Lecture 12

Formal Methods of Inference

The main idea:

There are 3 formal methods of inference:

Point Estimation →

Interval Estimation →

Hypothesis Testing →

Point Estimation

Intuitively:

Example – Taylor Swift
Researchers at the Universal Music Group are interested in gauging the U.S. market's reaction to Taylor Swift's latest 2017 album Reputation. She focused on a new audience and the company needs to know if she should keep with is. To do so, the researchers took a cluster sample of 500 males and got each person's average rating on a scale of 1-10 (1=worst, 10=best). The average rating was a 1.3. What does this say about what the U.S. thinks of her new album?

Is this an accurate estimate of the true U.S. opinion of the album?

Consider another cluster sample of 500 females aged 12-18. The average rating from this sample was 8.9.

What would you conclude about the true average U.S. rating?

Which one is correct?

Point estimation is simple, and deceptively intuitive. However, it is not robust. The key issue with point estimation is that it does not take into account or provide any information about the ***expected estimation error***.

Interval Estimation

Intuitively:

Example – Taylor Swift (cont.)
Now let's say that the researchers got a bit more cleaver about their sampling technique. This time they now take a simple random sample of size 500 and get that the mean rating was 7.6. The researchers take another simple random sample of size 500 and get a mean rating of 8.1.

What can we say?

The advantages of interval estimation over point estimation are clear. What is perhaps not so clear is how to actually construct a confidence interval and how we can just assign a probability to our interval estimate.

Confidence Intervals for the Population Mean μ

Let's consider the sampling distribution of sample means $(\bar{x})$. If the samples of size 100 are being taken from a population known to be normal with mean 40 and standard deviation 2, then we can say that:

Visually, the Sampling Distribution of Sample Means in this case is as follows:

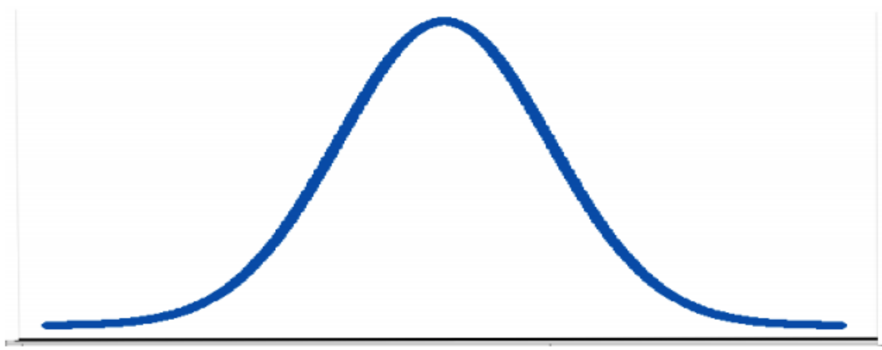

We can apply the Empirical Rule since

So, we can state that there is a 95% chance that the observed sample mean $\bar{x}$ falls within 2 standard deviations of the population mean μ.

So, what does this tell us?

Formulaically:

So, a 95% confidence interval (CI) for μ (when σ is known) is:

Similarly, from the Empirical Rule, we have that a 99.7% CI for μ is:

Some other intervals for key confidence levels include:

So, where is this coefficient coming from?

Generally, we have

How do we find z? This is the z-score from the Normal Table that carries a probability of $\frac{(1-CC)}{2} + CC$ where CC is the ***confidence coefficient***. This is the confidence level for the interval, in decimal form.

Example – IQ
The IQ level of students at a certain college has an unknown mean μ and standard deviation $\sigma = 15$. A simple random sample of 100 students was chosen and their average IQ score was $\bar{x} = 112$. Estimate μ with a 90%, 95%, and 99% C.I.

95% C.I.:

Interpretation:

Method 1: We are 95% confident that μ, the mean IQ level of *all* the college students, is covered by the interval (109, 115).

Method 2: We are 95% confident that when we estimate μ by $\bar{x} = 112$, we're "off" by no more than 3.

