Causal Inference II: Causal Graphs and Estimating Effects

In Causal Inference I we saw that the naive difference in means equals the true effect *plus selection bias*, and that **randomization** removes that bias. But we often cannot run an experiment:

Lecture Overview

Drawing our assumptions: causal graphs

A **causal graph** is a *directed acyclic graph* (DAG): nodes are variables, and an arrow $X \to Y$ means “X is a direct cause of Y” (Pearl et al. 2016).

Reading a causal graph

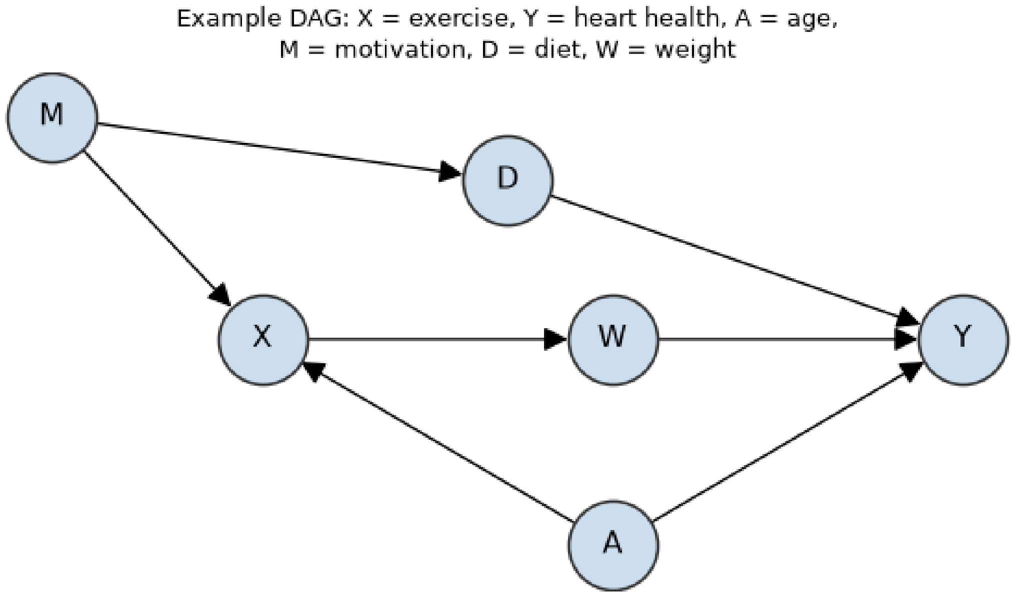

The three building blocks

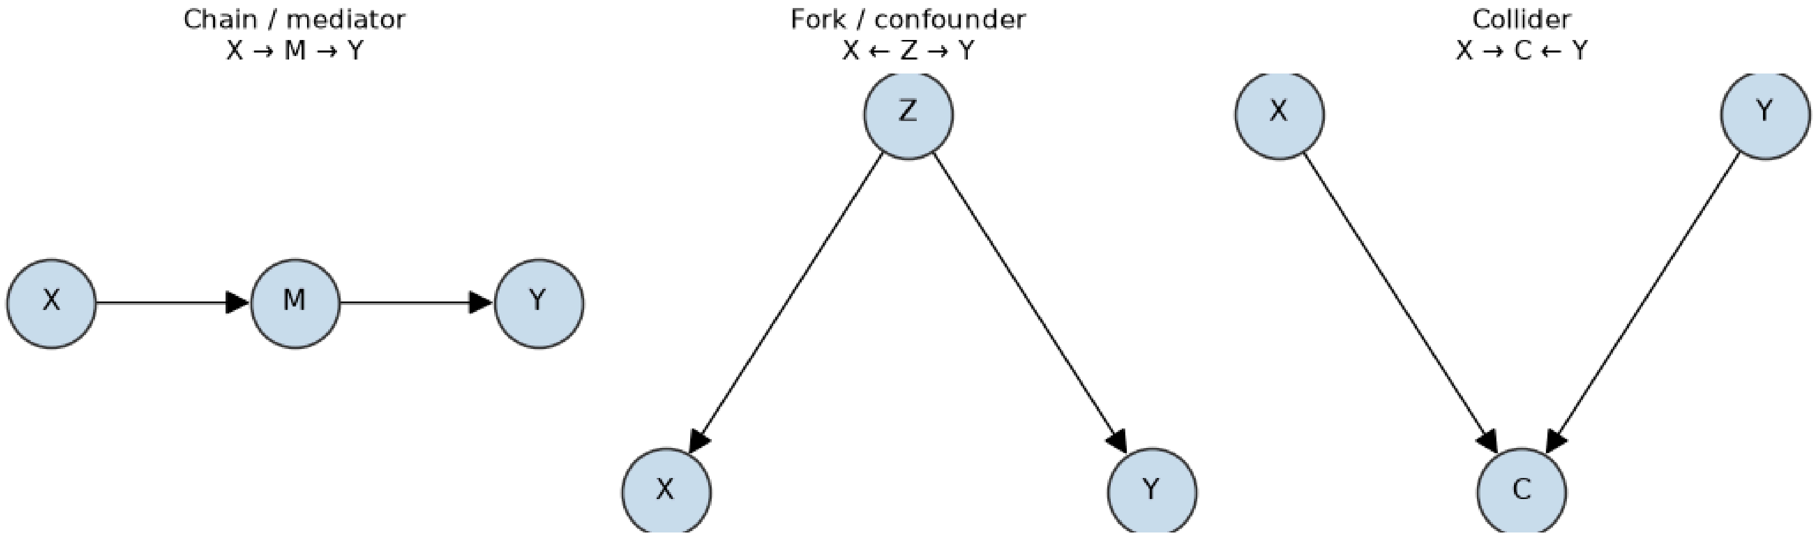

Open and blocked paths

A path carries association between its endpoints only if it is **open**. Two rules decide:

The golden rule: adjust for forks, not colliders

The three structures behave *oppositely* when you condition (adjust/stratify) on the middle node:

Structure	Middle node	Adjust for it?	Why
Fork $X \leftarrow Z \rightarrow Y$	confounder	**Yes**	blocks the spurious backdoor path
Chain $X \rightarrow M \rightarrow Y$	mediator	**No***	it's part of the effect you want
Collider $X \rightarrow C \leftarrow Y$	collider	**No!**	conditioning *creates* spurious association

* Adjust for a mediator only if you specifically want the *direct* (not total) effect.

Confounders: the backdoor

A **backdoor path** is a non-causal path from X to Y that starts with an arrow *into* X (e.g. $X \leftarrow Z \rightarrow Y$). Backdoor paths are the graph-language for confounding.

The backdoor procedure, by hand

Given a DAG, a treatment X and an outcome Y:

Worked example

Back to the exercise DAG (X = exercise, Y = heart health):

$$A \to X,\ A \to Y,\ M \to X,\ M \to D,\ D \to Y,\ X \to W,\ W \to Y$$

Your turn: pick the adjustment set

Tutoring (X) and final grade (Y), with: $\mathrm{GPA}_{\mathrm{prior}} \to X, \mathrm{GPA}_{\mathrm{prior}} \to Y$, $X \to \mathrm{Hours} \to Y$, and $X \to \mathrm{HonorRoll} \leftarrow Y$.

