Confidence Intervals

Confidence Intervals for Proportions

Imagine an experiment with two outcomes. Let p denote the unknown true probability of getting a successful outcome on any one trial. This probability is constant and does not change from experiment to experiment. If we repeat the experiment n times we can estimate p by $\hat{p}$ = total number of successes / total number of trials. For this type of situation we know the sampling distribution, the mean and the variance of $\hat{p}$ are given by

$$E(\hat{p}) = p$$
$$Var(\hat{p}) = \frac{p(1-p)}{n}$$
$$\hat{p} \approx N\left(p, \frac{p(1-p)}{n}\right)$$

Now we can do a little bit of algebraic manipulation. If we subtract from $\hat{p}$ the mean and divide by the square root of the variance we get a quantity which has approximately a standard Normal distribution.

$$\frac{\hat{p}-p}{\sqrt{\frac{p(1-p)}{n}}} \approx N(0,1)$$

The standard Normal distribution is a useful and familiar distribution. We can graph it as shown in Figure **??** and we can determine the probabilities in different regions. For example, the area under the curve between -1.96 and 1.96 is equal to 95% of the total area. If a random variable follows a standard Normal distribution only 5% of the time will an observation be less than -1.96 or greater than 1.96.

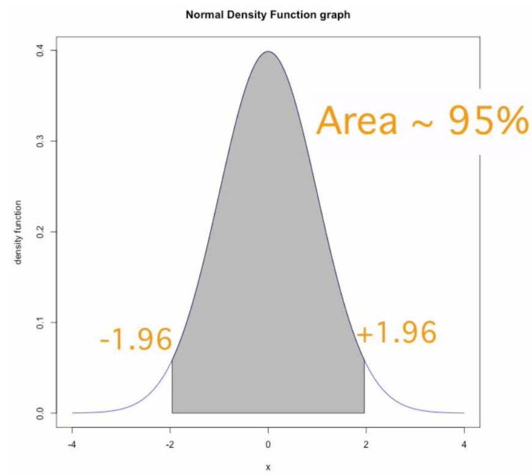

Figure 1: Critical values for 95% confidence interval

Now putting all these facts together

$$P\left(\left|\frac{\hat{p}-p}{\sqrt{\frac{p(1-p)}{n}}}\right| > 1.96\right) \doteq 0.05 = 5\%$$

$$P\left(\left|\frac{\hat{p}-p}{\sqrt{\frac{p(1-p)}{n}}}\right| \leq 1.96\right) \doteq 0.95 = 95\%$$

$$P\left(-1.96 \leq \left|\frac{\hat{p}-p}{\sqrt{\frac{p(1-p)}{n}}}\right| \leq 1.96\right) \doteq 0.95 = 95\%$$

$$P\left(\hat{p} - 1.96\sqrt{p(1-p)/n} \leq p \leq \hat{p} + 1.96\sqrt{p(1-p)/n}\right) \doteq 0.95 = 95\%$$

The 95% **confidence interval** for p is

$$\left[\hat{p} - 1.96\sqrt{p(1-p)/n}, \hat{p} + 1.96\sqrt{p(1-p)/n}\right]$$

The **margin of error** for a 95% confidence interval for p is $1.96\sqrt{p(1-p)/n}$ or half the width of the confidence interval. In the formula for the confidence interval above this quantity is subtracted and added to $\hat{p}$. You may wonder how the formula for the confidence interval is useful since it includes p which is unknown. There are two solutions we can apply.

