

6.2: Basic Confidence Intervals

As elsewhere in this chapter, we assume that we are working with some (large) population on which there is defined a quantitative RV X . The population mean μ_x is unknown, and we want to estimate it. world, unchanging but also probably unknown, simply because to compute it we would have to have access to the values of X for the entire population.

We continue also with our strange assumption that while we do not know μ_x , we do know the population standard deviation σ_x , of X .

Our strategy to estimate μ_x is to take an SRS of size n , compute the sample mean \bar{x} of X , and then to guess that $\mu_x \approx \bar{x}$. But this leaves us wondering how good an approximation \bar{x} is of μ_x .

The strategy we take for this is to figure how close μ_x must be to \bar{x} – or \bar{x} to μ_x , it's the same thing, and in fact to be precise enough to say what is the probability that μ_x is a certain distance from \bar{x} . That is, if we choose a target probability, call it L , we want to make an interval of real numbers centered on \bar{x} with the probability of μ_x being in that interval being L .

Actually, that is not really a sensible thing to ask for: probability, remember, is the fraction of times something happens in repeated experiments. But we are not repeatedly choosing μ_x and seeing if it is in that interval. Just the opposite, in fact: μ_x is fixed (although unknown to us), and every time we pick a new SRS – that's the repeated experiment, choosing new SRSs! – we can compute a new interval and hope that that new interval might contain μ_x . The probability L will correspond to what fraction of those newly computed intervals which contain the (fixed, but unknown) μ_x .

How to do even this?

Well, the Central Limit Theorem tells us that the distribution of \bar{x} as we take repeated SRSs – exactly the repeatable experiment we are imagining doing – is approximately Normal with mean μ_x and standard deviation σ_x/\sqrt{n} . By the 68-95-99.7 Rule:

1. 68% of the time we take samples, the resulting \bar{x} will be within σ_x/\sqrt{n} units on the number line of μ_x . Equivalently (since the distance from A to B is the same as the distance from B to A!), 68% of the time we take samples, μ_x will be within σ_x/\sqrt{n} of \bar{x} . In other words, 68% of the time we take samples, μ_x will happen to lie in the interval from $\bar{x} - \sigma_x/\sqrt{n}$ to $\bar{x} + \sigma_x/\sqrt{n}$.
2. Likewise, 95% of the time we take samples, the resulting \bar{x} will be within $2\sigma_x/\sqrt{n}$ units on the number line of μ_x . Equivalently (since the distance from A to B is still the same as the distance from B to A!), 95% of the time we take samples, μ_x will be within $2\sigma_x/\sqrt{n}$ of \bar{x} . In other words, 95% of the time we take samples, μ_x will happen to lie in the interval from $\bar{x} - 2\sigma_x/\sqrt{n}$ to $\bar{x} + 2\sigma_x/\sqrt{n}$.
3. Likewise, 99.7% of the time we take samples, the resulting \bar{x} will be within $3\sigma_x/\sqrt{n}$ units on the number line of μ_x . Equivalently (since the distance from A to B is still the same as the distance from B to A!), 99.7% of the time we take samples, μ_x will be within $3\sigma_x/\sqrt{n}$ of \bar{x} . In other words, 99.7% of the time we take samples, μ_x will happen to lie in the interval from $\bar{x} - 3\sigma_x/\sqrt{n}$ to $\bar{x} + 3\sigma_x/\sqrt{n}$.

Notice the general shape here is that the interval goes from $\bar{x} - z_L^* \sigma_x/\sqrt{n}$ to $\bar{x} + z_L^* \sigma_x/\sqrt{n}$, where this number z_L^* has a name:

[def:criticalvalue] The **critical value z_L^* with probability L** for the Normal distribution is the number such that the Normal distribution $N(\mu_x, \sigma_x)$ has probability L between $\mu_x - z_L^* \sigma_x$ and $\mu_x + z_L^* \sigma_x$.

Note the probability L in this definition is usually called the **confidence level**.

If you think about it, the 68-95-99.7 Rule is exactly telling us that $z_L^* = 1$ if $L = .68$, $z_L^* = 2$ if $L = .95$, and $z_L^* = 3$ if $L = .997$. It's actually convenient to make a table of similar values, which can be calculated on a computer from the formula for the Normal distribution.