Which of these adjustment sets are valid for the **total** effect of X on Y?

Colliders: adjusting can *create* bias

Conditioning on a collider (or its descendant) opens a spurious path. Classic example: suppose **talent** and **looks** are independent in the general population, but both help an actor get **famous**.

▶ Code

```
corr(talent, looks) overall        = +0.00
corr(talent, looks) among famous    = -0.55
```

Mediators: don’t “control away” your effect

If a job-training program raises earnings *by* teaching skills:

$$\text{Training} \to \text{Skills} \to \text{Earnings}$$

Estimating effects under “no unmeasured confounders”

Suppose we’ve drawn the DAG and found an adjustment set S that satisfies the backdoor criterion (the **ignorability / unconfoundedness** assumption). Several estimators then recover the effect:

Regression vs. stratification, side by side

▶ Code

```
True effect          = 3.00
Naive (ignore Z)     = 4.65   <- biased
Adjusted (+ Z)       = 3.00   <- recovers truth
```

Propensity scores in one slide

When S is high-dimensional, matching directly is hard. Rosenbaum & Rubin showed it is enough to match/weight on the scalar **propensity score** $e(S) = P(T{=}1 \mid S)$.

▶ Code

```
IPW estimate of ATE = 3.03   (true = 3.0)
```

When confounders are *unobserved*: natural experiments

Adjustment only works for confounders we can **measure**. When we can't, we look for "nature's randomization" — sources of variation in treatment that are *as good as random* (Angrist and Pischke 2014).

Three workhorse designs:

Difference-in-Differences

Compare the **change** in a treated group to the **change** in a control group over time. This differences out any *fixed* differences between the groups.

