Interval Estimates

Measures of location and dispersion such as the sample mean and standard deviation that consist of a single value are called **point estimators**.

Interval Estimates

Measures of location and dispersion such as the sample mean and standard deviation that consist of a single value are called **point estimators**.

A shortcoming of this type of estimator is that it provides no information about how precise or reliable the estimate is.

Interval Estimates

Measures of location and dispersion such as the sample mean and standard deviation that consist of a single value are called **point estimators**.

A shortcoming of this type of estimator is that it provides no information about how precise or reliable the estimate is.

An alternative type of estimate is an **interval estimate**, which instead of a single value provides a range of values.

Interval Estimates

Measures of location and dispersion such as the sample mean and standard deviation that consist of a single value are called **point estimators**.

A shortcoming of this type of estimator is that it provides no information about how precise or reliable the estimate is.

An alternative type of estimate is an **interval estimate**, which instead of a single value provides a range of values.

We will cover three approaches to constructing interval estimates:

- The classical or *frequentist* approach which produces a **confidence interval**
- The Bayesian approach with produces a **credible interval**
- The bootstrap procedure which produces a **bootstrap confidence interval**

Confidence Intervals

We start off by deciding the fraction we want to contain the parameter, usually by specifying the fraction of intervals that *will not* contain it, which is called α and is often chosen to be $.05$.

Confidence Intervals

We start off by deciding the fraction we want to contain the parameter, usually by specifying the fraction of intervals that *will not* contain it, which is called α and is often chosen to be $.05$.

The classical or frequentist approach then constructs a **confidence interval** in such a way that, regardless of the actual value of the target parameter, if we took a great many samples and constructed a confidence interval for each sample, the percentage of them that actually contain the parameter value would approach

$$100(1-\alpha)\% \quad \text{or} \quad 95\% \quad \text{if} \quad \alpha = .05$$

Confidence Intervals

We start off by deciding the fraction we want to contain the parameter, usually by specifying the fraction of intervals that *will not* contain it, which is called α and is often chosen to be $.05$.

The classical or frequentist approach then constructs a **confidence interval** in such a way that, regardless of the actual value of the target parameter, if we took a great many samples and constructed a confidence interval for each sample, the percentage of them that actually contain the parameter value would approach

$$100(1-\alpha)\% \quad \text{or} \quad 95\% \quad \text{if} \quad \alpha = .05$$

Of course, we actually construct only one interval, and we never know whether it contains the parameter or not.

Confidence Intervals

This description of a confidence interval sounds complicated, and it is.

Confidence Intervals

This description of a confidence interval sounds complicated, and it is.

One often hears a 95% confidence interval described as "the true parameter value falls into this interval with probability $.95$, but this is really not a correct statement since the true parameter value is not assumed to be a random variable in the frequentist approach.

Confidence Intervals

This description of a confidence interval sounds complicated, and it is.

One often hears a 95% confidence interval described as "the true parameter value falls into this interval with probability $.95$, but this is really not a correct statement since the true parameter value is not assumed to be a random variable in the frequentist approach.

One of the criticisms of the frequentist approach is that it produces interval estimates that are difficult to describe correctly.

Confidence Intervals

Like a point estimate, an interval estimate is based on a random sample.

Confidence Intervals

Like a point estimate, an interval estimate is based on a random sample.

It is constructed in a way that guarantees that, for a 95% confidence interval, over the long run if we construct many of them, the percentage of them that actually contain the true value of the parameter approaches 95%.

Confidence Intervals

Like a point estimate, an interval estimate is based on a random sample.

It is constructed in a way that guarantees that, for a 95% confidence interval, over the long run if we construct many of them, the percentage of them that actually contain the true value of the parameter approaches 95%.

The exact procedure for constructing a confidence interval varies depending on the target parameter and what we assume for the distribution and for other parameters.

Confidence Intervals

Like a point estimate, an interval estimate is based on a random sample.

It is constructed in a way that guarantees that, for a 95% confidence interval, over the long run if we construct many of them, the percentage of them that actually contain the true value of the parameter approaches 95%.

The exact procedure for constructing a confidence interval varies depending on the target parameter and what we assume for the distribution and for other parameters.

We will confine our discussion to the most common special cases.

Classical Large Sample (sigma known)

Case 1: Sampling from a normal population where we either know the actual population standard deviation, or have a large enough sample (at least 30, preferably 50 or more) so that we can reliably estimate it from the sample.

Classical Large Sample (sigma known)

Case 1: Sampling from a normal population where we either know the actual population standard deviation, or have a large enough sample (at least 30, preferably 50 or more) so that we can reliably estimate it from the sample.

It might seem unusual to know the standard deviation but not the mean of an underlying normal population, but this situation arises in the case of standardized measures such as IQ and SAT.

Classical Large Sample (sigma known)

Case 1: Sampling from a normal population where we either know the actual population standard deviation, or have a large enough sample (at least 30, preferably 50 or more) so that we can reliably estimate it from the sample.

It might seem unusual to know the standard deviation but not the mean of an underlying normal population, but this situation arises in the case of standardized measures such as IQ and SAT.

This is called a "large sample" or "sigma known" confidence interval because we treat the population standard deviation as a known quantity.

Classical Large Sample (sigma known)

As noted previously, the Cumulative Distribution Function $F(x)$ of the normal distribution cannot be expressed in a simple formula.

Classical Large Sample (sigma known)

As noted previously, the Cumulative Distribution Function $F(x)$ of the normal distribution cannot be expressed in a simple formula.

With statistical software readily available, our main tool evaluating the normal CDF $F(x)$ will be a computer or calculator.

Classical Large Sample (sigma known)

As noted previously, the Cumulative Distribution Function $F(x)$ of the normal distribution cannot be expressed in a simple formula.

With statistical software readily available, our main tool evaluating the normal CDF $F(x)$ will be a computer or calculator.

The actual code for producing the value of $F(x)$ is:

- $= NORMDIST(x, \mu, \sigma, TRUE)$ for a spreadsheet
- $pnorm(x, \mu, \sigma)$ for R

Classical Large Sample (sigma known)

As noted previously, the Cumulative Distribution Function $F(x)$ of the normal distribution cannot be expressed in a simple formula.

With statistical software readily available, our main tool evaluating the normal CDF $F(x)$ will be a computer or calculator.

The actual code for producing the value of $F(x)$ is:

- $= NORMDIST(x, \mu, \sigma, TRUE)$ for a spreadsheet
- $pnorm(x, \mu, \sigma)$ for R

Either method produces $F(x) = P(X \leq x)$ for a $N(\mu, \sigma)$ distribution.

Classical Large Sample (sigma known)

Sometimes we have the opposite problem: given a probability value p, we want to find x such that

$$P(X \leq x) = p \quad \text{when} \quad X \sim N(\mu, \sigma)$$

Classical Large Sample (sigma known)

Sometimes we have the opposite problem: given a probability value p, we want to find x such that

$$P(X \leq x) = p \quad \text{when} \quad X \sim N(\mu, \sigma)$$

We could try different values of x with $NORMDIST(x, \mu, \sigma, TRUE)$ or $pnorm(x, \mu, \sigma)$ but this would be time consuming.

Classical Large Sample (sigma known)

Sometimes we have the opposite problem: given a probability value p, we want to find x such that

$$P(X \le x) = p \quad \text{when} \quad X \sim N(\mu, \sigma)$$

We could try different values of x with $NORMDIST(x, \mu, \sigma, TRUE)$ or $pnorm(x, \mu, \sigma)$ but this would be time consuming.

A better way is to use the **inverse CDF** functions:

- $= NORMINV(p, \mu, \sigma)$ for a spreadsheet
- $qnorm(p, \mu, \sigma)$ for R

Classical Large Sample (sigma known)

Sometimes we have the opposite problem: given a probability value p, we want to find x such that

$$P(X \leq x) = p \quad \text{when} \quad X \sim N(\mu, \sigma)$$

We could try different values of x with $NORMDIST(x, \mu, \sigma, TRUE)$ or $pnorm(x, \mu, \sigma)$ but this would be time consuming.

A better way is to use the **inverse CDF** functions:

- $= NORMINV(p, \mu, \sigma)$ for a spreadsheet
- $qnorm(p, \mu, \sigma)$ for R

Either method produces x with the property that $F(x) = P(X \leq x) = p$ for a $N(\mu, \sigma)$ distribution.

Example

Find the value x with the property that a normal random variable with mean $\mu = 100$ and standard deviation $\sigma = 15$ takes a value less than or equal to x with probability $p = 0.05$.

Example

Find the value x with the property that a normal random variable with mean $\mu = 100$ and standard deviation $\sigma = 15$ takes a value less than or equal to x with probability $p = 0.05$.

Using either of the **inverse CDF** functions:

- $x = NORMINV(.05, 100, 15)$ for a spreadsheet
- $x = qnorm(.05, 100, 15)$ for R

we find that $x = 75.33$.

Example

Find the value x with the property that a normal random variable with mean $\mu = 100$ and standard deviation $\sigma = 15$ takes a value less than or equal to x with probability $p = 0.05$.

Using either of the **inverse CDF** functions:

- $x = NORMINV(.05, 100, 15)$ for a spreadsheet
- $x = qnorm(.05, 100, 15)$ for R

we find that $x = 75.33$.

As a check, entering $= NORMDIST(75.33, 100, 15)$ or $pnorm(75.33, 100, 15)$ should produce $p = .05$.

Example

Find the value x with the property that a normal random variable with mean $\mu = 500$ and standard deviation $\sigma = 100$ takes a value less than or equal to x with probability $p = 0.98$ (i.e., find the 98^{th} percentile SAT score).

Example

Find the value x with the property that a normal random variable with mean $\mu = 500$ and standard deviation $\sigma = 100$ takes a value less than or equal to x with probability $p = 0.98$ (i.e., find the 98^{th} percentile SAT score).

Using either of the **inverse CDF** functions:

- $x = NORMINV(.98, 500, 100)$ for a spreadsheet
- $x = qnorm(.98, 500, 100)$ for R

we find that $x = 705$.

Example

Find the value x with the property that a normal random variable with mean $\mu = 500$ and standard deviation $\sigma = 100$ takes a value less than or equal to x with probability $p = 0.98$ (i.e., find the 98^{th} percentile SAT score).

Using either of the **inverse CDF** functions:

- $x = NORMINV(.98, 500, 100)$ for a spreadsheet
- $x = qnorm(.98, 500, 100)$ for R

we find that $x = 705$.

As a check, entering $= NORMDIST(705, 500, 100)$ or $pnorm(705, 500, 100)$ should produce $p = .98$.

Classical Large Sample (sigma known)

Suppose we have a sample $X_1, X_2, \ldots, X_n$ from a normal population with **unknown** mean μ and **known** standard deviation σ.

Now we present a procedure for constructing an interval estimate (L, U) for the (unknown) mean μ.

Classical Large Sample (sigma known)

Suppose we have a sample $X_1, X_2, \ldots, X_n$ from a normal population with **unknown** mean μ and **known** standard deviation σ.

Now we present a procedure for constructing an interval estimate (L, U) for the (unknown) mean μ.

Suppose we want an interval (L, U) that contains the true value μ with probability $1 - \alpha$.

Classical Large Sample (sigma known)

Suppose we have a sample $X_1, X_2, \ldots, X_n$ from a normal population with **unknown** mean μ and **known** standard deviation σ.

Now we present a procedure for constructing an interval estimate (L, U) for the (unknown) mean μ.

Suppose we want an interval (L, U) that contains the true value μ with probability $1 - \alpha$.

The correct interpretation of this is that if we took a large number of samples and constructed an interval (L, U) for each sample, we would get many different intervals, and on average $100(1 - \alpha)$ percent of them would contain μ, that is,

$$P(L \leq \mu \leq U) = 1 - \alpha$$

Classical Large Sample (sigma known)

First we choose the level of confidence we want. Let's say this is 95%. Solve the following equation to get α:

$$100(1-\alpha) = 95 \quad \text{or} \quad 1 - \frac{95}{100} = .05 = \alpha$$

Classical Large Sample (sigma known)

First we choose the level of confidence we want. Let's say this is 95%. Solve the following equation to get α:

$$100(1-\alpha) = 95 \quad \text{or} \quad 1 - \frac{95}{100} = .05 = \alpha$$

Now we compute the endpoints (L, U) of the $100(1-\alpha)\%$ confidence interval using:

- The α value derived from the level of confidence
- $\overline{x}$, the sample mean
- σ the **known** population standard deviation.
- The sample size n

Classical Large Sample (sigma known)

Recall that in this situation, the sample mean has a normal distribution:

$$\overline{x} \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$$

Classical Large Sample (sigma known)

Recall that in this situation, the sample mean has a normal distribution:

$$\overline{x} \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$$

Now we compute the endpoints (L, U) of the $100(1-\alpha)\%$ confidence interval using:

- $L = NORMINV(\alpha/2, \overline{x}, \sigma/\sqrt{n})$ or
  $L = qnorm(\alpha/2, \overline{x}, \sigma/\sqrt{n})$
- $U = NORMINV(1-\alpha/2, \overline{x}, \sigma/\sqrt{n})$ or
  $U = qnorm(1-\alpha/2, \overline{x}, \sigma/\sqrt{n})$

Classical Large Sample (sigma known)

Recall that in this situation, the sample mean has a normal distribution:

$$\overline{x} \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$$

Now we compute the endpoints (L, U) of the $100(1-\alpha)\%$ confidence interval using:

- $L = NORMINV(\alpha/2, \overline{x}, \sigma/\sqrt{n})$ or
  $L = qnorm(\alpha/2, \overline{x}, \sigma/\sqrt{n})$
- $U = NORMINV(1-\alpha/2, \overline{x}, \sigma/\sqrt{n})$ or
  $U = qnorm(1-\alpha/2, \overline{x}, \sigma/\sqrt{n})$

Notice that once the values of α, n and σ are determined, L and U depend only on $\overline{x}$.

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 95% confidence interval for the mean SAT score in this school district.

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 95% confidence interval for the mean SAT score in this school district.

In this case, $\alpha = .05$, $n = 100$, $\overline{x} = 507$, and we assume that σ is known to be 100 because the SAT is standardized to have this value.

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 95% confidence interval for the mean SAT score in this school district.

In this case, $\alpha = .05$, $n = 100$, $\overline{x} = 507$, and we assume that σ is known to be 100 because the SAT is standardized to have this value.

Then:

$$L = NORMINV(.025, 507, 10)$$

or $L = qnorm(.025, 507, 10) = 487.4$

and

$$U = NORMINV(.975, 507, 10)$$

or $L = qnorm(.975, 507, 10) = 526.6$

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 99% confidence interval for the mean SAT score in this school district.

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 99% confidence interval for the mean SAT score in this school district.

In this case, $\alpha = .01$, $n = 100$, $\overline{x} = 507$, and we assume that σ is known to be 100 because the SAT is standardized to have this value.

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 99% confidence interval for the mean SAT score in this school district.

In this case, $\alpha = .01$, $n = 100$, $\overline{x} = 507$, and we assume that σ is known to be 100 because the SAT is standardized to have this value.

Then:

$$L = NORMINV(.005, 507, 10)$$
$$\text{or } L = qnorm(.005, 507, 10) = 481.24$$

and

$$U = NORMINV(.995, 507, 10)$$
$$\text{or } L = qnorm(.995, 507, 10) = 532.76$$

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 90% confidence interval for the mean SAT score in this school district.

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 90% confidence interval for the mean SAT score in this school district.

In this case, $\alpha = .10$, $n = 100$, $\overline{x} = 507$, and we assume that σ is known to be 100 because the SAT is standardized to have this value.

Example

A random sample of 100 students in a large school district are given the SAT. If the average score is 507, construct a 90% confidence interval for the mean SAT score in this school district.

In this case, $\alpha = .10$, $n = 100$, $\overline{x} = 507$, and we assume that σ is known to be 100 because the SAT is standardized to have this value.

Then:

$$L = NORMINV(.05, 507, 10)$$

or $L = qnorm(.005, 507, 10) = 490.55$

and

$$U = NORMINV(.95, 507, 10)$$

or $L = qnorm(.995, 507, 10) = 523.44$

Classical Large Sample (sigma known)

Generally the higher the level of confidence, the wider the interval. For the preceding examples,

90% confidence (L,U)=(490.5,523.4)

95% confidence (L,U)=(487.4,526.6)

99% confidence (L,U)=(481.2,532.8)

Classical Large Sample (sigma known)

Generally the higher the level of confidence, the wider the interval. For the preceding examples,

90% confidence (L,U)=(490.5,523.4)

95% confidence (L,U)=(487.4,526.6)

99% confidence (L,U)=(481.2,532.8)

Notice that our confidence interval is interval centered at $\overline{x}$ that would contain $100(1-\alpha)$ percent of the area under a normal curve with mean $\overline{x}$ and standard deviation $\sigma/\sqrt{n}$.

Example - R Simulations

There is an R program that simulates the behavior of classical large sample confidence intervals from the notes and handouts section of the web page. The link is named:

Classical large sample confidence intervals (R simulation)

Example - R Simulations

There is an R program that simulates the behavior of classical large sample confidence intervals from the notes and handouts section of the web page. The link is named:

Classical large sample confidence intervals (R simulation)

One way to execute this directly is to copy the link from the web page and paste it between the quotes in the command
`source("")`

Example - R Simulations

There is an R program that simulates the behavior of classical large sample confidence intervals from the notes and handouts section of the web page. The link is named:

Classical large sample confidence intervals (R simulation)

One way to execute this directly is to copy the link from the web page and paste it between the quotes in the command `source("")`

You can also download it and open it from the file menu in R, or run R in batch mode and specify it as input.