Firstly, if we assume that p is close to $\hat{p}$ we can substitute in $\hat{p}$ to get the following formula

$$\left[\hat{p} - 1.96\sqrt{\hat{p}(1-\hat{p})/n}, \hat{p} + 1.96\sqrt{\hat{p}(1-\hat{p})/n}\right]$$

Alternatively, we can substitute $\frac{1}{2}$ for p. This is a conservative approach and would result in the widest possible interval we could get for any set of observations:

$$\left[\hat{p} - 1.96\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}, \hat{p} + 1.96\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}\right] = [\hat{p} - 0.98/\sqrt{n}, \hat{p} + 0.98/\sqrt{n}]$$

EXAMPLE 1

Here is an excerpt from an opinion poll from a recent United States Presidential Election:

> *"President Barack Obama led Republican challenger Mitt Romney by one point in a poll released last night... The CBS News/New York Times poll... put Obama ahead, 48 percent to 47 percent...The poll of 563 likely voters taken Oct. 25-28 had a margin of error of plus or minus four percentage points."*

From the excerpt, $n = 563$. If we use the conservative assumption that $p = \frac{1}{2}$ we can calculate the margin of error.

$$\text{Margin of error} = 1.96\sqrt{\frac{1}{2}(1-\frac{1}{2})/563} \doteq 0.0414 = 4.13\% \approx 4\%$$

So the margin of error for the true proportion based on a sample size of 563 is just a bit greater than 4%. This is exactly what the opinion poll reported. Saying that the margin of error is "plus or minus 4 percentage points" means that half of the width of the confidence interval is approximately 4%.

You do not always have to use a 95% confidence interval. You can use any confidence level you want. For example, if you wanted to calculate a 90% confidence interval you would need to determine the cutoff points on a standard Normal distribution that contain 90% of the area. As shown in Figure **??** these cutoff points are -1.645 and 1.645.

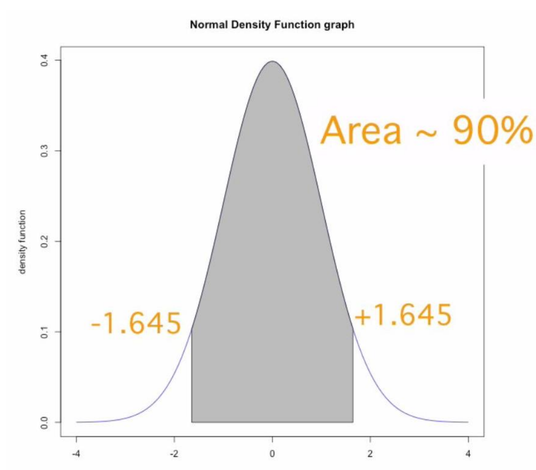

Figure 2: Critical values for 90% confidence interval

If you wanted to calculate a 99% confidence interval you would need to determine the cutoff points on a standard Normal distribution that contain 99% of the area. As shown in Figure **??** these cutoff points are -2.576 and 2.576.

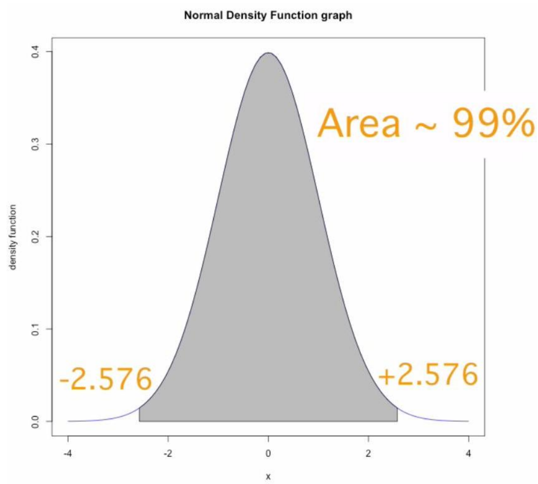

Figure 3: Critical values for 99% confidence interval

In general, for any value α, you can find a value $z_{\alpha/2}$ such that the area under the standard Normal curve which is less than $-z_{\alpha/2}$ is equal to $\alpha/2$. This is shown in Figure **??**.

