CONFIDENCE INTERVALS POPULATION MEANS

Unit 4A - Statistical Inference Part 1
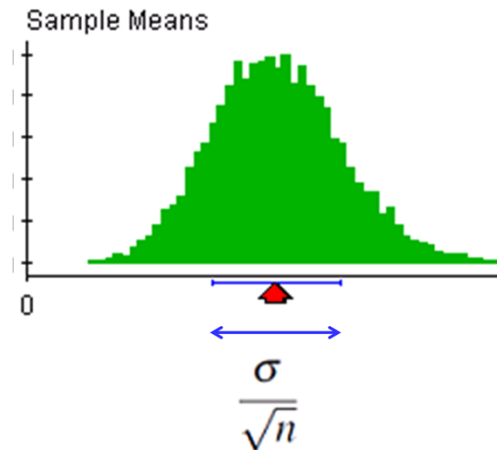
{Confidence Intervals for Population Means}

First we will develop the idea for the specific case of 95% confidence intervals, and then we will move on to the more general formula and some additional issues in practice.

How does this work?

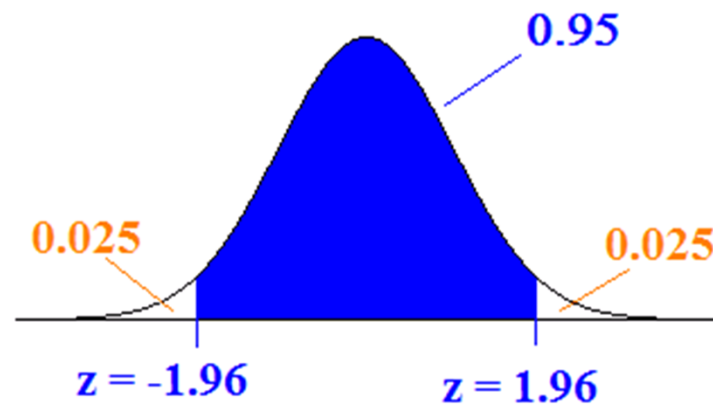Sample Means

$$\frac{\sigma}{\sqrt{n}}$$

Let's use the results of one of our earlier simulations to explain the idea.

Suppose we have a population, and we know the standard deviation of the population.

If we are working with a fixed sample size, then we can calculate the standard error of the sample mean using sigma divided by the square root of n.

How does this work?

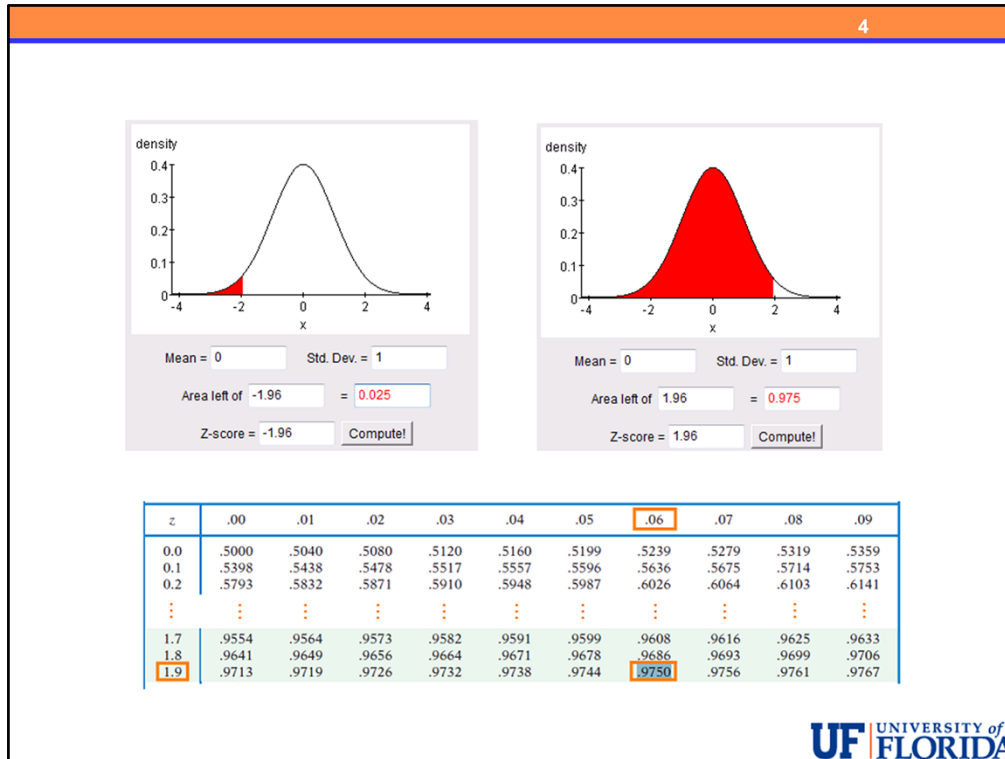0.95

0.025

0.025

z = -1.96

z = 1.96

Assuming we are willing to use a normal approximation for the sampling distribution, we can determine how far we need to go on either side of the population mean to contain 95% of the possible values of x-bar.

