

Stat 256: Causality

Introduction

Instructor: Chad Hazlett

UCLA

Spring 2026

Why causality?

Most interesting questions in science and life are causal.

- Does this drug / policy / event cause higher or lower Y ?
- What explains why Y happened?

Why causality?

Most interesting questions in science and life are causal.

- Does this drug / policy / event cause higher or lower Y ?
- What explains why Y happened?

And yet:

Why causality?

Most interesting questions in science and life are causal.

- Does this drug / policy / event cause higher or lower Y ?
- What explains why Y happened?

And yet:

- There is no causal inference without assumptions that are fundamentally unverifiable from data alone.

Why causality?

Most interesting questions in science and life are causal.

- Does this drug / policy / event cause higher or lower Y ?
- What explains why Y happened?

And yet:

- There is no causal inference without assumptions that are fundamentally unverifiable from data alone.
- Traditional statistics does not include a coherent language for causal effects, much less tools.

Why causality?

Most interesting questions in science and life are causal.

- Does this drug / policy / event cause higher or lower Y ?
- What explains why Y happened?

And yet:

- There is no causal inference without assumptions that are fundamentally unverifiable from data alone.
- Traditional statistics does not include a coherent language for causal effects, much less tools.
- Even “descriptive” analyses involve missingness, case selection, and generalization—problems rooted in causal processes.

Why causality?

Most interesting questions in science and life are causal.

- Does this drug / policy / event cause higher or lower Y ?
- What explains why Y happened?

And yet:

- There is no causal inference without assumptions that are fundamentally unverifiable from data alone.
- Traditional statistics does not include a coherent language for causal effects, much less tools.
- Even “descriptive” analyses involve missingness, case selection, and generalization—problems rooted in causal processes.
- Knowing what you can and cannot say, and protecting against abuse of your results, is vital to being an honest statistician.

Goals

- Equip you with tools for careful causal thinking

Goals

- Equip you with tools for careful causal thinking
- Equip you to read, critique, and (prepare to) write methodological work on causal inference

Goals

- Equip you with tools for careful causal thinking
- Equip you to read, critique, and (prepare to) write methodological work on causal inference
- Provide you with tools for applied research and the wisdom to
 - avoid abusing these tools
 - carefully assess the credibility of claims made by others

Goals

- Equip you with tools for careful causal thinking
- Equip you to read, critique, and (prepare to) write methodological work on causal inference
- Provide you with tools for applied research and the wisdom to
 - avoid abusing these tools
 - carefully assess the credibility of claims made by others
- What we focus *less* on:
 - some of the more esoteric theoretical directions
 - estimation horse-races while deeper problems go ignored

Theme 1: Bilingualism

Theme 1: Bilingualism

Two frameworks dominate work on causal inference:

Theme 1: Bilingualism

Two frameworks dominate work on causal inference:

1. **Structural causal models (SCMs) / directed acyclic graphs (DAGs)**

- Encode which variables cause which; derive what you can learn from observational data.
- Largely attributed to UCLA's own Judea Pearl.

Theme 1: Bilingualism

Two frameworks dominate work on causal inference:

1. **Structural causal models (SCMs) / directed acyclic graphs (DAGs)**

- Encode which variables cause which; derive what you can learn from observational data.
- Largely attributed to UCLA's own Judea Pearl.

2. **Potential outcomes (POs)**

- Each unit has potential outcomes under each treatment; write these down and make assumptions about how they relate to treatment assignment.
- Associated with Jerzy Neyman and Don Rubin.

Theme 1: Bilingualism

Two frameworks dominate work on causal inference:

1. **Structural causal models (SCMs) / directed acyclic graphs (DAGs)**

- Encode which variables cause which; derive what you can learn from observational data.
- Largely attributed to UCLA's own Judea Pearl.

2. **Potential outcomes (POs)**

- Each unit has potential outcomes under each treatment; write these down and make assumptions about how they relate to treatment assignment.
- Associated with Jerzy Neyman and Don Rubin.

These have sometimes been treated as rival camps, but that need not be the case. *We will develop fluency in both.*

Theme 2: Feasibility

- Great theoretical progress has been made in causal inference.

Theme 2: Feasibility

- Great theoretical progress has been made in causal inference.
- But in theory you can assert an assumption whether it is credible or not.

Theme 2: Feasibility

- Great theoretical progress has been made in causal inference.
- But in theory you can assert an assumption whether it is credible or not.
- We care about how investigators can *make or defend* causal claims in practice.

Theme 2: Feasibility

- Great theoretical progress has been made in causal inference.
- But in theory you can assert an assumption whether it is credible or not.
- We care about how investigators can *make or defend* causal claims in practice.
- “Identification strategies” link causal quantities to observable ones. For each, we ask:
 - What do the assumptions mean? What about nature or the research design supports them? How could you falsify them?
 - How do we avoid over-confidence built on unreliable assumptions?
 - How do we make progress when sharp assumptions aren’t defensible?

Road map: Part I — Core language and theory

Unit 1: DAGs and structural causal models

- Graphs, d-separation, back-door and front-door criteria, do-calculus
- Where do potential outcomes come from?

