Causal DAGs: the structural meaning of CI

PS 200C, Spring 2026

Chad Hazlett

UCLA

chazlett@ucla.edu

Where we are, and today's question

So far: **selection on observables**. We assumed CI:

$$\{Y_{1i}, Y_{0i}\} \perp\!\!\!\perp D_i \mid X_i$$

Then if we condition on the right X, treatment is as-if random within strata.

Today we talk about graphical causal models ("DAGs") that help show:

- **Which** X is the "right" X?
- Is any X good enough?
- What can **go wrong** if I pick certain X?

Core idea: Ask *why* D and Y could be associated. Three possibilities:

1. D causes Y — the effect we want.
2. A **common cause** of D and Y.
3. Analytic errors that create these associations.

DAGs visualize and tell us how to remove unwanted sources of association.

Confounding, with numbers

Suppose people self-select into a new drug. You find:

	Recovered (count / total)	
	Drug	No drug
Mild	81 of 87 (93%)	234 of 270 (87%)
Severe	192 of 263 (73%)	55 of 80 (69%)
Combined	273 of 350 (78%)	289 of 350 (83%)

Pooled, the drug "hurts" (78 vs 83), but within each severity stratum the drug helps (93 vs 87; 73 vs 69).

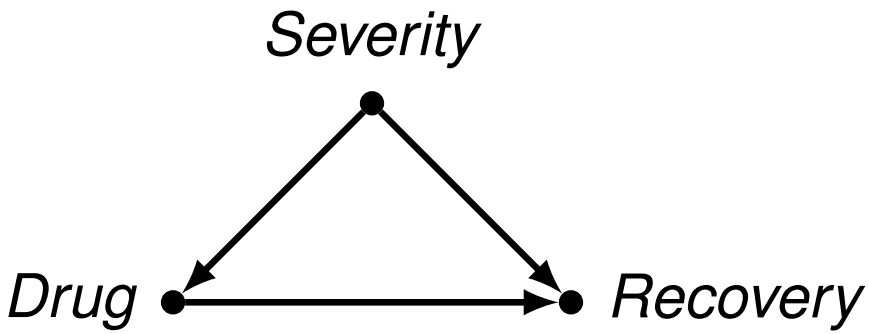
Why? The pooled comparison compares groups that differ in *severity*, not just the drug.

With POMs, we might have noted that $\mathbb{E}[Y(0)|D = 1] - \mathbb{E}[Y(0)|D = 0] < 0$ (negative bias)

Preview: Confounding, with a DAG

What we *believe* about the world:

- Severity increases drug-taking
- Severity causes reduced recovery
- The drug might cause recovery.



Now you can *see* the two reasons *Drug* and *Recovery* are associated:

1. *Drug* $\rightarrow$ *Recovery* (what we want),
2. *Drug* $\leftarrow$ *Severity* $\rightarrow$ *Recovery* (common cause — a spurious association).

Conditioning on *Severity* removes (2), leaving only (1).

What's a DAG?

A **Directed Acyclic Graph** (DAG):

- **Nodes** = variables.
- **Arrows, Edges** = direct causal relationships ($A \to B$: A is a direct cause of B).
- **No cycles**: you can't follow arrows and end up back where you started.
- **Missing arrows encode assumptions**: "A is *not* a direct cause of B." These are the assumptions doing the real work.
- **Dashed arrow** / **bidirected arc** = unobserved common cause.

Three things a DAG is *not*:

- Not a list of the variables in your dataset. (Many nodes may be unobserved.)
- Not a statistical model. (No parametric assumptions.)
- Not a complete theory of the world. (Just enough structure for the causal claim you want to make.)

Three path types

A "path" between two nodes is any sequence with edges connecting them (in any direction). Three elementary components of paths:

- **Chain**: $A \rightarrow B \rightarrow C$
- **Fork** (common cause): $A \leftarrow B \rightarrow C$
- **Collider** (common effect): $A \rightarrow B \leftarrow C$

Our task: figure out, for each, (a) when is A associated with C? (b) what happens when we condition on B?

Chain

$A \to B \to C$:

$$A \longrightarrow B \longrightarrow C$$

$A \perp\!\!\!\perp C$? No — A affects C through B.

$A \perp\!\!\!\perp C \mid B$? Yes — conditioning on B blocks the chain.

Moral: a chain transmits association; conditioning on the middle blocks it.

