



WHERE WE'VE BEEN

The first seven chapters of this book have given you the building blocks you will need to understand statistical inference and how it can be applied in practical situations. The first three chapters were concerned with using descriptive statistics, both graphical and numerical, to describe and interpret sets of measurements. In the next three chapters, you learned about probability and probability distributions—the basic tools used to describe *populations* of measurements. The binomial and the normal distributions were emphasized as important for practical applications. The seventh chapter provided the link between probability and statistical inference. Many statistics are either sums or averages calculated from sample measurements. The Central Limit Theorem states that, even if the sampled populations are not normal, the sampling distributions of these *statistics* will be approximately normal when the sample size n is large. These statistics are the tools you use for *inferential statistics*—making inferences about a population using information contained in a sample.



WHERE WE'RE GOING— STATISTICAL INFERENCE

Inference—specifically, decision making and prediction—is centuries old and plays a very important role in most peoples' lives. Here are some applications:

- The government needs to predict short- and long-term interest rates.
- A broker wants to forecast the behavior of the stock market.
- A metallurgist wants to decide whether a new type of steel is more resistant to high temperatures than the current type.
- A consumer wants to estimate the selling price of her house before putting it on the market.

There are many ways to make these decisions or predictions, some subjective and some more objective in nature. How good will your predictions or decisions be? Although you may feel that your own built-in decision-making ability is quite good, experience suggests that this may not be the case. It is the job of the mathematical statistician to provide methods of statistical inference making that are better and more reliable than just subjective guesses.

Statistical inference is concerned with making decisions or predictions about **parameters**—the numerical descriptive measures that characterize a population. Three parameters you encountered in earlier chapters are the population mean μ , the population standard deviation σ , and the binomial proportion p . In statistical inference, a practical problem is restated in the framework of a population with a specific parameter of interest. For example, the metallurgist could measure the *average* coefficients of expansion for both types of steel and then compare their values.

Methods for making inferences about population parameters fall into one of two categories:

- **Estimation:** Estimating or predicting the value of the parameter
- **Hypothesis testing:** Making a decision about the value of a parameter based on some preconceived idea about what its value might be



Parameter \leftrightarrow Population
Statistic \leftrightarrow Sample

EXAMPLE

8.1

The circuits in computers and other electronics equipment consist of one or more printed circuit boards (PCB), and computers are often repaired by simply replacing one or more defective PCBs. In an attempt to find the proper setting of a plating process applied to one side of a PCB, a production supervisor might *estimate* the average thickness of copper plating on PCBs using samples from several days of operation. Since he has no knowledge of the average thickness μ before observing the production process, his is an *estimation* problem.

EXAMPLE

8.2

The supervisor in Example 8.1 is told by the plant owner that the thickness of the copper plating must not be less than .001 inch in order for the process to be in control. To decide whether or not the process is in control, the supervisor might formulate a test. He could *hypothesize* that the process is in control—that is, assume that the average thickness of the copper plating is .001 or greater—and use samples from several days of operation to decide whether or not his hypothesis is correct. The supervisor's decision-making approach is called a *test of hypothesis*.

Which method of inference should be used? That is, should the parameter be estimated, or should you test a hypothesis concerning its value? The answer is dictated by the practical question posed and is often determined by personal preference. Since both estimation and tests of hypotheses are used frequently in scientific literature, we include both methods in this and the next chapter.

A statistical problem, which involves planning, analysis, and inference making, is incomplete without a measure of the **goodness of the inference**. That is, how accurate or reliable is the method you have used? If a stockbroker predicts that the price of a stock will be \$80 next Monday, will you be willing to take action to buy or sell your stock without knowing how reliable her prediction is? Will the prediction be within \$1, \$2, or \$10 of the actual price next Monday? Statistical procedures are important because they provide two types of information:

- Methods for making the inference
- A numerical measure of the goodness or reliability of the inference

8.3

TYPES OF ESTIMATORS

To estimate the value of a population parameter, you can use information from the sample in the form of an **estimator**. Estimators are calculated using information from the sample observations, and hence, by definition they are also *statistics*.

Definition An **estimator** is a rule, usually expressed as a formula, that tells us how to calculate an estimate based on information in the sample.

Estimators are used in two different ways:

- **Point estimation:** Based on sample data, a single number is calculated to estimate the population parameter. The rule or formula that describes this calculation is called the **point estimator**, and the resulting number is called a **point estimate**.

- **Interval estimation:** Based on sample data, two numbers are calculated to form an interval within which the parameter is expected to lie. The rule or formula that describes this calculation is called the **interval estimator**, and the resulting pair of numbers is called an **interval estimate** or **confidence interval**.

EXAMPLE

8.3

A veterinarian wants to estimate the average weight gain per month of 4-month-old golden retriever pups that have been placed on a lamb and rice diet. The *population* consists of the weight gains per month of all 4-month-old golden retriever pups that are given this particular diet. The veterinarian wants to estimate the unknown parameter μ , the average monthly weight gain for this *hypothetical* population. One possible *estimator* based on sample data is the sample mean, $\bar{x} = \sum x_i/n$. It could be used in the form of a single number or *point estimate*—for instance, 3.8 pounds—or you could use an *interval estimate* and estimate that the average weight gain will be between 2.7 and 4.9 pounds.

Both point and interval estimation procedures use information provided by the sampling distribution of the specific estimator you have chosen to use. We will begin by discussing *point estimation* and its use in estimating population means and proportions.

POINT ESTIMATION

In a practical situation, there may be several statistics that could be used as point estimators for a population parameter. To decide which of several choices is best, you need to know how the estimator behaves in repeated sampling, described by its *sampling distribution*.

By way of analogy, think of firing a revolver at a target. The parameter of interest is the bull's-eye, at which you are firing bullets. Each bullet represents a single sample estimate, fired by the revolver, which represents the estimator. Suppose your friend fires a single shot and hits the bull's-eye. Can you conclude that he is an excellent shot? Would you stand next to the target while he fires a second shot? Probably not, because you have no measure of how well he performs in repeated trials. Does he always hit the bull's-eye, or is he consistently too high or too low? Do his shots cluster closely around the target, or do they consistently miss the target by a wide margin? Figure 8.1 shows several target configurations. Which target would you pick as belonging to the best shot?



TIP
parameter — target's bull's-eye
estimator — bullet or arrow


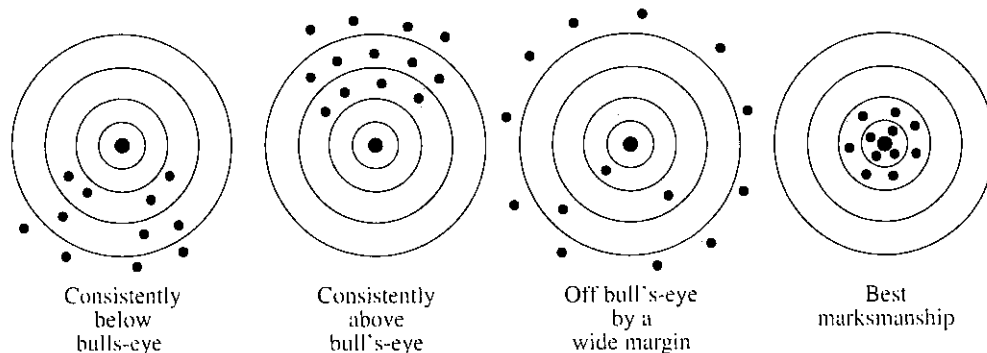


FIGURE 8.1
Which marksman is best?