We can ask: how many standard deviations away from the center of the sampling distribution do we need to go to contain 95% of the values of x-bar. The answer will be a z-score.

Although we don't know the true center of the sampling distribution, all we really need is the standard error of our statistic *(the standard deviation of the sampling distribution)* and the ability to use the normal distribution to find the appropriate z-score and we can find our confidence interval.

We could approximate this value using plus or minus 2 standard errors, as the standard deviation rule would give, however, now that we have discussed normal distributions in more depth, we can find the value more accurately.

To find the z-scores which contain the middle 95% of the standard normal distribution we need to note that there is 2.5% left over on each side.

density
0.4
0.3
0.2
0.1
0
-4 -2 0 2 4
x

Mean = 0   Std. Dev. = 1
Area left of -1.96 = 0.025
Z-score = -1.96   Compute!

density
0.4
0.3
0.2
0.1
0
-4 -2 0 2 4
x

Mean = 0   Std. Dev. = 1
Area left of 1.96 = 0.975
Z-score = 1.96   Compute!

z	.00	.01	.02	.03	.04	.05	.06	.07	.08	.09
0.0	.5000	.5040	.5080	.5120	.5160	.5199	.5239	.5279	.5319	.5359
0.1	.5398	.5438	.5478	.5517	.5557	.5596	.5636	.5675	.5714	.5753
0.2	.5793	.5832	.5871	.5910	.5948	.5987	.6026	.6064	.6103	.6141
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
1.7	.9554	.9564	.9573	.9582	.9591	.9599	.9608	.9616	.9625	.9633
1.8	.9641	.9649	.9656	.9664	.9671	.9678	.9686	.9693	.9699	.9706
1.9	.9713	.9719	.9726	.9732	.9738	.9744	.9750	.9756	.9761	.9767

UF UNIVERSITY of FLORIDA

Then we can use the online calculator or a normal table to find the appropriate values.

We really only need to find one of these values as we know these z-scores will be the same value with opposite signs when using the standard normal distribution.

For the tables and calculators we are using here, we need to find the area to the left of each of the unknown z-scores. The lower z-score is easiest, with 2.5% of the area below that value so we look up a probability of 0.025.

For the upper z-score, 2.5% of the area is TO THE RIGHT of the unknown value so the area to the left is 97.5% or a probability of 0.975.

When we type in 0.025 we find the z-score is negative 1.96 and when we enter 0.975, we get z = positive 1.96.

Using the table we need to find the closest area, in this case we have illustrated 0.975, and follow it both up and left – combining the values 1.9 with 0.06 we get 1.96 as well.

95% Confidence Interval for μ

$$\bar{x} \pm 1.96 * \frac{\sigma}{\sqrt{n}}$$

UF UNIVERSITY of FLORIDA

This leads to the following formula for the 95% confidence interval for a population mean whenever we feel the sampling distribution of x-bar is reasonably normal.

Currently, we are also assuming that we know the population standard deviation. This allows us to easily manage the probabilities required to understand the development of these results.

Later, we will discuss the more practical scenario where the population standard deviation is not known but instead must be estimated from the data.

Here is an example. A researcher is interested in estimating the mean score on the math portion of the SAT of all community college students in his state. The researcher collects data on a random sample of 650 community college students from his state and finds the average SAT-math score to be 475. It is known that scores roughly follow a normal distribution with a standard deviation of sigma = 100.

We can perform these calculations to find the 95% confidence interval for the population mean SAT math score of all community college students in the researcher's state.

Be sure to check these calculations yourself. We take 475 plus and minus our z-score of 1.96 times the standard error of x-bar, which here is 100 divided by the square root of 650. Following the correct order of operations, we start with multiplication and division by calculating 1.96 times 100 divided by the square root of 650 which is 7.7. Then we add and subtract 7.7 from the sample mean of 475.

The resulting confidence interval ranges from 467.3 to 482.7.

How can we interpret this result?

Interpretation

- We are 95% confident that the **mean** SAT-M score of **all** community college students in the researcher's state is covered by the interval (467.3,482.7)

- It is **NOT** correct to say "there is a 95% chance that μ is in the interval (467.3, 482.7)"!! **Why?**

UF UNIVERSITY of FLORIDA

We can interpret this interval by saying: **We are 95% confident that the <u>mean</u> SAT-M score of <u>all</u> community college students in the researcher's state is covered by the interval (467.3,482.7) .**

It is important to state the confidence level (95%).
                    AND
The parameter you are estimating in the words of the problem, here it is the mean SAT math score of all community college students in the researcher's state).

