Directed Acyclic Graphs

BUSS975: Causal Inference in Financial Research

Professor Ji-Woong Chung

Korea University

Outline

1. Why Graphs?

Where we left off

- Last lecture: a causal effect is a property of the data **plus the assignment mechanism**.
- The potential outcomes framework gave us:
  - precise **estimands** — ATE, ATT, ATU;
  - a decomposition of what goes wrong — selection bias and heterogeneity bias;
  - a gold standard — randomization, which makes $(Y^1, Y^0) \perp\!\!\!\perp D$.
- What it did **not** give us: a systematic way to answer the single most common question in empirical work.

"What should I control for?"

The question every referee asks

- A typical referee report: *"The authors should control for firm size, leverage, industry, and profitability."*
- Is that good advice? The honest answer in most papers is: **nobody knows**.
- Common (bad) heuristics:
  - "Control for everything you have." — Wrong, and this lecture shows why.
  - "Control for anything correlated with both D and Y." — Closer, but still wrong.
  - "More controls make the estimate more credible." — Emphatically wrong.
- Adding a control can **remove** bias, **add** bias, or **flip the sign** of your estimate.
- To tell which, you need a language for your assumptions about the data-generating process.

What a DAG is for

- A **directed acyclic graph** is a picture of your causal assumptions.
- Once drawn, it delivers — mechanically, by rules you can check —
  - which variables you **must** control for,
  - which variables you **must not** control for,
  - whether your causal effect is identified **at all** from observational data.
- The DAG does not tell you whether your assumptions are *true*. Nothing does.
- What it does is make them **explicit, communicable, and refutable** — so a referee can disagree with your *assumptions* rather than your *taste in controls*.

This lecture . . .

By the end you will be able to explain, precisely, why the following is true:

You can run a perfectly randomized experiment, add one perfectly measured control variable, and get an estimate with the wrong sign.

- If controlling for things can do that, “control for everything” is not a strategy — it is a coin flip.

2. DAG Notation

Nodes and arrows

- A DAG has **nodes** (random variables) and **directed edges** (arrows).
- An arrow $X \to Y$ means: X exerts a **direct causal effect** on Y — direct meaning *not* routed through any other variable in the graph.

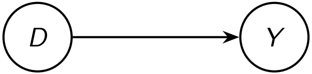

- Read: relationship lending (D) affects firm investment (Y).
- The graph is **non-parametric**: it says nothing about magnitude, sign, linearity, or functional form. Only *whether* an effect exists.

The missing arrow is the assumption

- Drawing an arrow is a weak claim: "X may affect Y."
- **Omitting** an arrow is a strong claim: "X has *no* effect on Y, for any unit, of any size."
- This is the opposite of how we usually read a regression table, and it is the most important thing to remember about DAGs.
- So the interesting content of a DAG is in the arrows you **did not draw**.
- Corollary: a DAG with arrows everywhere is honest but useless — nothing is identified. Progress requires exclusions, and exclusions require **institutional knowledge**.

Acyclicity

- "Acyclic" = no directed path leads from a variable back to itself.
- Forbidden:

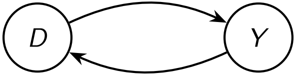

- This is a real restriction for finance, where equilibrium feedback is everywhere: leverage and performance, price and demand, coverage and returns.
- The usual fix is **time**: $\text{Leverage}_t \to \text{Perf}_{t+1} \to \text{Leverage}_{t+2}$ is acyclic.
- Genuine simultaneity is beyond what a DAG can express. We return to this in §8, What DAGs Cannot Do.

Observed and unobserved variables

- Solid circle = **observed**; dashed circle = **unobserved**.

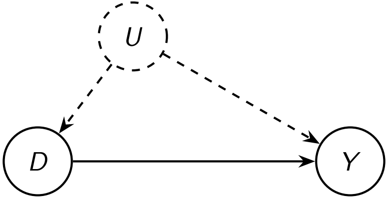

- U might be managerial talent, firm culture, private information, or local demand.
- The distinction matters enormously: a variable you cannot measure cannot be conditioned on, so it cannot be used to block anything.
- Much of this course is about what to do when a U like this exists.