Fork (common cause)

$A \leftarrow B \rightarrow C$:



$A \perp\!\!\!\perp C$? No — the common cause B induces association.

$A \perp\!\!\!\perp C \mid B$? Yes — conditioning on the common cause blocks it.

Moral: a common cause creates spurious association; conditioning on it removes it. **This is classical confounding.**

Collider (common effect)

$A \rightarrow B \leftarrow C$:



$A \perp\!\!\!\perp C$? Yes — no causal path, no common cause.

$A \perp\!\!\!\perp C \mid B$? **No!** — conditioning on the collider *opens* a path.

Moral: colliders are the reverse of the other two. They are *already blocked*, and conditioning on one *creates* an association that wasn't there.

Colliders are real

Pearl's favorite:

Rain $\rightarrow$ Wet lawn $\leftarrow$ Sprinkler was on

Observing a wet lawn, learning it did *not* rain raises the chance the sprinkler ran.

Hollywood:

Looks $\rightarrow$ Fame $\leftarrow$ Talent

Among famous actors, the good-looking ones are less talented (on average).

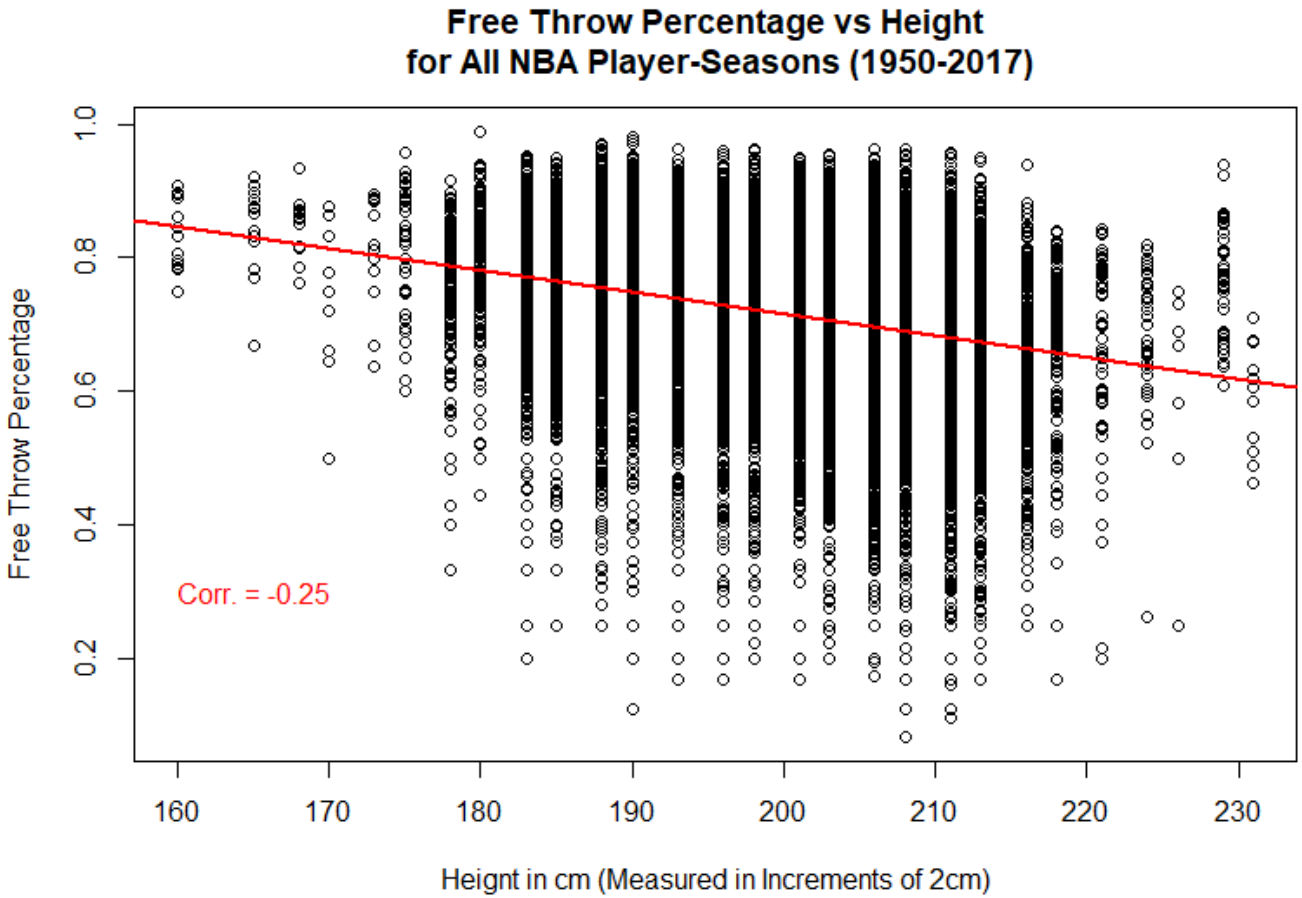
Academic tenure:

Productivity $\rightarrow$ Tenure $\leftarrow$ Originality

Among tenured faculty, the prolific ones are less original (on average).