Sampling distributions provide information that can be used to select the **best estimator**. What characteristics would be valuable? First, the **sampling distribution of the point estimator should be centered over the true value of the parameter to be estimated**. That is, the estimator should not consistently underestimate or overestimate the parameter of interest. Such an estimator is said to be **unbiased**.

Definition An estimator of a parameter is said to be **unbiased** if the mean of its distribution is equal to the true value of the parameter. Otherwise, the estimator is said to be **biased**.

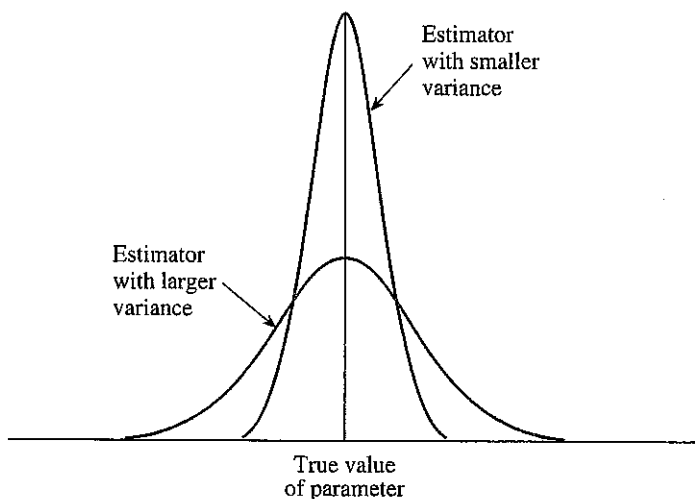
The sampling distributions for an unbiased estimator and a biased estimator are shown in Figure 8.2. The sampling distribution for the biased estimator is shifted to the right of the true value of the parameter. This biased estimator is more likely than an unbiased one to overestimate the value of the parameter.

FIGURE 8.2
Distributions for biased
and unbiased estimators



The second desirable characteristic of an estimator is that **the spread (as measured by the variance) of the sampling distribution should be as small as possible**. This ensures that, with a high probability, an individual estimate will fall close to the true value of the parameter. The sampling distributions for two unbiased estimators, one with a small variance[†] and the other with a larger variance, are shown in Figure 8.3.

FIGURE 8.3
Comparison of estimator
variability



[†]Statisticians usually use the term *variance of an estimator* when in fact they mean the variance of the sampling distribution of the estimator. This contractive expression is used almost universally.

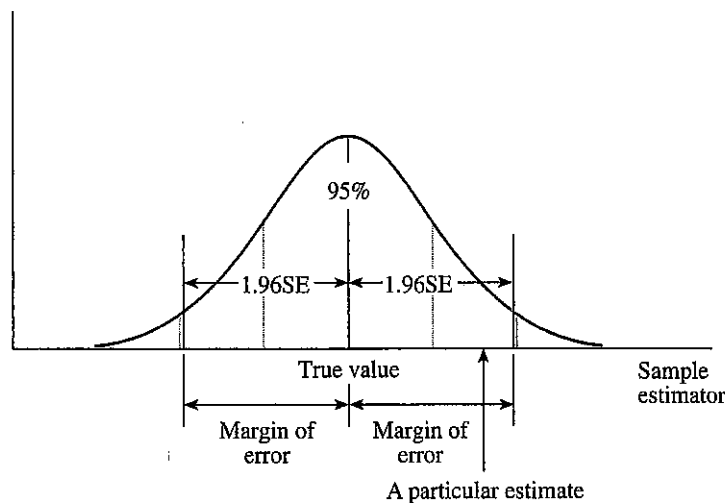
Naturally, you would prefer the estimator with the smaller variance because the estimates tend to lie closer to the true value of the parameter than in the distribution with the larger variance.

In real-life sampling situations, you may know that the sampling distribution of an estimator centers about the parameter that you are attempting to estimate, but all you have is the estimate computed from the n measurements contained in the sample. How far from the true value of the parameter will your estimate lie? How close is the marksman's bullet to the bull's-eye? The distance between the estimate and the true value of the parameter is called the **error of estimation**.

Definition The distance between an estimate and the estimated parameter is called the **error of estimation**.

In this chapter, you may assume that the sample sizes are always large and, therefore, that the *unbiased* estimators you will study have sampling distributions that can be approximated by a normal distribution (because of the Central Limit Theorem). Remember that, for any point estimator with a normal distribution, the Empirical Rule states that approximately 95% of all the point estimates will lie within two (or more exactly, 1.96) standard deviations of the mean of that distribution. For *unbiased* estimators, this implies that the difference between the point estimator and the true value of the parameter will be less than 1.96 standard deviations or 1.96 standard errors (SE). This quantity, called the 95% **margin of error** (or simply the “**margin of error**”), provides a practical upper bound for the error of estimation (see Figure 8.4). It is possible that the error of estimation will exceed this margin of error, but that is very unlikely.

FIGURE 8.4
Sampling distribution of
an unbiased estimator



95% Margin of error =
 $1.96 \times$ Standard error

POINT ESTIMATION OF A POPULATION PARAMETER

- Point estimator: a statistic calculated using sample measurements
- 95% Margin of error: $1.96 \times$ Standard error of the estimator

The sampling distributions for two unbiased point estimators were discussed in Chapter 7. It can be shown that both of these point estimators have the *minimum variability* of all unbiased estimators and are thus the *best estimators* you can find in each situation.

The variability of the estimator is measured using its standard error. However, you might have noticed that the standard error usually depends on unknown parameters such as σ or p . These parameters must be estimated using sample statistics such as s and \hat{p} . Although not exactly correct, experimenters generally refer to the estimated standard error as *the standard error*.

MY

PERSONAL TRAINER

How Do I Estimate a Population Mean or Proportion?

- To estimate the population mean μ for a quantitative population, the point estimator \bar{x} is *unbiased* with standard error estimated as

$$SE = \frac{s}{\sqrt{n}}^\dagger$$

The 95% margin of error when $n \geq 30$ is estimated as

$$\pm 1.96 \left(\frac{s}{\sqrt{n}} \right)$$

- To estimate the population proportion p for a binomial population, the point estimator $\hat{p} = x/n$ is *unbiased*, with standard error estimated as

$$SE = \sqrt{\frac{\hat{p}\hat{q}}{n}}$$

The 95% margin of error is estimated as

$$\pm 1.96 \sqrt{\frac{\hat{p}\hat{q}}{n}}$$

Assumptions: $n\hat{p} > 5$ and $n\hat{q} > 5$.

EXAMPLE

8.4

An environmentalist is conducting a study of the polar bear, a species found in and around the Arctic Ocean. Their range is limited by the availability of sea ice, which they use as a platform to hunt seals, the mainstay of their diet. The destruction of its habitat on the Arctic ice, which has been attributed to global warming, threatens the bear's survival as a species; it may become extinct within the century.¹ A random sample of $n = 50$ polar bears produced an average weight of 980 pounds with a standard deviation of 105 pounds. Use this information to estimate the average weight of all Arctic polar bears.

