PS200C Causal Inference, Spring 2026

Introduction

Chad Hazlett

UCLA

chazlett@ucla.edu

The stakes are real: Hormone Replacement Therapy

For decades, **40+ observational studies** consistently showed that postmenopausal women on hormone replacement therapy (HRT) had $\sim$30% lower rates of heart disease.

Panels of experts endorsed HRT. Millions of women were prescribed it.

But these were observational comparisons: women who *chose* HRT tended to be wealthier, healthier, and more engaged with the medical system—classic **healthy user bias**.

...what do we really learn from such observational evidence?

HRT: The danger of observational claims

The Women's Health Initiative (2002) randomized $\sim$16,000 women to HRT or placebo. It found a $\sim$29% *increase* in heart attacks and elevated cancer risk.

The trial was **halted early**.

- Physicians reversed course overnight
- Millions of women stopped HRT
- Even now people take guidance from that trial.

Here we go again

That interpretation of the trial result may have been hasty.

The WHI enrolled women who were on average **63 years old**—10–20 years past menopause.

- Later reanalysis showed that for women near menopause, HRT likely *is* protective.
- The "timing hypothesis" is now widely accepted, and guidelines have largely come back around.

Lesson: Even a large RCT doesn't necessarily tell you what you want.

- Design details (the treatment formulation, who is enrolled), stopping rules, interpretation can invite errors
- An experiment answers a causal question, but *which* causal question depends on details you cannot ignore.

More examples to keep in mind

- **Scared Straight programs:** Observational studies suggested exposing at-risk youth to prison life reduced delinquency. Multiple RCTs later showed it made outcomes *worse*. Programs were widely implemented before the evidence caught up.

- **Microfinance:** Celebrated as a path out of poverty based on observational comparisons of borrowers vs. non-borrowers. Six major RCTs found much more modest effects—helpful in some ways, but not transformative.

- **Smoking and lung cancer:** We will never have an RCT. Yet the causal conclusion is among the most trusted in all of science. Why? Not because of the volume of data, but because researchers could formalize *what it would take* for the result to be wrong—and the required confounding was implausible. (We will return to this when we cover sensitivity analysis.)

Modes of Statistical Inference

1. Descriptive Inference: summarizing and exploring data
    - means, correlations, regression coefficients
    - inferring "ideal points" from roll-call votes
    - inferring "topics" from texts and speeches
    - Note: the statistical inference you are used to—hypothesis tests, *p*-values, confidence intervals—is *descriptive* inference. It furthers the description of how variables look or move together. It does not, by itself, get you any closer to causal.
2. Predictive Inference: forecasting out-of-sample data points
    - predicting future state failures from past failures and covariates
    - predicting vote share from a survey of likely-voters
3. Causal Inference: engages *counterfactuals*
    - inferring whether the lack of war among democracies can be attributed to regime type
    - determining whether a medical treatment improved outcomes.
    - generally, estimating effect of some policy or institution or treatment on some outcome...

Is causal inference more difficult?

Version 1: Is it more difficult as a scientific problem?

- Yes.
- Why? Because the assumptions required to say things about unobservable counterfactuals are understandably onerous.

Version 2: Is it more difficult to *learn and do*?

- Yes and no.
- Yes: it is a different way of thinking that may be awkward and magical sounding at first!
- No: the challenges are conceptual, not technical or mathematical.

In fact, I hope the course helps to convince you that technical sophistication of methods often does not solve core problems...it's not about the latest estimator, machine learning, or more complicated standard errors.

Causal inference $\neq$ data analysis

The shift in thinking needed for this course begins by **getting away from the idea that data "speak for themselves".** Data $\neq$ evidence.

Rather, there are causal quantities you'd like to know, but *without further assumptions*, data are largely silent on them, no matter how much you have.

Stylized illustration:

- you have data on X and Y
- you want to learn some effect of X on Y
- you compute something with data, say, $coef(Y \sim X)$
- trust me (for now): *without further assumption*, any causal effect can produce any coefficient...and thus any coefficient you get is equally suggestive of any causal effect.