Paths

- A **path** is any sequence of edges connecting two nodes, *ignoring the direction of the arrows*.
- A **directed (causal) path** follows the arrows from start to finish.

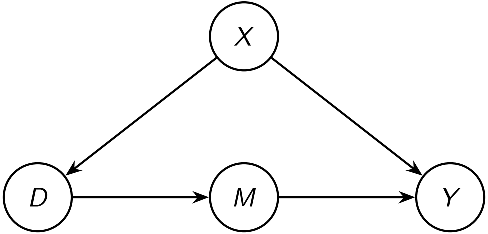

- $D \to M \to Y$ is a directed path (this is the causal effect).
- $D \leftarrow X \to Y$ is a path but **not** a directed path — yet it still generates correlation between D and Y.
- **Association flows along paths regardless of arrow direction. Causation flows only along directed paths.**

Where DAGs come from

- A DAG is **not** derived from the data. The data cannot tell you which way the arrows point.
- It comes from: economic theory, institutional detail, prior empirical work, and knowledge of how *your specific dataset* was assembled.
- Two researchers with different institutional knowledge will draw different DAGs and reach different conclusions — and that is the point. The disagreement is now **visible and arguable**.
- Practical advice: draw the DAG *before* you run the regression. If you cannot draw it, you do not yet understand your identification strategy.

Controlling for is not setting

- A recurring confusion, worth naming before we build the machinery.
- **Controlling for** X filters the data: "among firms that *happen* to have $X = x$." The population changes — we keep a subpopulation, which may be special in ways we cannot see.
- **Setting** X intervenes: "if we *assigned* $X = x$ to every firm." The population stays fixed; only X moves.
- Prediction needs only the first. Causal inference needs the second.
- In graph terms, setting X means **deleting every arrow into** X — X is no longer caused by its parents; it is caused by us.
    - That is exactly what a randomized experiment does physically.
    - It is why randomization needs no controls. There are no arrows left to block.

Ratings predict default; ratings do not prevent it

- Let $Y = 1$ if a firm defaults within five years, and let X be its credit rating.
- **Filtering**: among AAA issuers, default is rare. Ratings are superb *predictors* — selling this number is a profitable business.
- **Setting**: suppose a regulator relabelled every junk issuer AAA. Default rates would barely move. The label is not the cash flows.
- Filtering **selects** firms that were already creditworthy; setting **changes the label** and leaves creditworthiness untouched.[1]
- Ask of any empirical result: *is this a filter, or an intervention?*

[1] The causal effect is not exactly zero. Kisgen and Strahan 2010: when the SEC certified DBRS in 2003, yields moved 39bp per notch — the label alone, with fundamentals unchanged.

3. Three Elemental Structures

Every DAG is built from three pieces

Any path, however long, is a chain of three-node segments of exactly three types.

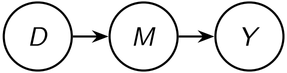

Chain (mediator)

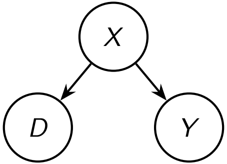

Fork (confounder)

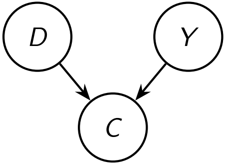

Collider

Chain: $D \to M \to Y$

- M is a **mediator**: the causal effect of D travels *through* it.
- The path is **open** by default — D and Y are correlated, and the correlation is causal.
- Conditioning on M **closes** the path — and thereby destroys the very effect you are trying to measure.
- Example: a credit shock raises investment, which raises employment.
  - Credit $\to$ Investment $\to$ Employment.
  - Controlling for investment when estimating the effect of credit on employment removes most of the answer.
- **Do not control for mediators** when you want the total effect.