Solution The random variable measured is weight, a quantitative random variable best described by its mean μ . The point estimate of μ , the average weight of all Arctic polar bears, is $\bar{x} = 980$ pounds. The margin of error is estimated as

$$1.96 SE = 1.96 \left(\frac{s}{\sqrt{n}} \right) = 1.96 \left(\frac{105}{\sqrt{50}} \right) = 29.10 \approx 29 \text{ pounds}$$


[†]When you sample from a normal distribution, the statistic $(\bar{x} - \mu)/(s/\sqrt{n})$ has a t distribution, which will be discussed in Chapter 10. When the sample is *large*, this statistic is approximately normally distributed whether the sampled population is normal or nonnormal.

You can be fairly confident that the sample estimate of 980 pounds is within ± 29 pounds of the population mean.

In reporting research results, investigators often attach either the sample standard deviation s (sometimes called SD) or the standard error s/\sqrt{n} (usually called SE or SEM) to the estimates of population means. You should always look for an explanation somewhere in the text of the report that tells you whether the investigator is reporting $\bar{x} \pm SD$ or $\bar{x} \pm SE$. In addition, the sample means and standard deviations or standard errors are often presented as “error bars” using the graphical format shown in Figure 8.5.

FIGURE 8.5

Plot of treatment means and their standard errors



EXAMPLE

8.5

In addition to the average weight of the Arctic polar bear, the environmentalist from Example 8.4 is also interested in the opinions of adults on the subject of global warming. In particular, he wants to estimate the proportion of adults who think that global warming is a very serious problem. In a random sample of $n = 100$ adults, 73% of the sample indicated that, from what they have heard or read, global warming is a very serious problem. Estimate the true population proportion of adults who believe that global warming is a very serious problem, and find the margin of error for the estimate.

Solution The parameter of interest is now p , the proportion of individuals in the population who believe that global warming is a very serious problem. The best estimator of p is the sample proportion \hat{p} , which for this sample is $\hat{p} = .73$. In order to find the margin of error, you can approximate the value of p with its estimate $\hat{p} = .73$:

$$1.96 \text{ SE} = 1.96 \sqrt{\frac{\hat{p}\hat{q}}{n}} = 1.96 \sqrt{\frac{.73(.27)}{100}} = .09$$

With this margin of error, you can be fairly confident that the estimate of .73 is within $\pm .09$ of the true value of p . Hence, you can conclude that the true value of p could be as small as .64 or as large as .82. This margin of error is quite large when compared to the estimate itself and reflects the fact that large samples are required to achieve a small margin of error when estimating p .

TABLE 8.1 Some Calculated Values of \sqrt{pq}

p	pq	\sqrt{pq}	p	pq	\sqrt{pq}
.1	.09	.30	.6	.24	.49
.2	.16	.40	.7	.21	.46
.3	.21	.46	.8	.16	.40
.4	.24	.49	.9	.09	.30
.5	.25	.50			

Table 8.1 shows how the numerator of the standard error of \hat{p} changes for various values of p . Notice that, for most values of p —especially when p is between .3 and .7—there is very little change in \sqrt{pq} , the numerator of SE, reaching its maximum value when $p = .5$. This means that the margin of error using the estimator \hat{p} will also be a maximum when $p = .5$. Some pollsters routinely use the maximum margin of error—often called the **sampling error**—when estimating p , in which case they calculate

$$1.96 \text{ SE} = 1.96 \sqrt{\frac{.5(.5)}{n}} \quad \text{or sometimes} \quad 2 \text{ SE} = 2 \sqrt{\frac{.5(.5)}{n}}$$

Gallup, Harris, and Roper polls generally use sample sizes of approximately 1000, so their margin of error is

$$1.96 \sqrt{\frac{.5(.5)}{1000}} = .031 \quad \text{or approximately } 3\%$$

In this case, the estimate is said to be within ± 3 percentage points of the true population proportion.

8.4 EXERCISES

BASIC TECHNIQUES

8.1 Explain what is meant by “margin of error” in point estimation.

8.2 What are two characteristics of the best point estimator for a population parameter?

8.3 Calculate the margin of error in estimating a population mean μ for these values:

a. $n = 30, \sigma^2 = .2$

b. $n = 30, \sigma^2 = .9$

c. $n = 30, \sigma^2 = 1.5$

8.4 Refer to Exercise 8.3. What effect does a larger population variance have on the margin of error?

8.5 Calculate the margin of error in estimating a population mean μ for these values:

a. $n = 50, s^2 = 4$

b. $n = 500, s^2 = 4$

c. $n = 5000, s^2 = 4$

8.6 Refer to Exercise 8.5. What effect does an increased sample size have on the margin of error?

8.7 Calculate the margin of error in estimating a binomial proportion for each of the following values of n . Use $p = .5$ to calculate the standard error of the estimator.

a. $n = 30$

b. $n = 100$

c. $n = 400$

d. $n = 1000$

8.8 Refer to Exercise 8.7. What effect does increasing the sample size have on the margin of error?

8.9 Calculate the margin of error in estimating a binomial proportion p using samples of size $n = 100$ and the following values for p :

a. $p = .1$

b. $p = .3$

c. $p = .5$

d. $p = .7$

e. $p = .9$

f. Which of the values of p produces the largest margin of error?

- Describe the sampled population(s).
- Find a point estimate for the average room rate for the Marriott hotel chain. Calculate the margin of error.
- Find a point estimate for the average room rate for the Radisson hotel chain. Calculate the margin of error.
- Find a point estimate for the average room rate for the Wyndham hotel chain. Calculate the margin of error.
- Display the results of parts b, c, and d graphically, using the form shown in Figure 8.5. Use this display to compare the average room rates for the three hotel chains.

8.19 “900” Numbers Radio and television stations often air controversial issues during broadcast time and ask viewers to indicate their agreement or disagreement with a given stand on the issue. A poll is conducted by asking those viewers who *agree* to call a certain 900 telephone number and those who *disagree* to call a second 900 telephone number. All respondents pay a fee for their calls.

- Does this polling technique result in a random sample?
- What can be said about the validity of the results of such a survey? Do you need to worry about a margin of error in this case?

8.20 Men On Mars? The Mars twin rovers, Spirit and Opportunity, which roamed the surface of Mars several years ago, found evidence that there was once water on Mars, raising the possibility that there was once life on the planet. Do you think that the United States should pursue a program to send humans to Mars? An opinion poll conducted by the Associated Press indicated that 49% of the 1034 adults surveyed think that we should pursue such a program.⁵

- Estimate the true proportion of Americans who think that the United States should pursue a program to send humans to Mars. Calculate the margin of error.
- The question posed in part a was only one of many questions concerning our space program that were asked in the opinion poll. If the Associated Press wanted to report one sampling error that would be valid for the entire poll, what value should they report?

8.21 Hungry Rats In an experiment to assess the strength of the hunger drive in rats, 30 previously trained animals were deprived of food for 24 hours. At the end of the 24-hour period, each animal was put into a cage where food was dispensed if the animal pressed a lever. The length of time the animal continued pressing the bar (although receiving no food) was recorded for each animal. If the data yielded a sample mean of 19.3 minutes with a standard deviation of 5.2 minutes, estimate the true mean time and calculate the margin of error.

