

Causal Inference Methods and Case Studies

STAT24630

Jingshu Wang

Lecture 2

Topic: Potential outcome framework

- Definition of concepts: intervention, potential outcomes, causal effects
- The SUTVA assumption
- Causal estimand and assignment mechanism
- Lord paradox
- An alternative language: Causal directed acyclic graph (DAG)
- Textbook chapters: Chapters 1 & 2

What is causality?

- Causality is tied to an action applied to a unit (intervention, treatment, manipulation)
- **“No causation without manipulation”** (Rubin, 1975)
- Manipulation need not be performed, but should be theoretically possible

Causal questions that we have encountered

- Will a new drug that increases HDL-C concentration reduce the risk of heart disease?
- If a smoking mother stops smoking , will the baby have a higher chance to survive?
- Will two doses of the Moderna vaccine stop the infection of COVID-19?
- If someone has a longer length of education, will his earnings in the future increase?

Other examples of causal questions:

- What is the effect of studying on test scores?
- What is the impact of “blind” auditions on job offers for female musicians?
- What’s the effect of taking aspirin on your headache?
- Does he have lung cancer because he smokes?

The potential outcome framework: basic concepts

- **Treatment:** An active intervention applied at a particular moment in time, whose effects we wish to assess relative to no intervention (the control)
 - Example: take the aspirin to treat headache
- **Unit:** A physical object at a particular moment in time
 - Example: a before-and-after comparison of my headache before and after deciding to take or not to take the aspirin involves two distinct units
- **Potential outcomes:** Outcomes that would be observed if active treatment is applied and that would be observed if control treatment is applied

Unit	Potential Outcomes		Causal Effect
	$Y(\text{Aspirin})$	$Y(\text{No Aspirin})$	
You	No Headache	Headache	Improvement due to Aspirin

The potential outcome framework: basic concepts

- How to define a causal effect?

1. Headache gone only with aspirin:

$$Y(\text{Aspirin}) = \text{No Headache}, Y(\text{No Aspirin}) = \text{Headache}$$

Positive causal effect of aspirin

2. No effect of aspirin, with a headache in both cases:

$$Y(\text{Aspirin}) = \text{Headache}, Y(\text{No Aspirin}) = \text{Headache}$$

No causal effect of aspirin

3. No effect of aspirin, with the headache gone in both cases:

$$Y(\text{Aspirin}) = \text{No Headache}, Y(\text{No Aspirin}) = \text{No Headache}$$

No causal effect of aspirin

4. Headache gone only without aspirin:

$$Y(\text{Aspirin}) = \text{Headache}, Y(\text{No Aspirin}) = \text{No Headache}$$

Negative causal effect of aspirin

- Treatment: $W_i = 1$ if treated or $W_i = 0$ if control

- Causal effect for unit i : $\tau_i = Y_i(1) - Y_i(0)$

The potential outcome framework: basic concepts

- **Pre-treatment covariates:** A background characteristic (measured or unmeasured) of a unit that could not have been affected by treatment assignment
 - Headache example: intensity of headache before making the decision to take aspirin or not
 - Evaluation of a job training program: age, previous educational achievement, family, socio-economic status, pre-training earnings, etc.
 - Pre-training earnings is a **pre-treatment covariate**, while post-training earnings are **potential outcomes**


Voters	Contact	Turnout		Age	Gender
i	W_i	$Y_i(1)$	$Y_i(0)$	X_{i1}	X_{i2}
1	1	1	?	20	M
2	0	?	0	55	F
3	0	?	1	40	F
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	1	0	?	62	M

- **Observed outcome:** $Y_i = Y_i(W_i)$. Only one of the potential outcomes is observed we can not return to a previous point in time to administer the “other” treatment

The potential outcomes are natural concepts




- *A Christmas Carol* (1843): a fiction by Charles Dickens
 - $Y(0)$ potential future if Scrooge continues his miserly ways
 - $Y(1)$ potential future with changed Scrooge



- *Sliding Doors* (1998): The film alternates between two storylines, showing two paths the central character's life could take depending on whether she catches a train.

The potential outcomes are natural concepts



Watch the YouTube video:

<https://www.youtube.com/watch?v=P9TShT3xn4Q>

But-for test in legal contexts

- The but-for test is a test commonly used in both tort law and criminal law to determine actual causation. The test asks, "but for the existence of X, would Y have occurred?"
- The Federal Judicial Center's "Reference Manual on Scientific Evidence" (1994, Chapter 3, p. 481) states:

*In most cases, the analysis considers the difference **between the plaintiff's economic position if the harmful event had not occurred and the plaintiff's actual economic position.** The damages study restates the plaintiff's position "but for" the harmful event; this part is often called but-for analysis. **Damages are the difference between the but-for value and the actual value.***

Fundamental problem of causal inference

- We can also use other quantities to define the causal effect:
ex. Log fold change $\log(Y_i(1)/Y_i(0))$, percentage change $\frac{Y_i(1)-Y_i(0)}{Y_i(0)} \times 100\%$.
- Any causal quantity is a function of potential outcomes
- Non-binary treatment:
 - Categorical: $Y_i(0), Y_i(1), \dots, Y_i(K - 1)$
 - Continuous: $Y_i(t)$ for any $t \in \mathbb{R}$
- Causal inference is a **missing data problem: only one potential outcome is observed for each unit**
 - potential outcomes are thought to be fixed for each unit
 - potential outcomes do have a distribution across units
 - treatment variable determines which potential outcome is observed
 - observed outcomes are random for each unit because the treatment is random

Underlying assumptions of the potential outcome framework: SUTVA

- The above notations implies three assumptions

1. **Causal ordering:** Y_i cannot causally affect W_i

Potential violations: feedback effects

2. **No interference** between units : each unit's potential outcomes remain the same no matter what treatments the other units receive.

$$Y_i(W_1, \dots, W_n) = Y_i(W_i)$$

Example: if we have no contact with each other, whether you take an aspirin has no effect on the status of my headache (not necessarily true if we are in the same room)

Counter-examples:

- vaccine effect, effect of job training programs
- **spillover:** In economics, a spillover is an economic event in one context that occurs because of something else in a seemingly unrelated context. For example, if consumer spending in the United States declines, it has spillover effects on the economies that depend on the U.S. as their largest export market.

3. **No different forms or versions** of each treatment level

Underlying assumptions of the potential outcome framework: SUTVA

- The above notations implies three assumptions
 1. **Causal ordering:** Y_i cannot causally affect W_i
Potential violations: feedback effects
 2. **No interference** between units : each unit's potential outcomes remain the same no matter what treatments the other units receive.
 3. **No different forms or versions** of each treatment level (consistency)
 - For example, by setting $W_i = 1$ for taking the aspirin, we assume that there is only one version of the tablets (old v.s. new) and at a particular dosage
 - If there are multiple versions (say different dosages), we need to redefine as sperate levels of treatments
 - Violated also if the effect defer depending on the method of administering the treatment. Assigned to receive a treatment v.s. choose to take the treatment
- **SUTVA:** Stable Unit Treatment Value Assumption 2 + 3
 - SUTVA can hold even if each unit's treatment options differ (e.g. each person's surgeon can differ)

Causal effects of immutable Characteristics

- **Immutable characteristics:** sex, race, age, etc.
- Can immutable characteristics have meaningful causal effects?
- Examples:
 - **Race:** Study race in the labor market by sending fictitious resumes to help-wanted ads in Boston and Chicago newspapers. To manipulate perceived race, resumes are randomly assigned African-American- or White-sounding names. Researchers find that white names receive 50 percent more callbacks for interviews.
[Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American economic review*, 2004]
 - **Gender:** Since the mid-1990's, one third of Village Council head positions in India have been randomly reserved for a woman: In these councils only women could be elected to the position of head. Researchers find that leaders invest more in infrastructure that is directly relevant to the needs of their own genders.
[Women as policy makers: Evidence from a randomized policy experiment in India. *Econometrica*, 2004]

Causal effects of immutable Characteristics

Extremely difficult with observational data

- Asian American Discrimination in Harvard Admissions
 - The organization Students for Fair Admissions and other plaintiffs filed a lawsuit against Harvard College in 2014, claiming that the college discriminates against Asian American applicants in its undergraduate admissions process.
 - Based on the admission data for Classes of 2014-2019, Economist Peter Arcidiacono from Duke developed a model to estimate the causal influence of Asian American status and concluded that “typical Asian American applicants would see their average admit rate rise by 19%, or approximately 1 percentage point, if they were treated as white applicants.”
 - Using the same dataset, Economist David Card from UC Berkeley argued that Arcidiacono’s models place too much emphasis on academic factors as predictors of admissions outcomes. By considering contextual factors including “high school, community and family background.” in his model, Card argued that the effect of considering racial and ethnic factors doesn’t result in a bias towards Asian-American students as Arcidiacono found.

Causal estimand

- Unit causal effects are difficult (impossible in most cases) to estimate
- Learning about causal effect typically requires multiple unit
- We can average the unit causal effects over n units
 - Average treatment effect : **SATE** = $\frac{1}{n} \sum_{i=1}^n \{Y_i(1) - Y_i(0)\}$
 - Average treatment effect for the treated: **SATT** = $\frac{1}{n_1} \sum_{i=1}^n W_i \{Y_i(1) - Y_i(0)\}$ where $n_1 = \sum_{i=1}^n W_i$
- If we treat the units as sampled from a population
 - Population average treatment effect: **PATE** = $\mathbb{E}(Y_i(1) - Y_i(0))$
 - Average treatment effect for the treated: **PATT** = $\mathbb{E}(Y_i(1) - Y_i(0) \mid W_i = 1)$
- Estimand: any value that can be calculated from Science
we want to define an estimand that is free from any model assumptions

Causal estimand

Other causal quantities of interest

- Heterogenous effects
 - Conditional average treatment effect (CATE)
$$\tau(\mathbf{x}) = \mathbb{E}(Y_i(1) - Y_i(0) \mid \mathbf{X}_i = \mathbf{x})$$
 - Applications in precision medicine and microtargeting
 - Quantile treatment effects
 - $Q_{Y(1)-Y(0)}(\alpha)$ where $Q_Z(\alpha)$ is the α th quantile of the random variable Z , e.x. $\alpha = 0.5$ corresponds to the median
 - (Easier to estimate) $Q_{Y(1)}(\alpha) - Q_{Y(0)}(\alpha)$

- Non-additive effect
 - Odds ratio for binary outcome

$$\frac{P(Y_i(1) = 1)/P(Y_i(1) = 0)}{P(Y_i(0) = 1)/P(Y_i(0) = 0)}$$

Example: vaccine efficacy = 1 – causal odds ratio

How to estimate the causal estimand?

How do we estimate the causal estimand from observed data?

Voters	Contact	Turnout		Age	Gender
i	T_i	$Y_i(1)$	$Y_i(0)$	X_{i1}	X_{i2}
1	1	1	?	20	M
2	0	?	0	55	F
3	0	?	1	40	F
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	1	0	?	62	M

- Causal inference is a missing data problem: only one potential outcome is observed for each unit
- **Understanding the assignment mechanism is important**
- Assignment mechanism: the process governing which units receive active treatment and which receive control

How to estimate the causal estimand?

- We can not simply compare between observed outcomes across units

Unit	Unknown			Known	
	Potential Outcomes		Causal Effect	Actual	Observed
	$Y(\text{Aspirin})$	$Y(\text{No Aspirin})$		Treatment W_i	Outcome Y_i^{obs}
You	No Headache	Headache	Improvement due to Aspirin	Aspirin	No Headache
I	No Headache	No Headache	None	No Aspirin	No Headache

- We should also not simply compare a before-after effect of the same individual (not the same unit)

Perfect doctor example

- Perfect doctor chooses the better treatment for each patient, i.e. the treatment under which the patient will live longer

The underlying fact: 0 (drug) v.s. 1(surgery)

Unit	Potential Outcomes		Causal Effect
	$Y_i(0)$	$Y_i(1)$	$Y_i(1) - Y_i(0)$
Patient #1	1	7	6
Patient #2	6	5	-1
Patient #3	1	5	4
Patient #4	8	7	-1
Average	4	6	2

Average causal effect: 2

- What is wrong?

The observed data with perfect doctors

Unit i	Treatment W_i	Observed Outcome Y_i^{obs}
Patient #1	1	7
Patient #2	0	6
Patient #3	1	5
Patient #4	0	8

Observed mean difference between two groups: -1

Perfect doctor example

Unit	W_i	$Y_i(0)$	$Y_i(1)$
1	1	?	7
2	0	6	?
3	1	?	5
4	0	8	?

Unit	W_i	$Y_i(0)$	$Y_i(1)$
1	1	7	7
2	0	6	6
3	1	7	5
4	0	8	6

- In the previous analysis, we implicitly impute the missing values by the group mean
- This is NOT reasonable for perfect doctors. We at least has the constraint that the missing potential outcomes (counterfactuals) should be no larger than the observed outcome
- More explicit answers relies on making assumptions

Experimental and observational studies

Two common types of studies

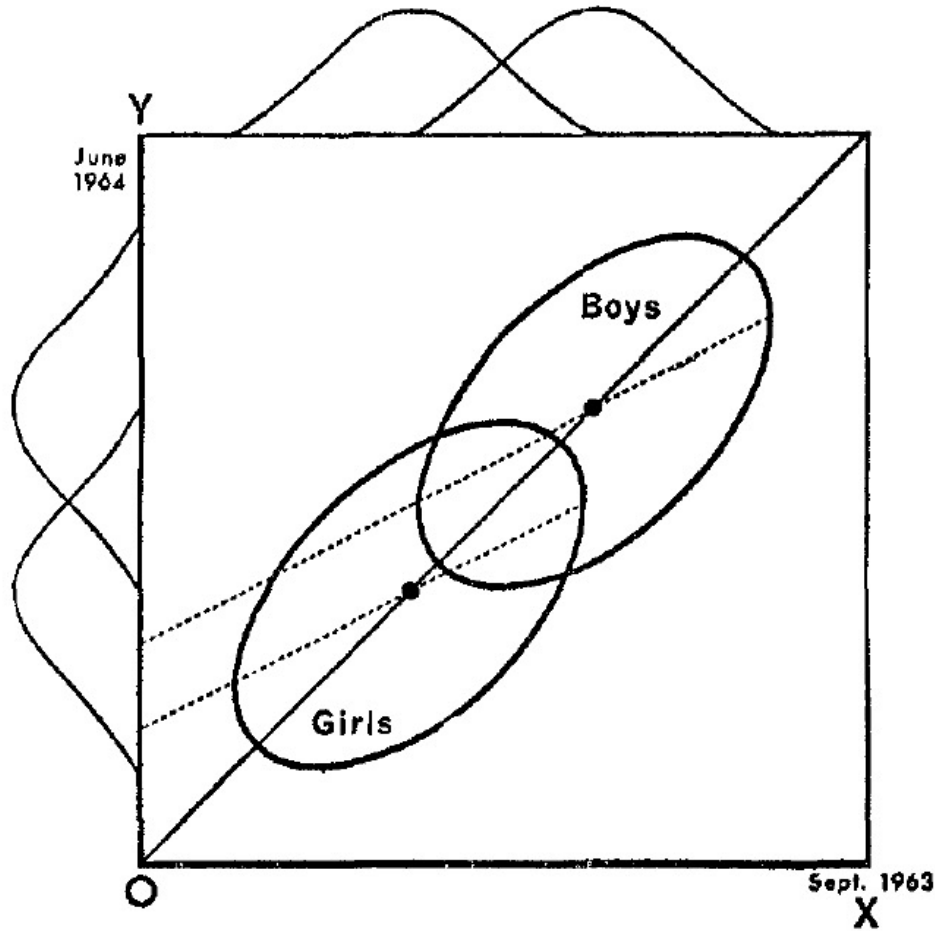
- **Randomized experiment:** the assignment probability $P(W|X)$ is known
 - Assignment can be completely random or depend on pre-treatment covariates of the individuals
 - Laboratory experiments, survey experiments, field experiments
- **Observational studies:** exact assignment probabilities may be unknown to the researcher
 - The researcher still has substantial information about the assignment mechanism
For instance, the researchers knows what X includes. In medical decisions in some situations are solely based on patients' medical records.
 - No intervention is actually performed
- Tradeoff between internal and external validity

