PSY 503: Foundations of Psychological Methods
Lecture 11: Estimation and Uncertainty I

Robin Gomila

Princeton

October 3, 2020

Congratulations!

You have learned so much

You have learned so much

- It's been less than 6 weeks!

You have learned so much

- It's been less than 6 weeks!
- Be proud of yourself!

You have learned so much

- It's been less than 6 weeks!
- Be proud of yourself!
- Your problem set answers are great!

You have learned so much

- It's been less than 6 weeks!
- Be proud of yourself!
- Your problem set answers are great!
- Crystallizing knowledge takes time!

Methodological plasticity

- Potential Outcomes / Experimental Design

Methodological plasticity

- Potential Outcomes / Experimental Design
- Probability theory

Methodological plasticity

- Potential Outcomes / Experimental Design
- Probability theory
- Bayes' rule

Methodological plasticity

- Potential Outcomes / Experimental Design
- Probability theory
- Bayes' rule
- Expectations

Methodological plasticity

- Potential Outcomes / Experimental Design
- Probability theory
- Bayes' rule
- Expectations
- R

Methodological plasticity

- Potential Outcomes / Experimental Design
- Probability theory
- Bayes' rule
- Expectations
- R
- R Markdown

Methodological plasticity

- Potential Outcomes / Experimental Design
- Probability theory
- Bayes' rule
- Expectations
- R
- R Markdown
- Simulation

Where are we going?

Where are we going?

- Statistical inference: right now!

Where are we going?

- Statistical inference: right now!
- Estimation (standard error, confidence intervals, statistical power)
- Hypothesis testing (p-values, null hypotheses)
- Regression

Where are we going?

- Statistical inference: right now!
- Estimation (standard error, confidence intervals, statistical power)
- Hypothesis testing (p-values, null hypotheses)
- Regression
- I know it's Monday morning. . .

Where are we going?

- Statistical inference: right now!
- Estimation (standard error, confidence intervals, statistical power)
- Hypothesis testing (p-values, null hypotheses)
- Regression
- I know it's Monday morning... but I hope that you are excited!

Statistical Inference

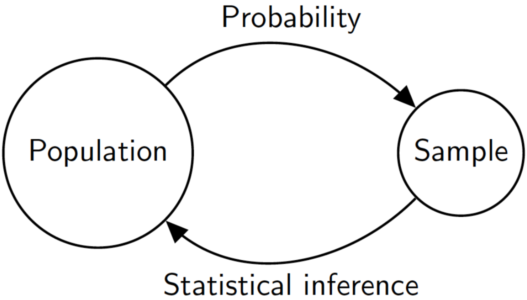

Estimation

Estimation

- **Parameter**
  - Quantity of Interest (e.g., ATE)

Estimation

- **Parameter**
  - Quantity of Interest (e.g., ATE)
- **Estimate**
  - Specific value obtained from a given set of observations (e.g., $\widehat{ATE}$)

Estimation

- **Parameter**
  - Quantity of Interest (e.g., ATE)
- **Estimate**
  - Specific value obtained from a given set of observations (e.g., $\widehat{ATE}$)
- **Estimator**
  - Procedure or formula used to make guesses about a parameter (e.g., difference in means between control and treatment groups)

Illustration: Presidential Election

- Suppose we want to estimate the the proportion of American voters who support Joe Biden

Illustration: Presidential Election

- Suppose we want to estimate the the proportion of American voters who support Joe Biden
- We randomly sample n American voters from the population of interest without replacement
    - Representative sample of American voters

Illustration: Presidential Election

- Suppose we want to estimate the the proportion of American voters who support Joe Biden
- We randomly sample n American voters from the population of interest without replacement
  - Representative sample of American voters
- We ask: Do you support Joe Biden?

Illustration: Presidential Election

- Suppose we want to estimate the the proportion of American voters who support Joe Biden
- We randomly sample n American voters from the population of interest without replacement
  - Representative sample of American voters
- We ask: Do you support Joe Biden?

Illustration: Presidential Election

- Let p denote the population proportion of Biden supporters.

Illustration: Presidential Election

- Let p denote the population proportion of Biden supporters.
- Let X represents response to survey question
  - $X_i = 1$ if voter i supports Biden
  - $X_i = 0$ if voter i does not support Biden

Illustration: Presidential Election

- Let p denote the population proportion of Biden supporters.
- Let X represents response to survey question
  - $X_i = 1$ if voter i supports Biden
  - $X_i = 0$ if voter i does not support Biden
- Random sampling: $\{X_1, X_2, ..., X_n\}$ are **independently and identically distributed** Bernoulli random variables with success probability p

Estimator: Sample mean (i.e., proportion)

Estimator: Sample mean (i.e., proportion)

$$\bar{X} = \sum_{i=1}^{n} \frac{X_i}{n}$$

We use this estimator to estimate the unknown parameter p.