INTERVAL ESTIMATION

8.5


long-term
probability
of the method,
not probability
of a numerical
single instance

An *interval estimator* is a rule for calculating two numbers—say, a and b —to create an interval that you are fairly certain contains the parameter of interest. The concept of “fairly certain” means “with high probability.” We measure this probability using the **confidence coefficient**, designated by $1 - \alpha$.

Definition The probability that a confidence interval will contain the estimated parameter is called the **confidence coefficient**.

MY TIP

like lariat roping:
parameter = fence post
interval estimate = lariat



For example, experimenters often construct 95% confidence intervals. This means that the confidence coefficient, or the probability that the interval will contain the estimated parameter, is .95. You can increase or decrease your amount of certainty by changing the confidence coefficient. Some values typically used by experimenters are .90, .95, .98, and .99.

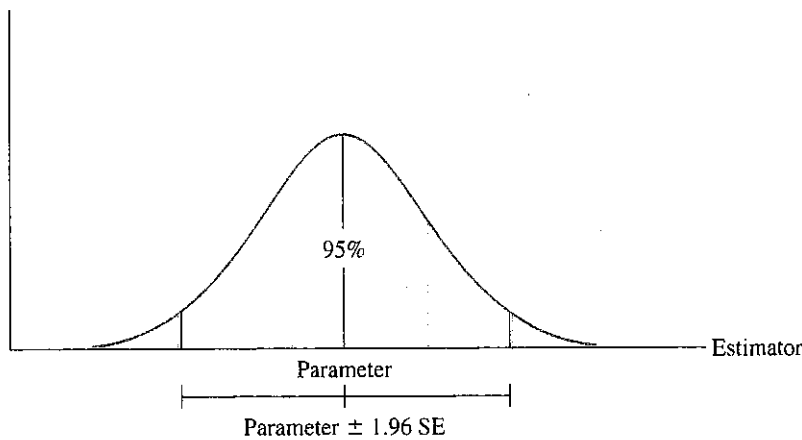
Consider an analogy—this time, throwing a lariat at a fence post. The fence post represents the parameter that you wish to estimate, and the loop formed by the lariat represents the confidence interval. Each time you throw your lariat, you hope to rope the fence post; however, sometimes your lariat misses. In the same way, each time

you draw a sample and construct a confidence interval for a parameter, you hope to include the parameter in your interval, but, just like the lariat, sometimes you miss. Your “success rate”—the proportion of intervals that “rope the post” in repeated sampling—is the confidence coefficient.

Constructing a Confidence Interval

When the sampling distribution of a point estimator is approximately normal, an interval estimator or **confidence interval** can be constructed using the following reasoning. For simplicity, assume that the confidence coefficient is .95 and refer to Figure 8.6.

FIGURE 8.6
Parameter ± 1.96 SE



- We know that, of all possible values of the estimator that we might select, 95% of them will be in the interval

$$\text{Parameter} \pm 1.96 \text{ SE}$$

shown in Figure 8.6.


- Since the value of the parameter is unknown, consider constructing the interval

$$\text{estimator} \pm 1.96 \text{ SE}$$

which has the same width as the first interval, but has a variable center.

- How often will this interval work properly and enclose the parameter of interest? Refer to Figure 8.7.

FIGURE 8.7
Some 95% confidence intervals






TIP

Like a game of ring toss:

parameter = peg

interval estimate = ring




The first two intervals work properly—the parameter (marked with a dotted line) is contained within both intervals. The third interval does not work, since it fails to enclose the parameter. This happened because the value of the estimator at the center of the interval was too far away from the parameter. Fortunately, values of the estimator only fall this far away 5% of the time—our procedure will work properly 95% of the time!

You may want to change the *confidence coefficient* from $(1 - \alpha) = .95$ to another confidence level $(1 - \alpha)$. To accomplish this, you need to change the value $z = 1.96$, which locates an area .95 in the center of the standard normal curve, to a value of z that locates the area $(1 - \alpha)$ in the center of the curve, as shown in Figure 8.8. Since the total area under the curve is 1, the remaining area in the two tails is α , and each tail contains area $\alpha/2$. The value of z that has “tail area” $\alpha/2$ to its right is called $z_{\alpha/2}$, and the area between $-z_{\alpha/2}$ and $z_{\alpha/2}$ is the confidence coefficient $(1 - \alpha)$. Values of $z_{\alpha/2}$ that are typically used by experimenters will become familiar to you as you begin to construct confidence intervals for different practical situations. Some of these values are given in Table 8.2.

FIGURE 8.8

Location of $z_{\alpha/2}$



A $(1 - \alpha)100\%$ LARGE-SAMPLE CONFIDENCE INTERVAL

$$(\text{Point estimator}) \pm z_{\alpha/2} \times (\text{Standard error of the estimator})$$

where $z_{\alpha/2}$ is the z -value with an area $\alpha/2$ in the right tail of a standard normal distribution. This formula generates two values; the **lower confidence limit (LCL)** and the **upper confidence limit (UCL)**.

TABLE 8.2 Values of z Commonly Used for Confidence Intervals

Confidence coefficient, $(1 - \alpha)$	α	$\alpha/2$	$z_{\alpha/2}$
.90	.10	.05	1.645
.95	.05	.025	1.96
.98	.02	.01	2.33
.99	.01	.005	2.58

Large-Sample Confidence Interval for a Population Mean μ

Practical problems very often lead to the estimation of μ , the mean of a population of quantitative measurements. Here are some examples:

- The average achievement of college students at a particular university
- The average strength of a new type of steel
- The average number of deaths per age category
- The average demand for a new cosmetics product

When the sample size n is large, the sample mean \bar{x} is the best point estimator for the population mean μ . Since its sampling distribution is approximately normal, it can be used to construct a confidence interval according to the general approach given earlier.

A $(1 - \alpha)100\%$ LARGE-SAMPLE CONFIDENCE INTERVAL FOR A POPULATION MEAN μ

$$\bar{x} \pm z_{\alpha/2} \frac{\sigma}{\sqrt{n}}$$

where $z_{\alpha/2}$ is the z -value corresponding to an area $\alpha/2$ in the upper tail of a standard normal z distribution, and

n = Sample size

σ = Standard deviation of the sampled population

If σ is unknown, it can be approximated by the sample standard deviation s when the sample size is large ($n \geq 30$) and the approximate confidence interval is

$$\bar{x} \pm z_{\alpha/2} \frac{s}{\sqrt{n}}$$

Another way to find the large-sample confidence interval for a population mean μ is to begin with the statistic

$$z = \frac{\bar{x} - \mu}{\sigma/\sqrt{n}}$$

which has a standard normal distribution. If you write $z_{\alpha/2}$ as the value of z with area $\alpha/2$ to its right, then you can write

$$P\left(-z_{\alpha/2} < \frac{\bar{x} - \mu}{\sigma/\sqrt{n}} < z_{\alpha/2}\right) = 1 - \alpha$$

You can rewrite this inequality as

$$-z_{\alpha/2} \frac{\sigma}{\sqrt{n}} < \bar{x} - \mu < z_{\alpha/2} \frac{\sigma}{\sqrt{n}}$$

$$-\bar{x} - z_{\alpha/2} \frac{\sigma}{\sqrt{n}} < -\mu < -\bar{x} + z_{\alpha/2} \frac{\sigma}{\sqrt{n}}$$

so that

$$P\left(\bar{x} - z_{\alpha/2} \frac{\sigma}{\sqrt{n}} < \mu < \bar{x} + z_{\alpha/2} \frac{\sigma}{\sqrt{n}}\right) = 1 - \alpha$$

Both $\bar{x} - z_{\alpha/2}(\sigma/\sqrt{n})$ and $\bar{x} + z_{\alpha/2}(\sigma/\sqrt{n})$, the lower and upper confidence limits, are actually random quantities that depend on the sample mean \bar{x} . Therefore, in repeated sampling, the random interval, $\bar{x} \pm z_{\alpha/2}(\sigma/\sqrt{n})$, will contain the population mean μ with probability $(1 - \alpha)$.