Lesson: Conditioning on something can give rise to a spurious association...this will matter when we start trying to isolate causal effects.

A collider in the wild



Why do you think that is? Answer with a DAG.

Source: Adam Rohde, via `basketball-reference.com` *and kaggle.com/datasets/drgilermo/nba-players-stats.*

d-separation

Putting the three patterns together:

Rule

*d**-separation.** A and C are d-separated by a set Z iff every path between them is blocked. A path is blocked when:*

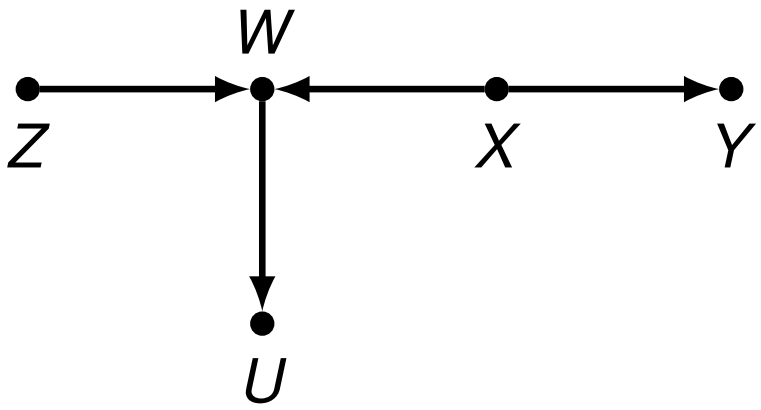
- *It contains a chain or fork at a node in Z, **or***
- *It contains a collider such that neither the collider nor any descendant of it is in Z.*

If A and C are d-separated by Z, then

$$A \perp\!\!\!\perp C \mid Z$$

In plain English: When A and C are d-separated by Z in this way, you can condition on Z as a way to remove the association between A and C due to those paths.

Practice: *d*-separation



Looking at Z and Y:

- d-separated by $\{\}$? Yes — the only path is $Z \rightarrow W \leftarrow X \rightarrow Y$, with W as collider; no conditioning $\Rightarrow$ blocked.
- d-separated by $\{W\}$? No — conditioning on the collider W *opens* the path.
- d-separated by $\{W, X\}$? Yes — W opens Z-Y, but X (a chain node on that path) re-blocks it.
- d-separated by $\{U\}$? No — U is a descendant of collider W, so conditioning on it also opens the path.

Back to the big picture

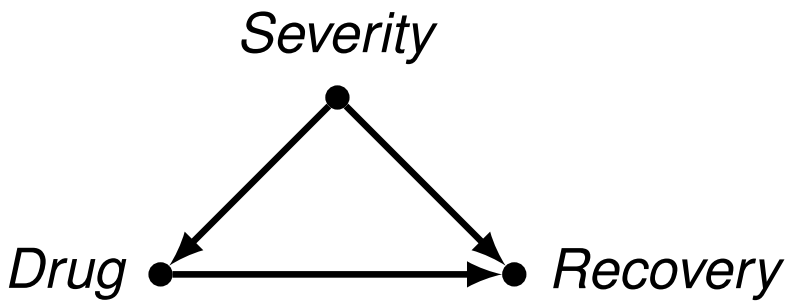
You do some analysis and see that D and Y associated. But is this down to $D \rightarrow Y$, or something else?

What we want: a rule for choosing set $\{X\}$ such that, after conditioning, the *only* remaining association between D and Y is the causal arrow $D \rightarrow Y$.

That rule is the **backdoor criterion**.

Front-door and back-door paths

- A **frontdoor path** from D to Y starts with an arrow *out of* D. These carry causation. We want to **leave them open**.
- A **backdoor path** from D to Y starts with an arrow *into* D. These carry spurious association. We want to **block them all**.



