

Stat 256: Causality

Spring 2026

Professor: Chad Hazlett

chazlett@ucla.edu

Class: Monday, Wednesday, 11:00–12:15, Haines A44

Office Hours: Wednesday 3:30–4:45, Bunche 3264

1 Goals

Very often, the goal of quantitative research is to establish or measure the causal effect of one variable on another. In many circumstances we cannot randomize which units receive a “treatment” of interest, complicating efforts to learn the effect of that treatment. Even when the treatment variable can be randomly manipulated (as in a randomized experiment), this does not enable us to answer all questions of interest, such as those relating to mechanism and mediation problems. Other times, even putatively “descriptive” statistical efforts involve problems such as case selection, missingness, attrition, and generalization—all of which actually require causal assumptions.

Fortunately, recent decades have seen an explosion of work on causality and turned what was previously the realm of philosophical debate into a mathematically tractable language (or really, two languages—more on that momentarily). Our goal is not to train you in every aspect of this quickly evolving area, but to provide you with a working understanding of the field, particularly as it applies to the challenges of real applied research. Specifically, this course aims to:

- equip you with tools for careful causal thinking,
- equip you to read causal inference papers on your own and to critique the conclusions they make, and
- provide you with tools for applied research, and the wisdom required to (i) avoid abusing these tools and (ii) carefully assess the credibility of claims made by other researchers in their applied work.

2 Approach

Our approach to this material focuses on two principles: bilingualism, and feasibility.

Bilingualism. Two main “languages” of causal reasoning dominate academic work on causality:

- An approach rooted in “structural causal models” (SCMs), often represented graphically and particularly with “directed acyclic graphs” (DAGs). This is most closely identified with UCLA’s own Judea Pearl.¹
- The “potential outcomes” (PO) language, built mainly on initial work by Jerzy Neyman as championed and elaborated by Don Rubin and many others.

This course takes the stance that both approaches are valuable and engage mostly the same underlying commitments, just emphasizing different perspectives. POs can be seen as arising from SCMs/DAGs. The

¹As well as foundational work by Peter Spirtes, Clark Glymour, Jin Tian, Thomas Verma, Jamie Robins, Sander Greenland (also at UCLA), Tyler VanderWeele, Miguel Hernan, and others.

causal estimands in which we are interested are often easy to write down using POs, as are the assumptions we must make in order to estimate these from data. DAGs are particularly effective at shedding light on what it would take to believe causal assumptions, and at showing you how they might be violated either by nature or by choices the analyst made. In some cases, POs naturally lead to identification ideas that are non-obvious from the DAG perspective, particularly where parametric restrictions are involved; other times it is the other way around. *A goal of this course is to build fluency in both languages to maximize our ability to see opportunities for making causal inferences, and to avoid errors in causal reasoning.*

Feasibility. A second major theme is a focus on practicality: (i) we focus on those aspects of causal inference that are useful to applied researchers being transparent and careful about what they can and cannot claim; and (ii) the credibility of inferences made using these approaches comes down to the assumptions on which they rest, so understanding those assumptions and what it would mean to believe them is of paramount importance. Practically, this means:

- Examining commonly relied-upon identification strategies—“clever ideas” for making causal claims from observational data. For each we must ask how one would justify or evaluate the credibility of the required assumptions, through both PO and DAG lenses.
- A focus on **partial identification and sensitivity analysis**: the assumptions required to point identify a causal quantity may not be defensible, but applied researchers should characterize how results might change under violations and what we must be prepared to believe in order to arrive at a particular causal conclusion.
- This course is *not* a deep dive into all frontier areas of causal inference: some areas in which recent advances have been made may remain too distant from empirical practice to be suitable here. But it will give you the notation and toolkit to understand and critically assess those proposals.

3 Organization

Flexibility and updating. *Everything below is subject to change.* If you have a good idea for something different to cover or a different type of assignment, we have the flexibility to consider it. This is still a relatively new class, and your input shapes what it will look like in the future.

The course website is definitive. When in doubt, trust the information, deadlines, etc. on the course website.