EXAMPLE 8.6

A scientist interested in monitoring chemical contaminants in food, and thereby the accumulation of contaminants in human diets, selected a random sample of $n = 50$ male adults. It was found that the average daily intake of dairy products was $\bar{x} = 756$ grams per day with a standard deviation of $s = 35$ grams per day. Use this sample information to construct a 95% confidence interval for the mean daily intake of dairy products for men.

Solution Since the sample size of $n = 50$ is large, the distribution of the sample mean \bar{x} is approximately normally distributed with mean μ and standard error estimated by s/\sqrt{n} . The approximate 95% confidence interval is

$$\begin{aligned} \bar{x} \pm 1.96 \left(\frac{s}{\sqrt{n}} \right) \\ 756 \pm 1.96 \left(\frac{35}{\sqrt{50}} \right) \\ 756 \pm 9.70 \end{aligned}$$

Hence, the 95% confidence interval for μ is from 746.30 to 765.70 grams per day.

TIP


A 95% confidence interval tells you that, if you were to construct many of these intervals (all of which would have slightly different endpoints), 95% of them would enclose the population mean.

Interpreting the Confidence Interval

What does it mean to say you are “95% confident” that the true value of the population mean μ is within a given interval? If you were to construct 20 such intervals, each using different sample information, your intervals might look like those shown in Figure 8.9. Of the 20 intervals, you might expect that 95% of them, or 19 out of 20, will perform as planned and contain μ within their upper and lower bounds.

FIGURE 8.9

Twenty confidence intervals for the mean for Example 8.6



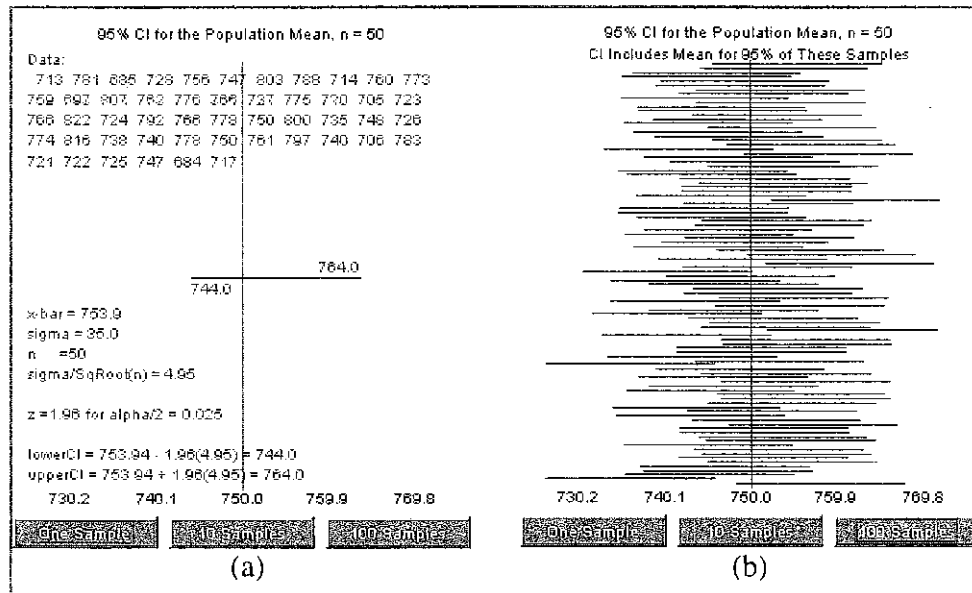
Remember that you cannot be absolutely sure that any one particular interval contains the mean μ . You will never know whether your particular interval is one of the 19 that “worked,” or whether it is the one interval that “missed.” Your confidence in the estimated interval follows from the fact that when repeated intervals are calculated, 95% of these intervals will contain μ .



APPLET

You can try this experiment on your own using the Java applet called **Interpreting Confidence Intervals**. The applet shown in Figure 8.10(a) shows the calculation of a 95% confidence interval for μ when $n = 50$ and $\sigma = 35$. For this particular confidence interval, we used the *One Sample* button. You can see the value of μ shown as a vertical green line on your monitor (gray in Figure 8.10). Notice that this confidence interval worked properly and enclosed the vertical line between its upper and lower limits. Figure 8.10(b) shows the calculation of 100 such intervals, using the *100 Samples* button. The intervals that fail to work properly are shown in red on your monitor (black in Figure 8.10). How many intervals fail to work? Is it close to the 95% confidence that we claim to have? You will use this applet again for the MyApplet Exercises section at the end of the chapter.

FIGURE 8.10
Interpreting Confidence Intervals applet



A good confidence interval has two desirable characteristics:

- It is as narrow as possible. The narrower the interval, the more exactly you have located the estimated parameter.
- It has a large confidence coefficient, near 1. The larger the confidence coefficient, the more likely it is that the interval will contain the estimated parameter.

EXAMPLE



Construct a 99% confidence interval for the mean daily intake of dairy products for adult men in Example 8.6.

Solution To change the confidence level to .99, you must find the appropriate value of the standard normal z that puts area $(1 - \alpha) = .99$ in the center of the curve. This value, with tail area $\alpha/2 = .005$ to its right, is found from Table 8.2 to be $z = 2.58$ (see Figure 8.11). The 99% confidence interval is then


$$\bar{x} \pm 2.58\left(\frac{s}{\sqrt{n}}\right)$$

$$756 \pm 2.58(4.95)$$

$$756 \pm 12.77$$

or 743.23 to 768.77 grams per day. This confidence interval is *wider* than the 95% confidence interval in Example 8.6.

FIGURE 8.11
Standard normal values
for a 99% confidence
interval



Right Tail Area	z-Value
.05	1.645
.025	1.96
.01	2.33
.005	2.58

The increased width is necessary to increase the confidence, just as you might want a wider loop on your lariat to ensure roping the fence post! The only way to *increase the confidence* without increasing the width of the interval is to *increase the sample size, n* .


The standard error of \bar{x} ,

$$SE = \frac{\sigma}{\sqrt{n}}$$

measures the variability or spread of the values of \bar{x} . The more variable the population data, measured by σ , the more variable will be \bar{x} , and the standard error will be larger. On the other hand, if you increase the sample size n , more information is available for estimating μ . The estimates should fall closer to μ and the standard error will be smaller. You can use the **Exploring Confidence Intervals** applet, shown in Figure 8.12, to see the effect of changing the sample size n , the standard deviation σ , and the confidence coefficient $1 - \alpha$ on the width of the confidence interval.