Another correct way of interpreting these results is:

Based upon our data, we estimate the mean SAT math score of all community college students in the researcher's state to be 475. The 95% confidence interval suggests this value could be as low as 467.3 to as high as 482.7.
Note that the confidence interval was obtained by taking 475 ± 7.7.

This means that we are 95% confident that by using the sample mean to estimate μ, our error is no more than 7.7 points.

It is **NOT** correct to say "there is a 95% chance that mu, the population mean, is in the interval (467.3, 482.7)"!! Why?
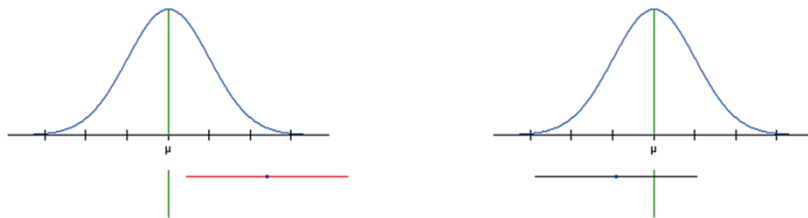
Once you are working with a particular sample, the interval either DOES or DOES NOT contain the parameter you just won't know which it is!!!

The probability is ONLY realized when we consider repeated sampling of x-bar and thus does not pertain to OUR PARTICULAR x-bar at the moment – the probability pertains to our process!

95% of the intervals we obtain in repeated samples will contain the true value, 5% will not – for a given sample, we don't know whether we have one of the 95% which work or one of the 5% which do not.

In statistics, we use the word "confidence" or "confident" to reflect this long run chance and we do NOT use the words "probability" or "chance" when interpreting the confidence interval from our data.

The illustrations show one of each. A sampling distribution is drawn centered at a particular population mean, mu. The vertical line represents the value of the population mean.

Explanations
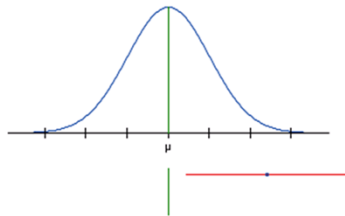
Suppose the population mean is really 460, then
P(467.3 < 460 < 482.7) = 0

Suppose the population mean is really 478, then
P(467.3 < 478 < 482.7) = 1

UF UNIVERSITY of FLORIDA
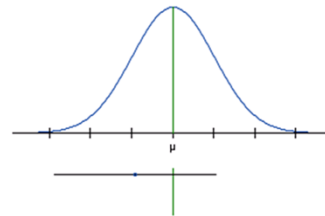
95% of our x-bars will fall with 1.96 standard errors of the true population mean. When this happens we will get an interval such as the one on the right where the interval, drawn as a black line, overlaps the target vertical line at mu.

However, 5% of the time, we will get an x-bar which falls in the extreme tail of the sampling distribution, and in these cases, we will obtain an interval which DOES NOT capture the true population mean – the interval does not overlap the target vertical line.

Let's return to our SAT math example where our confidence interval was (467.3,482.7).

Here is the question:

What is the probability that the true population mean is between 467.3 and 482.7?

This is a trick question!! Why? Because we don't KNOW the population mean. In reality we are working with three fixed numbers and we can only specify the correct probability once ALL of the values are known.

Suppose the true population mean is 460 then we are asking P(467.3 < 460 < 482.7) = 0

