CLARK UNIVERSITY

Confidence intervals

Math 218, Mathematical Statistics

D Joyce, Spring 2016

Introduction to confidence intervals. Although estimating a parameter θ by a particular number $\hat{\theta}$ may be the simplest kind of statistical inference, that often is not very satisfactory. Some indication of the spread of the likely values of θ explains a lot more. One way that's done is with confidence intervals. A typical confidence interval is a 95% confidence interval $[L, U]$ for θ and that's given by two statistics, L and U such that

$$P(L \leq \theta \leq U) = 0.95.$$

Other confidence levels besides 95% are defined similarly.

This concept is best explained with an example. Let's take a normal distribution with a known value for σ^2, but an unknown value for μ, and our job is to come up with a confidence interval for μ. The sample mean $\overline{X}$ is a point estimator for μ, and we know that $\overline{X}$ is a normal distribution with mean μ and variance σ^2/n. From the table for the standard normal distribution, the probability that a standard normal random variable Z lies between -1.96 and 1.96 is 95%. Therefore,

$$P\left(\mu - 1.96\,\frac{\sigma}{\sqrt{n}} \leq \overline{X} \leq \mu + 1.96\,\frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

We can rewrite the first inequality $\mu - 1.96\,\frac{\sigma}{\sqrt{n}} \leq \overline{X}$ as $\mu \leq \overline{X} + 1.96\,\frac{\sigma}{\sqrt{n}}$, and we can rewrite the second inequality $\overline{X} \leq \mu + 1.96\,\frac{\sigma}{\sqrt{n}}$ as $X - 1.96\,\frac{\sigma}{\sqrt{n}} \leq \mu$. Therefore, the statement of probability can be rewritten as

$$P\left(X - 1.96\,\frac{\sigma}{\sqrt{n}} \leq \mu \leq \overline{X} + 1.96\,\frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

We now have two statistics, $L = X - 1.96\,\frac{\sigma}{\sqrt{n}}$ and $U = X + 1.96\,\frac{\sigma}{\sqrt{n}}$, so that $P(L \leq \mu \leq U) = 0.95$.

We'll look at an example and discuss some of the difficulties of interpreting the meaning of confidence intervals and apply interval estimates to bent coins.

Interval estimates for Bernoulli distributions. Suppose we have a bent coin with unknown probability p of heads, so the unknown probability of tails is $q = 1 - p$. The mean of this distribution is $\mu = p$, and its variance is $\sigma^2 = pq$.

The sample mean $\overline{X}$, which is the fraction of heads that occur in n trials, has mean $\mu_{\overline{X}} = \mu = p$, variance $\sigma^2_{\overline{X}} = \sigma^2/n = pq/n$, and standard deviation $\sigma_{\overline{X}} = \sigma/\sqrt{n} = \sqrt{pq/n}$.

If n is large, then $\overline{X}$ is approximately normal, so we can apply the results of our last discussion on confidence intervals. We found that

$$P\left(\overline{X} - 1.96\,\frac{\sigma}{\sqrt{n}} \leq \mu \leq \overline{X} + 1.96\,\frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

So that a 95% confidence interval for μ is

$$\left[\overline{X} - 1.96\,\frac{\sigma}{\sqrt{n}}, \overline{X} + 1.96\,\frac{\sigma}{\sqrt{n}}\right].$$

We don't know what σ is, but we do know $\sigma^2 = pq = p(1-p)$. Since p is between 0 and 1, therefore $p(1-p)$ is between 0 and $\frac{1}{4}$, and the maximum $\frac{1}{4}$ occurs when $p = \frac{1}{2}$. Therefore, σ^2 is between 0 and $\frac{1}{4}$, so σ is between 0 and $\frac{1}{2}2$.

Thus, if we replace $\left[\overline{X} - 1.96\,\frac{\sigma}{\sqrt{n}}, \overline{X} + 1.96\,\frac{\sigma}{\sqrt{n}}\right]$, by $\left[\overline{X} - 1.96\,\frac{1}{2\sqrt{n}}, \overline{X} + 1.96\,\frac{1}{2\sqrt{n}}\right]$, we will have an interval that contains a 95% confidence interval no matter what the value of σ is. Since 1.96 is about 2, therefore

$$P(\overline{X} - 1/\sqrt{n} \leq \mu \leq \overline{X} + 1/\sqrt{n}) \geq 0.95,$$

so $[\overline{X} - 1/\sqrt{n}, \overline{X} + 1/\sqrt{n}]$ includes the unknown $p = \mu$ at least 95% of the time. Note that the length of this interval is $2/\sqrt{n}$.

Now suppose we have that bent coin with unknown p and we want to estimate p to one digit, with 95% confidence. The phrase "to within one digit" is usually interpreted to mean within 0.05, and that means the length of the interval is 0.1. How many times to we have to flip the coin? We want $2/\sqrt{n} = 0.1$, so that means $n = 400$. Thus, we've justified the rule of thumb that to get one digit of accuracy for the probability of success p, 400 trials are needed. To get two digits of accuracy, 40000 trials are needed, and that's an awful lot of coin flips.

Math 218 Home Page at
`http://math.clarku.edu/~djoyce/ma218/`