

Confidence intervals

Today, we're going to start talking about *confidence intervals*. We use confidence intervals as a tool in inferential statistics.

What this means is that *given some sample statistics*, we can say whether or not the actual population parameter is likely to be a certain value.

For example, if we have data from a survey that says 52% of voters prefer Al Gore and 47% of voters prefer George Bush, we can use confidence intervals to decide whether or not Al Gore is in fact more likely to win an election (i.e. if the true population parameter is greater than 50%).

Sample statistics are point estimates

To reinforce this idea, let's think about this hypothetical survey. From the survey we know that 52% of the sample respondents prefer Al Gore. This is our best available point estimate of the population parameter.

However, we know that a sample doesn't include the whole population—at best, we have several thousand out of around 200 million possible voters. It's possible, by *sheer chance*, that we have surveyed a sample that is atypical of the population, even if we have a completely random sample. So, in addition to our point estimates, we need to say how confident we are that we have a representative sample.

What this is about (continued)

That, in a nutshell, is what confidence intervals are about. We're mainly going to talk about the sample and population means, but this also works for other population parameters (such as the standard deviation).

Confidence intervals for a population

If we know the standard deviation of the population, confidence intervals for a sample are pretty easy to calculate.

$$\text{CI} = \bar{x} \pm (Z_{\alpha/2})\left(\frac{\sigma}{\sqrt{n}}\right)$$

This is the confidence interval for the population mean of x when we know the population standard deviation (σ).

There are four important parts of the formula; we'll talk about each of them in turn.

Standard error of the mean

One important part is σ/\sqrt{n} , which we call *the standard error of the mean*.

The standard error of the mean is a measurement of how much error results from the size of our sample. The bigger our sample, the less likely we are to accidentally end up with an unrepresentative sample, and therefore the standard error of the mean will be smaller. Conversely, if we have a small sample, we expect the chances of us having a “bad” sample to be higher, so the standard error will be bigger.

Standard error example

For example, if $\sigma = 1.3$ and $n = 144$, the standard error of the mean will be:

$$\frac{\sigma}{\sqrt{n}} = \frac{1.3}{\sqrt{144}} = \frac{1.3}{12} = 1.08\overline{33}$$

Alpha and confidence levels

The next part of the formula we need to discuss is the *confidence level*. The confidence level represents how willing we are to accidentally report a mistake. This is represented by $Z_{\alpha/2}$ in the equation.

The Greek letter α actually represents something called the *alpha level*. Normally, we use values of $\alpha = 0.05$ or $\alpha = 0.01$, which mean we are willing to be wrong five out of every hundred times, or one out of every hundred times, respectively.

Confidence levels

Normally, we instead of talking about α directly, we state our confidence level in terms of our chances of *not* making a mistake. $\alpha = .05$ corresponds to a “95% confidence level”, while $\alpha = .01$ corresponds to a “99% confidence level.”

Note that you can use other confidence levels, but people will probably look at you funny if you do. Anything lower than 95% is generally considered suspect, and above 99% is generally considered “showing off.”

Z scores for confidence levels

Now, the letter Z doesn't appear here by coincidence. In fact, we use the same Z tables that we learned how to use back in Chapter 6. However, since we normally only care about 95% and 99% confidence levels, we only have to worry about two Z values (we call them “critical values”):

$$Z_{.05/2} = Z_{.025} = 1.96 \quad \text{and} \quad Z_{0.01/2} = Z_{.005} = 2.58$$

You can look these up using the “small part” column in Appendix 1 if you like.

Confidence limits

The *confidence limits* are simply the product of the standard error of the mean and the Z score for a given confidence level. More formally:

$$e = (Z_{\alpha/2})\left(\frac{\sigma}{\sqrt{n}}\right)$$

e represents the confidence limit for a given α , σ and n .

Confidence limits (example)

From our example above, we can figure out the confidence limits for the sample mean, given $n = 144$ and $\sigma = 1.3$, as before, for $\alpha = .05$:

$$(Z_{\alpha/2})\left(\frac{\sigma}{\sqrt{n}}\right) = Z_{.025}\frac{1.3}{\sqrt{144}} = 1.96 \times \frac{1.3}{12} = 1.96 \times 1.083\bar{3} = 0.2133\bar{3}$$

Confidence limits (example - cont'd)

With $\alpha = .01$:

$$(Z_{\alpha/2})\left(\frac{\sigma}{\sqrt{n}}\right) = Z_{.005} \frac{1.3}{\sqrt{144}} = 2.58 \times \frac{1.3}{12} = 2.58 \times 1.083\bar{3} = 0.2795$$

Note that the interval is bigger with the smaller alpha.