The confidence intervals of Examples 8.6 and 8.7 are approximate because you substituted s as an approximation for σ . That is, instead of the confidence coefficient being .95, the value specified in the example, the true value of the coefficient may be .92, .94, or .97. But this discrepancy is of little concern from a practical point of view; as far as your “confidence” is concerned, there is little difference among these confidence coefficients. Most interval estimators used in statistics yield approximate confidence intervals because the assumptions upon which they are based are not satisfied exactly. Having made this point, we will not continue to refer to confidence intervals as “approximate.” It is of little practical concern as long as the actual confidence coefficient is near the value specified.

FIGURE 8.12
Exploring Confidence
Intervals applet



Large-Sample Confidence Interval for a Population Proportion p

Many research experiments or sample surveys have as their objective the estimation of the proportion of people or objects in a large group that possess a certain characteristic. Here are some examples:

- The proportion of sales that can be expected in a large number of customer contacts
- The proportion of seeds that germinate
- The proportion of “likely” voters who plan to vote for a particular political candidate

Each is a practical example of the binomial experiment, and the parameter to be estimated is the binomial proportion p .

When the sample size is large, the sample proportion,

$$\hat{p} = \frac{x}{n} = \frac{\text{Total number of successes}}{\text{Total number of trials}}$$

is the best point estimator for the population proportion p . Since its sampling distribution is approximately normal, with mean p and standard error $SE = \sqrt{pq/n}$, \hat{p} can be used to construct a confidence interval according to the general approach given in this section.

A $(1 - \alpha)100\%$ LARGE-SAMPLE CONFIDENCE INTERVAL FOR A POPULATION PROPORTION p

$$\hat{p} \pm z_{\alpha/2} \sqrt{\frac{pq}{n}}$$

where $z_{\alpha/2}$ is the z -value corresponding to an area $\alpha/2$ in the right tail of a standard normal z distribution. Since p and q are unknown, they are estimated using

the best point estimators: \hat{p} and \hat{q} . The sample size is considered large when the normal approximation to the binomial distribution is adequate—namely, when $n\hat{p} > 5$ and $n\hat{q} > 5$.

EXAMPLE

8.8

A random sample of 985 “likely” voters—those who are likely to vote in the upcoming election—were polled during a phone-athon conducted by the Republican Party. Of those surveyed, 592 indicated that they intended to vote for the Republican candidate in the upcoming election. Construct a 90% confidence interval for p , the proportion of likely voters in the population who intend to vote for the Republican candidate. Based on this information, can you conclude that the candidate will win the election?

Solution The point estimate for p is

$$\hat{p} = \frac{x}{n} = \frac{592}{985} = .601$$

and the standard error is

$$\sqrt{\frac{\hat{p}\hat{q}}{n}} = \sqrt{\frac{(.601)(.399)}{985}} = .016$$

The z -value for a 90% confidence interval is the value that has area $\alpha/2 = .05$ in the upper tail of the z distribution, or $z_{.05} = 1.645$ from Table 8.2. The 90% confidence interval for p is thus

$$\hat{p} \pm 1.645 \sqrt{\frac{\hat{p}\hat{q}}{n}}$$

$$.601 \pm .026$$

or $.575 < p < .627$. You estimate that the percentage of likely voters who intend to vote for the Republican candidate is between 57.5% and 62.7%. Will the candidate win the election? Assuming that she needs more than 50% of the vote to win, and since both the upper and lower confidence limits exceed this minimum value, you can say with 90% confidence that the candidate will win.

There are some problems, however, with this type of sample survey. What if the voters who consider themselves “likely to vote” do not actually go to the polls? What if a voter changes his or her mind between now and election day? What if a surveyed voter does not respond truthfully when questioned by the campaign worker? The 90% confidence interval you have constructed gives you 90% confidence only if you have selected a *random sample from the population of interest*. You can no longer be assured of “90% confidence” if your sample is biased, or if the population of voter responses changes before the day of the election!

You may have noticed that the point estimator with its 95% margin of error looks very similar to a 95% confidence interval for the same parameter. This close relationship exists for most of the parameters estimated in this book, but it is not true in general. Sometimes the best point estimator for a parameter *does not* fall in the middle of the best confidence interval; the best confidence interval may not even be a function of the best point estimator. Although this is a theoretical distinction, you should remember that there is a difference between point and interval estimation, and that the choice between the two depends on the preference of the experimenter.

ESTIMATING THE DIFFERENCE BETWEEN TWO POPULATION MEANS

8.6

A problem equally as important as the estimation of a single population mean μ for a quantitative population is the comparison of two population means. You may want to make comparisons like these:

- The average scores on the Medical College Admission Test (MCAT) for students whose major was biochemistry and those whose major was biology
- The average yields in a chemical plant using raw materials furnished by two different suppliers
- The average stem diameters of plants grown on two different types of nutrients

For each of these examples, there are two populations: the first with mean and variance μ_1 and σ_1^2 , and the second with mean and variance μ_2 and σ_2^2 . A random sample of n_1 measurements is drawn from population 1, and a second random sample of size n_2 is independently drawn from population 2. Finally, the estimates of the population parameters are calculated from the sample data using the estimators \bar{x}_1 , s_1^2 , \bar{x}_2 , and s_2^2 as shown in Table 8.3.

TABLE 8.3 Samples From Two Quantitative Populations

	Population 1	Population 2
Mean	μ_1	μ_2
Variance	σ_1^2	σ_2^2
	Sample 1	Sample 2
Mean	\bar{x}_1	\bar{x}_2
Variance	s_1^2	s_2^2
Sample Size	n_1	n_2

Intuitively, the difference between two sample means would provide the maximum information about the actual difference between two population means, and this is in fact the case. The best point estimator of the difference ($\mu_1 - \mu_2$) between the population means is $(\bar{x}_1 - \bar{x}_2)$. The sampling distribution of this estimator is not difficult to derive, but we state it here without proof.

PROPERTIES OF THE SAMPLING DISTRIBUTION OF $(\bar{x}_1 - \bar{x}_2)$, THE DIFFERENCE BETWEEN TWO SAMPLE MEANS

When independent random samples of n_1 and n_2 observations have been selected from populations with means μ_1 and μ_2 and variances σ_1^2 and σ_2^2 , respectively, the sampling distribution of the difference $(\bar{x}_1 - \bar{x}_2)$ has the following properties:

1. The mean of $(\bar{x}_1 - \bar{x}_2)$ is

$$\mu_1 - \mu_2$$

and the standard error is

$$SE = \sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}$$

which can be estimated as

$$SE = \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}} \text{ when the sample sizes are large.}$$

2. **If the sampled populations are normally distributed**, then the sampling distribution of $(\bar{x}_1 - \bar{x}_2)$ is **exactly** normally distributed, regardless of the sample size.
3. **If the sampled populations are not normally distributed**, then the sampling distribution of $(\bar{x}_1 - \bar{x}_2)$ is **approximately** normally distributed when n_1 and n_2 are both 30 or more, due to the Central Limit Theorem.