A brief history of the potential outcome approach

- Neyman (1923)
 - Define estimands in randomized experiment as functions of potential outcomes
 - While Neyman and others for half century restricted use of potential outcomes to randomized experiments
- Fisher (1925)
 - Proposing the necessity of physical randomization for credibly assessing causal effect
- Rubin (1974, 1975, 1977, 1978)
 - Define causal estimates with potential outcomes in all situations, not just randomized experiments
 - Discuss the assignment mechanism more extensively

Lord's paradox

[Lord 1967, 1969, 1975]



- A large university is interested in investigating the effects on the students of the diet provided in the university dining halls and any sex difference in these effects.
- The weight of each student at the time of his arrival in September (X) and his weight the following June (Y) are recorded.
- What is the differential causal effect of the diet on male and female weights

FIG. 1. Hypothetical scatterplots showing initial and final weight for boys and for girls.

Ellipses represent separate scatter plots

Lord's paradox

[Lord 1967, 1969, 1975]

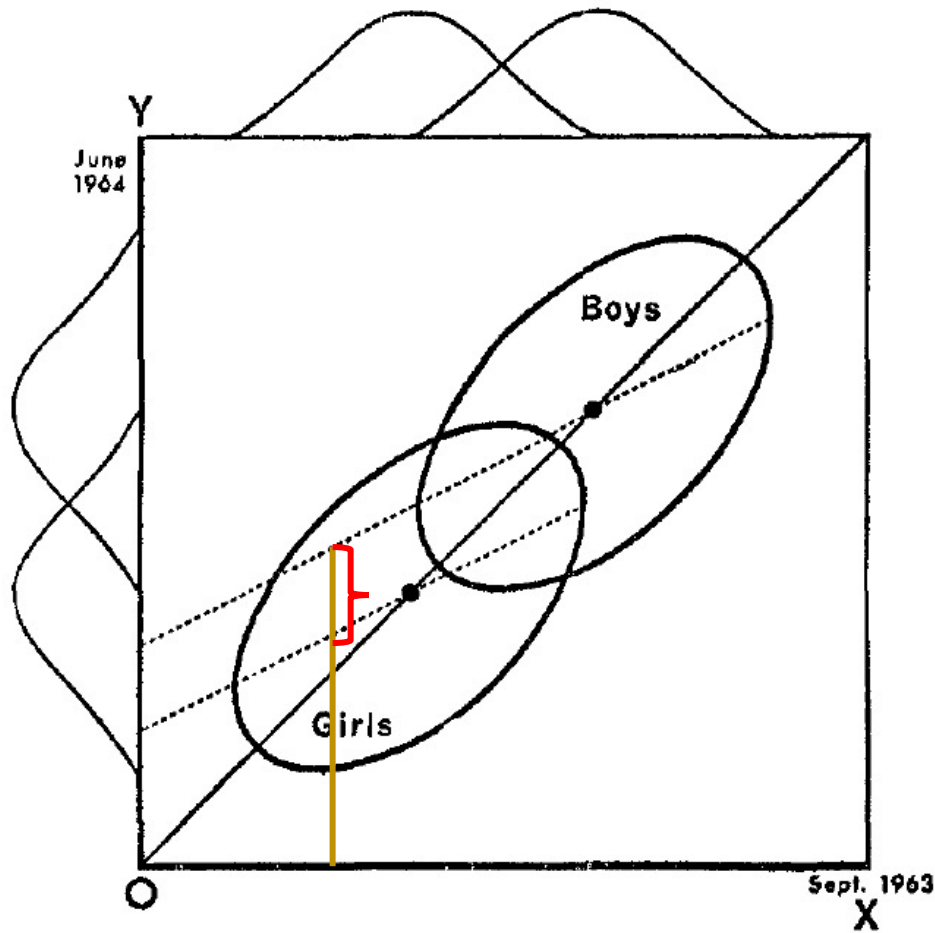


FIG. 1. Hypothetical scatterplots showing initial and final weight for boys and for girls.

Ellipses represent separate scatter plots

- **First statistician**
 - Distribution of weight of the girls at the beginning of the year and at the end of the year are identical. Same for the boys.
