

Week 9: Regression in the Social Sciences and Frameworks for Causal Inference

Brandon Stewart¹

Princeton

October 26–31, 2020

¹These slides are heavily influenced by Matt Blackwell, Justin Grimmer, Jens Hainmueller, Erin Hartman, Kosuke Imai and Ian Lundberg.

Where We've Been and Where We're Going...

- Last Week
 - ▶ diagnostics
- This Week
 - ▶ making an argument in social sciences
 - ▶ causal inference
 - ▶ two frameworks: potential outcomes and directed acyclic graphs
 - ▶ the experimental ideal
 - ▶ causation for non-manipulable variables
- Next Week
 - ▶ selection on observables
- Long Run
 - ▶ probability $\rightarrow$ inference $\rightarrow$ regression $\rightarrow$ causal inference

1 Making Arguments

- Regression
- Causal Inference
- Visualization

2 Core Ideas in Causal Inference

3 Potential Outcomes

- Framework
- Estimands
- Three Big Assumptions
- Average Treatment Effects
- What Gets to Be a Cause

4 Causal Directed Acyclic Graphs

5 Causation for Non-Manipulable Variables

Why Are We Doing All of This Again?

- We are all here because we are trying to do some **social science**, that is, we are in the business of knowledge production.
- Quantitative methods are an increasingly big part of that so whether you are **reading** or actively **doing** quantitative analysis it is going to be there.
- So why all the math? We are taking a **future-oriented** approach. We want to prepare you for the **next big thing**.
- Methods that became popular in the social sciences since I took the equivalent of this class: machine learning, text-as-data, Bayesian nonparametrics, design-based inference, DAG-based causal inference, deep learning.
- A **technical foundation** prepares you to learn new methods for the rest of your career. Trust me **now** is the time to invest.
- Knowing how methods work also makes you a better reader of work.

**DO ALL THE
MATH**



memegenerator.net

Quantitative Social Science

- Three components of quantitative social science:
 - 1 Argument
 - 2 Research Design
 - 3 Presentation
- This week we will focus on:
 - ▶ identification and causal inference (argument, design)
 - ▶ visualization and quantities of interest (argument, presentation)
- My core argument: to have a hope of success we need to be **clear about the estimand**. The implicit estimand is often (but not always) causal.

We will mostly talk about statistical methods here (it is a statistics class!) but the best work is a **combination** of substantive and statistical theory.

Regression as a Tool: A Review

- Regression is a tool for approximating a conditional expectation function, we can always think of it as saying ‘amongst the **subgroup** of units with covariates $X = x$ what is the average outcome.’
- This in turn is the best prediction of Y given X when ‘best’ is measured in terms of **mean squared error**.
- Confusion starts to creep in when we start talking about **marginal effects** in our prediction.
- Marginal effects are a really powerful way of summarizing differences across subgroups but they tend to lend themselves to **causal** interpretations that they don’t necessarily have.
- This is because they are about **different groups of units** not about **the same unit under intervention**.

A Concrete Example from Gerber, Green and Larimer

Non-parametric estimation

		Voted in 2002 General?	
		No	Yes
Voted in 2000 General?	No	.14	.34
	Yes	.21	.35

Additive regression

$$\beta = (0.16451, 0.03177, 0.15360)$$

		Voted in 2002 General?	
		No	Yes
Voted in 2000 General?	No	.16	.32
	Yes	.20	.35

Marginal Effects

Consider the model

$$Y = \beta_0 + X\beta_1 + Z\beta_2 + XZ\beta_3 + u$$

The **marginal “effect”** of X on Y is defined to be the association between X and Y holding the other variables constant. It is also the partial derivative:

$$\frac{\partial Y}{\partial X} = \beta_1 + Z\beta_3$$

If Z is binary, this says that,

- when $Z = 0$, the association between X and Y is β_1
- when $Z = 1$, the association between X and Y is $\beta_1 + \beta_3$

Marginal Effects

$$Y = \beta_0 + X\beta_1 + Z\beta_2 + XZ\beta_3 + u$$

$$\frac{\partial Y}{\partial X} = \beta_1 + Z\beta_3$$

What is the variance of the marginal effect?

$$\begin{aligned}\text{Var}\left(\frac{\partial Y}{\partial X}\right) &= \text{Var}(\hat{\beta}_1 + Z\hat{\beta}_3) \\ &= \text{Var}(\hat{\beta}_1) + Z^2 \text{Var}(\hat{\beta}_3) + 2Z\text{Cov}(\hat{\beta}_1, \hat{\beta}_3)\end{aligned}$$

If this model is fit using the `lm()` function, we can use `vcov(fit)` to extract the variance covariance matrix that has these variance and covariance elements.

Marginal Effects

Similarly, consider a model with a quadratic term:

$$Y = \beta_0 + X\beta_1 + X^2\beta_2 + u$$



What is the marginal “effect” of X ? What is its variance?

$$\frac{\partial Y}{\partial X} = \beta_1 + 2X\beta_2$$

$$\begin{aligned} \text{Var}\left(\frac{\partial Y}{\partial X}\right) &= \text{Var}(\hat{\beta}_1 + 2X\hat{\beta}_2) \\ &= \text{Var}(\hat{\beta}_1) + (2X)^2 \text{Var}(\hat{\beta}_2) + 2 * 2X * \text{Cov}(\hat{\beta}_1, \hat{\beta}_2) \end{aligned}$$

Plotting Marginal Effects

Given estimated coefficients, we could plot the marginal effect of X on Y as a function of X




Pursuing Single Number Summaries

- If you want to summarize marginal effect across all values of X when it depends on Z there are essentially two options:
 - ▶ calculate at the average observed value of Z .
 - ▶ average over the observed distribution setting Z to values observed in the dataset.
- More generally, we can always pose a specific question of our model and get the answer by plugging in the relevant predictions and averaging.
- You can see how this lends itself to improper causal thinking!

