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**, also known as a *confidence interval*.

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**, also known as a *confidence interval*.

An **interval estimate** consists of:

- An interval or range of values (a, b)
- A *confidence level*, usually expressed as a percentage, representing the probability that the interval contains the true population value.

Interval Estimates

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

Interval Estimates

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

Using the sample, an interval estimation procedure constructs an interval that has a specified probability, say $.95$, of containing the true population value. In this case, it would be called a 95% confidence interval.

Interval Estimates

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

Using the sample, an interval estimation procedure constructs an interval that has a specified probability, say $.95$, of containing the true population value. In this case, it would be called a 95% confidence interval.

Because it is based on a random sample, the interval itself is random.

Interval Estimates

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

Using the sample, an interval estimation procedure constructs an interval that has a specified probability, say $.95$, of containing the true population value. In this case, it would be called a 95% confidence interval.

Because it is based on a random sample, the interval itself is random.

This means that different samples from the same population are likely to produce different intervals.

Interval Estimates

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

Using the sample, an interval estimation procedure constructs an interval that has a specified probability, say $.95$, of containing the true population value. In this case, it would be called a 95% confidence interval.

Because it is based on a random sample, the interval itself is random.

This means that different samples from the same population are likely to produce different intervals.

The intervals are constructed in such a way that for each one, the probability that it contains the true population value is (in this example) $.95$

Interval Estimates

What follows is a procedure for using a random sample $X_1, X_2, \ldots, X_n$ to construct a confidence interval for the (unknown) mean of a normal population with **known** standard deviation σ.

Interval Estimates

What follows is a procedure for using a random sample $X_1, X_2, \ldots, X_n$ to construct a confidence interval for the (unknown) mean of a normal population with **known** standard deviation σ.

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.

Interval Estimates

What follows is a procedure for using a random sample $X_1, X_2, \ldots, X_n$ to construct a confidence interval for the (unknown) mean of a normal population with **known** standard deviation σ.

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.

The theory underlying the construction of confidence intervals from a random sample guarantees that the resulting interval will contain the true population mean with a certain probability or level of confidence, *regardless of what the actual population value* μ *is*.

Interval Estimates

What follows is a procedure for using a random sample $X_1, X_2, \ldots, X_n$ to construct a confidence interval for the (unknown) mean of a normal population with **known** standard deviation σ.

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.

The theory underlying the construction of confidence intervals from a random sample guarantees that the resulting interval will contain the true population mean with a certain probability or level of confidence, *regardless of what the actual population value* μ *is*.

Interval Estimates

The theory leads to different procedures for different statistics. A confidence interval for the standard deviation of a normal population is not constructed the same way as a confidence interval for the mean.

Interval Estimates

The theory leads to different procedures for different statistics. A confidence interval for the standard deviation of a normal population is not constructed the same way as a confidence interval for the mean.

Although the computational procedure will be different, a confidence interval for another population value such as the standard deviation still has the basic property that it contains the true population value σ with a certain probability or level of confidence.

Interval Estimates

The theory leads to different procedures for different statistics. A confidence interval for the standard deviation of a normal population is not constructed the same way as a confidence interval for the mean.

Although the computational procedure will be different, a confidence interval for another population value such as the standard deviation still has the basic property that it contains the true population value σ with a certain probability or level of confidence.

We will consider only confidence intervals for the population mean.

The Normal CDF

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

The Normal CDF

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 Normal CDF

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

The Normal CDF

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.

The Normal CDF Inverse

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

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

The Normal CDF Inverse

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.

The Normal CDF Inverse

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

The Normal CDF Inverse

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 $p = .98$.

Confidence Intervals for Means

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.

Confidence Intervals for Means

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 $1 - \alpha$.

Confidence Intervals for Means

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,

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

Constructing the Confidence Interval

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

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

Constructing the Confidence Interval

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

Constructing the Confidence Interval

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

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

Constructing the Confidence Interval

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

Constructing the Confidence Interval

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:

$$L = NORMINV(.005, 507, 10)$$

or $L = qnorm(.005, 507, 10) = 481.24$

and

$$U = NORMINV(.995, 507, 10)$$

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

Constructing the Confidence Interval

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)

Constructing the Confidence Interval

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 $\sigma/\sqrt{n}$.