 - no evidence of any interesting effect of the school diet on student weight. No evidence of any differential effect on the two sexes
- **Second statistician**
 - Linear regression of Y on both X and gender and coefficient of gender > 0
 - If initial weight a subgroup of boys and a subgroup of girls having identical, then the subgroup of boys is going to gain substantially more than the subgroup of girls.
- **Which statistician is correct?**

An explanation using potential outcomes

[Holland and Rubin, 1983]

- X_i : the initial weight in September
- $(Y_i(0), Y_i(1))$: potential weight in June if the student does not follow or follows the university's diet
- $G_i = g$: gender $g = 0, 1$
- Average effect of the university's diet: $\tau_g = E[Y_i(1) - Y_i(0) \mid G_i = g]$
- The effect difference between female and male: $\Delta = \tau_1 - \tau_0$
- Key issue leading to Lord paradox : $Y_i = Y_i(1)$ for all students, $Y_i(0)$ is never observed

An explanation using potential outcomes

[Holland and Rubin, 1983]

Unit	Sex (F)	September Weight (X_w)	Control: $Y(0)$	Treatment: $Y(1)$
1	0	170	?	160
2	0	180	?	200
3	0	175	?	175
4	0	160	?	160
5	1	120	?	115
6	1	130	?	125
7	1	115	?	120
8	1	125	?	130

- Either statistician can be correct, depending on the assumption
- Causal inference is a missing data problem, not a regression problem involving only observed values

Statistician 1:

If assuming $Y_i(0) = X_i$, then

$$Y_i(1) - Y_i(0) = Y_i - X_i$$

$$\Delta = 0 - 0 = 0$$

Statistician 2:

From the data we know $Y(1) = Y = \alpha_g + bX + \varepsilon_1$

If we assume $Y(0) = a + bX + \varepsilon_0$,

then $Y(1) - Y(0) = \alpha_g - a + \varepsilon_1 - \varepsilon_0$


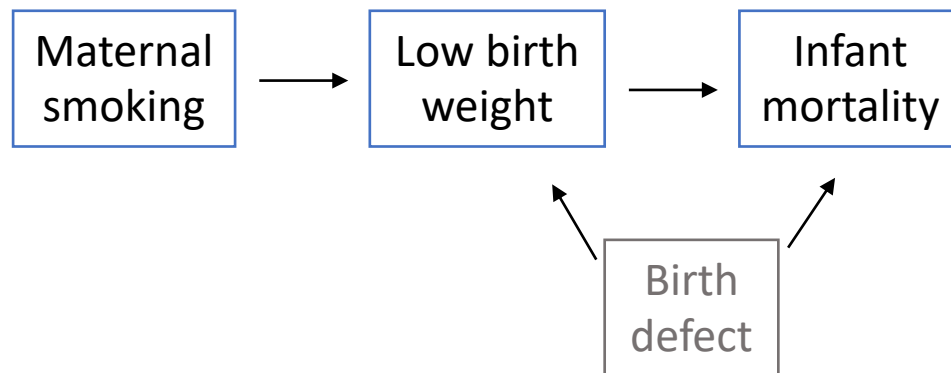
$$\Delta = \alpha_1 - \alpha_0 > 0$$

Causal directed acyclic graphs

- This course focuses on using the potential outcome language, but there is another language to quantitatively describe causal inference
- Causal DAG: a directed graph with no cycles and the arrows has a causal meaning
- Nodes are random variables
- PA_j : parents of X_j (nodes that have direct arrows to X_j)


$PA \rightarrow Y$
parent descendant

Examples:




What are implicitly assumed in a DAG?

- No directed cycle:




A can not cause itself

- The common causes of any pair of variables in the graph must be in the graph (either observed or unobserved)



- SUTVA: one version of treatment and no interference
- Intervention can be done on any node that has an arrow out

Connection with the potential outcome framework



For each X_j , assume existence of random error $E_{X_j} \perp PA_j$ and deterministic unknown function f_j for each node j

$$X_j(pa_j) = f_j(pa_j, E_{X_j})$$

Potential
outcomes for
joint intervention
on parents PA_j

Non-parametric
structural equation

Potential outcomes:


$$A_2(a_1) = f_2(a_1, E_{A_2})$$

$$Y(a_1, a_2) = f(a_1, a_2, E_Y)$$

$$Y(a_1) = f(a_1, A_2(a_1), E_Y)$$

Observed outcome:

$$Y = f(A_1, A_2, E_Y)$$



Comparison between the two languages

- **One reference:**
Imbens, G. W. (2020). Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature*, 58(4), 1129-79.
- As a graphical approach, DAG is superior in illustrating the causal relationships in a complex model and in clarifying some key assumptions
- With DAG, we can perform causal network discovery under additional assumptions
- DAG has more difficulties to capture and represent individual level heterogeneity
- Formal identification assumption for the causal estimand can be clearer using the potential outcome language
- In DAG, all variables are doable and the literature is silent about experiments

Reference papers to read

Causal effects of immutable Characteristics:

- Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American economic review*, 94(4), 991-1013.
<https://www.aeaweb.org/articles?id=10.1257/0002828042002561>
- Chattopadhyay, R., & Duflo, E. (2004). Women as policy makers: Evidence from a randomized policy experiment in India. *Econometrica*, 72(5), 1409-1443.
https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2004.00539.x?casa_token=1SVxHRxC5hQAAAAA:fz62pJ2QMeg41laiRMIERS5K1psbJVovA3-fgD5OTPjb2C-UkLuiQvv0CV2HSFq9z7r3WxxmH_p7NqE
- Arcidiacono, P., Kinsler, J., & Ransom, T. (2022). Asian American discrimination in Harvard admissions. *European Economic Review*, 104079. <https://www.iza.org/publications/dp/13172/asian-american-discrimination-in-harvard-admissions>

Lord's paradox:

- Holland, P. W., & Rubin, D. B. (1983). On Lord's paradox. *Principals of modern psychological measurement*, 3-25.
<https://onlinelibrary.wiley.com/doi/abs/10.1002/j.2333-8504.1982.tb01321.x>

Comparison between DAG and potential outcome framework:

- Imbens, G. W. (2020). Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature*, 58(4), 1129-79.
<https://www.aeaweb.org/articles?id=10.1257/jel.20191597>