What is Causal Inference?

- A causal inference is a statement about **counterfactuals** — it is a statement about the difference between what did and didn't happen
- The core puzzle of causal inference is how you get the information about what didn't happen
- The difference between prediction and causal inference is the **intervention** on the system under study
- Like it or not, social science theories are almost always expressed as causal claims: e.g. “an increase in X causes an increase in Y ” (or something more opaque meaning the same thing)
- The study of causal inference helps us understand the assumptions we need to make this kind of claim.
- Don't be casual about causal inference!
- This will be the subject of the rest of the week but for now let's change gears. . .

An Intro Motivation




Visualization

- Visualization is **hard** but ultimately extremely **important**
- It is absurd that we spend months collecting data, weeks analyzing it and five minutes slapping it into an unreadable table.
- Visualization can be used for many purposes
 - ▶ drawing people into a topic/dataset
 - ▶ presenting evidence
 - ▶ exploration/model checking
- Three steps involved
 - 1 clearly define the goal
 - 2 estimate quantities of interest
 - 3 present those quantities in a compelling way
- Good design involves thinking carefully about the **audience** (are you making the graph for yourself or someone else?)
- I **strongly recommend** Kieran Healy's visualization book — great summary of the fundamentals plus R code.


Examples

People spend more time with parents on Christmas



Source: Ian Lundberg


Examples



Source: Ian Lundberg


Examples

The test score gender gap in about 1,800 large school districts



Source: New York Times


Examples



Includes men who were ages 27 to 32 in 2010.


Source: New York Times

Examples




Source: New York Times

Examples



Source: Olivia Walch

Examples



Source: The Pudding

Examples





Figure 1. Average births per million people per day, 1930–1991. Each tile represents one month. The underlying count is number of births per month, standardized first by the total population for the period and then by the number of days in that month. Data for the United States are from the U.S. Census Bureau. Data for England and Wales are from the U.K. Office of National Statistics.

Source: Kieran Healy

Examples



Source: Kieran Healy

Case Study 1: Visualization in the New York Times

AMERICAS

How Stable Are Democracies? ‘Warning Signs Are Flashing

The Interpreter

By AMANDA TAUB NOV. 29, 2016

WASHINGTON — Yascha Mounk is used to being the most pessimistic person in the room. Mr. Mounk, a lecturer in government at Harvard, has spent the past few years challenging one of the bedrock assumptions of Western politics: that once a country becomes a liberal democracy, it will stay that way.

His research suggests something quite different: that liberal democracies around the world may be at serious risk of decline.

Mr. Mounk’s interest in the topic began rather unusually. In 2014, he published a book, “[Stranger in My Own Country](#).” It started as a memoir of his experiences growing up as a Jew in Germany, but became a broader investigation of how contemporary European nations were struggling to construct new, multicultural national identities.

Case Study 1: Visualization in the New York Times

The Danger of Deconsolidation


THE DEMOCRATIC DISCONNECT

Roberto Stefan Foa and Yascha Mounk

Roberto Stefan Foa is a principal investigator of the World Values Survey and fellow of the Laboratory for Comparative Social Research. His writing has appeared in a wide range of journals, books, and publications by the UN, OECD, and World Bank. Yascha Mounk is a lecturer on political theory in Harvard University's Government Department and a Carnegie Fellow at New America, a Washington, D.C.-based think tank. His dissertation on the role of personal responsibility in contemporary politics and philosophy will be published by Harvard University Press, and his essays have appeared in Foreign Affairs, the New York Times, and the Wall Street Journal.

Case Study 1: Visualization in the New York Times

Percentage of people who say it is “essential” to live in a democracy



Source: Yascha Mounk and Roberto Stefan Foa, “The Signs of Democratic Deconsolidation,” *Journal of Democracy* | By The New York Times


Case Study 1: Visualization in the New York Times



Ryan D. Enos @RyanDEnos · 19h




Lots of worried chatter a/b @amandataub
article on work of @Yascha_Mounk.
Important, but want to raise cautions 1/



Alternate Graphs

Percentage of people who say it is “essential” to live in a democracy





Source: Yascha Mounk and Roberto Stefan Foa, “The Signs of Democratic Deconsolidation,” *Journal of Democracy* | By The New York Times

Alternate Graphs


.@RyanDEnos Compare NYT/JoD (left) to the very same data analysed differently by Bartels and Achen (2016) (right). Extreme score vs means.

Across numerous countries, including Australia, Britain, the Netherlands, New Zealand, Sweden and the United States, the percentage of people who say it is "essential" to live in a democracy has plummeted, and it is especially low among younger generations.




Alternate Graphs

@RyanDEnos They also stop at the 80s cohort. The data has the 90's as well. I wonder why they would stop there...



Alternate Graphs

Percentage of people who say it is *extremely* important to live in a country that is governed democratically



Decade of birth

Source: ESS Wave 8

In reply to Ryan D. Enos




Benjamin Seck @bseck · 16h

@RyanDEnos Same analysis strategy with comparable data from @ESS_Survey (similar item, 0-10 scale) shows slightly different pattern, too.

Alternate Graphs

How important is it for you to live in a country that is governed democratically?




814 Bantam @jgbsch · 16h

@RyanDENos @bshor @nataliemjb @TomWCvdMeer this is a "quick and dirty" plot I did with WVS wave 6. Not quite so terrifying.

Alternate Graphs

How important is it for you to live in a country that is governed democratically? United States, 2011



Dimiter Toshkov @DToshkov · 31m

my take on the democratic deconsolidation graph that scared everyone yesterday. Blue is 1940s cohort, red is 1980s. First, United States

Thoughts

Two stories here:

- 1 Visualization and data coding choices are important
- 2 The internet is amazing (especially with replication data being available!)

