Confidence Intervals for Means - Known Variance

Gene Quinn

Interval Estimation

The maximum likelihood and method of moments estimates we have developed each consist of a **single value**

Interval Estimation

The maximum likelihood and method of moments estimates we have developed each consist of a **single value**

For this reason, they are called *point estimates*.

An obvious drawback of stating a single value is that it gives no information about the precision of the estimate.

Interval Estimation

The maximum likelihood and method of moments estimates we have developed each consist of a **single value**

For this reason, they are called *point estimates*.

An obvious drawback of stating a single value is that it gives no information about the precision of the estimate.

To address this shortcoming, a different type of estimate called an *interval estimate* can be used.

Unknown Variance

An **interval estimate** consists of:

- An **interval** (L, U), usually centered at a sample mean

Unknown Variance

An **interval estimate** consists of:

- An **interval** (L, U), usually centered at a sample mean
- A **level of confidence** that the interval contains the (unknown) parameter

Unknown Variance

An **interval estimate** consists of:

- An **interval** (L, U), usually centered at a sample mean
- A **level of confidence** that the interval contains the (unknown) parameter

The level of confidence is usually stated as a percentage, say 95%.

Unknown Variance

An **interval estimate** consists of:

- An **interval** (L, U), usually centered at a sample mean
- A **level of confidence** that the interval contains the (unknown) parameter

The level of confidence is usually stated as a percentage, say 95%.

The interpretation of this number is that, if a large number of samples are taken and a new confidence interval (L, U) is calculated for each of them, 95% of the time the interval will contain the unknown paramter value.

Computing the Interval

Suppose a random variable X has a normal distribution with mean $\mu = 120$ and variance $\sigma^2 = 16$

Computing the Interval

Suppose a random variable X has a normal distribution with mean $\mu = 120$ and variance $\sigma^2 = 16$

We denote this situation with the notation $X \sim (120, 16)$

Computing the Interval

Suppose a random variable X has a normal distribution with mean $\mu = 120$ and variance $\sigma^2 = 16$

We denote this situation with the notation $X \sim (120, 16)$

We are interested in finding an interval centered at $\mu = 120$ that has an area of 0.95.

That is, an interval $[L, U]$ such that

$$\int_L^U \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right) dx \ = \ 0.95$$

Computing the Interval

Suppose a random variable X has a normal distribution with mean $\mu = 120$ and variance $\sigma^2 = 16$

We denote this situation with the notation $X \sim (120, 16)$

We are interested in finding an interval centered at $\mu = 120$ that has an area of 0.95.

That is, an interval $[L, U]$ such that

$$\int_L^U \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right) dx \;=\; 0.95$$

Computing the Interval

It's impossible to find an antiderivative of

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right)$$

so we can't evaluate the integral directly.

Computing the Interval

It's impossible to find an antiderivative of

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right)$$

so we can't evaluate the integral directly.

However, spreadsheets and statistical packages provide functions to evaluate the integral numerically.

Computing the Interval

The spreadsheet function

=NORMINV(p, μ, σ)

returns the value x_p on the x-axis with the property that

$$\int_{-\infty}^{x_p} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right) dx \;=\; p$$

Computing the Interval

If we are interested in an interval centered at μ that contains 95% of the area under the bell curve associated with the $N(\mu, \sigma^2)$ distribution, the limits can be obtained as:

LOWER LIMIT: =NORMINV$(0.025, \mu, \sigma)$

UPPER LIMIT: =NORMINV$(0.975, \mu, \sigma)$

Computing the Interval

In the case $\mu = 120$ and $\sigma^2 = 16$,

LOWER LIMIT: =NORMINV$(0.025, 120, 4) = 112.16$

UPPER LIMIT: =NORMINV$(0.975, 120, 4) = 127.84$

Computing the Interval

The following graph depicts the interval $(112.16, 127.84)$ which contains 95% of the area under the bell curve:

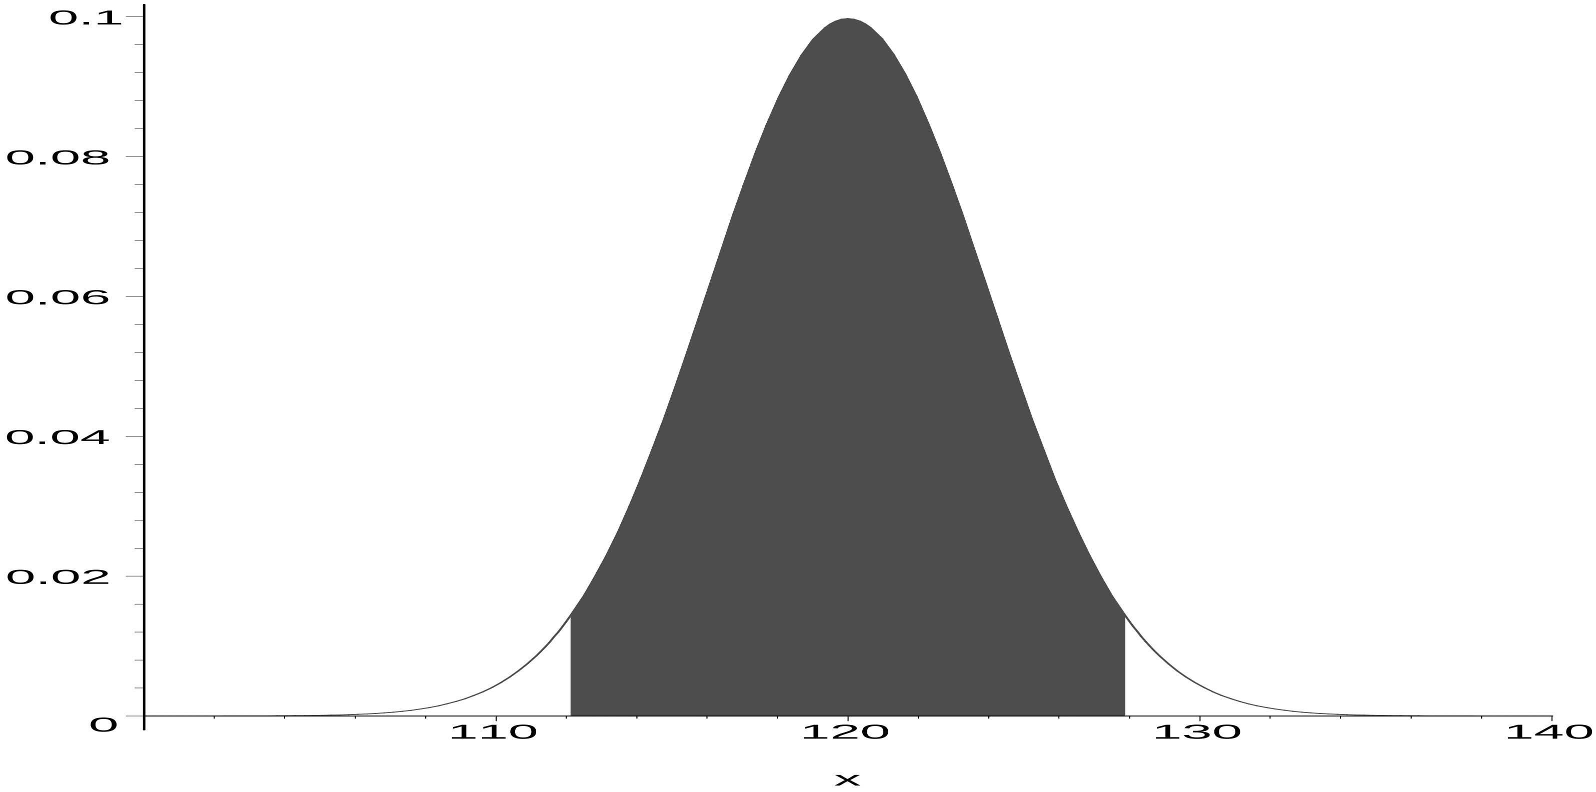

Sample Means

Now we consider how to use the mean $\overline{x}$ of a random sample to estimate the population mean μ.