Frontdoor: $Drug \rightarrow Recovery$.
Backdoor: $Drug \leftarrow Severity \rightarrow Recovery$. Condition on *Severity* to block it.

The backdoor criterion

Rule

Backdoor criterion. *$\{X\}$ satisfies the backdoor criterion relative to (D, Y) if:*

1. *No node in X is a descendant of D (don't block frontdoor paths),* ***and***
2. *X blocks every backdoor path from D to Y.*

For humans, find a set X that:

1. blocks all *spurious* paths (every backdoor),
2. leaves *causal* paths untouched (no descendants of D),
3. doesn't *open new* spurious paths (no colliders or their descendants unless paired with another blocker downstream).

Backdoor = CI: two languages, one object

Theorem (informal)[1]: if X satisfies the backdoor criterion in a DAG, then

$$\{Y_{1i}, Y_{0i}\} \perp\!\!\!\perp D_i \mid X_i$$

in any SCM consistent with that DAG.

We can then apply the "adjustment formula" if we wanted to recover $P(Y)$ as if we "made everyone have treatment $D = d$", i.e. $P(Y \mid do(d))$:

$$P(Y \mid do(d)) = \sum_x P(Y \mid D{=}d, X{=}x)\, P(X{=}x).$$

Or, since we have been working with expectations:

$$\mathbb{E}(Y \mid do(d)) = \sum_x \mathbb{E}(Y \mid D{=}d, X{=}x)\, P(X{=}x)$$

[1] For formal statements: Pearl, Glymour & Jewell (2016) *Primer* Thm 4.3.1; Shpitser, VanderWeele & Robins (2010) adjustment criterion.

Adjustment formula for ATE

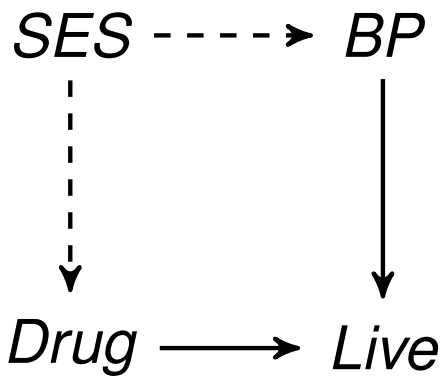
If you use this to get the ATE, $\mathbb{E}[Y \mid do(D = 1)] - \mathbb{E}[Y \mid do(D = 0)]$, you see exactly what we say with adjustment under SOO:

$$\sum_x (\mathbb{E}[Y \mid D=1, X=x] - \mathbb{E}[Y \mid D=0, X=x])\ P(X=x)$$

That is: figuring out what to condition on using the DAG and adjusting for it is the same thing we were doing by (i) claiming CI given those variables and (ii) conditioning on them.

As before, sub-stratification is the literal way to do so, but in practice we'll often need to use other approaches that give us an estimate as if $p(X|D = 1)$ and $p(X|D = 0)$.

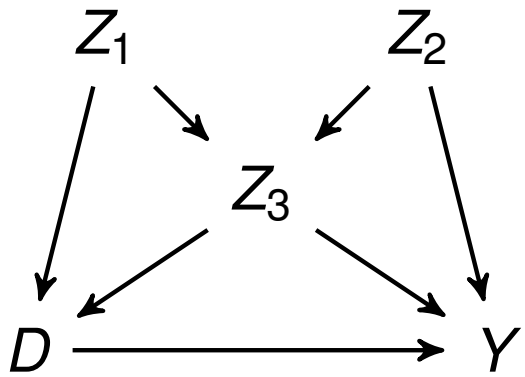
Example 1: a hidden confounder, but a workable adjustment set



SES unobserved (dashed). Can we identify the effect of *Drug* on *Live*?

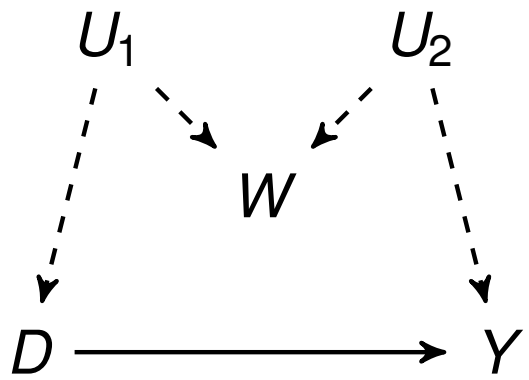
- Backdoor paths? $Drug \leftarrow SES \rightarrow BP \rightarrow Live$
- Observed set that blocks it? $\{BP\}$
- Equivalent CI statement: $\{Live_1, Live_0\} \perp\!\!\!\perp Drug \mid BP$.