90% C.I.:

99% C.I.:

Comments Regarding Confidence Intervals:

i. As the level of confidence required increases, the confidence interval gets w i d e r.

ii. We can only use the formulas for the C.I. for μ if
   - The sampling is random
   - The sampling distribution of $\bar{x}$ is normal (i.e. sample size n ≥ 30 (invoke the CLT) OR the population from which we sample is normal)

Margin of Error (m)

Formulaically:

→

What Does the Ideal C.I. Look Like?

Ideally:

So, this implies two key qualities that we look for in a confidence interval:

i.

ii.

How do we achieve this?

Let's look at the margin of error for the ideal C.I.:

The ideal is never practical, but it helps to give us an idea of what we're looking for. The way to obtain the 2 key qualities of a confidence interval is

How can we shrink the margin of error?

1. Compromise on the level of confidence (use a 90% confidence level instead of 99%, for example)

   Example – IQ (cont.)
   Compare the lengths of the 90% C.I. and 95% C.I. What do we notice?

   The implication → there is a trade-off between the level of confidence and how precise (narrow) the confidence interval is.

2. Take a larger sample size

   When you increase n, $z\frac{\sigma}{\sqrt{n}}$ gets smaller.

Example – Apples
We know the circumference of apples is normally distributed with unknown mean and standard deviation of .02 cm.

a. To provide inference, Austin took a simple random sample of 100 and got a mean of 1 cm. Calculate the 95% confidence interval for the true mean circumference of apples.

b. Austin then took another simple random sample of 200 and got a mean again of 1 cm. Calculate the 95% confidence interval for the true mean circumference of apples.

Note: In practice, a larger n (sample size) is not always available.

Sample Size Calculations

We can find the sample size needed in order to estimate μ with a certain level of confidence and a desired margin of error.

Recall the formula for the margin of error:

We can isolate for n in order to find the sample size needed:

Note: Always round UP for sample size calculations, no matter what your decimal is. This is sample size, so you need the whole experimental unit (doesn't make sense to take .3 of a person…)

Example – IQ (cont.)
What should my sample size be if I want to estimate μ with a 95% C.I. that has a margin of error m=2 (i.e, length =4)? (remember, $\sigma = 15$)

What if we don't know σ?

This happens most frequently. For this case, as long as the sample is "large", i.e. greater than 30, then we can use the sample standard deviation s instead of σ. All other calculations are the same.

Confidence Intervals for the Population Proportion p

The same logic applies here as did for the sampling distribution of sample means.

Let's consider the sampling distribution of sample proportions $(\hat{p})$. If the samples of size 100 are being taken from a population known to be Normal with p = .31, then we can say that

Visually, the Sampling Distribution of Sample Proportions in this case is as follows:

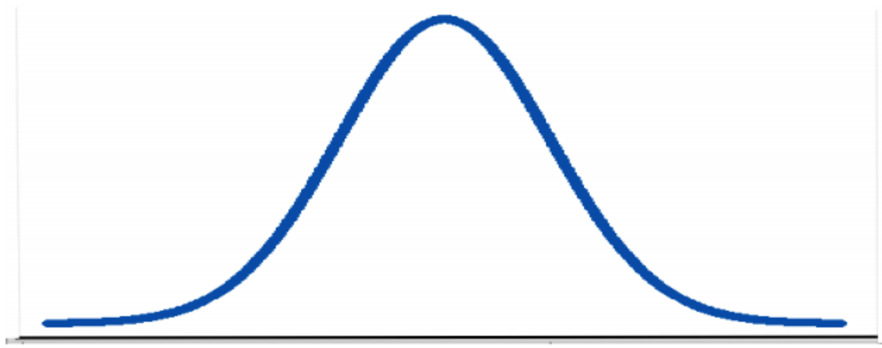

A 95% confidence interval for the parameter p is:

In general, the confidence interval for p is:

As always, the level of confidence dictates the z-score used:

This formula is safe to use when

i. The sample is random

ii. Population is large relative to the sample

iii. Normal sampling distribution of $\hat{p}$: $n\hat{p} \geq 10, n(1-\hat{p}) \geq 10$

Example – Polls

We would like to estimate p, the proportion of US adults who support legalizing the use of marijuana.