Since $(\mu_1 - \mu_2)$ is the mean of the sampling distribution, it follows that $(\bar{x}_1 - \bar{x}_2)$ is an unbiased estimator of $(\mu_1 - \mu_2)$ with an approximately normal distribution when n_1 and n_2 are large. That is, the statistic

$$z = \frac{(\bar{x}_1 - \bar{x}_2) - (\mu_1 - \mu_2)}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$$

has an approximately standard normal z distribution, and the general procedures of Section 8.5 can be used to construct point and interval estimates. Although the choice between point and interval estimation depends on your personal preference, most experimenters choose to construct confidence intervals for two-sample problems. The appropriate formulas for both methods are given next.

LARGE-SAMPLE POINT ESTIMATION OF $(\mu_1 - \mu_2)$

Point estimator: $(\bar{x}_1 - \bar{x}_2)$

$$95\% \text{ Margin of error: } \pm 1.96 SE = \pm 1.96 \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$$

A $(1 - \alpha)100\%$ LARGE-SAMPLE CONFIDENCE INTERVAL FOR $(\mu_1 - \mu_2)$

$$(\bar{x}_1 - \bar{x}_2) \pm z_{\alpha/2} \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$$

MY TIP

Right Tail Area	z-Value
.05	1.645
.025	1.96
.01	2.33
.005	2.58

EXAMPLE

8.9

The wearing qualities of two types of automobile tires were compared by road-testing samples of $n_1 = n_2 = 100$ tires for each type. The number of miles until wearout was defined as a specific amount of tire wear. The test results are given in Table 8.4. Estimate $(\mu_1 - \mu_2)$, the difference in mean miles to wearout, using a 99% confidence interval. Is there a difference in the average wearing quality for the two types of tires?

TABLE 8.4

Sample Data Summary for Two Types of Tires

Tire 1	Tire 2
$\bar{x}_1 = 26,400$ miles	$\bar{x}_2 = 25,100$ miles
$s_1^2 = 1,440,000$	$s_2^2 = 1,960,000$

Solution The point estimate of $(\mu_1 - \mu_2)$ is

$$(\bar{x}_1 - \bar{x}_2) = 26,400 - 25,100 = 1300 \text{ miles}$$

and the standard error of $(\bar{x}_1 - \bar{x}_2)$ is estimated as

$$SE = \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}} = \sqrt{\frac{1,440,000}{100} + \frac{1,960,000}{100}} = 184.4 \text{ miles}$$

The 99% confidence interval is calculated as

$$(\bar{x}_1 - \bar{x}_2) \pm 2.58 \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$$

$$1300 \pm 2.58(184.4)$$

$$1300 \pm 475.8$$

or $824.2 < (\mu_1 - \mu_2) < 1775.8$. The difference in the average miles to wearout for the two types of tires is estimated to lie between LCL = 824.2 and UCL = 1775.8 miles of wear.

Based on this confidence interval, can you conclude that there is a difference in the average miles to wearout for the two types of tires? If there were no difference in the two population means, then μ_1 and μ_2 would be equal and $(\mu_1 - \mu_2) = 0$. If you look at the confidence interval you constructed, you will see that 0 is not one of the possible values for $(\mu_1 - \mu_2)$. Therefore, it is not likely that the means are the same; you can conclude that there is a difference in the average miles to wearout for the two types of tires. The confidence interval has allowed you to *make a decision* about the equality of the two population means.

**TIP**

If 0 is not in the interval you can conclude that there is a difference in the population means.

EXAMPLE

8.10

The scientist in Example 8.6 wondered whether there was a difference in the average daily intakes of dairy products between men and women. He took a sample of $n = 50$ adult women and recorded their daily intakes of dairy products in grams per day. He did the same for adult men. A summary of his sample results is listed in Table 8.5. Construct a 95% confidence interval for the difference in the average daily intakes of dairy products for men and women. Can you conclude that there is a difference in the average daily intakes for men and women?

TABLE 8.5

Sample Values for Daily Intakes of Dairy Products

	Men	Women
Sample Size	50	50
Sample Mean	756	762
Sample Standard Deviation	35	30

Solution The confidence interval is constructed using a value of z with tail area $\alpha/2 = .025$ to its right; that is, $z_{.025} = 1.96$. Using the sample standard deviations to approximate the unknown population standard deviations, the 95% confidence interval is

$$\begin{aligned}(\bar{x}_1 - \bar{x}_2) \pm 1.96 \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}} \\(756 - 762) \pm 1.96 \sqrt{\frac{35^2}{50} + \frac{30^2}{50}} \\-6 \pm 12.78\end{aligned}$$

or $-18.78 < (\mu_1 - \mu_2) < 6.78$. Look at the possible values for $(\mu_1 - \mu_2)$ in the confidence interval. It is possible that the difference $(\mu_1 - \mu_2)$ could be negative (indicating that the average for women exceeds the average for men), it could be positive (indicating that men have the higher average), or it could be 0 (indicating no difference between the averages). Based on this information, you *should not be willing to conclude* that there is a difference in the average daily intakes of dairy products for men and women.

Examples 8.9 and 8.10 deserve further comment with regard to using sample estimates in place of unknown parameters. The sampling distribution of

$$\frac{(\bar{x}_1 - \bar{x}_2) - (\mu_1 - \mu_2)}{\sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}}$$

has a standard normal distribution for all sample sizes when both sampled populations are normal and an *approximate* standard normal distribution when the sampled populations are not normal but the sample sizes are large (≥ 30). When σ_1^2 and σ_2^2 are not known and are estimated by the sample estimates s_1^2 and s_2^2 , the resulting statistic will still have an approximate standard normal distribution when the sample sizes are large. The behavior of this statistic when the population variances are unknown and the sample sizes are small will be discussed in Chapter 10.

8.6 EXERCISES

BASIC TECHNIQUES

8.39 Independent random samples were selected from populations 1 and 2. The sample sizes, means, and variances are as follows:

	Population	
	1	2
Sample Size	35	49
Sample Mean	12.7	7.4
Sample Variance	1.38	4.14

a. Find a 95% confidence interval for estimating the difference in the population means $(\mu_1 - \mu_2)$.

b. Based on the confidence interval in part a, can you conclude that there is a difference in the means for the two populations? Explain.

8.40 Independent random samples were selected from populations 1 and 2. The sample sizes, means, and variances are as follows:

	Population	
	1	2
Sample Size	64	64
Sample Mean	2.9	5.1
Sample Variance	0.83	1.67

ESTIMATING THE DIFFERENCE BETWEEN TWO BINOMIAL PROPORTIONS

8.7

A simple extension of the estimation of a binomial proportion p is the estimation of the difference between two binomial proportions. You may wish to make comparisons like these:

- The proportion of defective items manufactured in two production lines
- The proportion of female voters and the proportion of male voters who favor an equal rights amendment
- The germination rates of untreated seeds and seeds treated with a fungicide

These comparisons can be made using the difference $(p_1 - p_2)$ between two binomial proportions, p_1 and p_2 . Independent random samples consisting of n_1 and n_2 trials are drawn from populations 1 and 2, respectively, and the sample estimates \hat{p}_1 and \hat{p}_2 are calculated. The unbiased estimator of the difference $(p_1 - p_2)$ is the sample difference $(\hat{p}_1 - \hat{p}_2)$.

PROPERTIES OF THE SAMPLING DISTRIBUTION OF THE DIFFERENCE $(\hat{p}_1 - \hat{p}_2)$ BETWEEN TWO SAMPLE PROPORTIONS

Assume that independent random samples of n_1 and n_2 observations have been selected from binomial populations with parameters p_1 and p_2 , respectively. The sampling distribution of the difference between sample proportions

$$(\hat{p}_1 - \hat{p}_2) = \left(\frac{x_1}{n_1} - \frac{x_2}{n_2} \right)$$

has these properties:

1. The mean of $(\hat{p}_1 - \hat{p}_2)$ is

$$p_1 - p_2$$

and the standard error is

$$SE = \sqrt{\frac{p_1 q_1}{n_1} + \frac{p_2 q_2}{n_2}}$$

which is estimated as

$$SE = \sqrt{\frac{\hat{p}_1 \hat{q}_1}{n_1} + \frac{\hat{p}_2 \hat{q}_2}{n_2}}$$

2. The sampling distribution of $(\hat{p}_1 - \hat{p}_2)$ can be approximated by a normal distribution when n_1 and n_2 are large, due to the Central Limit Theorem.

Although the range of a single proportion is from 0 to 1, the difference between two proportions ranges from -1 to 1 . To use a normal distribution to approximate the distribution of $(\hat{p}_1 - \hat{p}_2)$, both \hat{p}_1 and \hat{p}_2 should be approximately normal; that is, $n_1 \hat{p}_1 > 5$, $n_1 \hat{q}_1 > 5$, and $n_2 \hat{p}_2 > 5$, $n_2 \hat{q}_2 > 5$.

The appropriate formulas for point and interval estimation are given next.

LARGE-SAMPLE POINT ESTIMATION OF $(p_1 - p_2)$ Point estimator: $(\hat{p}_1 - \hat{p}_2)$

$$95\% \text{ Margin of error: } \pm 1.96 \text{ SE} = \pm 1.96 \sqrt{\frac{\hat{p}_1 \hat{q}_1}{n_1} + \frac{\hat{p}_2 \hat{q}_2}{n_2}}$$

A $(1 - \alpha)100\%$ LARGE-SAMPLE CONFIDENCE INTERVAL FOR $(p_1 - p_2)$

$$(\hat{p}_1 - \hat{p}_2) \pm z_{\alpha/2} \sqrt{\frac{\hat{p}_1 \hat{q}_1}{n_1} + \frac{\hat{p}_2 \hat{q}_2}{n_2}}$$

Assumption: n_1 and n_2 must be sufficiently large so that the sampling distribution of $(\hat{p}_1 - \hat{p}_2)$ can be approximated by a normal distribution—namely, if $n_1 \hat{p}_1$, $n_1 \hat{q}_1$, $n_2 \hat{p}_2$, and $n_2 \hat{q}_2$ are all greater than 5.

EXAMPLE**8.11**

A bond proposal for school construction will be submitted to the voters at the next municipal election. A major portion of the money derived from this bond issue will be used to build schools in a rapidly developing section of the city, and the remainder will be used to renovate and update school buildings in the rest of the city. To assess the viability of the bond proposal, a random sample of $n_1 = 50$ residents in the developing section and $n_2 = 100$ residents from the other parts of the city were asked whether they plan to vote for the proposal. The results are tabulated in Table 8.6.

TABLE 8.6**Sample Values for Opinion on Bond Proposal**

	Developing Section	Rest of the City
Sample Size	50	100
Number Favoring Proposal	38	65
Proportion Favoring Proposal	.76	.65

1. Estimate the difference in the true proportions favoring the bond proposal with a 99% confidence interval.
2. If both samples were pooled into one sample of size $n = 150$, with 103 in favor of the proposal, provide a point estimate of the proportion of city residents who will vote for the bond proposal. What is the margin of error?

Solution

1. The best point estimate of the difference $(p_1 - p_2)$ is given by

$$(\hat{p}_1 - \hat{p}_2) = .76 - .65 = .11$$

and the standard error of $(\hat{p}_1 - \hat{p}_2)$ is estimated as

$$\sqrt{\frac{\hat{p}_1 \hat{q}_1}{n_1} + \frac{\hat{p}_2 \hat{q}_2}{n_2}} = \sqrt{\frac{(.76)(.24)}{50} + \frac{(.65)(.35)}{100}} = .0770$$

For a 99% confidence interval, $z_{.005} = 2.58$, and the approximate 99% confidence interval is found as

$$\begin{aligned}
 (\hat{p}_1 - \hat{p}_2) \pm z_{.005} \sqrt{\frac{\hat{p}_1 \hat{q}_1}{n_1} + \frac{\hat{p}_2 \hat{q}_2}{n_2}} \\
 .11 \pm (2.58)(.0770) \\
 .11 \pm .199
 \end{aligned}$$

or $(-.089, .309)$. Since this interval contains the value $(p_1 - p_2) = 0$, it is possible that $p_1 = p_2$, which implies that there may be no difference in the proportions favoring the bond issue in the two sections of the city.

2. If there is no difference in the two proportions, then the two samples are not really different and might well be combined to obtain an overall estimate of the proportion of the city residents who will vote for the bond issue. If both samples are pooled, then $n = 150$ and

$$\hat{p} = \frac{103}{150} = .69$$

Therefore, the point estimate of the overall value of p is .69, with a margin of error given by

$$\pm 1.96 \sqrt{\frac{(.69)(.31)}{150}} = \pm 1.96(.0378) = \pm .074$$

Notice that $.69 \pm .074$ produces the interval .62 to .76, which includes only proportions greater than .5. Therefore, if voter attitudes do not change adversely prior to the election, the bond proposal should pass by a reasonable majority.

BASIC TECHNIQUES

8.50 Independent random samples of $n_1 = 500$ and $n_2 = 500$ observations were selected from binomial populations 1 and 2, and $x_1 = 120$ and $x_2 = 147$ successes were observed.

- What is the best point estimator for the difference $(p_1 - p_2)$ in the two binomial proportions?
- Calculate the approximate standard error for the statistic used in part a.
- What is the margin of error for this point estimate?

Independent random samples of $n_1 = 800$ and $n_2 = 640$ observations were selected from binomial populations 1 and 2, and $x_1 = 337$ and $x_2 = 374$ successes were observed.

- Find a 90% confidence interval for the difference $(p_1 - p_2)$ in the two population proportions. Interpret the interval.
- What assumptions must you make for the confidence interval to be valid? Are these assumptions met?

8.52 Independent random samples of $n_1 = 1265$ and $n_2 = 1688$ observations were selected from binomial populations 1 and 2, and $x_1 = 849$ and $x_2 = 910$ successes were observed.

- Find a 99% confidence interval for the difference $(p_1 - p_2)$ in the two population proportions. What does "99% confidence" mean?
- Based on the confidence interval in part a, can you conclude that there is a difference in the two binomial proportions? Explain.

APPLICATIONS

8.53 **CANDIES** Does Mars, Incorporated use the same proportion of red candies in its plain and peanut varieties? A random sample of 56 plain M&M'S contained 12 red candies, and another random sample of 32 peanut M&M'S contained 8 red candies.

- Construct a 95% confidence interval for the difference in the proportions of red candies for the plain and peanut varieties.