

Fact Sheet – Confidence Intervals (Ch 9)

A confidence interval for an unknown population parameter (p or μ) consists of an interval of numbers based upon the sample statistic (\hat{p} or μ).

The **level of confidence** represents the expected percentage of intervals that will contain the unknown population parameter. (The unknown population parameter lies between the lower bound and upper bound of the interval.)

For example, a 95% confidence interval for a population proportion p implies that if 100 different intervals are constructed, each based on a different sample from the same population, then we expect 95 of the intervals to include the population proportion p and 5 of the intervals to not include the population proportion p .

Give It a Try

- Open StatCrunch.
- StatCrunch > Applets > Confidence Intervals > for a proportion
- Leave Proportion as 0.50 and First Level as 0.95, click Create Applet.
- The intervals that do not contain the population proportion ($p = 0.50$) are shown in red, those that do contain p are shown in green. The “proportion contained” should be close to 0.95.
- Keep clicking “Simulate 100”, and notice that the proportion contained will stay close to 0.95.

Basic Format

We take a sample, and compute the sample proportion (\hat{p}) or sample mean (\bar{x}).

That sample statistic is called a **point estimate**, and will be the center of the confidence interval.

From there, we compute a **margin of error (E)**. This is based on the level of confidence and the standard deviation for that sampling distribution.

The level of confidence is related to the z-value (or t-value) that separates the middle 90/95/99% of the values from the rest.

Lower Bound

point estimate – margin of error

Upper Bound

point estimate + margin of error