Interval Estimation

Javier E. Flores

February 13, 2019

Introduction

Today we're going to be talking about confidence (the statistical kind).

So in the spirit of this topic, here is a picture of a statistician about 95% confident he would die.

Introduction

On a more serious note, in order to properly frame our discussion of statistical confidence we need to first introduce the idea of *interval estimation*.

You might recall from our last discussion on sampling distributions that μ is the most likely value of $\bar{x}$ for a single sample.

This considered, if our goal was to *estimate* μ, our best estimate would be $\bar{x}$.

However, we also learned that the reliability of this estimate is dependent on the standard error (variability) of the sampling distribution.

Due to the variability of our sampling distribution, chances are that this **point estimate** $\bar{x}$ will never be exactly μ.

Introduction

As a solution to this problem, rather than rely on a single point estimate, statisticians often use **interval estimates** which provide a range of plausible values for the parameter of interest.

We'll explore the utility of this idea with a game of trivia:

- A series of 16 trivia questions covering a range of topics will be asked.
- Your group must come to a consensus on an answer and (in addition) provide an interval of values you believe would capture the truth 80% of the time.
- To prevent cheating, only 15 randomly chosen questions will count.
- The group whose intervals capture the truth 12/15 (80%) of the time will win a bag of candy.

Trivia

Q1) How many licks does it really take to get to the center of a Tootsie Pop?

Q2) How many times does Paul McCartney sing the phrase ”let it be” in The Beatles’ song *Let it Be*?

Q3) On average, how many gallons of beer are consumed on Super Bowl Sunday?

Q4) What percent of U.S. 1$ bills are contaminated with disease-causing bacteria?

Trivia

Q5) In dollars, what is the highest selling price for a painting done by elephants?

Q6) In pounds, what is the weight of an average dairy cow?

Q7) What percent of American adults have a tattoo?

Q8) How many years separated the creation of the hamburger and the creation of the hamburger bun?

Trivia

Q9) How many slices of pizza are eaten in the U.S. every second?

Q10) The Hercules beetle can carry an object that is how many times its own weight?

Q11) In the 2001 UK census, how many people registered "Jedi" as their religion?

Q12) On average, how many people attended the SPAMARAMA festival each year to celebrate SPAM canned meat?

Trivia

Q13) In years, how old was the oldest woman to parachute from a plane?

Q14) How many words does the average American speak per day?

Q15) In feet, what is the length of the longest recorded human beard?

Q16) What percent of Americans believe that aliens have visited Earth?

Trivia Answers

A1) 252 (Tootsie Pop Official Website)

A2) 41 (*Let it Be*)

A3) 325 million (Wikipedia, 2013)

A4) 94 (Milwaukee Journal Sentinel)

A5) 39000 (*Guinness World Records, 2013*)

A6) 1400 (Agricultural Council of America)

A7) 21 (Harris Poll, 2013)

A8) 31 (What's Cooking America Website)

Trivia Answers

A9) 350 (Agricultural Council of America, 2011)

A10) 850 (Institute of Physics)

A11) 390000 (UK Office for National Statistics)

A12) 10000 (SPAMARAMA Website)

A13) 100 (*Guinness World Records, 2013*)

A14) 15942 (ABC News)

A15) 17.5 (Smithsonian Institution, 2013)

A16) 48 (Roper Poll)

Interval Estimation

If not obvious from how few of you correctly answered these questions, point estimates are almost always wrong.

Fun Fact: These questions all came from the game *Wits and Wagers*. In all the times I've played, I've only known one person to answer a question exactly right.

A much better option is to use interval estimates, which (should) contain your point estimate along with some **margin of error**.

Even with the option to construct intervals, this trivia activity also served to demonstrate how difficult it is to come up with an interval providing a specified degree of **coverage** (i.e. 80%).

Worry not, because we'll learn ways to create informative intervals of a specified coverage before the end of this lecture.

Interval Estimation

In general, an interval estimate of a population parameter has the form:

$$\text{Point Estimate} \;\pm\; \text{Margin of Error}$$

In order to create meaningful intervals, the margin of error should specify some degree of precision (e.g. 80% of the time these intervals will contain the true population parameter).

Question: Why would we want to specify a coverage less than 100%?

Confidence Interval

In statistics, an interval estimate's success rate in capturing the true parameter value is known as the **confidence level**.

Confidence intervals are simply interval estimates constructed with a specified confidence level.

Be mindful of the following clarifications regarding confidence intervals.

- Any one specific interval will either contain the parameter or not (i.e. 100% coverage or 0% coverage). Either way, the confidence level for the interval remains what it was when initially specified (e.g. 80%).
- Since the confidence interval is a statistic, it will vary from sample to sample despite being constructed the same way each time.
- This considered, the correct way to understand the confidence level is as the proportion of similarly constructed intervals from each (hypothetical) sample that contain the true parameter value.

Constructing Confidence Intervals

Suppose the following distribution corresponded to the sampling distribution of the sample mean (but we could use any statistic):

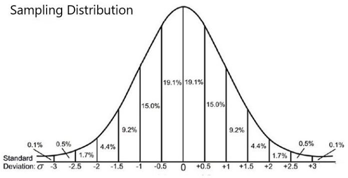

From our previous discussions we know that for a single sample, our most likely point estimate is the center of the sampling distribution.