Fork: $D \leftarrow X \rightarrow Y$

- X is a **confounder**: a common cause of both treatment and outcome.
- The path is **open** by default — D and Y are correlated, but the correlation is **not** causal.
- Conditioning on X **closes** the path and removes the bias.
- This is exactly omitted variable bias from last lecture, drawn as a picture.
- Example: firm quality drives both whether a firm has a lending relationship and how much it invests.
- **Do control for confounders** — if you can measure them.

Collider: $D \to C \leftarrow Y$

- C is a **collider**: a common *effect* of treatment and outcome.
- The path is **closed** by default — D and Y are *not* correlated through it.
- Conditioning on C **opens** the path and *creates* a correlation that did not exist.
- This is the one with no counterpart in your regression training, and it is where most damage is done.
- Example: skill and luck are independent, but both determine whether a fund survives. Study only survivors and skill and luck become correlated.
- **Do not control for colliders** — and note that "control for" includes *sample selection*.

The one table to memorize

Structure	Shape	**Association flows already?**	**Conditioning on middle node**	**Should you?**
Chain (mediator)	$D \to M \to Y$	Yes, *causal*	**blocks** it	**No**
Fork (confounder)	$D \leftarrow X \to Y$	Yes, *spurious*	**blocks** it	**Yes**
Collider	$D \to C \leftarrow Y$	No	**opens** it	**No**

- Column 3 asks whether D and Y are *already* correlated through this structure; column 4, what conditioning does to that.
- Only the **fork** is worth conditioning on: it is the only structure where association flows and is *not* the effect you want.
- Mediators and colliders are both "leave alone" — for opposite reasons. Blocking a mediator deletes the answer; opening a collider invents one.
- Conditioning on a **descendant** of a collider partially opens it too. A proxy for a collider is still a collider.

4. Backdoor Paths and the Backdoor Criterion

Backdoor paths

- A **backdoor path** is a path connecting D and Y that begins with an arrow pointing *into* D.
- Compare: $D \to Y$ leaves D through an arrow pointing *out*, so it is a front-door — the causal path.
- Backdoor paths generate correlation without causation — "back doors" through which association leaks.
- Identification strategy in one sentence:

Leave every directed path from D to Y open.
Close every backdoor path.

- The **backdoor criterion**: a set of variables $\mathbf{Z}$ identifies the causal effect of D on Y if
  1. $\mathbf{Z}$ blocks every backdoor path from D to Y, and
  2. $\mathbf{Z}$ contains no descendant of D.
- Condition 2 is what rules out mediators and post-treatment colliders.

A DAG: Example

Effect of a bank **relationship** (D) on firm **investment** (Y).

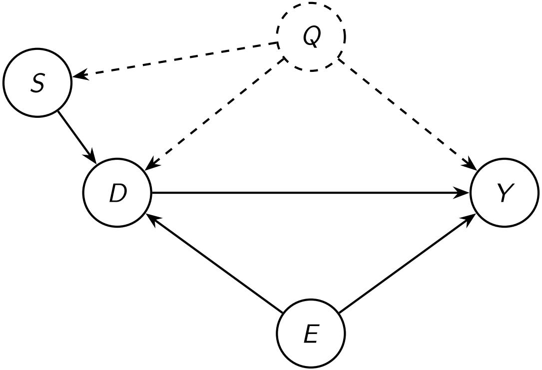

- Q = firm quality (**unobserved**), S = firm size, E = local economic conditions.

Enumerate the paths

From the previous DAG:

1. $D \rightarrow Y$ — **directed**. Leave open; this is the estimand.
2. $D \leftarrow Q \rightarrow Y$ — backdoor, open (fork at Q).
3. $D \leftarrow S \leftarrow Q \rightarrow Y$ — backdoor, open (chain then fork).
4. $D \leftarrow E \rightarrow Y$ — backdoor, open (fork at E).

- Conditioning on $\{Q, E\}$ blocks paths 2, 3, and 4. Identified.
- But Q is **unobserved**. So the achievable set is $\{S, E\}$:
  - blocks path 3 (at S) and path 4 (at E),
  - leaves path 2 **open**.
- Conclusion: the effect is **not identified** by controlling for observables.
- That conclusion — reached before running anything — is the DAG earning its keep.