Sample Means

Now we consider how to use the mean $\overline{x}$ of a random sample to estimate the population mean μ.

Note that $\overline{x}$ is a *random variable* and as such it usually has a density function, expected value, variance, cdf, and any of the other constructs associated with random variables.

Sample Means

Now we consider how to use the mean $\overline{x}$ of a random sample to estimate the population mean μ.

Note that $\overline{x}$ is a *random variable* and as such it usually has a density function, expected value, variance, cdf, and any of the other constructs associated with random variables.

Suppose

$$x = \{x_1, \ldots, x_n\}$$

represents a random sample from some population.

Sample Means

By the definition of a random sample, the individual random variables x_i that make up x are:

- Independent
- Identically distributed (i.e., they all have the same pdf)

Sample Means

By the definition of a random sample, the individual random variables x_i that make up x are:

- Independent
- Identically distributed (i.e., they all have the same pdf)

The previous two conditions are sometimes abreviated "IID" or "Independent, identically distributed".

Sample Means

Suppose the random variables x_i have:

$$\mathsf{E}(x_i) \;=\; \mu \quad \text{and} \quad \mathsf{Var}(x_i) \;=\; \sigma^2, \quad i = 1, \ldots, n$$

Sample Means

Suppose the random variables x_i have:

$$\mathsf{E}(x_i) = \mu \quad \text{and} \quad \mathsf{Var}(x_i) = \sigma^2, \quad i = 1, \ldots, n$$

Then the vector of expected values and variance-covariance matrix associated with the joint distribution of the x_i is:

$$\vec{\mu} = \begin{bmatrix} \mu \\ \mu \\ \vdots \\ \mu \end{bmatrix} \quad \text{and} \quad V = \begin{bmatrix} \sigma^2 & & & \\ & \sigma^2 & & \\ & & \ddots & \\ & & & \sigma^2 \end{bmatrix} = \sigma^2 I$$

Sample Means

The sample mean $\overline{x}$ is a linear combination $t'x$ of the x_i, with

$$t = \begin{bmatrix} \frac{1}{n} \\ \frac{1}{n} \\ \vdots \\ \frac{1}{n} \end{bmatrix}$$

Sample Means

The sample mean $\overline{x}$ is a linear combination $t'x$ of the x_i, with

$$t = \begin{bmatrix} \frac{1}{n} \\ \frac{1}{n} \\ \vdots \\ \frac{1}{n} \end{bmatrix}$$

From an earlier result, the expected value of $\overline{x}$ is:

$$\mathsf{E}(\overline{x}) = t'\vec{\mu} = \sum_{i=1}^{n} \frac{1}{n}\mu = \mu$$

Sample Means

The variance of $\overline{x}$ is:

$$\text{Var}(\overline{x}) = t'Vt = \sum_{i=1}^{n} \frac{\sigma^2}{n} \cdot n = \frac{\sigma^2}{n}$$

Sample Means

The variance of $\overline{x}$ is:

$$\text{Var}(\overline{x}) \;=\; t'Vt \;=\; \sum_{i=1}^{n} \frac{\sigma^2}{n} \cdot n \;=\; \frac{\sigma^2}{n}$$

These results hold as long as the expected values and variances of the x_i exist.

They **do not** require any assumptions about their distributions.