Example 2: multiple valid adjustment sets



- Backdoor paths? $D \leftarrow Z_1 \rightarrow Z_3 \rightarrow Y$; $D \leftarrow Z_3 \rightarrow Y$; $D \leftarrow Z_3 \leftarrow Z_2 \rightarrow Y$.
- Several valid sets: $\{Z_3\}$, $\{Z_1, Z_2\}$, $\{Z_3, Z_1\}$, ...
- *Invariance*: all give you the same estimand in the non-parametric sense. Different estimators in practice, though.

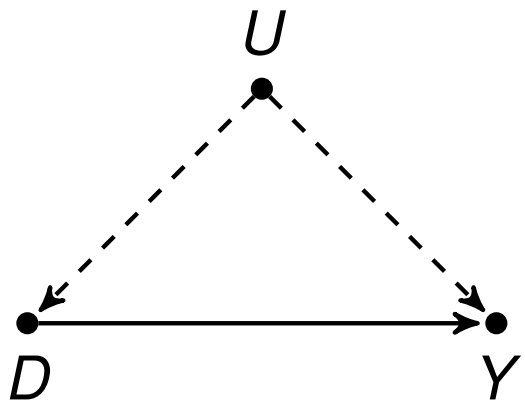
Example 3: How about now?



- Backdoor paths? The only candidate is $D \leftarrow U_1 \rightarrow W \leftarrow U_2 \rightarrow Y$.
- W is a collider on that path, so the path is already blocked. Empty set suffices!
- What if your colleague says, "let's play it safe and condition on W"? You open it. Now D and Y are spuriously associated through the (now unblocked) U_1, U_2.

Lessons: 1. DAGs help you see if CI holds; and 2. Conditioning on extra variables can *break* an identification that was fine without them.

Example 4: Failure of identification



- Backdoor path: $D \leftarrow U \rightarrow Y$. U unobserved.
- No observed variable blocks it. No set satisfies the backdoor criterion.

DAGs tell you when you're stuck — not just when you're fine.

What to do? SOO+sensitivity, or find another ID strategy.

Complaints (and responses)

"I can't really know the true causal structure."

True. But SOO isn't any different...you can't support it with claiming ausal structure. Drawing the DAG forces you to be *explicit* about what you're assuming. Different people can then disagree productively, not just pick different controls silently.

"I've always worked with potential outcomes, not DAGs."

They're two languages for the same object. CI = backdoor; the adjustment formula = stratification. DAGs are especially good at spotting *traps* (colliders, post-treatment conditioning) that are hard to see in PO notation.

"Different people draw different DAGs."

This is a feature not a bug: the DAG is surfacing key disagreement that needs to be resolved or what-if'd out.

Resources: going deeper

Today was a quick practical tour. For more practice or details:

Books:

- *What If* (Hernán and Robins, 2024). Chapters 6, 7. Free online.
- *Causal Inference in Statistics: A Primer* (Pearl, Glymour, Jewell, 2016). Chapters 1.4, 1.5, 2, 3.
- *Causality* (Pearl, 2009). The full treatment.

Online courses / lectures:

- *Causal Diagrams: Draw Your Assumptions Before Your Conclusions* — Miguel Hernán, EdX.
- *A Crash Course in Causality* — Jason Roy, Coursera.
- *Richard McElreath, "Science Before Statistics: Causal Inference." `https://www.youtube.com/watch?v=KNPYUVmY3NM`
- *Richard McElreath, *Statistical Rethinking* lecture series on YouTube (more on DAGs + Bayesian modeling).

Takeaways

- **"Why are *D* and *Y* associated?"** A DAG enumerates the answers: causal, common cause, or something else (e.g. a collider trap).
- **Paths come in three flavors**: chains and forks transmit association; colliders block it (and conditioning flips them).
- ***d*-separation** is the bookkeeping: does your conditioning set kill every path you don't want?
- **Backdoor criterion = CI**. Same statement, two languages. The DAG gives CI its structural meaning.
- **Conditioning can backfire**. Collider conditioning, post-treatment conditioning. More *X* is not always better.
- **The graph tells you when you're stuck**. Example 4 is the honest payoff: identification isn't guaranteed.