CNN/ORC Poll. Jan. 3-5, 2014. N=1,010 adults nationwide. Margin of error ± 3.

"Do you think the use of marijuana should be made legal, or not?"

	Should be made legal %	Should not be made legal %	Unsure %
ALL	55	44	1
Democrats	62	37	1
Independents	59	41	1
Republicans	36	61	2

Calculate a 90% and 95% confidence interval for the true proportion p of all US adults who support legalizing the use of marijuana.

90% C.I.:

95% C.I.:

Interpretation:

Method 1: We are 95% confident that the proportion of all U.S. adults who believe that the use of marijuana should be legal is covered by the interval (.52, .58). ***(HENCE, THE MAJORITY)***

Method 2: We are 95% confident that when we estimate p (the proportion of all U.S. adults who believe that the use of marijuana should be legal) by $\hat{p} = .55$, we are "off" by no more than .03.

Let's consider some more recent data:

Pew Research Center. Oct. 15-20, 2014. N=2,003 adults nationwide. Margin of error ± 2.5.

"Do you think the use of marijuana should be made legal, or not?"

	Yes, legal %	No, illegal %	Unsure/ Refused %
10/15-20/14	52	45	3
2/12-26/14	54	42	3
3/13-17/13	52	45	3
2/22 - 3/1/11	45	50	5
3/10-14/10	41	52	7

Calculate a 95% confidence interval for the true proportion p of all US adults who support legalizing the use of marijuana.

Interpretation: We are 95% confident that the proportion of all U.S. adults who believe that the use of marijuana should be legal is covered by the interval (.495, .545). ***(THIS TIME WE CANNOT BE CERTAIN AT THE CLAIMED CONFIDENCE LEVEL THAT THE MAJORITY SUPPORT THIS LEGISLATURE).***

Sample Size Calculations

Recall the formula for the confidence interval for the parameter p:

Similar to when we were estimating the mean, we can look at the margin of error and isolate for n:

It's very common, especially in polling, to determine the sample size based on a desire margin of error (m). However, the margin of error depends on the sample result $\hat{p}$. We would like to be able to control the margin of error with the sample size ***BEFORE*** the poll is conducted.

Some mathematical intuition will allow us to see that for the 95% confidence interval:

In practice, polls are often designed so that they will have a margin of error of 3% at the 95% confidence level. So, the "target" sample size is:

Example – Legalization of Marijuana (cont.)

CNN/ORC Poll. Jan. 3-5, 2014. N=1,010 adults nationwide. Margin of error ± 3.

"Do you think the use of marijuana should be made legal, or not?"

	Should be made legal %	Should not be made legal %	Unsure %
ALL	55	44	1
Democrats	62	37	1
Independents	59	41	1
Republicans	36	61	2

Did they hit their goal?

Example – Presidential Election
During the last days leading into the previous presidential elections, Gallup wanted to conduct a poll of likely voters that would have a margin of error of only 1.5% at the 95% confidence level. What is the "target" sample size?

Note how we had to have 95% confidence!

Using Confidence Intervals

Since a confidence interval is the set of plausible values for the parameter of interest, we can test any candidate parameter value by looking to see whether it is among the set of plausible values:

1. If a supposed parameter value lies inside the confidence interval, then it is a plausible value for the true parameter value.

2. If a supposed parameter value lies outside the confidence interval, then it is not likely to be the true parameter value.

Example – Legalization of Marijuana (cont.)

a. Based on the data, is it plausible that the proportion of all U.S. adults who are in favor of legalizing marijuana is as low as 50%?

   95% C.I.:

b. Based on the data, is it plausible that the proportion of all U.S. adults who are in favor of legalizing marijuana is as high as 57%?

   95% C.I.: