-square test for goodness of fit is $X^2 = \sum \frac{(Observed\ count - Expected\ count)^2}{Expected\ count}$ with $degrees\ of\ freedom = number\ of\ categories - 1.$ The distribution of the test statistic assuming the null hypothesis is true (null distribution) can be either a randomization distribution or, when a probability model is assumed to be true, a theoretical distribution (chi-square).
	Determine the *p*-value for chi-square test for goodness of fit significance test.	The *p*-value for a chi-square test for goodness of fit for a number of degrees of freedom is found using the appropriate table or computer generated output.
	Interpret the *p*-value for the chi-square test for goodness of fit.	An interpretation of the *p*-value for the chi-square test for goodness of fit is the probability, given the null hypothesis and probability model are true, of obtaining a test statistic as, or more, extreme than the observed value.
	Justify a claim about the population based on the results of a chi-square test for goodness of fit.	A decision to either reject or fail to reject the null hypothesis is based on comparison of the p-value to the significance level, α. The results of a chi-square test for goodness of fit can serve as the statistical reasoning to support the answer to a research question about the population that was sampled.
Expected Counts in Two-Way Tables (8.4)	Calculate expected counts for two-way tables of categorical data.	The expected count in a particular cell of a two-way table of categorical data can be calculated using the formula: $expected\ count = \frac{(row\ total)(column\ total)}{table\ total}$
Setting Up a Chi-Square Test for Homogeneity or Independence (8.5)	Identify the null and alternative hypotheses for a chi-square test for homogeneity or independence.	The appropriate hypotheses for a chi-square test for homogeneity are: H_0: There is no difference in distributions of a categorical variable across populations or treatments. H_a: There is a difference in distributions of a categorical variable across populations or treatments.

§ 8: Inference for Categorical Data: Chi-Square ★

§ 6 introduced inference for proportions of categorical data. § 8 introduces chi-square tests, which can be used when there are two or more categories. Students need to understand how to select from the following tests: the chi-square test for goodness of fit (for a distribution of proportions of one categorical variable in a population), the chi-square test for independence (for associations between categorical variables within a single population), or the chi-square test for homogeneity (for comparing distributions of a categorical variable across populations or treatments). To integrate conceptual understanding, teachers can make connections between frequency tables, conditional probability, and calculating expected counts. The chi-square statistic is introduced to measure the distance between observed and expected counts relative to expected counts.

		The appropriate hypotheses for a chi-square test for independence are: H_0: There is no association between two categorical variables in a given population or the two categorical variables are independent. H_a: Two categorical variables in a population are associated or dependent.
	Identify an appropriate testing method for comparing distributions in two-way tables of categorical data.	When comparing distributions to determine whether proportions in each category for categorical data collected from different populations are the same, the appropriate test is the chi-square test for homogeneity. To determine whether row and column variables in a two-way table of categorical data might be associated in the population from which the data were sampled, the appropriate test is the chi-square test for independence.
	Verify the conditions for making statistical inferences when testing a chi-square distribution for independence or homogeneity.	In order to make statistical inferences for a chi-square test for two-way tables (homogeneity or independence), we must verify the following: a. To check for independence: i. For a test for independence: Data should be collected using a simple random sample. ii. For a test for homogeneity: Data should be collected using a stratified random sample or randomized experiment. iii. When sampling without replacement, check that $n \leq 10\% \ N$. b. The chi-square tests for independence and homogeneity become more accurate with more observations, so large counts should be used (shape). i. A conservative check for large counts is that all expected counts should be greater than 5.

§ 8: Inference for Categorical Data: Chi-Square ★

§ 6 introduced inference for proportions of categorical data. § 8 introduces chi-square tests, which can be used when there are two or more categories. Students need to understand how to select from the following tests: the chi-square test for goodness of fit (for a distribution of proportions of one categorical variable in a population), the chi-square test for independence (for associations between categorical variables within a single population), or the chi-square test for homogeneity (for comparing distributions of a categorical variable across populations or treatments). To integrate conceptual understanding, teachers can make connections between frequency tables, conditional probability, and calculating expected counts. The chi-square statistic is introduced to measure the distance between observed and expected counts relative to expected counts.