Case Study 2: Sean Taylor's Night Off

 **Sean J. Taylor**
@seanjtaylor

I think this is an interesting topic but found this visualization hard to follow (no surprise if you've been reading my complaints about animated plots).

I have nothing to do tonight so i'm going to try to re-visualize this data. Starting a THREAD I'll keep updated as I go.

 **Nathan Yau** @flowingdata · Oct 16, 2019
How commuting is too much? It depends on where you live.
flowingdata.com/2019/10/16/how...

COMMUTE THRESHOLDS OF AMERICA
AT LEAST 170 MINUTES ROUNDTRIP



GIF

12:39 AM · Oct 19, 2019 · Twitter Web App


<https://twitter.com/seanjtaylor/status/1185415182761254912>

Case Study 2: Sean Taylor's Night Off

```
df %>%
```


```
  ggplot(aes(x = TRANTIME)) + geom_histogram(binwidth=5)
```

Don't know how to automatically pick scale for object of type haven_l




<https://twitter.com/seanjtaylor/status/1185415182761254912>

Case Study 2: Sean Taylor's Night Off



<https://twitter.com/seanjtaylor/status/1185415182761254912>

Case Study 2: Sean Taylor's Night Off



<https://twitter.com/seanjtaylor/status/1185415182761254912>

We Covered

- Thoughts about making a non-causal argument.
- Regression and marginal effects.
- Visualization.


Next Time: Core Ideas in Causal Inference

Where We've Been and Where We're Going...

- Last Week
 - ▶ diagnostics
- This Week
 - ▶ making an argument in social sciences
 - ▶ causal inference
 - ▶ two frameworks: potential outcomes and directed acyclic graphs
 - ▶ the experimental ideal
 - ▶ causation for non-manipulable variables
- Next Week
 - ▶ selection on observables
- Long Run
 - ▶ probability $\rightarrow$ inference $\rightarrow$ regression $\rightarrow$ causal inference

- 1 Making Arguments
 - Regression
 - Causal Inference
 - Visualization
- 2 Core Ideas in Causal Inference
- 3 Potential Outcomes
 - Framework
 - Estimands
 - Three Big Assumptions
 - Average Treatment Effects
 - What Gets to Be a Cause
- 4 Causal Directed Acyclic Graphs
- 5 Causation for Non-Manipulable Variables

Causation



Fundamental Problem of Causal Inference

- Causal inference is the study of counterfactuals.
- The hard thing about counterfactuals is that we never get to see all of them: **Fundamental Problem of Causal Inference** Holland (1986)).
- Assumptions and careful design are the only way out of this problem because we never get to see the truth.
- When it works though it can be a powerful view into the things that we care the most about.
- By convention we often care the counterfactual levels we care about **treated** and **control** and we often consider only binary **treatment variables** because continuous variables are often even more complicated!

Causal Workflow

- 1) **Question** $\leftarrow$ the thing we care about
- 2) **Estimand** $\leftarrow$ the causal quantity of interest
- 3) **Ideal Experiment** $\leftarrow$ what's the counterfactual we care about
- 4) **Identification Strategy** $\leftarrow$ how we connect features of a probability distribution of observed data to causal estimand.
- 5) **Estimation** $\leftarrow$ how we estimate a feature of a probability distribution from observed data.
- 6) **Inference/Uncertainty** $\leftarrow$ what would have happened if we observed a different treatment assignment? (and possibly sampled a different population)

Identification

- A quantity of interest is **identified** when (given stated assumptions) access to **infinite** data would result in the estimate taking on only a single value.
- For example, having all dummy variables in a linear model is not statistically **identified** because they cannot be distinguished from the intercept.
- **Causal identification** is what we can learn about a causal effect from available data.
- If an effect is not identified, no estimation method will recover it.
- ‘What’s your identification strategy?’ means ‘what are the assumptions that allow you to claim that the **association** you’ve estimated has a **causal interpretation**?’
- Identification depends on **assumptions** not statistical models.
- As we will see this is **not** a conversation about estimation: in other words, if someone answers “regression” they have made a **category error**

Identification vs. Estimation

- **Identification**: How much can you learn about the estimand if you have an infinite amount of data?
- **Estimation**: How much can you learn about the estimand from a finite sample?
- Identification precedes estimation

The role of assumptions:

- Often identification requires (hopefully minimal) assumptions
- Even when identification is possible, estimation may impose additional assumptions (i.e. that the linear approximation to the CEF is good enough)
- **Law of Decreasing Credibility (Manski)**: The credibility of inference decreases with the strength of the assumptions maintained

Confounding: The Threat to Identification

- **Confounding** is the bias caused by common causes of the treatment and outcome.
 - ▶ Leads to “spurious correlation.”
- In observational studies, the goal is to avoid confounding inherent in the data.
- Pervasive in the social sciences:
 - ▶ effect of income on voting (confounding: age)
 - ▶ effect of job training program on employment (confounding: motivation)
 - ▶ effect of political institutions on economic development (confounding: previous economic development)
- No unmeasured confounding assumes that we’ve measured all sources of confounding.

Mostly Harmless Econometrics Frequently Asked Questions

- What is the causal relationship of interest?
- What is the experiment that could ideally be used to capture the causal effect of interest?
- What is your identification strategy?
- What is your mode of statistical inference?

Avoiding Common Areas of Confusion

- **contribution not attribution**: we care about a difference which doesn't make it the main reason, nor does it imply a morality claim, it doesn't make T the reason it happened, it doesn't mean that T is "responsible" for Y
- T can 'cause' Y if it is neither necessary nor sufficient
- If you know that on average A causes B and B causes C this doesn't mean you know that A causes C (example $A \rightarrow B$ for one subgroup, $B \rightarrow C$ for second subgroup, still no $A \rightarrow C$)
- estimation of causal effects does not require identical treatment and control groups
- you need a **clear counterfactual** to have a well-defined causal effect. For example of 'the recession was caused by Wall Street' may make intuitive sense but is it well-defined?