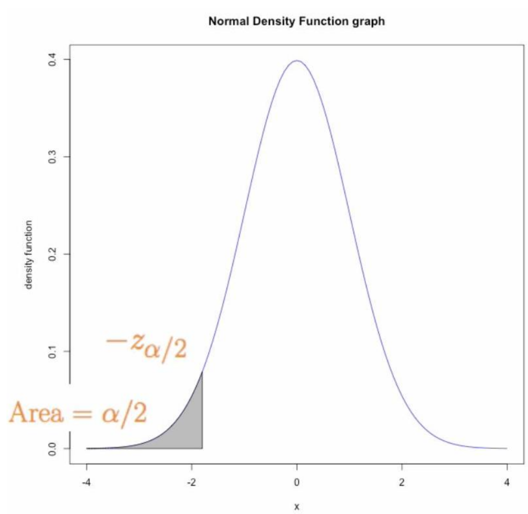

Figure 4: Critical values for $(1 - \alpha)\%$ confidence interval

If you take the interval from $-z_{\alpha/2}$ to $+z_{\alpha/2}$ this will contain $(1 - \alpha)\%$ of the area under the curve. So the confidence interval will miss the target only $\alpha\%$ of the time. This situation is shown in Figure **??**.

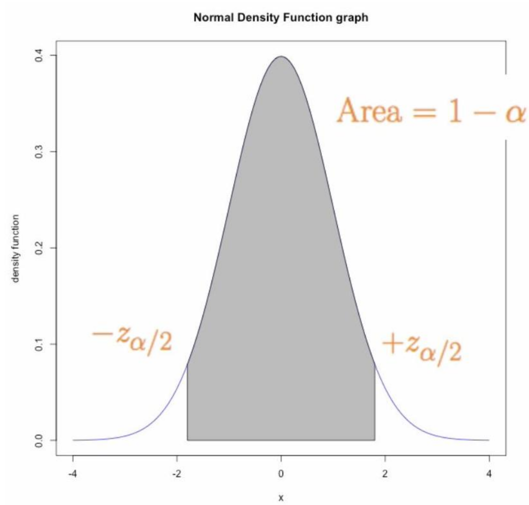

Figure 5: Critical values for $(1-\alpha)\%$ confidence interval

EXAMPLE 2

Let's return to the beer bottle cap example. We flipped a cap 1,000 times and observed 576 reds. In this case $n = 1000$ and $\hat{p} = 0.576$. We can now calculate the 95% confidence interval for the true probability, p, of the bottle cap landing on the red side.

$$\begin{aligned}
95\% \text{ CI for } p &= \left[\hat{p} - 1.96\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}, \hat{p} + 1.96\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}\right] \\
&= [0.545, 0.607] \\
&= [54.5\%, 60.7\%]
\end{aligned}$$

We can also calculate the 90% confidence interval for p. Note that this interval is narrower than the 95% confidence interval.

$$\begin{aligned}
90\% \text{ CI for } p &= \left[\hat{p} - 1.645\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}, \hat{p} + 1.645\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}\right] \\
&= [0.550, 0.602] \\
&= [55.0\%, 60.2\%]
\end{aligned}$$

If we want to be more certain that the confidence interval covers p we can calculate a 99% confidence interval. Note that this interval is wider than both the previous two intervals.

$$\begin{aligned}
99\% \text{ CI for } p &= \left[\hat{p} - 2.576\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}, \hat{p} + 2.576\sqrt{\frac{1}{2}\left(1-\frac{1}{2}\right)/n}\right] \\
&= [0.535, 0.617] \\
&= [53.5\%, 61.7\%]
\end{aligned}$$

This example demonstrates that there is a tradeoff between how certain we want to be and how accurate (narrow) the confidence interval is. If we chose $\alpha = 0.05$ we can be 95% confident that the true probability of observing red on a flip of the beer cap is anywhere from 54.5% to 60.7%. But if we want to be more certain by constructing a 99% interval, then the interval [53.5%, 61.7%] will be wider.