Carrying Out a Chi-Square Test for Homogeneity or Independence (8.6)	Calculate the appropriate statistic for a chi-square test for homogeneity or independence.	The appropriate test statistic for a chi-square test for homogeneity or independence is the chi-square statistic: $X^2 = \sum \frac{(Observed\ count - Expected\ count)^2}{Expected\ count}$ with degrees of freedom equal to: $(number\ of\ rows - 1)(number\ of\ columns - 1)$.
	Determine the p-value for a chi-square significance test for independence or homogeneity.	The *p*-value for a chi-square test for independence or homogeneity for a number of degrees of freedom is found using the appropriate table or technology. – – – For a test of independence or homogeneity for a two-way table, the *p*-value is the proportion of values in a chi-square distribution with appropriate degrees of freedom that are equal to or larger than the test statistic.
	Interpret the *p*-value for the chi-square test for homogeneity or independence.	An interpretation of the *p*-value for the chi-square test for homogeneity or independence is the probability, given the null hypothesis and probability model are true, of obtaining a test statistic as, or more, extreme than the observed value.
	Justify a claim about the population based on the results of a chi-square test for homogeneity or independence.	A decision to either reject or fail to reject the null hypothesis for a chi-square test for homogeneity or independence is based on comparison of the p-value to the significance level, α. – – – The results of a chi-square test for homogeneity or independence can serve as the statistical reasoning to support the answer to a research question about the population that was sampled (independence) or the populations that were sampled (homogeneity).

§ 9: Inference for Quantitative Data: Slopes ★

Students may be surprised to learn that there is variability in slope. In their experience in previous courses, the slope of the line of best fit does not vary for a particular set of bivariate quantitative data. However, suppose that every student in a university physics course collects data on spring length for 10 different hanging masses and calculates the least-squares regression line for their sample data. The students' slopes would likely vary as part of an approximately normal sampling distribution centered at the (true) slope of the population regression line relating spring length to hanging mass. In this unit, students will learn how to construct confidence intervals for and perform significance tests about the slope of a population regression line when appropriate conditions are met.

Topic	Learning Objective	Essential Knowledge
Introducing Statistics: Do Those Points Align? (9.1)	Identify questions suggested by variation in scatter plots.	Variation in points' positions relative to a theoretical line may be random or non-random.
Confidence Intervals for the Slope of a Regression Model (9.2)	Identify an appropriate confidence interval procedure for a slope of a regression model.	Consider a response variable, y, that is normally distributed with standard deviation, σ. The standard deviation σ can be estimated using the standard deviation of the residuals, $$s = \sqrt{\frac{\sum(y_i - \hat{y}_i)^2}{n-2}}$$ For a simple random sample of *n* observations, let *b* represent the slope of a sample regression line. Then the mean of the sampling distribution of *b* equals the population mean slope: $\mu_b = \beta$. The standard deviation of the sampling distribution for β is $\sigma_b = \frac{\sigma}{\sigma_x \sqrt{n}}$, where $\sigma_x = \sqrt{\frac{\sum(x_i - \bar{x})^2}{n}}$. The appropriate confidence interval for the slope of a regression model is a *t*-interval for the slope.
	Verify the conditions to calculate confidence intervals for the slope of a regression model.-1	In order to calculate a confidence interval to estimate the slope of a regression line, we must check the following: a. The true relationship between *x* and *y* is linear. Analysis of residuals may be used to verify linearity. b. The standard deviation for *y*, σ_y, does not vary with *x*. Analysis of residuals may be used to check for approximately equal standard deviations for all *x*. c. To check for independence: i. Data should be collected using a random sample or a randomized experiment. ii. When sampling without replacement,

§ 9: Inference for Quantitative Data: Slopes ★

Students may be surprised to learn that there is variability in slope. In their experience in previous courses, the slope of the line of best fit does not vary for a particular set of bivariate quantitative data. However, suppose that every student in a university physics course collects data on spring length for 10 different hanging masses and calculates the least-squares regression line for their sample data. The students' slopes would likely vary as part of an approximately normal sampling distribution centered at the (true) slope of the population regression line relating spring length to hanging mass. In this unit, students will learn how to construct confidence intervals for and perform significance tests about the slope of a population regression line when appropriate conditions are met.

		check that $n \leq 10\% \ N$. d. For a particular value of *x*, the responses (*y*-values) are approximately normally distributed. Analysis of graphical representations of residuals may be used to check for normality. i. If the observed distribution is skewed, *n* should be greater than 30.
	Determine the given margin of error for the slope of a regression model.	For the slope of a regression line, the margin of error is the critical value (t^*) times the standard error (*SE*) of the slope. The standard error for the slope of a regression line with sample standard deviation, *s*, is $SE = \frac{s}{s_x\sqrt{n-1}}$, where *s* is the estimate of σ and s_x is the sample standard deviation of the *x* values.
	Calculate an appropriate confidence interval for the slope of a regression model.	The point estimate for the slope of a regression model is the slope of the line of best fit, *b*. For the slope of a regression model, the interval estimate is $b \pm t^*(SE_b)$.
Justifying a Claim About the Slope of a Regression Model Based on a Confidence Interval (9.3)	Interpret a confidence interval for the slope of a regression model.	In repeated random sampling with the same sample size, approximately C% of confidence intervals created will capture the slope of the regression model, i.e., the true slope of the population regression model. An interpretation for a confidence interval for the slope of a regression line should include a reference to the sample taken and details about the population it represents.
	Justify a claim based on a confidence interval for the slope of a regression model.	A confidence interval for the slope of a regression model provides an interval of values that may provide sufficient evidence to support a particular claim in context.
	Identify the effects of sample size on the width of a confidence interval for the slope of a regression model.	When all other things remain the same, the width of the confidence interval for the slope of a regression model tends to decrease as the sample size increases.