<http://egap.org/methods-guides/10-things-you-need-know-about-causal-inference>

We Covered

- Identification vs. Estimation in Causal Inference
- What Causal Inference is Broadly

Next Time: Potential Outcomes

Where We've Been and Where We're Going...

- Last Week
 - ▶ diagnostics
- This Week
 - ▶ making an argument in social sciences
 - ▶ causal inference
 - ▶ two frameworks: potential outcomes and directed acyclic graphs
 - ▶ the experimental ideal
 - ▶ causation for non-manipulable variables
- Next Week
 - ▶ selection on observables
- Long Run
 - ▶ probability $\rightarrow$ inference $\rightarrow$ regression $\rightarrow$ causal inference

- 1 Making Arguments
 - Regression
 - Causal Inference
 - Visualization
- 2 Core Ideas in Causal Inference

- 3 Potential Outcomes
 - Framework
 - Estimands
 - Three Big Assumptions
 - Average Treatment Effects
 - What Gets to Be a Cause

- 4 Causal Directed Acyclic Graphs

- 5 Causation for Non-Manipulable Variables

The Potential Outcomes Framework

- Potential Outcomes is one of two major frameworks that we will consider for doing causal inference.
- It is a way of thinking about counterfactuals and the assumptions required to make statements about them.
- We will first step through the **framework**, then discuss **estimands**, **three big assumptions** and finally what **counts as a cause**.

Potential Outcomes

Definitions:

T_i : Dichotomous Treatment assignment for unit i (multi-valued treatments are possible too—just more potential outcomes for each unit)

$$T_i = \begin{cases} 1 & \text{Unit is assigned to treatment} \\ 0 & \text{Unit is not assigned to treatment} \end{cases}$$

Y_i : Outcome for unit i

Potential outcomes for unit i :

$$Y_i(T_i) = \begin{cases} Y_i(1) & \text{Potential outcome for unit } i \text{ with treatment} \\ Y_i(0) & \text{Potential outcome for unit } i \text{ without treatment} \end{cases}$$

Pre-treatment covariates X_i

τ_i : The treatment effect

$$\tau_i = Y_i(1) - Y_i(0)$$

Potential Outcomes – Aspirin Example

Definitions:

T_i : Unit assigned to:

$$T_i = \begin{cases} 1 & \text{Receive Aspirin} \\ 0 & \text{Receive Placebo} \end{cases}$$

$(T_i = 1)$



$(Y_i = 1)$



$(T_i = 0)$



$(Y_i = 0)$



Y_i : Outcome for unit i – Patient has headache, or not

Potential outcomes for unit i :

$$Y_i(T_i) = \begin{cases} Y_i(1) & \text{Headache (or not) for unit } i \text{ with Aspirin} \\ Y_i(0) & \text{Headache (or not) for unit } i \text{ with placebo} \end{cases}$$

Pre-treatment covariates X_i

Illustrated potential outcomes here and later courtesy of Erin Hartman

What is random in the potential outcomes framework?

Note that potential outcomes are thought of as **fixed**, and that they, and the difference between them, can vary by arbitrary amounts for each unit i . There is some true distribution of potential outcomes across the population.

Treatment assignment is the source of randomness

Causal Inference is a Missing Data Problem

Definition: Observed Outcome

$$Y_i = T_i * Y_i(1) + (1 - T_i) * Y_i(0)$$

Inherently, since we cannot observe both treatment and control for unit i , thus we only observe Y_i , causal inference suffers from a **missing data problem**.

No methodology allows us to simultaneously observe both potential outcomes, $Y_i(1)$ and $Y_i(0)$, making τ_i unobservable—and unidentifiable without additional assumptions (**Fundamental Problem of Causal Inference** Holland (1986))

Causal Inference is a Missing Data Problem

Example: Aspirin's Impact on Headaches

Patient		Pill	Headache Status			Age	Academic
i		T_i	$Y_i(0)$	$Y_i(1)$	Y_i	X_{1i}	X_{2i}
1		1	0	0	0	25	Y
2			0	1		55	N
3			1	1		62	Y
4		0	1	1	1	80	N
5		1	0	1	1	32	Y
6			1	0		45	N
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
n		0	0	0	0	71	N

Some Estimands of Interest

- **Sample average treatment effect (SATE)**

$$\frac{1}{n} \sum_{i=1}^n (Y_i(1) - Y_i(0))$$

- **Population average treatment effect (PATE)**

$$\frac{1}{N} \sum_{i=1}^N (Y_i(1) - Y_i(0))$$

- **Population average treatment effect for the treated (PATT)**

$$\mathbb{E}(Y_i(1) - Y_i(0) \mid T_i = 1)$$

- **Population conditional average treatment effect (CATE)**

$$\mathbb{E}(Y_i(1) - Y_i(0) \mid X_i = x)$$

- **Treatment effect heterogeneity:** Zero ATE doesn't mean zero effect for everyone

- 1 Making Arguments
 - Regression
 - Causal Inference
 - Visualization
- 2 Core Ideas in Causal Inference

- 3 Potential Outcomes
 - Framework
 - Estimands
 - Three Big Assumptions
 - Average Treatment Effects
 - What Gets to Be a Cause

- 4 Causal Directed Acyclic Graphs

- 5 Causation for Non-Manipulable Variables

Built in Assumptions

The notation implies three related assumptions:

- No simultaneity
- No interference
 - ▶ We are implicitly stating that the potential outcomes for that unit are unaffected by the treatment status of other units
 - ▶ If this is not true, the number of potential outcomes for unit i grows
 - ▶ Ex: in an experiment with 3 units, if the potential outcomes for unit i depend on the treatment assignment of units j and k , the potential outcomes for unit i are defined by $Y(i, j, k)$:

$$\begin{array}{ll} Y(1, 0, 0) & Y(0, 0, 0) \\ Y(1, 1, 0) & Y(0, 1, 0) \\ Y(1, 0, 1) & Y(0, 0, 1) \\ Y(1, 1, 1) & Y(0, 1, 1) \end{array}$$

- Same version of the treatment

How do we proceed?

