Constructing the Confidence Interval

Recall that the mean of a random sample of size n from a $N(\mu, \sigma)$ population also has a normal distribution:

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

Constructing the Confidence Interval

Recall that the mean of a random sample of size n from a $N(\mu, \sigma)$ population also has a normal distribution:

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

This is the basic fact that allowed us to construct $100(1-\alpha)\%$ confidence intervals for μ, *as long as we knew the value of the population standard deviation* σ

Constructing the Confidence Interval

Recall that the mean of a random sample of size n from a $N(\mu, \sigma)$ population also has a normal distribution:

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

This is the basic fact that allowed us to construct $100(1-\alpha)\%$ confidence intervals for μ, *as long as we knew the value of the population standard deviation* σ

More often than not, in practice we cannot assume that we know σ.

Constructing the Confidence Interval

Since we need to come up with some estimate of the population standard deviation, one approach is to simply use the sample standard deviation s in place of the population standard deviation σ.

Constructing the Confidence Interval

Since we need to come up with some estimate of the population standard deviation, one approach is to simply use the sample standard deviation s in place of the population standard deviation σ.

Prior to about 1915, this is exactly what was done.

Constructing the Confidence Interval

Since we need to come up with some estimate of the population standard deviation, one approach is to simply use the sample standard deviation s in place of the population standard deviation σ.

Prior to about 1915, this is exactly what was done.

While this seems innocent enough, the situation is much more complicated because s is a random variable while σ is a constant.

Constructing the Confidence Interval

Around the time of the first world war, a chemist at the Guiness Brewery named William Gosset derived the complicated distribution of the statistic

$$T = \frac{\overline{x} - \mu}{s/\sqrt{n}}$$

Constructing the Confidence Interval

Around the time of the first world war, a chemist at the Guiness Brewery named William Gosset derived the complicated distribution of the statistic

$$T = \frac{\overline{x} - \mu}{s/\sqrt{n}}$$

Recall that a sum of the squares of standard normal random variables has a distribution called *Chi-square* with a parameter called the "degrees of freedom" that is the number of independent random variables squared and summed.

Constructing the Confidence Interval

Around the time of the first world war, a chemist at the Guiness Brewery named William Gosset derived the complicated distribution of the statistic

$$T = \frac{\overline{x} - \mu}{s/\sqrt{n}}$$

Recall that a sum of the squares of standard normal random variables has a distribution called *Chi-square* with a parameter called the "degrees of freedom" that is the number of independent random variables squared and summed.

As it turns out, s^2/n has a Chi-square distribution with $n-1$ degrees of freedom.

Constructing the Confidence Interval

This means that the random variable T is the ratio of a normal random variable to the square root of an (independent) Chi-square random variable:

$$T = \frac{\overline{x} - \mu}{s/\sqrt{n}}$$

Constructing the Confidence Interval

This means that the random variable T is the ratio of a normal random variable to the square root of an (independent) Chi-square random variable:

$$T = \frac{\overline{x} - \mu}{s/\sqrt{n}}$$

The brewery did not allow employees to publish articles, so Gosset published the result under the name "student"

Constructing the Confidence Interval

This means that the random variable T is the ratio of a normal random variable to the square root of an (independent) Chi-square random variable:

$$T = \frac{\overline{x} - \mu}{s/\sqrt{n}}$$

The brewery did not allow employees to publish articles, so Gosset published the result under the name "student"

The t distribution was known for many years as "student's t distribution".

Constructing the Confidence Interval

The t distribution appears in Table A5 in the text, although we will generally use a computer.

Constructing the Confidence Interval

The t distribution appears in Table A5 in the text, although we will generally use a computer.

Directions for obtaining thresholds for the t distribution for two-tailed and one-tailed regions are documented in the technology pages (linked from the web page).

Constructing the Confidence Interval

The t distribution appears in Table A5 in the text, although we will generally use a computer.

Directions for obtaining thresholds for the t distribution for two-tailed and one-tailed regions are documented in the technology pages (linked from the web page).

Note carefully that spreadsheets implement the t distribution functions $TDIST$ and $TINV$ differently from $NORMDIST$ and $NORMINV$.

Constructing the Confidence Interval

The t distribution values can also be obtained using R.

Unlike spreadsheets, R is consistent in its implementation across distributions.

Constructing the Confidence Interval

The t distribution values can also be obtained using R.

Unlike spreadsheets, R is consistent in its implementation across distributions.

Numerical values for a t distribution with n degrees of freedom:

$P(X \leq x)$	$pt(x, n)$
x such that $P(X \leq x) = q)$	$qt(x, n)$

Constructing the Confidence Interval

The two-sided $100(1-\alpha)\%$ confidence interval when both the mean and standard deviation are estimated from the sample is based on:

- the sample mean $\overline{x}$
- the sample standard deviation s
- then sample size n
- the value of α

Constructing the Confidence Interval

The two-sided $100(1-\alpha)\%$ confidence interval when both the mean and standard deviation are estimated from the sample is based on:

- the sample mean $\overline{x}$
- the sample standard deviation s
- then sample size n
- the value of α

The lower and upper limits of the two-sided confidence interval are:

Lower limit $\quad \overline{x} + qt(\alpha/2, n-1) \cdot s/\sqrt{n}$

Upper limit $\quad \overline{x} + qt(1-\alpha/2, n-1) \cdot s/\sqrt{n}$

Constructing the Confidence Interval

The upper one-sided $100(1-\alpha)\%$ confidence interval when both the mean and standard deviation are estimated from the sample is:

Lower limit	$-\infty$
Upper limit	$\overline{x} + qt(1-\alpha, n-1) \cdot s/\sqrt{n}$

Constructing the Confidence Interval

The upper one-sided $100(1-\alpha)\%$ confidence interval when both the mean and standard deviation are estimated from the sample is:

Lower limit	$-\infty$
Upper limit	$\overline{x} + qt(1-\alpha, n-1) \cdot s/\sqrt{n}$

The lower one-sided $100(1-\alpha)\%$ confidence interval when both the mean and standard deviation are estimated from the sample is:

Lower limit	$\overline{x} + qt(\alpha, n-1) \cdot s/\sqrt{n}$
Upper limit	∞

Example

A random sample of 100 students have their heights measured. If the average score is 69.3 inches with a sample standard deviation of 3.4, construct a two-sided 95% confidence interval for the mean height of students.

