IMS 12 (ish): Bootstrap Confidence Intervals

Recall: Data Distributions vs Sampling Distributions

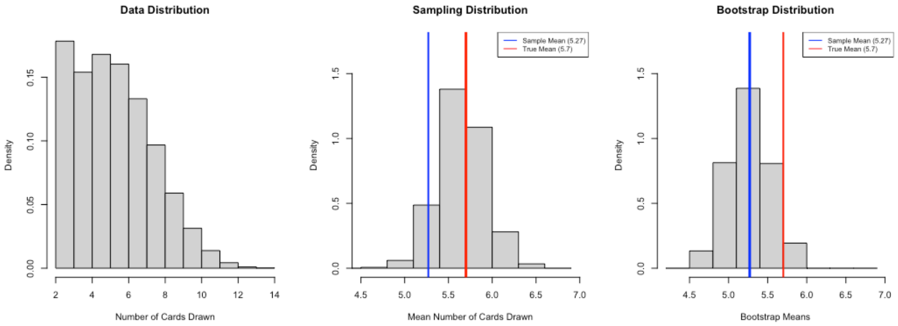

Recall: The Bootstrap

A **bootstrap sample** mimics repeatedly drawing samples from the population, to approximate characteristics of the true sampling distribution for a given statistic ($\hat{p}$ or $\overline{x}$) Recall, to obtain the bootstrap distribution, we repeatedly:

- Generate a bootstrap sample of size n from the original observations (i.e., we draw, with replacement, observations from the original sample).
- Compute the statistic of interest from the bootstrap sample.
- Repeat previous steps B times (B is typically in the thousands).

Warm Up Activity

Pin the tail on the Donkey

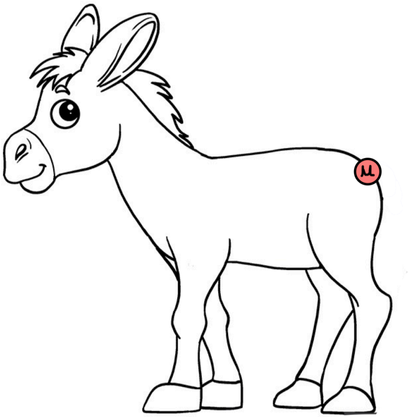

Estimation

Definition

- A **point estimate** is a single number that is our *best guess* for the parameter.
- An **interval estimate** is an interval of numbers that is believed to contain the actual value of the parameter.

- The sample mean $\bar{x}$ is a point estimate for population mean μ.
- The sample proportion $\hat{p}$ is a point estimate for the population proportion p.
- An interval estimate is often more useful, it helps us to gauge the accuracy of the point estimate.

Estimation

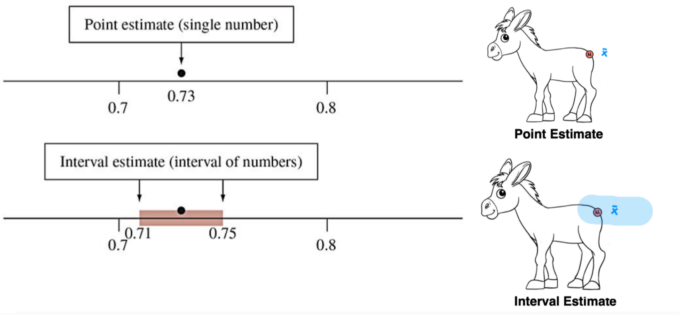
Point estimate (single number)
0.7
0.73
0.8
Interval estimate (interval of numbers)
0.7
0.71
0.75
0.8
Point Estimate
Interval Estimate

Interval Estimate

- An **interval estimate** indicates precision by giving an interval of numbers around the point estimate.
- The interval is made up of numbers that are the most believable values for the unknown parameter, based on the data observed.
- An **interval estimate** is designed to contain the parameter with a certain degree of confidence, say 95, they are referred to as **confidence intervals**.

Confidence Intervals

Confidence Interval

- A **confidence interval** is an interval containing the most plausible values for a parameter.
- The largest and smallest values in the intervals are called the **upper bound** and **lower bound**
- The chance that the interval produced *captures the parameter* is called the **confidence level**.

Common confidence levels are 90%, 95%, 99%

Confidence Intervals

What method do we use to construct a confidence interval?

- The key is the *sampling distribution* of the point estimate.
- The sampling distribution tells us the *probability* that the point estimate will fall within any certain distance of the parameter.
- For example: *If I have ($n =$) 60 simulations the probability I will observe a mean number flips of between 5.2 and 6.2 is about 95%*

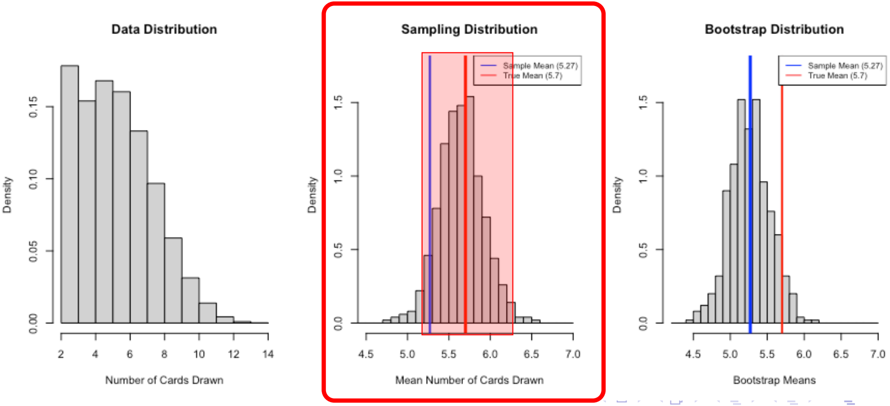

Confidence Intervals

However, in practice most of the time we do NOT have access to the sampling distribution.

