

Chapter 8: Estimating with Confidence

Section 8.1 Confidence Intervals: The Basics

A **point estimator** is a statistic that provides an estimate of a population parameter. Ideally, a point estimate is our “best guess” at the value of an unknown parameter.

If you were estimating a population mean μ , you would probably use \bar{x} . If you were estimating a population proportion p , you might use \hat{p} . In both cases, you would be providing a point estimate of the parameter of interest.

Example 1: In each of the following settings, determine the point estimator you would use and calculate the value of the point estimate.

a) Quality control inspectors want to estimate the mean lifetime μ of the AA batteries produced in an hour at a factory. They select a random sample of 50 batteries during each hour of production and then drain them under conditions that mimic normal use. Here are the lifetimes (in hours) of the batteries from one such sample:

16.73 15.60 16.31 17.57 16.14 17.28 16.67 17.28 17.27 17.50 15.46 16.50 16.19
15.59 17.54 16.46 15.63 16.82 17.16 16.62 16.71 16.69 17.98 16.36 17.80 16.61
15.99 15.64 17.20 17.24 16.68 16.55 17.48 15.58 17.61 15.98 16.99 16.93 16.01
17.54 17.41 16.91 16.60 16.78 15.75 17.31 16.50 16.72 17.55 16.46

Use the sample mean \bar{x} as a point estimator for the population mean μ . For these data, our point estimate is $\bar{x} = 16.718$ hours.

b) What proportion p of U.S. high school students smoke? The 2011 Youth Risk Behavioral Survey questioned a random sample of 15,425 students in grades 9 to 12. Of these, 2792 said they had smoked cigarettes at least one day in the past month.

Use the sample proportion \hat{p} as a point estimator for the population proportion p . For this survey, our point estimate is

$$\hat{p} = \frac{2792}{15,425} = 0.181$$

c) The quality control inspectors in part (a) want to investigate the **variability** in battery lifetimes by estimating the population variance σ^2 .

Use the sample variance s_x^2 as a point estimator for the population variance σ^2 . For the battery life data, our point estimate is $s_x^2 = 0.441$ hours².

Activity p. 476


The Idea of a Confidence Interval

Recall the “Mystery Mean” Activity. Is the value of the population mean μ exactly 240.80? Probably not.

However, since the sample mean is 240.80, we could guess that μ is “somewhere” around 240.80. **How close to 240.80 is μ likely to be?**

To answer this question, we must ask another:

How would the sample mean \bar{x} vary if we took many SRSs of size 16 from the population?



Example 2: When Mr. Schiel’s class discussed the results of the mystery mean Activity, students used the following logic to come up with an “interval estimate” for the unknown population mean μ .

1. The sample mean $\bar{x} = 240.80$ is our point estimate for μ . We don’t expect \bar{x} to be exactly equal to μ , so we want to say how precise this estimate is.


2. In repeated samples, the values of \bar{x} follow a Normal distribution with mean μ and standard deviation 5, as in the figure to the right.

3. The 95 part of the 68-95-99.7 rule for Normal distributions says that \bar{x} is within $2(5) = 10$ (that’s 2 standard deviations) of the population mean μ in about 95% of all samples of size $n = 16$. See the figure to the right.

4. Whenever \bar{x} is within 10 points of μ , μ is within 10 points of \bar{x} . This happens in about 95% of all possible samples. So the interval from $\bar{x} - 10$ to $\bar{x} + 10$ “captures” the population mean μ in about 95% of all samples of size 16.

5. If we estimate that μ lies somewhere in the interval from $\bar{x} - 10 = 240.80 - 10 = 230.80$ to $\bar{x} + 10 = 240.80 + 10 = 250.80$

we’d be calculating an interval using a method that captures the true μ in about 95% of all possible samples of this size.



A **C% confidence interval** gives an interval of plausible values for a parameter. The interval is calculated from the data and has the form

$$\text{point estimate} \pm \text{margin of error}$$

The difference between the point estimate and the true parameter value will be less than the **margin of error**.