Simulation: the fork, and closing it

$Q \to D$, $Q \to Y$, $D \to Y$ with a true effect of 2.

```
set.seed(975)
n    <- 100000
qual <- rnorm(n)                       # firm quality (the confounder)
D    <- 0.6 * qual + rnorm(n)          # relationship lending
Y    <- 2 * D + 1.5 * qual + rnorm(n)    # TRUE effect of D on Y is 2

b <- function(f) unname(coef(lm(f))[2])
round(c(naive = b(Y ~ D), control_Q = b(Y ~ D + qual)), 3)
```

```
naive control_Q
2.662     2.001
```

- Leaving the backdoor open: the estimate is too large by about 0.66.
- Blocking it at Q: we recover the truth.
- Here "blocking" just means putting Q in the regression.

Multiple valid strategies

- Often several different sets satisfy the backdoor criterion.
- Any of them gives an **unbiased** estimate — so choose among them on **precision**.
- Rule of thumb: prefer the set that includes strong predictors of Y.
  - They do not change identification; they shrink residual variance and tighten standard errors.
- This is the legitimate version of “add controls to make the estimate sharper”.

5. Collider Bias

Colliders create correlation out of nothing

- The claim that surprises people: conditioning on a collider makes **independent** variables dependent.
- Intuition. Suppose a fund survives only if skill $+$ luck is high enough.
  - Among survivors, a manager with low skill *must* have had high luck — otherwise they would not be in the sample.
  - So within survivors, skill and luck are **negatively** correlated, though they are independent in the population.
- The correlation is manufactured entirely by **who is in the sample**.

Survivorship as a collider

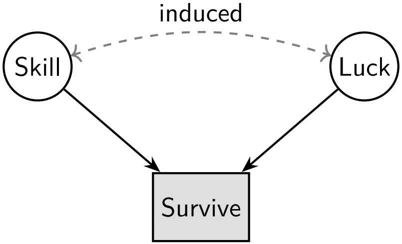

- Grey box = we conditioned on it (here, by studying only surviving funds).
- Skill $\rightarrow$ Survive $\leftarrow$ Luck is a collider. It was closed. We opened it.

Simulation: the surviving-fund fallacy

```r
set.seed(975)
n     <- 100000
skill <- rnorm(n)                          # independent...
luck  <- rnorm(n)                          # ...by construction
perf  <- skill + luck
surv  <- perf > quantile(perf, 0.80)       # only the top 20% survive

round(c(cor_all       = cor(skill, luck),
        cor_survivors = cor(skill[surv], luck[surv])), 3)
```

```
      cor_all cor_survivors
       -0.003        -0.640
```

```r
round(c(slope_all       = unname(coef(lm(perf ~ skill))[2]),
        slope_survivors = unname(coef(lm(perf[surv] ~ skill[surv]))[2])), 3)
```

```
      slope_all slope_survivors
          0.997           0.358
```

- Skill and luck: independent in the population, correlated -0.64 among survivors.
- The return to skill looks like 0.36 instead of 1.00 — understated.

This is not a hypothetical

- The same structure appears throughout empirical finance, usually unnoticed:
  - **Fund performance studies** on surviving funds.
  - **Compustat / Execucomp**: a firm appears only if it listed and survived. Listing is caused by size, quality, and opportunity — a collider.
  - **M&A studies** conditioned on **completed** deals: completion is caused by bidder synergies *and* target resistance.
  - **Rated firms only**: a firm is rated when it chooses to issue public debt.
  - **IPO samples**: going public is a common effect of quality and market conditions.
- Empirical finance treats these as a list of separate warnings — survivorship bias, sample coverage, deal completion. They are **one structural error**: conditioning on a collider.

Selection *is* conditioning

- You do not need to type a variable into a regression to condition on it.
- You condition on a variable whenever you
  - include it as a control,
  - stratify or match on it,
  - **drop observations** based on it,
  - or use a dataset that was **assembled** using it.
- The last one is the dangerous case, because it is invisible in your code.

