Recap: Causal Inference I

DS701 Session 16 — Wed Oct 28, 2026

Today’s plan

Quiz 2 is five days away (Mon Nov 2) and covers sessions 9–15 — *not* today. Causal inference (this lecture and lecture 15) is examined on **quiz 3**. Today is not quiz review; use office hours for that.

Knowledge-Check Review

KC 1: Two readings of one coefficient

A regression coefficient $\hat{\beta}_1$ on X has two possible readings — an associational one and a causal one. State both in one sentence each, say which one least squares always gives you, and give the definition of a ***confounder*** *from the lecture.*

KC 2: The selection-bias term

The lecture decomposed the naive difference in means, $\mathbb{E}[Y \mid T{=}1] - \mathbb{E}[Y \mid T{=}0]$, into a causal part plus a ***selection-bias*** *term. In words, what does the selection-bias term measure, why is it nonzero in the hospital-visits example, and why does randomly assigning T make it zero?*

KC 3: Kidney stones, then ice cream

In the kidney-stone study, Treatment A had the higher recovery rate for small stones ***and*** *for large stones, yet Treatment B had the higher rate overall. Which comparison should a doctor trust, and why? Then connect it to the ice cream regression: what plays the role of stone size there, and what did "adjusting for it" look like in code?*

Highlights

1. A regression coefficient is an association, not an effect

A regression coefficient is an association, not an effect. OLS estimates $P(Y \mid X)$; a causal claim is about $P(Y \mid do(X))$, and the two differ whenever there is confounding.

Prediction — seeing

$$P(Y \mid X = x)$$

"Among people I *observe* with $X = x$, what is Y?"

Every model in this course so far. Rung 1 of Pearl's ladder.

Intervention — doing

$$P(Y \mid do(X = x))$$

"If I *set* $X = x$ for everyone, what is Y?"

What a decision-maker needs. Rung 2.

2. A confounder is a common cause

A confounder is a common cause of treatment and outcome. Adjusting for it (holding it fixed, or adding it to the regression) removes the spurious association — but only if you measured it and knew to include it.