Sample Means

However, if the x_i are normally distributed, that is,

$$x_i \sim N(\mu, \sigma^2), \quad i = 1, \ldots, n$$

then

$$\overline{x} \sim N(\mu, \frac{\sigma^2}{n})$$

Sample Means

However, if the x_i are normally distributed, that is,

$$x_i \sim N(\mu, \sigma^2), \quad i = 1, \ldots, n$$

then

$$\overline{x} \sim N(\mu, \frac{\sigma^2}{n})$$

This distribution is the basis for constructing confidence intervals for the parameter μ.

Sample Means

Theorem: If

$$x_i \sim N(\mu, \sigma^2), \quad i = 1, \ldots, n$$

then 95% of the time the interval

$$\left[\overline{x} - 1.96\sqrt{\frac{\sigma^2}{n}}, \overline{x} + 1.96\sqrt{\frac{\sigma^2}{n}}\right]$$

will contain the (unknown) population mean μ.

Sample Means

Theorem: More generally, for $\alpha \in (0,1)$, if

$$x_i \sim N(\mu, \sigma^2), \quad i = 1, \ldots, n$$

then $100(1-\alpha)\%$ of the time the interval

$$\left[\overline{x} - z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}, \overline{x} + z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}\right]$$

will contain the (unknown) population mean μ.

Sample Means

Theorem: More generally, for $\alpha \in (0,1)$, if

$$x_i \sim N(\mu, \sigma^2), \quad i = 1, \ldots, n$$

then $100(1-\alpha)\%$ of the time the interval

$$\left[\overline{x} - z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}, \overline{x} + z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}\right]$$

will contain the (unknown) population mean μ.

Here

$$z_{\alpha/2} \quad = \text{NORMSINV}(1 - \alpha/2)$$

Sample Means

Example: Suppose x is a random sample of size 16 from a normal population with *known* variance Var$(x_i) = 4, \quad i = 1, \ldots, n$.

If the sample mean $\overline{x}$ is 13.5, construct a 95% confidence interval for the population mean μ.

Sample Means

Example: Suppose x is a random sample of size 16 from a normal population with *known* variance
Var$(x_i) = 4, \quad i = 1, \ldots, n$.

If the sample mean $\overline{x}$ is 13.5, construct a 95% confidence interval for the population mean μ.

Here $\alpha = 0.05$ and

$$z_{\alpha/2} \quad \text{=NORMSINV}(1 - 0.05/2) \;=\; 1.96$$

Sample Means

The 95% confidence interval for the population mean μ is:

$$\left[\overline{x} - z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}, \overline{x} + z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}\right]$$

Sample Means

The 95% confidence interval for the population mean μ is:

$$\left[\overline{x} - z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}, \overline{x} + z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}\right]$$

$$= \left[13.5 - 1.96\sqrt{\frac{4}{16}}, 13.5 + 1.96\sqrt{\frac{4}{16}}\right]$$

Sample Means

The 95% confidence interval for the population mean μ is:

$$\left[\overline{x} - z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}, \overline{x} + z_{\alpha/2}\sqrt{\frac{\sigma^2}{n}}\right]$$

$$= \left[13.5 - 1.96\sqrt{\frac{4}{16}}, 13.5 + 1.96\sqrt{\frac{4}{16}}\right]$$

$$= \left[13.5 - \frac{1.96}{2}, 13.5 + \frac{1.96}{2}\right] = [12.52, 14.48]$$

Sample Means

The interpretation is that we are 95% sure that the interval [12.52,14.48] contains the unknown population mean μ.

Sample Means

The interpretation is that we are 95% sure that the interval [12.52,14.48] contains the unknown population mean μ.

More precisely, if we repeated the experiment of drawing a sample of 16 many times and calculated a confidence interval each time, 19 out of 20 of those intervals would contain μ.

Sample Means

The interpretation is that we are 95% sure that the interval [12.52,14.48] contains the unknown population mean μ.

More precisely, if we repeated the experiment of drawing a sample of 16 many times and calculated a confidence interval each time, 19 out of 20 of those intervals would contain μ.

For an individual sample mean, we have no idea whether the interval contains μ or not.