The **confidence level C** gives the overall success rate of the method for calculating the confidence interval. That is, in C% of all possible samples, the method would yield an interval that captures the true parameter value.

Interpreting Confidence Intervals and Confidence Levels

Confidence interval: To interpret a C% confidence interval for an unknown parameter, say, "We are C% confident that the interval from ____ to ____ captures the [parameter in context]."

Example 3: Two weeks before a presidential election, a polling organization asked a random sample of registered voters the following question: "If the presidential election were held today, would you vote for candidate A or candidate B?" Based on this poll, the 95% confidence interval for the population proportion who favor candidate A is (0.48, 0.54).

a) Interpret the confidence interval.

We are 95% confident that the interval from 0.48 to 0.54 captures the true proportion of all registered voters who favor candidate A in the election.

b) What is the point estimate that was used to create the interval? What is the margin of error?

A confidence interval has the form

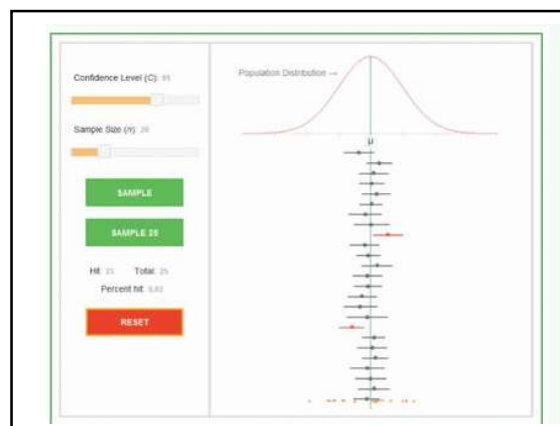
$$\text{point estimate} \pm \text{margin of error}$$

The point estimate is at the **midpoint** of the interval. Here, the point estimate is $\hat{p} = 0.51$. The margin of error gives the **distance from the point estimate** to either end of the interval. So the margin of error for this interval is 0.03.

c) Based on this poll, a political reporter claims that the majority of registered voters favor candidate A. Use the confidence interval to evaluate this claim.

Any value from 0.48 to 0.54 is a plausible value for the population proportion p that favors candidate A. Because there are plausible values of p less than 0.50, the confidence interval does not give convincing evidence to support the reporter's claim that a majority (more than 50%) of registered voters favor candidate A.

Activity:



Confidence level: To say that we are 95% confident is shorthand for "If we take many samples of the same size from this population, about 95% of them will result in an interval that captures the actual parameter value."

The confidence level tells us how likely it is that the method we are using will produce an interval that captures the population parameter *if we use it many times*. However, in practice we tend to calculate only a single confidence interval for a given situation.

The confidence level does not tell us the chance that a particular confidence interval captures the population parameter.

Instead, the confidence interval gives us a set of plausible values for the parameter.

Example 4: Interpreting a confidence level.

The confidence level in the mystery mean example—roughly 95%—tells us that if we take many SRSs of size 16 from Mr. Schiel's mystery population, the Interval $\bar{x} \pm 10$ will capture the population mean μ for about 95% of those samples.

```

MATH101.FX20.MATH101.BASIS.G1
mean(randNorm(μ, 20, 16))
240.80

```

Be sure you understand the basis for our confidence. There are only two possibilities:

1. The interval from 230.80 to 250.80 contains the population mean μ .
2. The interval from 230.80 to 250.80 does *not* contain the population mean μ . Our SRS was one of the few samples for which \bar{x} is not within 10 points of the true μ . Only about 5% of all samples result in a confidence interval that fails to capture μ .

What's the probability that our 95% confidence interval captures the parameter?

It's *not* 95%! Before we execute the command `mean(randNorm(M,20,16))`, we have a 95% chance of getting a sample mean that's within $2\sigma_{\bar{x}}$ of the mystery μ , which would lead to a confidence interval that captures μ . Once we have chosen a random sample, the sample mean \bar{x} either is or isn't within $2\sigma_{\bar{x}}$ of μ . And the resulting confidence interval either does or doesn't contain μ . After we construct a confidence interval, the probability that it captures the population parameter is either 1 (it does) or 0 (it doesn't).

