Ch 8. Confidence Intervals — Modules

33. Normal Mean CI's (variance known)

34. Normal Mean CI's (variance unknown)

35. CI's for Other Parameters

33. Normal Mean Confidence Intervals (variance known)

Intro to CI's

CI's for Normal Mean (variance known)

Sample-Size Calculation

CI's for Difference of Two Normal Means (var's known)

Introduction to Confidence Intervals

Instead of estimating a parameter by a point estimator alone, give a (random) **interval** that contains the unknown parameter with a certain probability.

Example: $\bar{X}$ is a point estimator for the parameter μ. A 95% **confidence interval** for μ might look like

$$\mu \in \left[\bar{X} - z_{.025}\sqrt{\sigma^2/n},\ \ \bar{X} + z_{.025}\sqrt{\sigma^2/n}\right],$$

where $z_{.025}$ is the Nor$(0,1)$'s 0.975 quantile. This means that μ is in the interval with probability 0.95.

Definition: A $100(1-\alpha)\%$ **confidence interval** for an unknown parameter θ is given by two random variables L and U satisfying

$$\Pr(L \leq \theta \leq U) = 1 - \alpha.$$

L is the **lower** confidence limit (and is a RV).

U is the **upper** confidence limit (and is a RV).

$1 - \alpha$ is the **confidence coefficient**, specified in advance. There is a $1-\alpha$ chance that θ actually lies between L and U.

Example: We're 95% sure that President Bush's popularity is $56\% \pm 3\%$.

Since $L \leq \theta \leq U$, we call $[L, U]$ a **two-sided** CI for θ.

If L is such that $\Pr(L \leq \theta) = 1 - \alpha$, then $[L, \infty)$ is a $100(1 - \alpha)\%$ **one-sided lower** CI for θ.

Similarly, if U is such that $\Pr(\theta \leq U) = 1 - \alpha$, then $(-\infty, U]$ is a $100(1 - \alpha)\%$ **one-sided upper** CI for θ.

Example: Here are some results from 10 independent samples, each consisting of 100 different observations. From each sample, we use the 100 obs'ns to re-calculate L and U. Is the unknown θ in $[L, U]$?

Sample #	L	U	θ	CI covers θ?
1	1.86	2.23	2	Yes
2	1.90	2.31	2	Yes
3	3.21	3.86	2	No
4	1.75	2.10	2	Yes
5	1.72	2.03	2	Yes
⋮	⋮	⋮	⋮	
10	1.62	1.98	2	No

As the number of samples gets large, the proportion of CI's that cover the unknown θ approaches $1-\alpha$.

Sometimes CI's miss θ too high, i.e., $\theta > U$.

Sometimes CI's miss θ too low, i.e., $\theta < L$.

But $1-\alpha$ of the time, they're OK.

CI's for Normal Mean (variance known)

Set-up: Sample from a normal distribution with unknown mean μ and *known* variance σ^2.

Goal: Obtain a CI for μ.

Remark: This is an unrealistic case, since if we didn't know μ in real life, then we probably wouldn't know σ^2 either. But it's a good place to start the discussion.

Details. . .

Suppose $X_1, \ldots, X_n \stackrel{\text{iid}}{\sim} \text{Nor}(\mu, \sigma^2)$, where σ^2 is *known*.

Use $\bar{X} = \sum_{i=1}^{n} X_i/n$ as our point estimator. Recall

$$\bar{X} \sim \text{Nor}(\mu, \sigma^2/n) \Rightarrow Z \equiv \frac{\bar{X} - \mu}{\sqrt{\sigma^2/n}} \sim \text{Nor}(0, 1).$$

The quantity Z is called a **pivot**. It's a "starting point" for us.

The definition of Z implies that

$$\begin{aligned}
1-\alpha &= \Pr\Big(-z_{\alpha/2} \le Z \le z_{\alpha/2}\Big) \\
&= \Pr\left(-z_{\alpha/2} \le \frac{\bar{X}-\mu}{\sqrt{\sigma^2/n}} \le z_{\alpha/2}\right) \\
&= \Pr\Big(-z_{\alpha/2}\sqrt{\sigma^2/n} \le \bar{X}-\mu \le z_{\alpha/2}\sqrt{\sigma^2/n}\Big) \\
&= \Pr\Big(\underbrace{\bar{X}-z_{\alpha/2}\sqrt{\sigma^2/n}}_{L} \le \mu \le \underbrace{\bar{X}+z_{\alpha/2}\sqrt{\sigma^2/n}}_{U}\Big) \\
&= \Pr\Big(L \le \mu \le U\Big).
\end{aligned}$$

Remarks:

Notice how we used the pivot to "isolate" μ all by itself to the middle of the inequalities.

After you observe $X_1, \ldots, X_n$, you calculate L and U. Nothing is unknown, since L and U don't involve μ.

Sometimes we'll write the CI as $\mu \in \bar{X} \pm H$, where the **half-width** is

$$H \equiv z_{\alpha/2}\sqrt{\sigma^2/n}.$$

Example: Suppose we take $n = 25$ i.i.d. obs'ns from a Nor(μ, σ^2) distribution where we somehow *know* that $\sigma = 30$. Further suppose that $\bar{X}$ turns out to be 278. Let's find a $1 - \alpha = 0.95$ CI for μ.

$$\begin{aligned} \mu &\in \bar{X} \pm z_{\alpha/2}\sqrt{\sigma^2/n} \\ &= 278 \pm z_{.025}(30/5) \\ &= 278 \pm 11.76 \quad (z_{.025} = 1.96). \end{aligned}$$

So a 95% CI for μ is $266.24 \le \mu \le 289.76$.