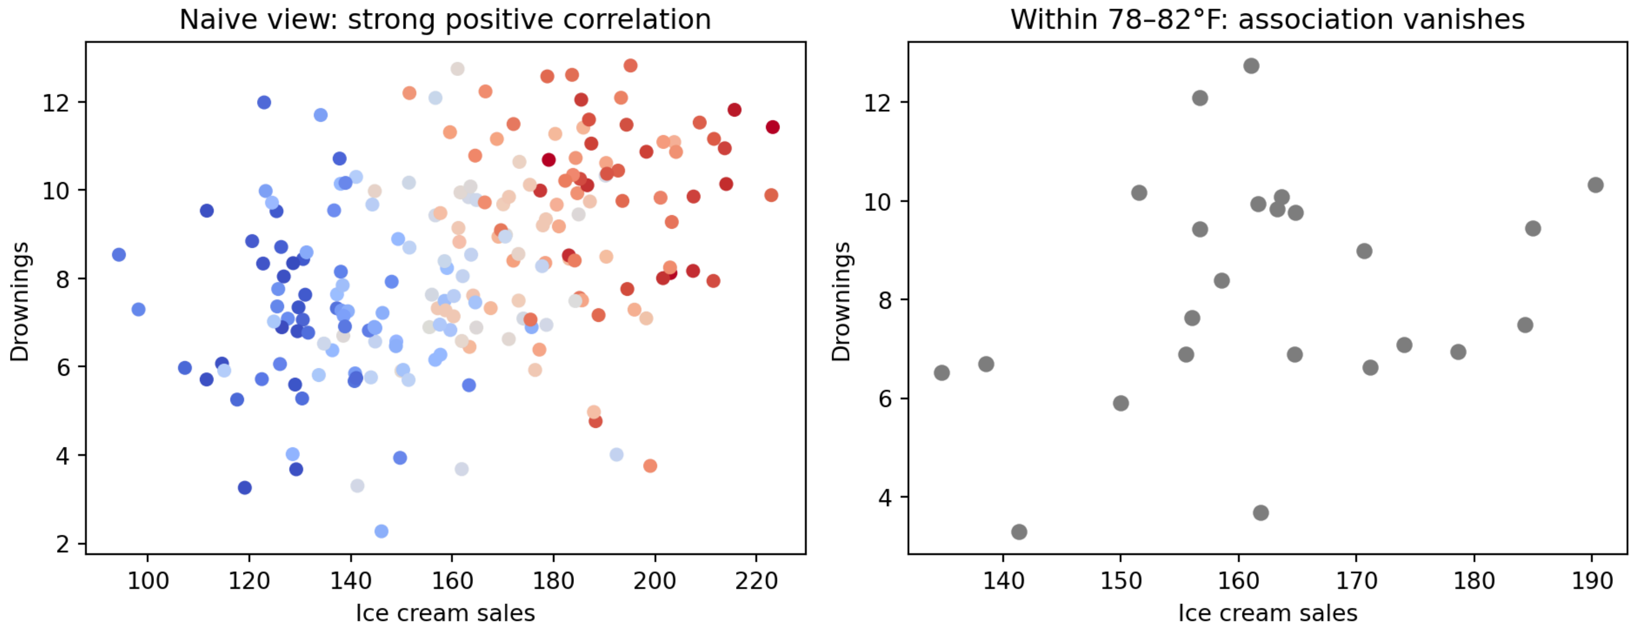

3. Potential outcomes define the effect — you observe one

Potential outcomes $Y(1), Y(0)$ define the effect; you only ever observe one. This is the fundamental problem of causal inference, so we estimate averages like the $\mathrm{ATE} = \mathbb{E}[Y(1) - Y(0)]$.

Unit	T_i	$Y_i(0)$	$Y_i(1)$	τ_i
Ann	1	**?**	7	?
Bob	0	4	**?**	?
Cara	1	**?**	9	?
Dan	0	3	**?**	?

$$Y_i^{\text{obs}} = T_i\, Y_i(1) + (1 - T_i)\, Y_i(0)$$

4. Naive difference $=$ causal effect $+$ selection bias

Naive difference in means $=$ causal effect $+$ selection bias. Simpson's paradox is what selection bias looks like in a table: within-group and overall comparisons can point in opposite directions.

$$\underbrace{\mathbb{E}[Y(1) \mid T{=}1] - \mathbb{E}[Y(0) \mid T{=}0]}_{\text{naive difference}} = \underbrace{\text{ATT}}_{\text{causal}} + \underbrace{\mathbb{E}[Y(0) \mid T{=}1] - \mathbb{E}[Y(0) \mid T{=}0]}_{\text{selection bias}}$$

5. Randomization zeroes the bias term

Randomization makes treatment independent of potential outcomes, so the selection-bias term is zero and the naive difference is unbiased for the ATE.

$$\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp T_i$$

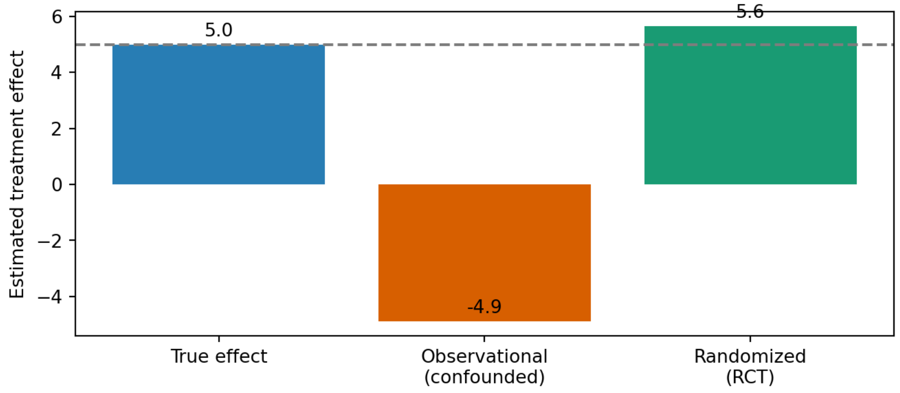

True effect $+5$; sicker patients likelier to be treated ⇒ observational -4.9; coin flip ⇒ $+5.6$.

In-Class Activity

Activity: confounding scenarios, then watch regression lie

Goal: classify real-world comparisons (confounded / reverse / Simpson / randomized) and name the confounder; recompute a Simpson's-paradox table by hand and *adjust* it; then simulate a confounder with a known true effect β, watch `Y ~ X` get it wrong, `Y ~ X + Z` recover it, and randomization make the naive fit right — and finally condition on a **collider** and watch a spurious effect appear from nothing.

- Work in groups of 2–3. Open **your section's** notebook (different scenarios, table, and β).
- Parts marked **(autograded)** are submitted to Gradescope; the rest is participation.
- Staff will circulate — be ready to explain any part of your work.
- Dependencies: `numpy`, `pandas`, `matplotlib`, `statsmodels`.

Section A1 Open in Colab Open on GitHub

Section B1 Open in Colab Open on GitHub

Going deeper

Full lecture notes: Causal Inference I: Correlation, Causation, and Counterfactuals

- What does a coefficient mean? — the two readings, side by side
- A regression you would recognize — the ice cream / drowning fit and its data-generating code
- Confounding — definition; adjusting by adding `temp`
- Two different questions — $P(Y \mid X)$ vs. $P(Y \mid do(X))$; the ladder of causation
- Where do spurious associations come from? — the five sources
- The potential outcomes framework, the fundamental problem, the ATE
- Selection bias — the decomposition of the naive difference
- Kidney stones and Berkeley admissions — Simpson's paradox in real tables
- The gold standard: randomized experiments — independence kills the bias term; the simulated drug study
- Next: Causal Inference II — causal graphs, what to adjust for and what *not* to