In other words, "looking at the data" is not a sufficient strategy for learning about the world...

So what *does* it take? That's what this course is about.

Identification: Assumptions + Data $\rightarrow$ Causal Conclusions

Causal inference is about that "further assumption" bit:

- What assumptions let us connect the causal thing we want to know to something in the data?
- Can we identify our quantity of interest with something we can estimate?
- We speak of identification assumptions or identification strategies that provide such bridges.

Key implications:

- Most identifying assumptions cannot be proven correct using data. It's not about "knowing the right test" to run.
- We will always be appealing to information, assumptions, or context from *outside* the data.
- Credibility does NOT derive from "what method you used." There are not "causal methods"! You can't say "it's causal because I did xyz."
- This is one of the hardest things for many students to absorb, in part because we are surrounded by mistakes and false advertising.

You cannot escape causal reasoning

You might think causal inference is only for people making “clearly causal claims,” and only when they can’t run an experiment. Both are wrong.

- Many “comparisons” are only interesting because the audience will read them as suggestive of a causal relationship. If you compare outcomes across groups, you are implicitly in causal territory whether you say so or not.
- Even genuinely descriptive tasks invoke causal reasoning. Handling **missingness** requires assumptions about *why* data are missing—a causal question. Choosing what to **condition on** requires knowing the causal structure, or you risk conditioning on a collider and creating bias that wasn’t there before.
- Randomized experiments don’t make causal reasoning unnecessary—they make it *more* necessary. What **mechanism** produced the total effect? How do you handle **attrition** (which is post-treatment)? How do you choose covariates for **precision** without introducing bias? These all require understanding causal structure.
- Even **generalization**—asking whether a result applies beyond the study sample—is a causal question about what would happen under different conditions.

Questions Answered in This Course

- Why do we accept (certain) results from randomized experiments as causal? Are they so special?
- We often cannot run randomized experiments or want to know the effect of something that "really happened". Can we still make a causal inference?
- How can we understand how vulnerable a given estimate might be to threats such as confounding?
- **How can you conduct top quality work that seeks to make causal claims in tricky situations, while accurately characterizing what can and cannot be claimed?**

Approach for 2026

This year's course reflects a deliberate shift:

- **No traditional problem sets.** AI tools can now handle much of the mechanical work—estimation, coding, debugging—that used to dominate your time outside class.
- Instead, we invest that time where it matters most: **applying causal reasoning to real research problems.**
- Starting early in the quarter, you develop a **final project**: formulate a causal question, identify an appropriate strategy, carry out an empirical analysis.
- In lecture, the emphasis is on **intuition and assumptions**—truly understanding the (heroic) assumptions required by each identification strategy.
- You are expected and encouraged to **use AI tools** (Claude, ChatGPT, etc.) as tutors. We want you to learn to *think* about causal inference; the tools can help you *do* it.

Topics

#	Topic
1	Introduction and the causal inference problem
2	Potential outcomes and the identification problem
3	Randomized experiments
4	Selection on observables
5	DAGs and structural causal models
6	Sensitivity analysis
7	Difference-in-differences
8	Instrumental variables
9	Regression discontinuity
10	Special topics (time permitting)

Topics do not map exactly to weeks. Detailed readings for each topic will be posted on the course website.

Is This Class Right for You?

Vitally important:

1. probability theory
2. statistical inference
3. (linear) regression
4. familiarity with statistical software R and hopefully LaTeX

Nice to have:

1. multivariate calculus (e.g. recognizing objects like $\frac{\partial y}{\partial x}, \frac{\partial^2 y}{\partial x x'}$)
2. linear algebra (test: $\boldsymbol{X}\beta$, $x_i^\top \beta$, derive β_{OLS})

For students who have taken PS200A & B, you are all set.

For others you may have obtained this background elsewhere.