Example. Your sample is the CRSP–Compustat merged file. A firm appears only if it (i) listed on a US exchange and (ii) filed accounts that got matched. Both are consequences of size, age, and quality — the same things driving your outcome. You conditioned on a collider without writing a single line of code that mentions it.

- Practical advice: write *how the sample was constructed* into the DAG, as nodes. Sample-construction filters are conditioning statements.

Discrimination and collider bias

- Gender discrimination is often denied like this: *condition on occupation and the wage gap shrinks or disappears* — within the same job, women are paid the same. **But what if occupation is itself a consequence of discrimination?**
- Suppose every woman is "treated" with discrimination and no man is. Firms respond in **two ways**: they pay her less, *and* they do not promote her — so she is paid less again.

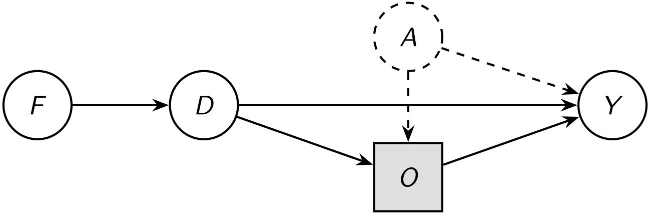

F = female, D = discrimination, O = occupation, A = ability (unobserved), Y = wage

- Men and women are **equally productive**, so there is no arrow $F \to Y$. We are not after the effect of *female* on earnings; we are after the effect of **discrimination**.

What controlling for occupation actually does

Three paths run from D to Y:

1. $D \to Y$ — the **direct** effect. The causal path of interest.
2. $D \to O \to Y$ — **mediation**, not a backdoor path: discrimination worsens the job, and the worse job pays less.
3. $D \to O \leftarrow A \to Y$ — contains the **collider** O, so it is *closed* to begin with.

- Regress Y on D alone and you get the **total effect**: the direct effect plus the part mediated by occupational sorting.
- Now add occupation, to "compare men and women in similar jobs." Two things happen at once:
  - it **closes** the mediation channel $D \to O \to Y$ — reasonable, if the direct effect is what you want;
  - it **opens** path 3, because conditioning on a collider opens it. The control ironically introduces a new pattern of bias.
- To shut path 3 again you would have to control for occupation **and ability** — and ability is unobserved. So this design **does not satisfy the backdoor criterion**.

The simulation

```r
set.seed(975)
n              <- 10000
female         <- as.integer(runif(n) >= 0.5)   # half the population is female
ability        <- rnorm(n)                      # independent of gender; UNOBSERVED
discrimination <- female                        # all women experience discrimination
occupation     <- 1 + 2*ability - 2*discrimination + rnorm(n)
wage           <- 1 - 1*discrimination + 1*occupation + 2*ability + rnorm(n)

b <- function(f) unname(round(coef(lm(f))["discrimination"], 3))
c(unconditional = b(wage ~ discrimination),
  add_occupation = b(wage ~ discrimination + occupation),
  add_occ_ability = b(wage ~ discrimination + occupation + ability))
```

```
 unconditional  add_occupation add_occ_ability
        -2.983           0.591          -0.996
```

- Discrimination is **hard-coded**: -1 directly on wages, and -2 on occupation. True direct effect $= -1$; true total effect $= -1 + 1 \times (-2) = -3$.

Reading the three regressions

- **Unconditional ≈ -3.** A large negative effect — the combination of the direct effect and the indirect effect through occupation. This is the *total* effect, and it is correct if that is what you want.
- **Add occupation $\approx +0.6$. The sign flips: it literally becomes positive.** We know this is wrong, because the direct effect was hard-coded at -1. Conditioning on the collider O opened $D \to O \leftarrow A \to Y$, and the spurious path is strong enough to pervert the entire relationship.
- **Add occupation and ability ≈ -1.** Now we recover the direct effect, because conditioning on A closes the path that conditioning on O opened.
- The rescue only worked because ability was **observable in the simulation**. In real data it is not — which is the whole point. "Control for occupation" is not the fix; it is the bug.