- Instead we can use the *bootstrap distribution*.
- This does NOT give us a *probability*, because in these situations we do NOT know the parameter (true population mean/proportion).
- Instead we say *confident*, to reflect that we *feel good about our estimate*.

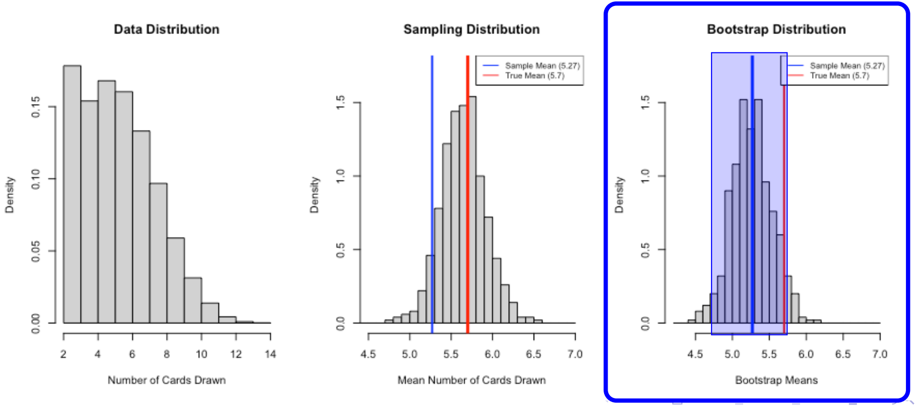

Confidence Intervals with Bootstrap Distributions

- To obtain a 95% confidence interval using a bootstrap distribution we want to find a lower bound and upper bound of the the middle 95% of the bootstrap estimates.
- This can be done using *percentiles*.

Definition

A *k*-**th percentile**, is a value at which a given percentage *k* of values in its frequency distribution falls.

For example:

- The 2.5^{th} percentile of the bootstrap means for number of cards is 4.77
- The 97.5^{th} percentile of the bootstrap means for number of cards is 5.77

Therefore: The middle 95% of the bootstrap distribution is between 4.77 and 5.77.

Interpretation

What is a bootstrap confidence interval? Is it a probability? Is it an estimate?

- It is not a probability because we are not using the sampling distribution.
- We are using an estimate of the sample distribution (the bootstrap distribution), which should be *very similar* to the sampling distribution in the long run.
- We say *confident* to reflect that we are estimating the parameter *and* the distribution.

Interpretation

Why say 95% confident?

- This 95% confidence refers to a long-run probability.
- The **long-run-probability** is the probability whether an interval so constructed captures the parameter over many random samples (of the given size) taken from the population.
- For example: take 100 samples, estimate bootstrap distributions for each sample, then obtain 95% confidence interval for each bootstrap distribution, then 95 are expected to capture the (true population) parameter and 5 are not.

There is a subtle but important distinction between **long-run-probability** and **probability**. The former refers to what happens if this whole process is repeated over and over. The later refers to the behavior of the process right now (which we do not know because we do not know the true parameter value).

Interpretation

Different ways to interpret the confidence interval.

Interpreting a Confidence Interval

1. The statistic is $\bar{x}$ or $\hat{p}$ with a 95% confidence interval of [lower bound, upper bound].
2. We are 95% confident that the confidence interval [lower bound, upper bound] captures the parameter.

Words/values in orange should be adjusted based on the setting.

Interpretation Examples

Interpretation Option 1

- The estimated proportion of those who could identify a song based on tapping is $\hat{p} = .025$ with a 95% confidence interval of [0.008, 0.041].
- The estimated mean height of Smith students is $\bar{x} = 65.2$ inches with a 95% confidence interval of [64, 66.4].

Interpretation Option 2

- We are 95% confident that the confidence interval [0.008, 0.041] captures the true proportion (p) of those who could identify a song based on tapping.
- We are 95% confident that the confidence interval [64, 66.4] captures the true mean (μ) height of Smith students.

Interpretation Cautions

Confidence intervals:

- **are only about the population parameter**. A confidence interval says nothing about individual observations.
- **are not a 95% probability**. Confidence interval do not capture the population parameter with a certain probability.

Confidence Intervals

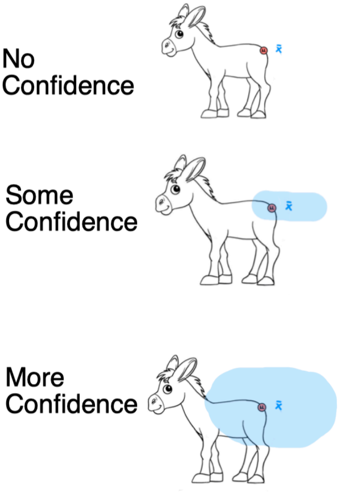

- Remember we are trying to capture the mean! The mean is not *falling into our interval*.
- The bigger the interval, the higher the confidence.