Suppose the population mean is really 478, then P(467.3 < 478 < 482.7) = 1
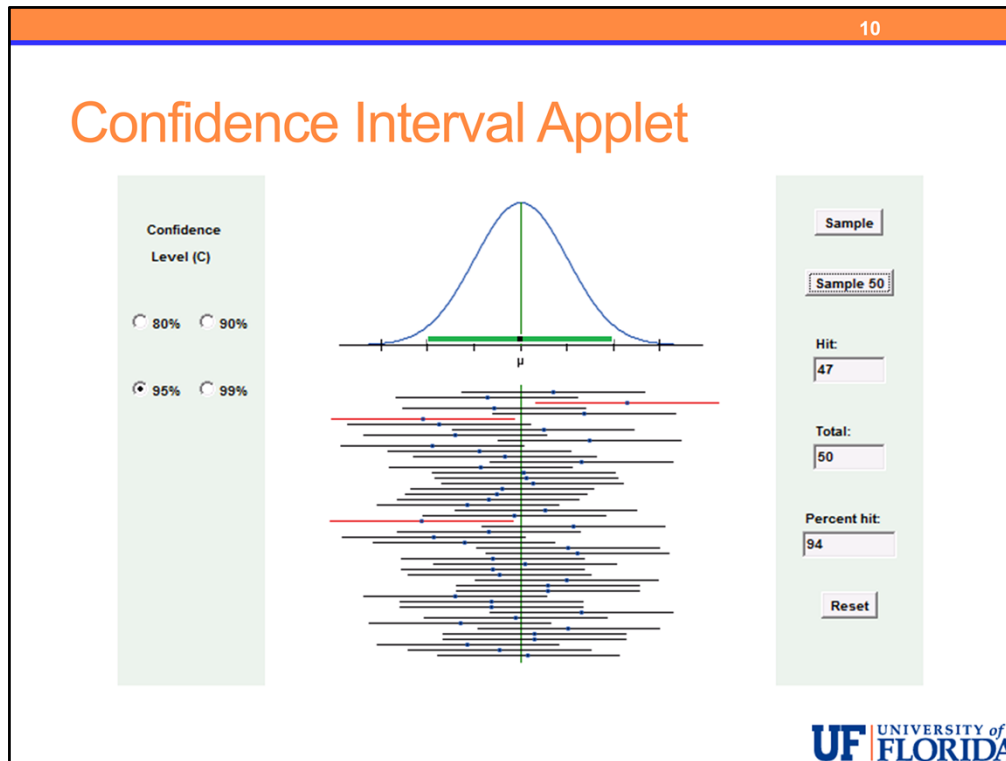
http://content.bfwpub.com/webroot_pubcontent/Content/BCS_4/IPS7e/Student/Statistical%20Applets/confidenceinterval.html

What we can say, is that if we repeat this PROCESS, 95% of the time, our interval will contain the true value of the parameter and 5% of the time it will not! In practice, we simply do not know which case we have and the phrase "confidence" implies this process of repeated sampling.

The words probability or chance do not and should not be used when interpreting the results of confidence intervals in practice.

The probability is associated with the randomness of our statistic so that for a particular interval we only speak of being "95% confident" which translates into an understanding about the process not a statement about our particular sample.

This applet also illustrates nicely how the sampling distribution gives us the needed width of our interval. Even though we don't know the true center of the distribution, we can find or at least estimate, the width needed to capture 95% of sample means.

Using this width and centering it at our estimate, the sample mean from our data, we find our confidence interval.
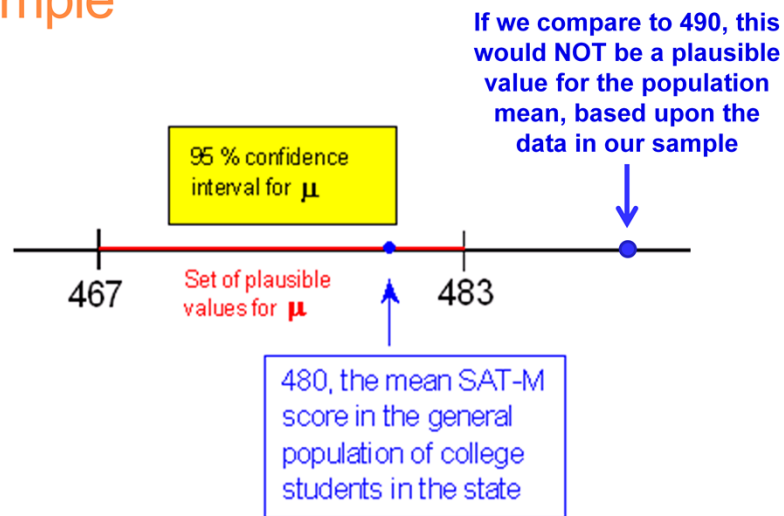
Here is a way that we can use the confidence interval:

Do the results of this study provide evidence that μ, the mean SAT-M score of community college students, is lower than the mean SAT-M score in the general population of college students in that state (which is 480)?

The 95% confidence interval for μ was found to be (467.3,482.7)

Note that 480, the mean SAT-M score in the general population of college students in that state, falls inside the interval, so it is one of the plausible values for μ

Thus we cannot claim that the mean SAT math score of community college students is lower than 480. In fact, it could be higher than 480 based upon our confidence interval.

If on the other hand, the mean SAT-math score of the general population had been 490, then we would be able to claim that the mean SAT math score of community college students is lower than 490 since all of the plausible values of mu are less than that value.

In this way, we will see that confidence intervals can be used to conduct hypothesis tests.

CONFIDENCE INTERVALS POPULATION MEANS

We have seen how to calculate a 95% confidence interval for a population mean when the sampling distribution of x-bar is normally distributed.

And we have discussed correct and incorrect interpretations for this confidence interval.

Next we will discuss the general form of the confidence interval which can be used for any confidence level.