Combined, the previous assumptions give us

- **Stable Unit Treatment Value Assumption** (SUTVA)
- Potential violations:
 - ▶ feedback effects
 - ▶ spill-over effects, carry-over effects
 - ▶ different treatment administration

We also need to assume **Positivity** $0 < P(T_i = 1) < 1 \forall i$ with probability 1.

Ignorability

Identification by randomization:

- If treatment is randomized, then treatment is unrelated to any and all underlying characteristics, observed and unobserved (and even unknown)
- Randomization therefore means treatment assignment is **independent of the potential outcomes** $Y_i(1)$ and $Y_i(0)$, i.e.

$$\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp T_i$$

- This is sometimes called **unconfoundedness** or **ignorability**
- Another way of thinking of it: The distributions of the potential outcomes $(Y_i(1), Y_i(0))$ are the same for the treatment and control group.
- Yet another way of thinking of it: The treatment and control group are exchangeable, or balanced (on observables and unobservables) on average

How do we proceed?

Identification by conditional independence:

- If treatment is not randomized, then treatment may be related to underlying characteristics, observed and unobserved, which are related to the potential outcomes
- Therefore, we need to assume that treatment assignment is independent of the potential outcomes $Y_i(1)$ and $Y_i(0)$, conditional on some pre-treatment characteristics X , i.e.

$$\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp T_i \mid X_i$$

- Conditioning set should yield $Y_i(0)$, $Y_i(1)$ and T_i conditionally independent. (This is next week's topic).
- This is **conditional ignorability**.

The Selection Problem

- Why is this difficult? **selection bias**
- The core idea is that the people who get treatment might look different from those who get control and thus they are not good **counterfactuals** for each other.
- Let's look at what we get from a naive difference in means with a binary treatment:

$$\begin{aligned} E[Y_i | T_i = 1] - E[Y_i | T_i = 0] \\ &= E[Y_i(1) | T_i = 1] - E[Y_i(0) | T_i = 0] \\ &= E[Y_i(1) | T_i = 1] - E[Y_i(0) | T_i = 1] + E[Y_i(0) | T_i = 1] - E[Y_i(0) | T_i = 0] \\ &= \underbrace{E[Y_i(1) - Y_i(0) | T_i = 1]}_{\text{Average Treatment Effect on Treated}} + \underbrace{E[Y_i(0) | T_i = 1] - E[Y_i(0) | T_i = 0]}_{\text{selection bias}} \end{aligned}$$

- Naive estimator = Average Treatment Effect on Treated + Selection Bias
- Selection bias: how different the treated and control groups are in terms of their potential outcome under control.

Selection Makes Us Care About Assignment Mechanisms

Assignment Mechanism

“The process that determines which units receive which treatments, hence which potential outcomes are realized and thus can be observed, and, conversely, which potential outcomes are missing.”

(Imbens and Rubin, 2015, p. 31)

Key Assumptions:

- **Individualistic assignment:** Limits the dependence of a particular unit's assignment probability on the values of the covariates and potential outcomes for other units
- **Probabilistic assignment:** Requires the assignment mechanism to imply a non-zero probability for each treatment value, for every unit
- **Unconfounded assignment:** Disallows dependence of the assignment mechanism on the potential outcomes

The Assignment Mechanism

- Since missing potential outcomes are unobservable we must make assumptions to fill in, i.e. **estimate** missing potential outcomes.
- In the causal inference literature, we typically make assumptions about the **assignment mechanism** to do so.

Types of Assignment Mechanisms

- random assignment
- selection on observables
- selection on unobservables

Most statistical models of causal inference attain identification of treatment effects by restricting the assignment mechanism in some way.

Three Big Assumptions

To review, we've talked about three big assumptions

- 1 SUTVA
- 2 Positivity
- 3 (Conditional) Ignorability

1 Making Arguments

- Regression
- Causal Inference
- Visualization

2 Core Ideas in Causal Inference

3 Potential Outcomes

- Framework
- Estimands
- Three Big Assumptions
- Average Treatment Effects
- What Gets to Be a Cause

4 Causal Directed Acyclic Graphs

5 Causation for Non-Manipulable Variables

Average Treatment Effects

Suppose we have N observations in population ($i = 1, \dots, N$)

$$\begin{aligned}\text{ATE} &= \frac{1}{N} \sum_{i=1}^N (Y_i(1) - Y_i(0)) \\ &= E[Y(1) - Y(0)] \text{ Average over population!!!}\end{aligned}$$

- Population parameter
- It is fixed and unchanging

Estimating ATE under Random Assignment

Estimator for ATE:

$$\begin{aligned}\widehat{\text{ATE}} &= \text{Average (Treated Units)} - \text{Average (Control Units)} \\ &= \frac{\sum_{i=1}^N Y_i(1) T_i}{\sum_{i=1}^N T_i} - \frac{\sum_{i=1}^N Y_i(0)(1 - T_i)}{\sum_{i=1}^N (1 - T_i)} \\ &= \sum_{i=1}^N \left[\frac{Y_i(1) T_i}{n_t} - \frac{Y_i(0)(1 - T_i)}{n_c} \right] \\ &= E[Y(1) | T = 1] - E[Y(0) | T = 0]\end{aligned}$$

Average Treatment Effect

Imagine a study population with 4 units:

i	T_i	$Y_i(1)$	$Y_i(0)$	τ_i
1	1	10	4	6
2	1	1	2	-1
3	0	3	3	0
4	0	5	2	3

What is the ATE?

$$E[Y_i(1) - Y_i(0)] = 1/4 \times (6 + -1 + 0 + 3) = 2$$

Note: Average effect is positive, but τ_i are negative for some units!

Average Treatment Effect on the Treated

Imagine a study population with 4 units:

i	T_i	$Y_i(1)$	$Y_i(0)$	τ_i
1	1	10	4	6
2	1	1	2	-1
3	0	3	3	0
4	0	5	2	3

What is the ATT and ATC?

$$E[Y_i(1) - Y_i(0) | T_i = 1] = 1/2 \times (6 + -1) = 2.5$$

$$E[Y_i(1) - Y_i(0) | T_i = 0] = 1/2 \times (0 + 3) = 1.5$$

Naive Comparison: Difference in Means

Comparisons between observed outcomes of treated and control units can often be misleading.

- units which select treatment may not be like units which select control.
- i.e. selection into treatment is often associated with the potential outcomes
- this means we have violated the assumption of unconfoundness $(Y(1), Y(0)) \perp T$

What Gets to Be a Cause?

We can imagine a world where individual i is assigned to treatment and control conditions

What is the Hypothetical Experiment?

Problem: Immutable (or difficult to change) characteristics

- Effect of gender on promotion
- Effect of race on traffic stops

Consider causal effect of race on traffic stops:

- Do we mean effect of officer perceiving a certain race?
- Do we mean randomly assigning race at birth?
- manipulating perceptions is a lot different from manipulating the characteristic

No Causation Without Manipulation

Caveats and Implications

- Does not dismiss claims of discrimination on immutable characteristics as legitimate
 - Pervasive effects of racism/sexism in society
 - Suggests: we need a different empirical strategy to evaluate claims
 - What facet of institutionalized racism (or its consequences) causes racial disparities?
- Correlation problem :
 - Regression models can estimate **coefficients** for immutable characteristics
 - But are necessarily imprecise: what do scholars have in mind in models?
- Design Principle:
 - Pretend you're God designing experiment
 - If that experiment does not exist, be concerned about interpretation

No causation without manipulation?

Always ask:
what is the experiment I would run if I had infinite resources and power?

Summing Up: Neyman-Rubin causal model

- Useful for studying the “effects of causes”, less so for the “causes of effects”.
- No assumption of homogeneity, allows for causal effects to vary unit by unit
 - ▶ No single “causal effect”, thus the need to be precise about the target estimand. (This is true even for perfect experiments.)
- Distinguishes between observed outcomes and potential outcomes.
- Causal inference is a missing data problem: we typically make assumptions about the assignment mechanism to go from descriptive inference to causal inference.

Neyman-Rubin Potential Outcomes Model



Figure: Neyman



Figure: Rubin

Brief History of Potential Outcomes and Causal Inference

- Introduction of potential outcomes in randomized experiments by Neyman (1923)
 - ▶ Super-population inference and confidence intervals
- Introduction of randomization as the “reasoned basis” for inference by Fisher (1925)
 - ▶ p -values and permutation inference
- Causal effects defined at the unit level, allowing for effects to be defined without a known assignment mechanism by Rubin (1974)
- Potential outcomes expanded to observational studies by Rubin (1974)
- Formalization of the assignment mechanism in potential outcomes by Rubin (1975, 1978)
- Pearl (1995) develops graphical models for causal inference

For more detailed see Morgan and Winship.

We Covered

- Potential Outcomes!
- Estimands!
- Three Big Assumptions!
- Treatment Effects!
- No Causation without Manipulation!

Next Time: Causal Directed Acyclic Graphs (Causal DAGs)

Where We've Been and Where We're Going...


- Last Week
 - ▶ diagnostics
- This Week
 - ▶ making an argument in social sciences
 - ▶ causal inference
 - ▶ two frameworks: potential outcomes and directed acyclic graphs
 - ▶ the experimental ideal
 - ▶ causation for non-manipulable variables
- Next Week
 - ▶ selection on observables
- Long Run
 - ▶ probability $\rightarrow$ inference $\rightarrow$ regression $\rightarrow$ causal inference

- 1 Making Arguments
 - Regression
 - Causal Inference
 - Visualization
- 2 Core Ideas in Causal Inference
- 3 Potential Outcomes
 - Framework
 - Estimands
 - Three Big Assumptions
 - Average Treatment Effects
 - What Gets to Be a Cause
- 4 Causal Directed Acyclic Graphs
- 5 Causation for Non-Manipulable Variables

Graphical Models


- A general framework for representing causal relationships based on directed acyclic graphs (DAG)
- The work we discuss here comes out of developments by Judea Pearl and others
- Particularly useful for thinking through issues of identification.
- Provides a graphical representation of the models and a set of rules (do-calculus) for identifying the causal effect.
- Nice software that takes the graph and returns an identification strategy: **DAGitty** at <http://dagitty.net>

Components of a DAG




- nodes represent variables (unobserved typically called U or V)
- (directed) arrows represent **causal** effects
- absence of nodes represents **no common causes** of any pair of variables
- absence of arrows represents **no causal effect**
- positioning conveys no mathematical meaning but often is oriented left-to-right with causal ordering for readability.
- dashed lines are used in context dependent ways
- all relationships are **non-parametric**

Relationships in a DAG




- Parents (Children): directly causing (caused by) a node
- Ancestors (Descendants): directly or indirectly causing (caused by) a node
- Path: a route that connects the variables (path is causal when all arrows point the same way)
- **Acyclic** implies that there are no cycles and a variable can't cause itself
- Causal Markov assumption: condition on its **direct causes**, a variable is independent of its non-descendants.
- We will talk in depth about two types of relationships: **confounders** and **colliders**

Confounders



- X is a **confounder** (or common cause)
- Even without a **causal** effect or directed edge between T and Y they will have a **marginal** associational relationship
- **Conditional** on X , T and Y are unrelated in this graph.
- We can think of conditioning on a confounder as blocking the flow of association.