Example

A random sample of 100 students have their heights measured. If the average score is 69.3 inches with a sample standard deviation of 3.4, construct a two-sided 95% confidence interval for the mean height of students.

In this case, $\alpha = .05$, $n = 100$, $\overline{x} = 69.3$ and $s = 3.4$. Since we have no reason to assume we know the standard deviation of heights, we will use a t distribution to construct the confidence interval.

Example

A random sample of 100 students have their heights measured. If the average score is 69.3 inches with a sample standard deviation of 3.4, construct a two-sided 95% confidence interval for the mean height of students.

In this case, $\alpha = .05$, $n = 100$, $\overline{x} = 69.3$ and $s = 3.4$. Since we have no reason to assume we know the standard deviation of heights, we will use a t distribution to construct the confidence interval.

Lower limit	$\overline{x} + qt(\alpha/2, n-1) \cdot s/\sqrt{n}$
Lower limit	$69.3 + qt(.025, 99) * 3.4/10$=$68.625$
Upper limit	$\overline{x} + qt(1-\alpha/2, n-1) \cdot s/\sqrt{n}$
Upper limit	$69.3 + qt(0.975, 99) * 3.4/10$=$69.975$

Example

The speed of 250 cars in the right hand lane of an interstate highway is measured with radar. If the average speed is $60.2mph$ and the sample standard deviation is 5.3, construct a 99% confidence interval for the mean vehicle speed.

Example

The speed of 250 cars in the right hand lane of an interstate highway is measured with radar. If the average speed is $60.2mph$ and the sample standard deviation is 5.3, construct a 99% confidence interval for the mean vehicle speed.

In this case, $\alpha = .01$, $n = 250$, $\overline{x} = 60.2$ and $s = 5.3$. Since we have no reason to assume we know the true standard deviation of speeds, we will use a t distribution to construct the confidence interval.

Example

The speed of 250 cars in the right hand lane of an interstate highway is measured with radar. If the average speed is $60.2mph$ and the sample standard deviation is 5.3, construct a 99% confidence interval for the mean vehicle speed.

In this case, $\alpha = .01$, $n = 250$, $\overline{x} = 60.2$ and $s = 5.3$. Since we have no reason to assume we know the true standard deviation of speeds, we will use a t distribution to construct the confidence interval.

Lower limit	$\overline{x} + qt(\alpha/2, n-1) \cdot s/\sqrt{n}$
Lower limit	$60.2 + qt(.005, 249) * 5.3/\sqrt{250}$=$59.330$
Upper limit	$\overline{x} + qt(1-\alpha/2, n-1) \cdot s/\sqrt{n}$
Upper limit	$60.2 + qt(0.995, 99) * 5.3/\sqrt{250}$=$61.070$

Example - Technology Page

Viscosity of oil from 50 test vehicles after $10,000$ miles of driving averaged 43.1 with a standard deviation of 10.3. Find a two-sided 95% confidence interval for the mean viscosity of oil after $10,000$ miles.

Example - Technology Page

Viscosity of oil from 50 test vehicles after $10,000$ miles of driving averaged 43.1 with a standard deviation of 10.3. Find a two-sided 95% confidence interval for the mean viscosity of oil after $10,000$ miles.

Download the spreadsheet from the Technology Page in the section:

- Confidence Intervals for Population Means
- Sigma unknown

Example - Technology Page

A sample of 80 lobsters caught in Narragansett Bay averaged 12.4 ounces in weight with a standard deviation of 3.1. Find a two-sided 99% confidence interval for the mean weight of lobsters in Narragansett Bay.

Example - Technology Page

A sample of 80 lobsters caught in Narragansett Bay averaged 12.4 ounces in weight with a standard deviation of 3.1. Find a two-sided 99% confidence interval for the mean weight of lobsters in Narragansett Bay.

Download the spreadsheet from the Technology Page in the section:

- Confidence Intervals for Population Means
- Sigma unknown

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 as if we knew sigma
- use the sample standard deviation in place of the true population standard deviation
- 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 as if we knew sigma
- use the sample standard deviation in place of the true population standard deviation
- 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 $H3$:

=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 $H3$:

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

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

Monte Carlo Experiment

Now replicate cell H3 horizontally 9 times, filling cells I3-Q3.

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

Monte Carlo Experiment

Now replicate cell H3 horizontally 9 times, filling cells I3-Q3.

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

Now carefully enter the following in cell $G3$:

=AVERAGE(H3:Q3)

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

Monte Carlo Experiment

Now replicate cell H3 horizontally 9 times, filling cells I3-Q3.

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

Now carefully enter the following in cell $G3$:

=AVERAGE(H3:Q3)

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

Now carefully enter the following in cell $F3$:

=STDEV(H3:Q3)

Monte Carlo Experiment

Now carefully enter the following in cell $B3$:

=NORMINV($B$1/2,G3,F3/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,G3,F3/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,G3,F3/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 **lower** than the expected $1000(1-\alpha)$

This illustrates that we need to use the t distribution when the sample size is small and the standard deviation is estimated from the sample.