6. Bad Controls

The bad control problem

- A **bad control** is a variable that is itself affected by the treatment.
- Last lecture we previewed this and deferred the explanation. Here it is:
  - if $D \rightarrow X$, then X is a descendant of D;
  - conditioning on it violates part 2 of the backdoor criterion;
  - it either blocks part of the causal effect (if X is a mediator) or opens a collider path (if X has another parent).
- The second case is the nasty one, because it can happen even when D was **randomly assigned**.

Randomization does not protect you

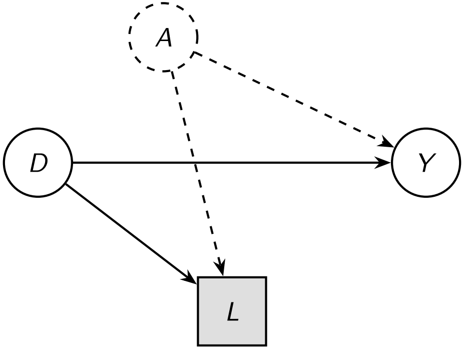

- D = credit supply shock (**randomly assigned**), Y = investment, L = leverage, A = firm quality (unobserved).
- Nothing points into D: there are **no backdoor paths**. The naive regression is already correct.
- But L is a collider on $D \to L \leftarrow A \to Y$. Control for leverage and you open it.

Simulation: how to ruin an experiment

```r
set.seed(975)
n <- 100000
D <- rnorm(n)                       # credit shock -- RANDOMLY ASSIGNED
A <- rnorm(n)                       # firm quality -- unobserved
L <- D + A + rnorm(n)               # leverage -- affected by treatment
Y <- 1 * D + 3 * A + rnorm(n)       # TRUE effect of D on Y is +1

b <- function(f) unname(coef(lm(f))[2])
round(c(no_controls     = b(Y ~ D),
        control_L       = b(Y ~ D + L),
        control_L_and_A = b(Y ~ D + L + A)), 3)
```

```
    no_controls       control_L control_L_and_A
          0.993          -0.493           1.008
```

- Treatment was randomized. The true effect is $+1$. One control variable turns it into -0.5.

Reading the damage

- Without controls: D is randomized, no backdoor exists, estimate ≈ 1. Correct.
- Adding L: opens $D \to L \leftarrow A \to Y$, and the estimate drops to -0.5.
- Adding A as well: closes the newly opened path, restoring $+1$ — but A is unobserved in any real study.
- The lesson is not "leverage is a bad control." It is:

Whether a variable is a good control is not a property of the variable.
It is a property of its position in the graph.

Bad controls in finance

Common post-treatment controls, and what they do:

- **Leverage / firm size / Tobin's Q** in a study of a shock that plausibly moves them.
- **Analyst coverage** when studying disclosure — disclosure changes coverage.
- **Deal completion** in M&A — completion is post-announcement.
- **Survival to $t+1$** in any panel with attrition.
- **Post-treatment industry classification** for firms that change industries after treatment.

- The operational rule: **controls must be determined before treatment**.
- When in doubt, report the specification without the questionable control. If the estimate moves a lot, you have learned something — and you must explain which specification the DAG supports.

7. Neutral Controls and Precision

Not every control matters

- Some variables are neither confounders nor colliders. Conditioning on them changes nothing about bias.
- Two useful cases:
  - $X \to Y$ only (a predictor of the outcome, unrelated to D): **harmless and helpful** — it reduces residual variance and tightens standard errors.
  - $X \to D$ only (a predictor of treatment, unrelated to Y): **harmless but harmful in practice** — it absorbs variation in D and *inflates* standard errors.
- The second is the seed of the weak-instrument and "overcontrolling" problems you will meet in the IV lecture.