Estimation error

- How good is an estimate?

$$\text{estimation error} = \text{estimate - truth} = \bar{X}_n - p$$

Estimation error

- How good is an estimate?

$$\text{estimation error} = \text{estimate - truth} = \bar{X}_n - p$$

- How do we compute estimator error when we do statistical inference?

Estimation error

- How good is an estimate?

$$\text{estimation error} = \text{estimate - truth} = \bar{X}_n - p$$

- How do we compute estimator error when we do statistical inference?
  - We can never compute estimation error because we do not know p

Revisiting bias

- Any estimator (e.g., sample mean $\bar{X}$, difference in means between treatment and control) can be considered a random variable that has its own distribution over the repeated use of random sampling

Revisiting bias

- Any estimator (e.g., sample mean $\bar{X}$, difference in means between treatment and control) can be considered a random variable that has its own distribution over the repeated use of random sampling
- Presidential Election example: sample mean/proportion $\bar{X}_n$ is a binomial random variable, divided by n, with success probability p and size n.

Revisiting bias

- Any estimator (e.g., sample mean $\bar{X}$, difference in means between treatment and control) can be considered a random variable that has its own distribution over the repeated use of random sampling
- Presidential Election example: sample mean/proportion $\bar{X}_n$ is a binomial random variable, divided by n, with success probability p and size n. We have

$$\text{bias} = \mathbb{E}[\text{estimation error}] = \mathbb{E}[\text{estimate} - \text{truth}] = \mathbb{E}[\bar{X}_n] - p$$

Revisiting bias

- Any estimator (e.g., sample mean $\bar{X}$, difference in means between treatment and control) can be considered a random variable that has its own distribution over the repeated use of random sampling
- Presidential Election example: sample mean/proportion $\bar{X}_n$ is a binomial random variable, divided by n, with success probability p and size n. We have

$$\text{bias} = \mathbb{E}[\text{estimation error}] = \mathbb{E}[\text{estimate} - \text{truth}] = \mathbb{E}[\bar{X}_n] - p$$

- In this example, random sampling implies that

$$\mathbb{E}[\bar{X}_n] - p = p - p = 0$$

General result for the sample mean

- Regardless of the distribution of the random variables, random sampling provides a way to use the sample average as an unbiased estimator of the population mean. We have

$$\mathbb{E}[\bar{X_n}] = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[X_i] = \mathbb{E}[X]$$

General result for the sample mean

- Regardless of the distribution of the random variables, random sampling provides a way to use the sample average as an unbiased estimator of the population mean. We have

$$\mathbb{E}[\bar{X}_n] = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[X_i] = \mathbb{E}[X]$$

- Random sampling eliminates bias

Convergence of Estimators

Law of Large Numbers

Law of Large Numbers

As the sample size increases, the sample average converges to the expectation or population average.

Law of Large Numbers

Let's open R Studio...

Law of Large Numbers

Law of Large Numbers

Suppose that we obtain a random sample of n i.i.d. observations X_1, X_2, ..., X_n, from a random variables with expectation $\mathbb{E}[X]$. The **law of large numbers** states

Law of Large Numbers

Suppose that we obtain a random sample of n i.i.d. observations X_1, X_2, ..., X_n, from a random variables with expectation $\mathbb{E}[X]$. The **law of large numbers** states

$$\bar{X_n} = \frac{1}{n}\sum_{i=1}^{n} X_i \to \mathbb{E}[X]$$

where we use $\to$ as shorthand for convergence.

Law of Large Numbers

- Implies that sample average $\bar{X}_n$ will better approximate $\mathbb{E}[X]$ as sample size increases

Law of Large Numbers

- Implies that sample average $\bar{X}_n$ will better approximate $\mathbb{E}[X]$ as sample size increases
- Super powerful: can be applied in most settings without knowledge of the underlying probability distribution

Law of Large Numbers and Monte Carlo Simulation

Law of Large Numbers and Monte Carlo Simulation

- The Law of Large number justifies the use of *Monte Carlo* simulations
  - i.e., repeat simulation trials many many times and take the mean of these trials
  - e.g., we used Monte Carlo simulation for the trendy restaurant in NY and pregnancy probabilities

Law of Large Numbers and estimation error

Law of Large Numbers and estimation error

- Estimation error becomes smaller as the sample size increases

Consistency of estimators

An **estimator** is said to be **consistent** is it converges to the parameter as the sample size goes to infinity.