Sample-Size Calculation

If we had taken more obs'ns, the CI would have gotten shorter, since $H = z_{\alpha/2}\sqrt{\sigma^2/n}$.

In fact, how many obs'ns should be taken to make the half-length (or "error") $\leq \epsilon$?

$$z_{\alpha/2}\sqrt{\sigma^2/n} \leq \epsilon \;\text{ iff }\; \sigma^2/n \leq (\epsilon/z_{\alpha/2})^2$$

iff

$$n \geq (\sigma z_{\alpha/2}/\epsilon)^2.$$

Example: Suppose, in the previous example, that we want the half-length to be $\leq$ 10, i.e., $\mu \in \bar{X} \pm 10$. What should n be?

$$n \geq (\sigma z_{\alpha/2}/\epsilon)^2 = \Big((30)(1.96)/10\Big)^2 = 34.57.$$

Just to make n an integer, round up to $n = 35$.

Remark: We can similarly obtain one-sided CI's for μ (if we're just interested in one bound)...

$100(1-\alpha)\%$ upper CI for μ:

$$\mu \;\leq\; \bar{X} + z_\alpha \sqrt{\sigma^2/n}.$$

$100(1-\alpha)\%$ lower CI for μ:

$$\mu \;\geq\; \bar{X} - z_\alpha \sqrt{\sigma^2/n}.$$

Note that we use the $1-\alpha$ quantile (not $1-\alpha/2$).

CI's for Diff. of Two Normal Means (var's known)

Idea: Compare the means of two competing alternatives or processes by getting a CI for their difference.

Example: Do Georgia Tech students have higher avg. IQ's than Univ. of Georgia students? Answer: Yes.

We'll again assume that the variances of the two competitors are somehow known. We'll do the more-realistic unknown variance case in the next module.

Suppose we have samples of sizes n_x and n_y from the two competing populations.

$$
\begin{aligned}
X_1, X_2, \ldots, X_{n_x} &\overset{\text{iid}}{\sim} \text{Nor}(\mu_x, \sigma_x^2) \quad \text{(population 1)} \\
Y_1, Y_2, \ldots, Y_{n_y} &\overset{\text{iid}}{\sim} \text{Nor}(\mu_y, \sigma_y^2) \quad \text{(population 2)},
\end{aligned}
$$

where the means μ_x and μ_y are *unknown*, while σ_x^2 and σ_y^2 are somehow *known*.

Also assume that the X_i's are *indep* of the Y_i's.

Let's find a CI for the difference in means, $\mu_x - \mu_y$.

Define the sample means from pop'ns 1 and 2,

$$\bar{X} \equiv \frac{1}{n_x} \sum_{i=1}^{n_x} X_i \quad \text{and} \quad \bar{Y} \equiv \frac{1}{n_y} \sum_{i=1}^{n_y} Y_i.$$

Obviously,

$$\bar{X} \sim \text{Nor}(\mu_x, \sigma_x^2/n_x) \text{ and } \quad \bar{Y} \sim \text{Nor}(\mu_y, \sigma_y^2/n_y),$$

so that

$$\bar{X} - \bar{Y} \sim \text{Nor}\left(\mu_x - \mu_y, \frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}\right).$$

This implies that

$$Z \equiv \frac{\bar{X} - \bar{Y} - (\mu_x - \mu_y)}{\sqrt{\frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}}} \sim \text{Nor}(0, 1),$$

so that

$$\Pr(-z_{\alpha/2} \le Z \le z_{\alpha/2}) = 1 - \alpha.$$

Using the same manipulations as in the single-population case, we get a two-sided $100(1-\alpha)\%$ CI for $\mu_x - \mu_y$:

$$\mu_x - \mu_y \in \bar{X} - \bar{Y} \pm z_{\alpha/2}\sqrt{\frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}}.$$

Similarly,

One-sided upper CI:

$$\mu_x - \mu_y \;\leq\; \bar{X} - \bar{Y} + z_\alpha \sqrt{\frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}}.$$

One-sided lower CI:

$$\mu_x - \mu_y \;\geq\; \bar{X} - \bar{Y} - z_\alpha \sqrt{\frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}}.$$

Example (H&M): A manufacturer looks at muzzle velocities of two types of ammo. He assumes that the velocities are normally distributed with $\sigma_x = 1.10$ meters/sec and $\sigma_y = 1.50$ m/s, respectively. He takes random samples of 10 Type 1 shells and 20 Type 2's.

He observes $\bar{X} = 500$ m/s and $\bar{Y} = 494$ m/s, resp.

Find a 90% two-sided CI for $\mu_x - \mu_y$.

Plugging into the two-sided equation with $z_{\alpha/2} = z_{.05}$, we get

$$\begin{aligned} \mu_x - \mu_y &= \bar{X} - \bar{Y} \pm z_{\alpha/2}\sqrt{\frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}} \\ &= 500 - 494 \pm 1.645\sqrt{\frac{1.21}{10} + \frac{2.25}{20}}, \end{aligned}$$

implying that $5.21 \le \mu_x - \mu_y \le 6.79$.

In other words, we're 90% sure that $\mu_x - \mu_y$ lies in this interval.

Remark: Let's assume that both sample sizes $n_x = n_y = n$. To obtain a half-length $\leq \epsilon$, we require

$$n \;\geq\; \frac{z^2_{\alpha/2}(\sigma_x^2 + \sigma_y^2)}{\epsilon^2}.$$

The next module will examine the more-realistic case in which the variance of the underlying obs'ns is unknown. . .