Materials. As we come to each unit, we will share the relevant references and reading expectations via the course website. *Often there will not be readings. When there are, we keep them minimal, so please do them.*

You should obtain a copy of:

- Pearl, J., Glymour, M., & Jewell, N. P. (2016). *Causal Inference in Statistics: A Primer*. John Wiley & Sons.

The following is freely available online and will be used throughout:

- Hernán, M.A. & Robins, J.M. (2024). *Causal Inference: What If*. Chapman & Hall/CRC. Available at: <https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/>

Additional recommended references (not required):

- Imbens, G.W. & Rubin, D.B. (2015). *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press.

- Cunningham, S. (2021). *Causal Inference: The Mixtape*. Yale University Press. Freely available at <https://mixtape.scunning.com/>.
- Huntington-Klein, N. (2022). *The Effect: An Introduction to Research Design and Causality*. Chapman & Hall/CRC. Freely available at <https://theeffectbook.net/>.

4 Prerequisites

Students are expected to have proficiency in basic R coding, probability theory, linear algebra, multivariate calculus, and statistics up through inference and regression. For students with strong mathematical and statistics backgrounds, this could be a first course in causal inference. For those less sure of these prerequisites, PS 200C would be a good introduction to both the mathematical requirements and causal inference.

5 Course Outline

Part I: Core Language and Theory

Unit 1: The graphical (DAG) approach.

- Graph basics: nodes, edges, and missing edges.
- Relation to structural causal models (SCMs).
- D-separation and conditional independence.
- Back-door criterion and adjustment.
- Front-door criterion.
- The *do*-calculus.
- Origins of potential outcomes from SCMs.

Primary reference: Pearl, Glymour, & Jewell (2016), *Primer*.

Unit 2: The potential outcomes (PO) approach.

- Basic definitions, potential outcomes, and causal estimands.
- Randomized experiments: identification, estimation, variance.
- Conditional ignorability (unconfoundedness) and selection on observables.
- Bridging POs and DAGs: equivalences and complementarities.

Part II: Identification Strategies

Unit 3: Selection on observables (SOO) in practice. The most commonly invoked identification strategy amounts to some version of assuming no unobserved confounders and then adjusting for observed covariates. The number of approaches for doing this has exploded. They differ only in *how* they approximate statistical conditioning, not in their causal sophistication. We cover the main families:

- Sub-classification, matching, and regression
- Propensity score methods (matching, weighting)

- Covariate-balancing weights
- Brief survey of combined/double-robust approaches (augmented IPW, double-ML, TMLE)
- Sensitivity analysis

Unit 4: Instrumental variables

- IV assumptions in PO and DAG frameworks
- LATE and complier interpretation
- Sensitivity analysis

Unit 5: Regression discontinuity.

- Sharp and fuzzy RD
- Bandwidth selection, local polynomial estimation
- Diagnostics: manipulation testing, covariate balance
- Generalizability concerns

Unit 6: Difference-in-differences/ no time-varying confounding)

- The canonical two-period, two-group case
- Parallel trends: meaning, plausibility, falsification efforts
- DID $\neq$ two-way fixed effects. Staggered adoption and heterogeneous treatment timing: why classical TWFE fails and modern alternatives (Callaway & Sant'Anna; Sun & Abraham; de Chaisemartin & D'Haultfoeuille), event studies and pre-trend testing
- Sensitivity analysis

Unit 7: Time-varying confounding

- Synthetic control
- Interactive fixed effects and related approaches
- When parallel trends is too strong: relaxing it

Unit 8: Ideas in safe inference and partial identification.

- Stability-controlled quasi-experiments (SCQE).
- Placebo-based adjustment and negative controls.
- Bounds and partial identification more broadly.

Unit 9: Mediation and TRACE

- Why mediation is a fundamentally harder problem.
- Direct and indirect effects: definitions and assumptions.
- DAG and PO perspectives on mediation.
- Marginal structural models
- Separating units rather than paths: the TRACE

Part III: You teach!