Unbiasedness and Consistency of the sample mean

- Good estimators: **unbiased** and **consistent**!
- The sample mean is a good estimator of the population mean

Unbiasedness and Consistency of the sample mean

- Good estimators: **unbiased** and **consistent**!
- The sample mean is a good estimator of the population mean

$$\mathbb{E}[\bar{X}_n] = \mathbb{E}[X] \quad \text{and} \quad \bar{X}_n \rightarrow \mathbb{E}[X]$$

Biased and not consistent estimator of the population mean

Biased and not consistent estimator of the population mean

$$\frac{1}{n}\sum_{i=1}^{n} X_i + n$$

Biased but consistent estimator of the population mean

Biased but consistent estimator of the population mean

$$\frac{1}{n}\sum_{i=1}^{n} X_i + \frac{1}{n}$$

Unbiased but not consistent estimator of the population mean

Unbiased but not consistent estimator of the population mean

$$X_n$$

Difference-in-means estimators of the average treatment effect in experiments

SATE

SATE

- The **difference-in-means** between treatment and control groups in an experiment is an **unbiased** and **consistent** estimator of the sample average treatment effect (SATE) IF:

SATE

- The **difference-in-means** between treatment and control groups in an experiment is an **unbiased** and **consistent** estimator of the sample average treatment effect (SATE) IF:
    - Participants randomly assigned to an experimental condition

PATE

- The **difference-in-means** between treatment and control groups in an experiment is an **unbiased** and **consistent** estimator of the population average treatment effect (PATE) IF:

PATE

- The **difference-in-means** between treatment and control groups in an experiment is an **unbiased** and **consistent** estimator of the population average treatment effect (PATE) IF:
  1. Participants were randomly sampled from the population
  2. Participants randomly assigned to an experimental condition

Understanding and Quantifying Uncertainty of Estimates

Quantifying uncertainty

- Law of Large Number cannot quantify how good an estimate becomes as n increases

Quantifying uncertainty

- Law of Large Number cannot quantify how good an estimate becomes as n increases

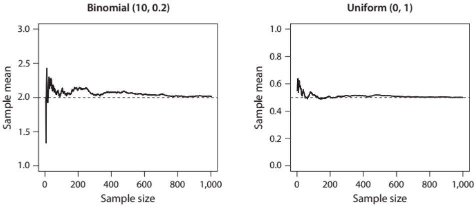

- e.g., convergence seems to occur faster for the uniform distribution on this figure, compared to binomial distribution

Quantifying uncertainty

- Law of Large Number cannot quantify how good an estimate becomes as n increases

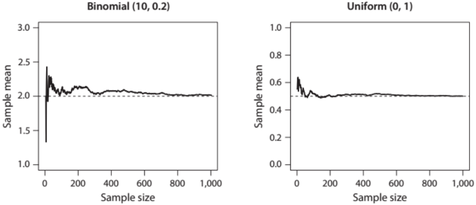

- e.g., convergence seems to occur faster for the uniform distribution on this figure, compared to binomial distribution
- In practice, we only observe one sample mean and do not know the expectation
- We need a different tool to know how well our sample mean approximates the expectation

Quantifying uncertainty

- Let V be the proportion of Princeton undergrads who would report that they violated a social norm last week
- Suppose that we know that $V \sim Bern(0.25)$
- V returns 1 if student respond yes, no otherwise

Sampling distribution of mean estimator

Back to R Studio...

Standard Error of the sample mean

Suppose that we have a sample of n i.i.d. random variables $\{X_1, X_2, ..., X_n\}$. The standard error of the sample mean $\bar{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i$ is given by

$$\text{standard error of the sample mean} = \sqrt{\widehat{\mathbb{V}[\bar{X}_n]}}$$

Standard Error of the sample mean

- Let's take a closer look at: $\mathbb{V}[\bar{X}_n]$:

$$\begin{aligned}\mathbb{V}[\bar{X}_n] &= \mathbb{V}\left[\frac{1}{n}\sum_{i=1}^{n} X_i\right] \\ &= \frac{1}{n^2}\mathbb{V}\left[\sum_{i=1}^{n} X_i\right] \\ &= \frac{1}{n}\mathbb{V}[X_i] \\ &= \frac{\mathbb{V}[X]}{n}\end{aligned}$$

Standard Error of the sample mean

$$\text{standard error of the sample mean} = \sqrt{\widehat{\mathbb{V}[\bar{X}_n]}} = \sqrt{\frac{\widehat{\mathbb{V}[X]}}{n}}$$

- Therefore, we can compute the standard error of the estimator by estimating the variance of the estimator $\mathbb{V}[\bar{X}]$ using the sample variance.

Standard Error of the difference-in-means estimator

Suppose that we have a sample of n i.i.d. random variables $\{T_i, T_2, ..., T_n\}$. And that we also have another sample of of m independently and i.i.d random variables $\{C_i, C_2, ..., C_n\}$.

Then the **standard error** of the difference-in-means estimator $\sum_{i=1}^{n} \frac{T_i}{n} - \sum_{i=1}^{m} \frac{C_i}{m}$ is given by

$$\text{standard error of the difference-in-means estimator} = \sqrt{\frac{\widehat{\mathbb{V}[T]}}{n} + \frac{\widehat{\mathbb{V}[C]}}{m}}$$

Central Limit Theorem

The distribution of the sample mean approaches the normal distribution as the sample size increases.

Central Limit Theorem

- Kicks in regardless of the distribution of the random variable
- This result is incredibly useful because the normal distribution is a parametric distribution
- This allows us to quantify the uncertainty of our estimates!