Colliders



- X is now a **collider** because two arrows point into it
- In this scenario T and Y are **not marginally associated**
- If we control for X they become associated and create a connection between T and Y

Colliders are scary because you can induce dependence



Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable

Felix Elwert¹ and Christopher Winship²

¹Department of Sociology, University of Wisconsin, Madison, Wisconsin 53706;
email: elwert@wisc.edu

²Department of Sociology, Harvard University, Cambridge, Massachusetts 02138;
email: cwinship@hs.harvard.edu

Annu. Rev. Sociol. 2014. 40:31–51

First published online as a Review in Advance on
June 2, 2014

The Annual Review of Sociology is online at
soc.annualreviews.org

This article doi:
10.1215/03612083-1245453

Copyright © 2014 by Annual Reviews.
All rights reserved.


Keywords

causality, directed acyclic graphs, identification, confounding,
selection

Abstract


Endogenous selection bias is a central problem for causal inference. Recognizing the problem, however, can be difficult in practice. This article introduces a purely graphical way of characterizing endogenous selection bias and of understanding its consequences (Elwert et al. 2009). We use causal graphs (direct acyclic graphs, or DAGs) to highlight that endogenous selection bias stems from conditioning (e.g., controlling, stratifying, or selecting) on a so-called collider variable, i.e., a variable that is itself caused by two other variables, one that is (or is associated with) the treatment and another that is (or is associated with) the outcome. Endogenous selection bias can result from direct conditioning on the outcome variable, a post-outcome variable, a post-treatment variable, and even a pre-treatment variable. We highlight the difference between endogenous selection bias, common-cause confounding, and overcontrol bias and discuss numerous examples from social stratification, cultural sociology, social network analysis, political sociology, social demography, and the sociology of education.

From Confounders to Back-Door Paths



- Identify causal effect of T on Y by conditioning on X , Z or X and Z
- We can formalize this logic with the idea of a **back-door** path
- A back-door path is “a path between any causally ordered sequence of two variables that begins with a directed edge that points to the first variable.” (Morgan and Winship 2013)
- Two paths from T to Y here:
 - 1 $T \rightarrow Y$ (directed or causal path)
 - 2 $T \leftarrow X \rightarrow Z \rightarrow Y$ (back-door path)
- Observed marginal association between T and Y is a composite of these two paths and thus does not identify the causal effect of T on Y
- We want to **block** the back-door path to leave only the causal effect

Colliders and Back-Door Paths



- Z is a **collider** and it lies along a back-door path from T to Y
- Conditioning on a collider on a back-door path does not help and in fact causes new associations
- Here we are fine unless we condition on Z which opens a path $T \leftarrow V \leftrightarrow U \rightarrow Y$ (this particular case is called *M*-bias)
- So how do we know which back-door paths to block?

D-Separation


- Graphs provide us a way to think about conditional independence statements. Consider disjoint subsets of the vertices A , B and C
- A is **D-separated** from B by C if and only if C **blocks** every path from a vertex in A to a vertex in B
- A path p is said to be blocked by a set of vertices C if and only if at least one of the following conditions holds:
 - 1 p contains a **chain** structure $a \rightarrow c \rightarrow b$ or a **fork** structure $a \leftarrow c \rightarrow b$ where the node c is in the set C
 - 2 p contains a **collider** structure $a \rightarrow y \leftarrow b$ where **neither** y nor its descendents are in C
- If A is not D -separated from B by C we say that A is **D-connected** to B by C

Backdoor Criterion

- Generally we want to know if we can **nonparametrically** identify the average effect of T on Y given a set of possible conditioning variables X
- Backdoor Criterion for X
 - ① No node in X is a descendent of T
(i.e. don't condition on post-treatment variables!)
 - ② X D -separates every path between T and Y that has an incoming arrow into T (backdoor path)
- In essence, we are trying to **block** all non-causal paths, so we can estimate the **causal** path.
- Backdoor criterion is just one way to identify the effect: but its the most popular approach in the social sciences and what we are trying to do 99% of the time.
- We will see some other approaches late in the semester.

Blocking backdoor paths: College and earnings


What do we need to include to block all backdoor paths between college and earnings?



Ability, parents' income, parents' education, extended family who pay for college and help you find a job, neighborhood characteristics that affect high school quality and also the availability of local jobs, ... **lots of things!**


Non-causal paths: Part 2

Now consider this graph. Is there an unblocked backdoor path from T to Y ?



No need to condition! X_2 already blocks this path. it is a **collider**.

Colliders: Be careful!



Y is a **collider**. X_1 and X_2 are not associated, but they are when we hold Y constant.

What situations might produce this?

- X_1 being in a car accident. X_2 is having cancer. Y is being in a hospital.
- X_1 is living in a warm climate. X_2 is being an elite swimmer. Y is going swimming in January.
- X_1 is family income. X_2 is religiosity. Y attendance at a Catholic high school.

Colliders: When drawing a DAG helps

Example extended from Elwert & Winship 2014

Hypothetical substantive question:

Does acting ability causally affect the probability of marriage?

Hypothetical approach: Estimate on a sample of Hollywood actors and actresses.

We want to estimate:



Acting ability $\longrightarrow$ Marriage

Should we worry about this design? It depends on our **theory** about how these variables are related. We can argue about identification with a DAG.

Colliders: When drawing a DAG helps

Example extended from Elwert & Winship 2014

Suppose working in Hollywood is a function of two factors: acting ability and beauty. In the general population, these two are uncorrelated. However, among those who work in Hollywood, those who are bad at acting must be beautiful.




Colliders: When drawing a DAG helps

Example extended from Elwert & Winship 2014

This is an example of conditioning on a collider! We induce a negative association between acting ability and beauty.

Acting ability



Under the assumptions above, our results are driven by collider conditioning!