Also recall that the mean of our sampling distribution (0 in this case) is unbiased for the population mean.

Constructing Confidence Intervals

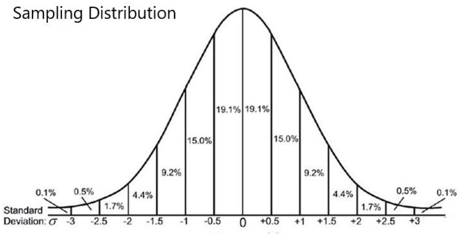

Shown on this distributions are various percentages which describe the probabilities of obtaining estimates in the specified ranges.

Considering this and the two previously mentioned points, discuss with your groups what you feel would be a reasonable thing to do to construct a confidence interval with, say, 95% coverage.

Solution

We know that the most probable values of (any) sample mean is at or near the center of the sampling distribution.

Therefore, any sample mean (in this case) is likely close to the true population mean.

The most reasonable way to construct a 95% confidence interval would then be:

$$\text{Sample Statistic} \;\pm\; 2{*}\text{SE}$$

Why 2∗SE? Because 95% of the distribution values are within 2 standard errors of the center. Let's see exactly how well a confidence interval constructed this way performs.

Confidence Interval Coverage

To assess the performance of this interval, we'll use a sampling distribution constructed from the class survey data. Specifically, our interest is in estimating the average time spent on social media (μ). Our sampling distribution will be made by taking several samples of 20 students and computing this mean for each sample:

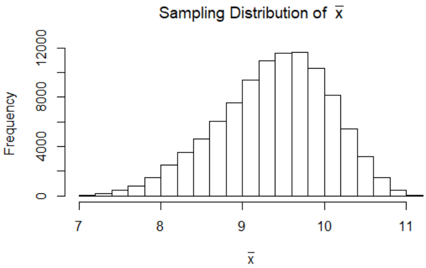

Confidence Interval Coverage

The standard error of $\bar{x}$ is 0.677. The first random sample had a sample mean of $\bar{x} = 9.3$. Therefore, the 95% confidence interval for μ is:

$$9.3 \ \pm \ 2 * 0.677 = (7.946, 10.654)$$

The true average time spent on social media in our class (pretending for a moment that the two people who did not take the survey are not in our class) is 9.4, so this interval captures the true μ.

Confidence Interval Coverage

To determine the coverage, we repeat this procedure for each random sample:

ID	Sample Mean	Calculation	95% CI
1	9.3	$9.3 \pm 2*0.677$	(7.946, 10.654)
2	8.25	$8.25 \pm 2*0.677$	(6.896, 9.604)
⋮	⋮	⋮	⋮
100k	9.65	$9.65 \pm 2*0.677$	(8.296, 11.004)

Confidence Interval Coverage

For the first 100 samples, 95 out of 100 samples contain the true population mean.

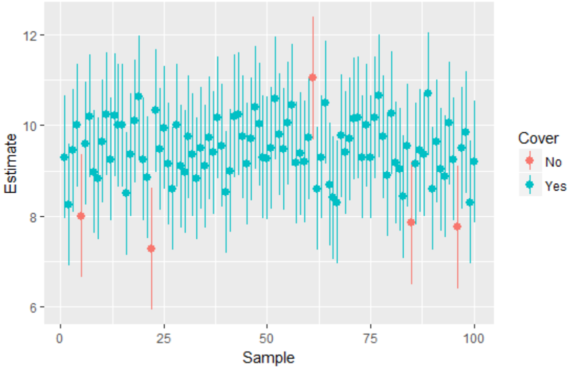

The coverage over all samples was about 96% (still very close to 95%)!

Interpreting Confidence Intervals

Practically speaking, we will usually only have a single sample and single confidence interval.

As I mentioned previously, this interval will either contain the parameter or not. In other words, that specific interval will have either a 100% chance or 0% chance of containing the population parameter.

For this reason, we **avoid** saying things like: "There is a 95% chance that μ is contained in this interval."

The appropriate statement would be to use the word "confident": "We are 95% confident that the interval contains the population parameter μ."

(Practical) Confidence Interval Construction

The past few slides have demonstrated how to construct a confidence interval using the sampling distribution. We've also demonstrated how well this interval performs.

However, the sampling distribution is only obtained through repeated sampling. How is it that we would have access to it with a single sample? (Answer: we don't)

However, there are ways that different (very clever) statisticians came up with that allow us to reconstruct the sampling distribution for various situations despite this limitation.

The first of these methods that we will explore (in the next lab) is known as the **bootstrap**.

Wrap-Up

Right now, you should...

- Know the difference between point and interval estimates, and understand the advantages of using interval estimates.
- Know how to construct a confidence interval when given a sample statistic and standard error.
- Correctly interpret confidence intervals.
- Explain the relationship between interval coverage and confidence level.

These notes cover section 3.2 of the textbook. Please read through the section and its examples along with any links provided in this lecture.