Monte Carlo Experiment

Now we will perform a simulation experiment in which we:

- generate random samples of size 10 from a normal population with known mean and standard deviation
- construct $100(1-\alpha)\%$ confidence intervals for each sample
- determine the proportion of these intervals that contain the true population mean

Monte Carlo Experiment

Now we will perform a simulation experiment in which we:

- generate random samples of size 10 from a normal population with known mean and standard deviation
- construct $100(1-\alpha)\%$ confidence intervals for each sample
- determine the proportion of these intervals that contain the true population mean

Naturally, we expect that approximately $100(1-\alpha)$ percent of the intervals will contain the true mean.

Monte Carlo Experiment

Starting with a blank spreadsheet, enter the following values:

Value	Cell Address
ALPHA	A1
0.05	B1
MU	C1
10	D1
SIGMA	E1
3	F1

Monte Carlo Experiment

Starting with a blank spreadsheet, enter the following values:

Value	Cell Address
ALPHA	A1
0.05	B1
MU	C1
10	D1
SIGMA	E1
3	F1

Now carefully enter the following in cell $E3$:

=NORMINV(RAND(),$D$1,$F$1)

Monte Carlo Experiment

Starting with a blank spreadsheet, enter the following values:

Value	Cell Address
ALPHA	A1
0.05	B1
MU	C1
10	D1
SIGMA	E1
3	F1

Now carefully enter the following in cell $E3$:

=NORMINV(RAND(),$D$1,$F$1)

This generates a $N(\mu, \sigma)$ random variable.

Monte Carlo Experiment

Now replicate cell E3 horizontally 9 times, filling cells F3-N3.

This gives us a simulated random sample of size $n = 10$ from a $N(\mu, \sigma)$ population.

Monte Carlo Experiment

Now replicate cell E3 horizontally 9 times, filling cells F3-N3.

This gives us a simulated random sample of size $n = 10$ from a $N(\mu, \sigma)$ population.

Now carefully enter the following in cell $D3$:

=AVERAGE(E3:N3)

This computes the sample mean $\overline{x}$.

Monte Carlo Experiment

Now carefully enter the following in cell $B3$:

=NORMINV($B$1/2,D3,$F$1/SQRT(10))

This gives us the lower limit of the $100(1-\alpha)\%$ confidence interval for μ.

Monte Carlo Experiment

Now carefully enter the following in cell $B3$:

=NORMINV($B$1/2,D3,$F$1/SQRT(10))

This gives us the lower limit of the $100(1-\alpha)\%$ confidence interval for μ.

Now carefully enter the following in cell $C3$:

=NORMINV(1-$B$1/2,D3,$F$1/SQRT(10))

This gives us the upper limit of the $100(1-\alpha)\%$ confidence interval for μ.

Monte Carlo Experiment

All that remains is to set up a counter. Enter the following in cell $A3$:

=IF(AND(B3<$D$1,$D$1<C3),1,0)

This will code a value of one if the true population mean, in cell $D1$, lies in the confidence interval, and a value of zero if it does not.

Monte Carlo Experiment

All that remains is to set up a counter. Enter the following in cell $A3$:

=IF(AND(B3<$D$1,$D$1<C3),1,0)

This will code a value of one if the true population mean, in cell $D1$, lies in the confidence interval, and a value of zero if it does not.

Of course we need more than one trial. Replicate cells $A3 - N3$ down to row 1003.

Monte Carlo Experiment

Finally, set up a count of the number of ones in column A: In cell $A2$, enter:

=SUM(A3:A1002)

Monte Carlo Experiment

Finally, set up a count of the number of ones in column A: In cell $A2$, enter:

=SUM(A3:A1002)

The result should be approximately $1000(1-\alpha)$

Monte Carlo Experiment

To recap, in this numerical experiment we:

- Generated 1,000 random samples of size 10 from a $N(\mu, \sigma)$ distribution
- Computed the sample mean for each of them
- Constructed a $100(1 - \alpha)\%$ confidence interval for μ from each sample
- Counted the number of confidence intervals that actually contained μ

Monte Carlo Experiment

To recap, in this numerical experiment we:

- Generated 1,000 random samples of size 10 from a $N(\mu, \sigma)$ distribution
- Computed the sample mean for each of them
- Constructed a $100(1-\alpha)\%$ confidence interval for μ from each sample
- Counted the number of confidence intervals that actually contained μ

You can experiment with the spreadsheet by replicating the experiment (F9), and changing values of α, μ, and σ.