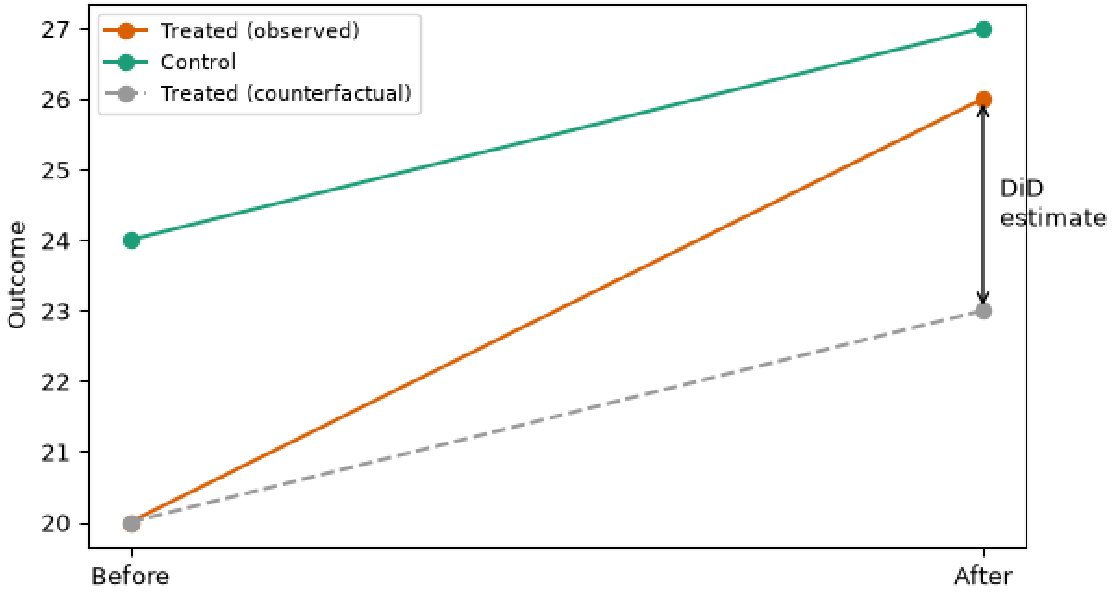

Instrumental variables & regression discontinuity

Instrumental variable (IV)

An **instrument** Z affects the outcome *only through* the treatment:

$$Z \to T \to Y, \quad Z \not\to Y \text{ directly.}$$

- Vietnam **draft lottery** as an instrument for military service.
- **Distance to college** as an instrument for years of schooling.

Z acts like the coin flip we wish we had run.

Regression discontinuity (RD)

Treatment switches on at a **cutoff** of a running variable. Units just above vs. just below are comparable.

- Scholarship awarded at a **test-score threshold**.
- Class-size caps that trigger a new section.

Compare outcomes in a narrow window around the cutoff.

Assumptions you can't skip

Every observational causal estimate relies on assumptions (Hernán and Robins 2020):

Tools and the modern workflow

Summary

Small-Group Activity

🎯 Activity: Be the causal analyst (≈30 min)

Format: groups of 3–4. **Time:** ~20 min of group work + ~10 min report-out.

You will get a scenario. For it, your group must produce four deliverables:

Choose a scenario

Pick **one** (or your instructor will assign one):

A. Product / tech

Does adding a **“free shipping” badge** on product pages increase **purchases**? You have historical logs; badges were shown more often on popular items.

B. Health

Does a new **fitness-tracking app** reduce users’ **resting heart rate**? People who download it are already more health-conscious.

C. Education

Does an **optional tutoring program** raise **final grades**? Students self-select into tutoring, and struggling students are more likely to sign up.

D. Policy

Does a city’s **new bike-lane network** reduce **traffic accidents**? It was built in denser, wealthier neighborhoods first.

Report out

Each group takes **~2 minutes** to share:

References

- *Causal Inference: What If* — Hernán & Robins (free PDF)
- *Causal Inference: A Primer* — Pearl, Glymour & Jewell
- "A Crash Course in Good and Bad Controls" — Cinelli, Forney & Pearl
- DoWhy documentation and dagitty.net

Bibliography

Angrist, Joshua D., and Jörn-Steffen Pischke. 2014. *Mastering 'Metrics: The Path from Cause to Effect*. Princeton University Press.

Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *American Economic Review* 84 (4): 772–93.

Cinelli, Carlos, Andrew Forney, and Judea Pearl. 2024. "A Crash Course in Good and Bad Controls." *Sociological Methods & Research* 53: 1071–104. https://doi.org/10.1177/00491241221099552.

Hernán, Miguel A., and James M. Robins. 2020. *Causal Inference: What If*. Chapman & Hall/CRC. https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/.

Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge University Press.

Pearl, Judea, Madelyn Glymour, and Nicholas P. Jewell. 2016. *Causal Inference in Statistics: A Primer*. Wiley.

Sharma, Amit, and Emre Kiciman. 2020. *DoWhy: An End-to-End Library for Causal Inference*. https://arxiv.org/abs/2011.04216.