In this section, students will work in groups of **2–3** to cover advanced topics, presenting them to the class in the form of developed lectures. See “Presentations” below for details.

6 Class Requirements and Grading

Problem sets (25%). Problem sets are the primary tool for learning. Expect 4–5 over the quarter. Late submissions will not be accepted; if you need an extension, ask **before** the deadline. Students are encouraged to discuss difficulties, but every step of every problem must be produced by the individual student and all work must be written up independently. *Neither code nor written solutions may be copied.* For analytical questions, include intermediate steps and reasoning. For data analysis, include annotated code. Your submission must list any students with whom you discussed the problems.

Advanced topic presentation (25%). Working in groups of 2–3, you will choose a topic from a provided list and deliver a **25-minute lecture plus 5 minutes of Q&A**. You will also submit, by the end of the quarter, a brief tutorial on the topic that includes a clear explanation, worked examples, and 1–2 problem-set-style questions for the class.

Topics will include (but are not limited to):

- Causal inference in networks/interference
- Causal attribution, necessity, sufficiency (causes of effects)
- Generalization, transportability, and external validity
- Advanced heterogeneous effect estimation (CATE, causal forests)
- Negative outcome controls, placebos, and falsification tests
- Bounds and partial identification beyond sensitivity analysis
- Other identification strategies not covered in class

Expect the following deadlines:

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

Critical review memo (25%). (NOTE, THIS IS SUBJECT TO CHANGE/REFINEMENT). An individual written assignment in which you select either (i) a published empirical paper that uses one of the identification strategies covered in the course; or (ii) examine a complicated new method, critique it, offer your proposed refinements, show simulation results, failure modes, or new applications.

In the case of an empirical paper,

- Explain the paper’s identification strategy using both PO and DAG frameworks.
- Articulate the key assumptions and evaluate their credibility.
- Propose a sensitivity analysis or partial identification exercise: what would you want to see to be convinced (or to know the bounds of what we can learn)?
- Offer constructive suggestions for strengthening the causal claims.

In the case of a methods paper,

- Explain the paper’s idea both rigorously and with intuition, giving comparisons to other methods and involving any new insights you’ve acquired. Put it in the landscape of things we learned about in class.
- Articulate the key assumptions and argue whether they are workable, or what situations they might be best suited to.
- Do new simulations, different from any the authors proposed; demonstrate failure modes worth thinking about.
- If you’re lucky and ambitious, propose an extension, improvement, or alternative!

Due **Week 10**. We will discuss suitable papers early in the quarter; you are also welcome to propose your own. We will also have meetings about it along the way, to develop ideas but also as part of the evaluation mechanism.

Active-learning activities (25%). This course is a seminar and your active engagement is essential. To encourage this, facilitate learning, and to motivate (and record) attendance we will have almost daily in-class activities...yes, a little quiz. These give an immediate way to rehearse and reinforce what you just saw. We will do these for my lectures, and this will be the comprehension questions for the student-led questions. In my other classes, I have replaced problem sets with these entirely. In this class however, we need problem sets for more in-depth skill-development...but we need the in-class activities to facilitate learning, engagement, and because it’s the only time to see if you are learning anything without an AI fallback. They will be graded lightly (0, 1 for an attempt that is entirely wrong, 2 for a good attempt, 3 for a great attempt). You can miss two without worrying. There will be no make-up opportunities. If you do not plan to attend regularly, this is not the right course for you!

7 Use of AI Tools

You are welcome to use AI tools (e.g., ChatGPT, Claude, GitHub Copilot) as learning aids—to check your understanding, debug code, or explore ideas. However:

- Problem set write-ups must reflect *your own* understanding. Using AI to generate solutions you do not understand defeats the purpose and will be apparent, believe it or not!
- Your critical review memo must be your own analytical work, but it’s okay (and maybe good) to “discuss” your ideas with AI to help you develop your response and look for holes etc.
- If you use AI tools substantively on any assignment, briefly note how (e.g., “used Claude to debug my R code for Q3”). It’s totally okay, and I want to get a sense of usage.

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