Example 5: The Pew Internet and American Life Project asked a random sample of 2253 U.S. adults, "Do you ever...use Twitter or another service to share updates about yourself or to see updates about others?" Of the sample, 19% said "Yes." According to Pew, the resulting 95% confidence interval is (0.167, 0.213). Interpret the confidence interval and the confidence level.


Confidence interval: We are 95% confident that the interval from 0.167 to 0.213 captures the true proportion p of all U.S. adults who use Twitter or another service for updates.

Confidence level: If many samples of 2253 U.S. adults were taken, the resulting confidence intervals would capture the true proportion of all U.S. adults who use Twitter or another service for updates for about 95% of those samples.

Constructing a Confidence Interval

Why settle for 95% confidence when estimating a parameter? The price we pay for greater confidence is a wider interval.

When we calculated a 95% confidence interval for the mystery mean μ , we started with **estimate \pm margin of error**



Our estimate came from the sample statistic \bar{x} . Since the sampling distribution of \bar{x} is Normal, about 95% of the values of \bar{x} will lie within 2 standard deviations ($2\sigma_{\bar{x}}$) of the mystery mean μ . That is, our interval could be written as:

$$240.80 \pm 2 \cdot 5 = \bar{x} \pm 2\sigma_{\bar{x}}$$

This leads to a more general formula for confidence intervals:

$$\text{statistic} \pm (\text{critical value}) \cdot (\text{standard deviation of statistic})$$

The **critical value** is a multiplier that makes the interval wide enough to have the stated capture rate. The critical value depends on both the confidence level C and the sampling distribution of the statistic.

Calculating a Confidence Interval

The confidence interval for estimating a population parameter has the form

$$\text{statistic} \pm (\text{critical value}) \cdot (\text{standard deviation of statistic})$$

where the statistic we use is the point estimator for the parameter.

Properties of Confidence Intervals:

*The user chooses the confidence level, and the margin of error follows from this choice.

*We would like high confidence and also a small margin of error. High confidence says that our method almost always gives correct answers. A small margin of error says that we have pinned down the parameter quite precisely.

*Our general formula for a confidence interval is

$$\text{Statistic} \pm (\text{critical value}) \cdot (\text{standard deviation of statistic})$$

Margin of error


Properties of Confidence Intervals Cont.:

*We can see that the margin of error depends on the critical value and the standard deviation of the statistic. The critical value is tied directly to the confidence level: greater confidence requires a larger critical value.

*The standard deviation of the statistic depends on the sample size n : larger samples give more precise estimates, which means less variability in the statistic.

*The margin of error gets smaller when:

- ✓ The confidence level decreases. There is a trade-off between the confidence level and the margin of error. To obtain a smaller margin of error from the same data, you must be willing to accept lower confidence.



- ✓ The sample size n increases. Increasing the sample size n reduces the margin of error for any fixed confidence level.

Using Confidence Intervals Wisely

Here are two important cautions to keep in mind when constructing and interpreting confidence intervals.

*Our method of calculation assumes that the data come from an SRS of size n from the population of interest. Other types of random samples (stratified or cluster, say) might be preferable to an SRS in a given setting, but they require more complex calculations than the ones we'll use.

*The margin of error in a confidence interval covers only chance variation due to random sampling or random assignment. The margin of error is obtained from the sampling distribution. It indicates how close our estimate is likely to be to the unknown parameter if we repeat the random sampling or random assignment process many times. Practical difficulties, such as undercoverage and nonresponse in a sample survey, can lead to additional errors that may be larger than this chance variation. Remember this unpleasant fact when reading the results of an opinion poll or other sample survey. The way in which a survey or experiment is conducted influences the trustworthiness of its results in ways that are not included in the announced margin of error.