On to the confidence interval

Once we have our confidence limits, all we need to do is figure out our confidence interval:

$$\text{CI} = \bar{x} \pm (Z_{\alpha/2})\left(\frac{\sigma}{\sqrt{n}}\right)$$

Examples

If $\bar{x} = 0.20$, we can find the confidence intervals for $\alpha = .05$:

$$\text{CI} = \bar{x} \pm (Z_{\alpha/2})\left(\frac{\sigma}{\sqrt{n}}\right) = 0.20 \pm 0.168 = (0.032 < \mu < 0.368)$$

When $\alpha = .01$:

$$\text{CI} = \bar{x} \pm (Z_{\alpha/2})\left(\frac{\sigma}{\sqrt{n}}\right) = 0.20 \pm 0.2795 = (-0.0795 < \mu < 0.4795)$$

So what does this mean?

What this means is that the *population mean* is somewhere between 0.032 and 0.368, with a 95% confidence level, and between -0.0795 and 0.4795, with a 99% confidence level. To put it another way, we believe, with the given margin of error, that the mean of the population is somewhere between those values, and that we'd get a result between those values if we looked at the entire population 95% (or 99%) of the time.

By the way, *margin of error* is another way to say “confidence limit.” So when you see on the news that “President Bush’s approval rating is 68%, with a 4% margin of error,” this means that they are confident that Bush’s true approval rating is between 64% and 72%.

Minimum sample sizes

From the confidence limit formula, we can derive the minimum n (sample size) required to have a given confidence level. That is, how many items out of the population must we look at to be confident that the population parameter is close to the sample statistic?

To do this, we rearrange the confidence limit formula and get:

$$n = \left(\frac{(Z_{\alpha/2})\sigma}{e} \right)^2$$

Example

So, for example, if we want to know the average age of Ole Miss students within a month (1/12 of a year, or .00833 years), with a 95% confidence level ($\alpha = .05$), and we know that the standard deviation $\sigma = 1.8$:

$$n = \left(\frac{(Z_{\alpha/2})\sigma}{e} \right)^2 = \left(\frac{1.96 \times 1.8}{0.0833} \right)^2 = \left(\frac{3.528}{0.0833} \right)^2 = (32.566)^2 = 1060.6$$

This means we would need to sample 1061 students to determine this information with a 95% confidence level. (Note that sometimes $n > N!$).

What if we don't know the population std. deviation?

Most of the time, we don't already know the population standard deviation (σ); in fact, to figure it out, we would have collected enough information to get the population mean μ as well.

However, there is an analogous formula we can use with the *sample* standard deviation, s :

$$\text{CI} = \bar{x} \pm (t_{\alpha/2})\left(\frac{s}{\sqrt{n}}\right)$$

Note that as well as changing σ to s , we've also changed Z into t . This means we have to learn a new distribution function.

Student's t distribution

The new distribution is called the *Student's t distribution*. (It was actually invented by William Gosset and employed to monitor the production of Guinness.)

Student's t distribution is a lot like the standard normal curve, except that it is designed specifically to deal with the problem of small samples. (For large samples, where $n > 200$ or so, the t distribution and Z distribution become the same.)

A new table; degrees of freedom

The t distribution is illustrated in Appendix 2 of Schacht. Note that this table is different than the table in Appendix 1 in two ways: only certain alpha levels for t are shown across the top, and down the left-hand column are a set of values called “df” or *degrees of freedom*.

Degrees of freedom have to do with what mathematicians call a “free parameters problem” (basically, it has to do with how many of the parameters we can estimate based on the data). Since we’re not mathematicians, all we need to really worry about is that $df = n - 1$ for the t distribution.

Finding t values

So, for any given α and df (or n), we can find the correct t value. For example, let's find the t value for $\alpha = .05$ and 18 degrees of freedom. (Look under the “two-tailed test” column.) $t_{\text{crit}} = 2.101$.

Note that if the number of degrees of freedom doesn't appear in the table, use the next-lowest number that appears. So, for example, the critical value for $\alpha = .01$ and $df = 53$ is 2.704.

Of course, if you have more than about 200 degrees of freedom, you should use the ∞ row, which has the Z scores for the appropriate alpha levels.

Let's work an example

Now that we know the t distribution, we can figure out some confidence intervals.

Let's assume we have a sample of 20 political science undergrads at Ole Miss. Their mean IQ $\bar{x} = 107$ with a standard deviation $s = 17$. What is the confidence interval for all political science undergrads at Ole Miss, with a 95% confidence level?

Since $n = 20$, we have $n - 1 = 19$ degrees of freedom. The critical value for t with $\alpha = .05$ and $df = 19$ is 2.093. Now, all we have to do is substitute into the formula:

Plug and chug

$$\begin{aligned}\text{CI} &= \bar{x} \pm (t_{\alpha/2})\left(\frac{s}{\sqrt{n}}\right) = 107 \pm (2.093)\left(\frac{17}{\sqrt{20}}\right) \\ &= 107 \pm (2.093)\frac{17}{4.47} \\ &= 107 \pm (2.093)(3.081) = 107 \pm 7.956 \\ &= 99.04 < \mu < 114.96\end{aligned}$$

Homework

Answer questions 1 and 3 on pages 136–7.

Prepare for a quiz next Wednesday on the material from chapters 6–8.