Defining Causal Effects

- Pearl's graphical model framework comes with a handy operator called the $\text{do}()$ operator.
- $P(Y|\text{do}(T = t))$ is distinct from $P(Y|T = t)$ with the former being the outcome under **intervention** and the latter being an observed value.
- This can often be helpful for distinguishing data as it exists in the world and data as it might exist in the counterfactual world.
- The do-calculus is actually a much broader set of rules that operate on the DAG structure to help us calculate causal effects (or learn when we can't!).

Thoughts on DAGs and Potential Outcomes

- Two very different languages for talking about and thinking about causal inferences.
- Potential outcomes is very focused on thinking about the **treatment assignment** mechanism and helpful for heterogeneity of treatment effects.
- Potential outcomes is also less of a “foreign language” for most statisticians, but in my experience lumps together a lot of identification assumptions in opaque ignorability conditions.
- Graphical Models with DAGs are very visually appealing but the operations on the graph can be challenging
- DAGs very helpful for thinking through **identification** and the entire **causal process**
- Note that both are about **non-parametric identification** and not **estimation**. This is good and bad.
 - ▶ Good: provides a very general framework that applies in non-linear scenarios and interactions
 - ▶ Bad: identification results for identification only holds when variable is completely controlled for (which may be difficult!)

We Covered

- How to read DAGs.
- We got a hint of what is coming next week with blocking backdoor paths.

Next Time: Causation for Non-Manipulable Variables

Where We've Been and Where We're Going...

- Last Week
 - ▶ diagnostics
- This Week
 - ▶ making an argument in social sciences
 - ▶ causal inference
 - ▶ two frameworks: potential outcomes and directed acyclic graphs
 - ▶ the experimental ideal
 - ▶ causation for non-manipulable variables
- Next Week
 - ▶ selection on observables
- Long Run
 - ▶ probability $\rightarrow$ inference $\rightarrow$ regression $\rightarrow$ causal inference

- 1 Making Arguments
 - Regression
 - Causal Inference
 - Visualization
- 2 Core Ideas in Causal Inference
- 3 Potential Outcomes
 - Framework
 - Estimands
 - Three Big Assumptions
 - Average Treatment Effects
 - What Gets to Be a Cause
- 4 Causal Directed Acyclic Graphs
- 5 Causation for Non-Manipulable Variables

No Causation Without Manipulation

- One of the difficulties that students and practitioners have with causal inference is the need for **manipulation** or an **ideal experiment**.
- In many areas the key variables are arguably **immutable** such as race or gender.
- Sen and Wasow argue that we can improve our empirical work on this by seeing race/ethnicity as a **composite** variable or 'a bundle of sticks' which can be manipulated separately.
- Lundberg offers a perspective where the non-manipulable variable defines social categories but is not the treatment itself.
- More broadly there is a need to define what the **proposed intervention is** because even cases that can be manipulated can be very opaque (e.g. obesity).

Bundle of Sticks


Sen and Wasow (2016) “Race as a Bundle of Sticks: Designs that Estimate Effects of Seemingly Immutable Characteristics” *Annual Review of Political Science*.

The Trouble with Race As Treatment

There are three problems with race as a treatment in the causal inference sense

- ① Race cannot be **manipulated**
 - ▶ without the capacity to manipulate the question is arguably ill-posed and the estimand is unidentified
- ② Everything else is **post-treatment**
 - ▶ everything else comes after race which is perhaps unsatisfying
 - ▶ this also presumes we are only interested in the total effect
- ③ Race is **unstable**
 - ▶ there is substantial variance across treatments which is a SUTVA violation

The Bundle of Sticks



Design 1: Exposure Studies

- Approach

- a) “one or more elements of race is identified as a relevant cue”
- b) “subjects are treated by exposure to the racial cue”
- c) “unit of analysis is the individual or institution being exposed”

- Examples

- ▶ Psychology (Steele 1997 on stereotype threat)
- ▶ Audit/Correspondence Studies (Pager 2003, Bertrand and Mullainathan 2004)
- ▶ Survey Experiments with Racial Cues (Mendelberg 2001)
- ▶ Field Experiments with Racial Cues (Green 2004, Enos 2011)
- ▶ Observational Studies (Greiner and Rubin 2010, Wasow 2012)

Design 2: Within-Group Studies

- Approach: identify variation within the racial group along constitutive element.
- Example: Sharkey (2010) exploiting temporal variation in local homicides in Chicago to identify a significant neighborhood effect of proximity to violence on cognitive performance of African-American children

Concluding Thoughts

We can study race with causal inference, it just takes very **careful design**.


Table 2 Overview of exposure and within-group research designs



	Exposure	Within-Group
Unit	Individuals or institutions, potentially from any group	Members of a particular group
Typical treatment	Racial cue or signal (e.g., include distinctively ethnic names on a resume)	Constitutive element of the composite of race (e.g., address anxiety about social belonging in college)
Role of element of race	One "stick" is a proxy for the bundle (e.g., in a phone call with a landlord, dialect signals many traits associated with race)	One "stick" explains part of the bundle (e.g., Middle Passage might partly explain high rates of hypertension among African-Americans)
Examples	Correspondence and audit studies Implicit Association Tests	Experimental manipulation of a constitutive psychological dimension of race Within-race matching

Gap Closing Estimands

Lundberg (2020) “The gap-closing estimand: A causal approach to study interventions that close disparities across social categories” Working Paper.

Thanks to Ian Lundberg for the slides that follow!





Categories

Treatment

Counterfactual Disparity


Professional

Occupational Stress and Commitment

Offspring of some e

Class Origin

Class Destination



This Week in Review

- We talked about what regression is doing and how we go about making an argument.
- This week we began our journey into **causal inference**.
- The next few weeks we are going to talk about how to use these frameworks to estimate causal effects across a wide variety of scenarios.
- We will make liberal use of **both** frameworks based on whatever is the most convenient to communicate the point.
- You want to have some familiarity with the core concepts of the frameworks—but don't worry, we will review them more in coming weeks.

Next Time: Causality with Measured Confounding