§ 9: Inference for Quantitative Data: Slopes ★

Students may be surprised to learn that there is variability in slope. In their experience in previous courses, the slope of the line of best fit does not vary for a particular set of bivariate quantitative data. However, suppose that every student in a university physics course collects data on spring length for 10 different hanging masses and calculates the least-squares regression line for their sample data. The students' slopes would likely vary as part of an approximately normal sampling distribution centered at the (true) slope of the population regression line relating spring length to hanging mass. In this unit, students will learn how to construct confidence intervals for and perform significance tests about the slope of a population regression line when appropriate conditions are met.

Setting Up a Test for the Slope of a Regression Model (9.4)	Identify the appropriate selection of a testing method for a slope of a regression model.	The null hypothesis for a t-test for a slope is: $H_0: \beta = \beta_0$, where β_0 is the hypothesized value from the null hypothesis. The alternative hypothesis is $H_a: \beta < \beta_0$ or $H_a: \beta > \beta_0$, or $H_a: \beta \neq \beta_0$.
	Verify the conditions for the significance test for the slope of a regression model.	In order to make statistical inferences when testing for the slope of a regression model, we must check the following: a. The true relationship between i and *y* is linear. Analysis of residuals may be used to verify linearity. b. The standard deviation for *y*, σ_y, does not vary with *x*. Analysis of residuals may be used to check for approximately equal standard deviations for all *x*. c. To check for independence: i. Data should be collected using a random sample or a randomized experiment. ii. When sampling without replacement, check that $n \leq 10\% \ N$. d. For a particular value of *x*, the responses (*y*-values) are approximately normally distributed. Analysis of graphical representations of residuals may be used to check for normality. i. If the observed distribution is skewed, *n* should be greater than 30. ii. If the sample size is less than 30, the distribution of the sample data should be free from strong skewness and outliers.
Carrying Out a Test for the Slope of a Regression Model (9.5)	Calculate an appropriate test statistic for the slope of a regression model.	The distribution of the slope of a regression model assuming all conditions are satisfied and the null hypothesis is true (null distribution) is a *t*-distribution. For simple linear regression when random sampling from a population for the response that can be modeled with a normal distribution for each value of the explanatory variable, the sampling distribution of $t = \frac{b-\beta}{SE_b}$ has a

§ 9: Inference for Quantitative Data: Slopes ★

Students may be surprised to learn that there is variability in slope. In their experience in previous courses, the slope of the line of best fit does not vary for a particular set of bivariate quantitative data. However, suppose that every student in a university physics course collects data on spring length for 10 different hanging masses and calculates the least-squares regression line for their sample data. The students' slopes would likely vary as part of an approximately normal sampling distribution centered at the (true) slope of the population regression line relating spring length to hanging mass. In this unit, students will learn how to construct confidence intervals for and perform significance tests about the slope of a population regression line when appropriate conditions are met.

		t-distribution with degrees of freedom equal to $n-2$. When testing the slope in a simple linear regression model with one parameter, the slope, the test for the slope has $df = n-1$.
	Interpret the p-value of a significance test for the slope of a regression model.	An interpretation of the p-value of a significance test for the slope of a regression model should recognize that the p-value is computed by assuming that the null hypothesis is true, i.e., by assuming that the true population slope is equal to the particular value stated in the null hypothesis.
	Justify a claim about the population based on the results of a significance test for the slope of a regression model.	A formal decision explicitly compares the p-value to the significance α. If the p-value $\leq \alpha$, then reject the null hypothesis, $H_0: \beta = \beta_0$. If the p-value $> \alpha$, then fail to reject the null hypothesis. The results of a significance test for the slope of a regression model can serve as the statistical reasoning to support the answer to a research question about that sample.

§ 10: Final Project ★

A short description of the culminating final project, in narrative form, goes here.

Topic	Learning Objective	Essential Knowledge	Strategies/Resources