Unit 2: Potential outcomes

- Definitions, estimands, randomized experiments
- Conditional ignorability / selection on observables
- Bridging POs and DAGs

Road map: Part II — Identification strategies

Unit 3: Selection on observables in practice

- Sub-classification, matching, regression, propensity scores, balancing weights
- Brief survey of double-robust / ML-augmented approaches
- Sensitivity analysis for SOO

Unit 4: Instrumental variables

- IV in PO and DAG frameworks
- Practical applications and critiques, sensitivity analysis

Unit 5: Regression discontinuity

- Sharp and fuzzy RD

Unit 6: Difference-in-differences

- Parallel trends; DID $\neq$ TWFE; staggered adoption; sensitivity

Road map: Part II (continued)

Unit 7: Time-varying confounding / panel methods

- Synthetic control, LDV, interactive fixed effects, relaxing parallel trends

Unit 8: Sensitivity analysis and partial identification

- OVB for bounding; sensitivity for IV and other strategies
- Placebo treatments and outcomes; stability-controlled quasi-experiments

Unit 9: Mediation and TRACE

- Why mediation is fundamentally harder; direct and indirect effects
- Marginal structural models; separating units rather than paths
- The TRACE: separating units rather than effects to address post-treatment problems

Road map: Part III — You teach!

- Groups of 2–3 choose an advanced topic
- 25-minute lecture + 5 min Q&A
- Submit a tutorial with worked examples and problems

Road map: Part III — You teach!

- Groups of 2–3 choose an advanced topic
- 25-minute lecture + 5 min Q&A
- Submit a tutorial with worked examples and problems

Example topics:

- Causal inference in networks / interference
- Causal attribution, necessity, sufficiency (causes of effects)
- Generalization, transportability, external validity
- Advanced heterogeneous effect estimation (CATE, causal forests)
- Negative outcome controls, placebos, falsification
- Other ID strategies not covered in class
- *Custom topics* (by permission)

More on presentations

Why do this?

- Mimics the experience of a smaller advanced seminar
- Learning something well enough to present it produces far better understanding
- You are well situated to explain new material to your peers

More on presentations

Why do this?

- Mimics the experience of a smaller advanced seminar
- Learning something well enough to present it produces far better understanding
- You are well situated to explain new material to your peers

Common pitfalls:

- Not enough attention to what the identification assumptions really are, how to understand them, what it would take to believe them.
- Splitting up the topic so that nobody understands the big picture.

Requirements and grading

- **Problem sets (25%):** Expect 4–5 over the quarter. Posted and submitted online. Ask for extensions *before* the deadline.

Requirements and grading

- **Problem sets (25%):** Expect 4–5 over the quarter. Posted and submitted online. Ask for extensions *before* the deadline.
- **Advanced topic presentation (25%):** Groups of 2–3; lecture + tutorial.

Requirements and grading

- **Problem sets (25%):** Expect 4–5 over the quarter. Posted and submitted online. Ask for extensions *before* the deadline.
- **Advanced topic presentation (25%):** Groups of 2–3; lecture + tutorial.
- **Critical review memo (25%):** Individual. Either critique an empirical paper's identification strategy, or take a deep dive into a methods paper (simulations, failure modes, extensions). Due Week 10.

Requirements and grading

- **Problem sets (25%):** Expect 4–5 over the quarter. Posted and submitted online. Ask for extensions *before* the deadline.
- **Advanced topic presentation (25%):** Groups of 2–3; lecture + tutorial.
- **Critical review memo (25%):** Individual. Either critique an empirical paper's identification strategy, or take a deep dive into a methods paper (simulations, failure modes, extensions). Due Week 10.
- **Active-learning activities (25%):** Near-daily in-class exercises to rehearse and reinforce the material. Graded lightly (0/1/2/3). Lowest two dropped. No make-ups.

Timeline (preliminary)

- **Week 3:** Sign up for presentation groups and topics.
- **Week 6:** Check-in meeting with instructor on presentation and memo.
- **Weeks 9–10:** Student-led lectures.
- **Week 10:** Critical review memo due.
- **Finals week:** All presentation materials due.

Prerequisites

We want to accommodate various fields and training backgrounds, but the skills we count on are:

- Essential: probability theory, statistics including inference and regression.
- Useful: linear algebra, multivariate calculus.
- Essential: enough R for AI to help you, and a way to write in $\text{T}_{\text{E}}\text{X}$ (e.g. Rmarkdown).

Books

We will share references and reading expectations as needed.

We will share references and reading expectations as needed.

For Unit 1 on DAGs (starting Week 1) you should have:

- Pearl, J., Glymour, M., & Jewell, N. P. (2016). *Causal Inference in Statistics: A Primer*. Wiley.
- Hernán, M.A. & Robins, J.M. (2024). *Causal Inference: What If*. Chapman & Hall/CRC. (Free online)

We will share references and reading expectations as needed.

For Unit 1 on DAGs (starting Week 1) you should have:

- Pearl, J., Glymour, M., & Jewell, N. P. (2016). *Causal Inference in Statistics: A Primer*. Wiley.
- Hernán, M.A. & Robins, J.M. (2024). *Causal Inference: What If*. Chapman & Hall/CRC. (Free online)

Other useful references:

- Pearl, J. (2009). *Causality*. Cambridge University Press. 2nd ed.
- Cunningham, S. (2021). *Causal Inference: The Mixtape*. (Free online)
- Huntington-Klein, N. (2022). *The Effect*. (Free online)

AI tools

You are welcome to use AI tools (ChatGPT, Claude, Copilot, etc.) as learning aids.

- Problem set write-ups must reflect *your own* understanding.
- Your critical review memo must be your own analytical work—but discussing ideas with AI to develop your thinking is fine and encouraged.
- If you use AI substantively, briefly note how.

The goal is to learn the material. AI can help with that if you use it as leverage rather than a way out.

Other logistics

- Lectures: Monday, Wednesday, 11:00–12:15, Haines A44. In-person only.
- Slides will be posted on the course website.
- My office hours: Wednesday 3:30–4:45, Bunche 3264

Questions?

Introductions

- Your name
- Your department / program / year
- What brings you to this course?
- One causal question you find interesting (any domain)