[fact:critvaltable] Here is a useful table of critical values for a range of possible confidence levels:

L	.5	.8	.9	.95	.99	.999
z_L^*	.674	1.282	1.645	1.960	2.576	3.291

Note that, oddly, the z_L^* here for $L = .95$ is not 2, but rather 1.96! This is actually more accurate value to use, which you may choose to use, or you may continue to use $z_L^* = 2$ when $L = .95$, if you like, out of fidelity to the 68-95-99.7 Rule.

Now, using these accurate critical values we can define an interval which tells us where we expect the value of μ_x to lie.

[def:confint] For a probability value L , called the **confidence level**, the interval of real numbers from $\bar{x} - z_L^* \sigma_x / \sqrt{n}$ to $\bar{x} + z_L^* \sigma_x / \sqrt{n}$ is called the **confidence interval for μ_x with confidence level L** .

The meaning of *confidence* here is quite precise (and a little strange):

[fact:confidence] Any particular confidence interval with confidence level L might or might not actually contain the sought-after parameter μ_x . Rather, what it means to have confidence level L is that if we take repeated, independent SRSs and compute the confidence interval again for each new \bar{x} from the new SRSs, then a fraction of size L of these new intervals will contain μ_x .

Note that any particular confidence interval might or might not contain μ_x not because μ_x is moving around, but rather the SRSs are different each time, so the \bar{x} is (potentially) different, and hence the interval is moving around. The μ_x is fixed (but unknown), while the confidence intervals move.

Sometimes the piece we add and subtract from the \bar{x} to make a confidence interval is given a name of its own:

[def:marginoferror] When we write a confidence interval for the population mean μ_x of some quantitative variable X in the form $\bar{x} - E$ to $\bar{x} + E$, where $E = z_L^* \sigma_x / \sqrt{n}$, we call E the **margin of error** [or, sometimes, the **sampling error**] of the confidence interval.

Note that if a confidence interval is given without a stated confidence level, particularly in the popular press, we should assume that the implied level was .95.

Cautions

The thing that most often goes wrong when using confidence intervals is that the sample data used to compute the sample mean \bar{x} and then the endpoints $\bar{x} \pm E$ of the interval is not from a good SRS. It is hard to get SRSs, so this is not unexpected. But we nevertheless frequently assume that some sample is an SRS, so that we can use it to make a confidence interval, even of that assumption is not really justified.

Another thing that can happen to make confidence intervals less accurate is to choose too small a sample size n . We have seen that our approach to confidence intervals relies upon the CLT, therefore it typically needs samples of size at least 30.

[eg:CI1] A survey of 463 first-year students at Euphoria State University [ESU] found that the [sample] average of how long they reported studying per week was 15.3 hours. Suppose somehow we know that the population standard deviation of hours of study per week at ESU is 8.5. Then we can find a confidence interval at the 99% confidence level for the mean study per week of all first-year students by calculating the margin of error to be

$$E = z_L^* \sigma_x / \sqrt{n} = 2.576 \cdot 8.5 / \sqrt{463} = 1.01759 \quad (6.2.1)$$

and then noting that the confidence interval goes from

$$\bar{x} - E = 15.3 - 1.01759 = 14.28241 \quad (6.2.2)$$

to

$$\bar{x} + E = 15.3 + 1.01759 = 16.31759. \quad (6.2.3)$$

Note that for this calculation to be doing what we want it to do, we must assume that the 463 first-year students were an SRS out of the entire population of first-year students at ESU.

Note also that what it means that we have 99% confidence in this interval from 14.28241 to 16.31759 (or $[14.28241, 16.31759]$ in interval notation) is not, in fact, that we any confidence at all in those particular numbers. Rather, we have confidence in the *method*, in the sense that if we imagine independently taking many future SRSs of size 463 and recalculating new confidence intervals, then 99% of these future intervals will contain the one, fixed, unknown μ_x .

This page titled [6.2: Basic Confidence Intervals](#) is shared under a [CC BY-SA 4.0](#) license and was authored, remixed, and/or curated by [Jonathan A. Poritz](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.