We will distribute a short **diagnostic self-assessment** (ungraded) to help you identify any gaps. The Resources page has materials keyed to each topic.

Books and Readings

- The slides are detailed and, together with the in-class activities, form the core material.
- Required readings will be posted on the course website for each topic.
- The book you should buy:
  - Angrist, Joshua D. and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. Princeton University Press.
- We will also use:
  - Hernán MA, Robins JM. *Causal Inference: What If*.

  Freely available at: `https://miguelhernan.org/whatifbook`
- A more comprehensive and annotated set of resources is on the course website's **Resources** page.

Course Requirements and Grades

- **In-class activities (20%)**
  - Short (5–10 min) pen-and-paper exercises at end of most lectures
  - Graded on completion and good-faith effort
  - Lowest two dropped (accommodates absences)
  - Also serves as the attendance mechanism
- **Project milestones (30%)**
  - Milestone 1 (Week 3): Topic and group formation
  - Milestone 2 (Week 5): Identification strategy proposal
  - Milestone 3 (Week 8): Data and preliminary analysis
- **Final project (50%)**
  - Research paper ($\sim$10–15 pages) in groups of 1–3
  - Presentations in the last week of class
  - Due during finals week

The Final Project

The project is the heart of the course.

- Formulate a causal question, choose an identification strategy, carry out an analysis.
- Be honest about limitations—discuss which assumptions are most vulnerable.
- You may use AI tools freely for coding, data wrangling, and exploring ideas. Include a short section describing how you used AI in the research process.
- What AI *cannot* do for you: choose an identification strategy, evaluate whether its assumptions are credible, or interpret what results do and do not tell you. That is your job.

AI as a Learning Tool

You are expected to use AI throughout this course—not as a shortcut, but as a **tutor**.

- When something from lecture is unclear, ask Claude or ChatGPT to explain it differently, work through an example, or connect it to something you already know.
- Learning to ask good questions of an AI is closely related to learning to think clearly about a problem.
- Each project milestone includes a brief **AI learning log**: one or two exchanges where you used an AI tool to clarify a concept, with a sentence about what you learned.
- This is not assessed for quality—it normalizes the practice and gives us insight into what topics are causing confusion.

Section

- **Your TA:** Barney Chen — `barneychen@ucla.edu`
- Section: Thursday 9:00–9:50 AM (Zoom — see course website for link)
- Section serves two roles:
  1. **Review and supplement**: technical details, estimation, and coding that we defer from lecture.
  2. **Project check-ins**: discuss your project with Barney and other groups, get feedback, troubleshoot.
- Barney's office hours: Mondays 5 PM, Bunche 3244 (and by appointment)

Logistics

- **Class:** MW 2:00–3:15 PM, Bunche 2160
- **Course website:**
  `https://chadhazlett.github.io/ps200c-2026/`
  Syllabus, slides, resources, project guidelines, section materials.
- **Announcements:** BruinLearn (make sure you are receiving them).
- **My office hours:** Wednesday 3:30–4:45 PM and by appointment, Bunche 3264.
- Come to class, engage, ask questions. If you aren't going to attend regularly, don't take the class.

Some Resources

1. CAPS: https://www.counseling.ucla.edu. They also have a 24/7 crisis support at 310-825-0768.
2. The Bruin Fund offers funding for students with computing needs.
3. I'm always happy to help you get any services you need or to figure out what's available.

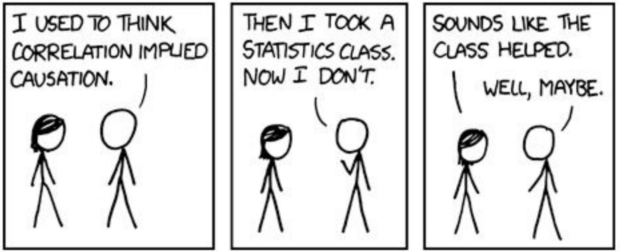

- Questions?
- Introductions!