Precision: the setup

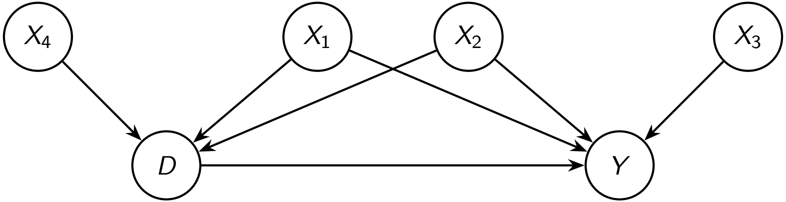

- X_1, X_2 are **confounders** — they point at both D and Y. Any valid set must include them.
- X_3 points at Y only. X_4 points at D only. Neither sits on a backdoor path, so **neither is needed** for identification.
- All three sets below therefore satisfy the backdoor criterion. They differ only in precision.

Precision: simulating three valid strategies

```r
set.seed(975)
sim <- function() {
  n  <- 1000
  x1 <- rnorm(n); x2 <- rnorm(n)                  # confounders: -> D and -> Y
  x3 <- rnorm(n)                                  # predicts Y only
  x4 <- rnorm(n)                                  # predicts D only
  D  <- 0.5*x1 + 0.5*x2 + 1.5*x4 + rnorm(n)
  Y  <- 2*D + 2*x1 + 2*x2 + 4*x3 + rnorm(n)       # TRUE effect = 2
  b  <- function(f) unname(coef(lm(f))[2])
  c(minimal = b(Y ~ D + x1 + x2),
    plus_x3 = b(Y ~ D + x1 + x2 + x3),
    plus_x4 = b(Y ~ D + x1 + x2 + x4))
}
r <- replicate(1000, sim())
round(rbind(mean = apply(r, 1, mean), sd = apply(r, 1, sd)), 4)
```

```
     minimal plus_x3 plus_x4
mean  1.9999  2.0002  1.9962
sd    0.0735  0.0172  0.1299
```

Reading the precision results

- All three are **unbiased** — every mean is 2.000. Identification is settled; only variance differs.
- Adding X_3, which predicts **only** Y, cuts the standard deviation from 0.074 to 0.017 — four times sharper. It soaks up residual variance without touching identification.
- Adding X_4, which predicts **only** D, raises it to 0.130. It absorbs the variation in D you were relying on.
- **Practical guidance**:
  1. First use the DAG to find sets that identify the effect. That is about *bias*.
  2. Then, among those sets, prefer strong predictors of Y and avoid strong predictors of D alone. That is about *variance*.
- Never reverse the order. Choosing controls by what tightens the standard error is how sign flips get published.

8. What DAGs Cannot Do

Honest limitations

1. **DAGs do not test assumptions.** A DAG is an argument, not evidence. The data cannot tell you which way an arrow points.
2. **DAGs are non-parametric.** They tell you *whether* an effect is identified, never how large it is, nor the functional form, nor the estimator.
3. **DAGs are acyclic.** Genuine simultaneity — the equilibrium feedback that pervades finance — cannot be drawn. Time-indexing helps; it does not always suffice.
4. **DAGs say nothing about magnitude.** A technically open backdoor path with a tiny effect may matter less than a closed one. The graph is binary; the world is not.
5. **DAGs do not choose estimands.** ATE versus ATT is a potential outcomes question.

DAGs and potential outcomes are complements

Potential outcomes is better for	DAGs are better for
Defining estimands (ATE, ATT, ATU)	Choosing what to control for
Heterogeneous effects, δ_i	Spotting colliders and bad controls
Deriving estimators and their bias	Communicating a design in one picture
Inference and standard errors	Proving an effect is *not* identified

- Use the DAG to decide your specification; use potential outcomes to say what the coefficient means.

How to use this in your own work

1. Draw the DAG **before** the regression — including how the sample was built.
2. Mark unobserved variables with dashes. Be honest about what you cannot measure.
3. List every path from D to Y. Classify each: directed, or backdoor.
4. Find a set satisfying the backdoor criterion. If none exists using observables, **say so** — and go find a research design (IV, RDD, DiD) instead of more controls.
5. Among valid sets, pick for precision.

Wrapping Up

Summary [Part 1]

- A DAG is a picture of your causal assumptions; the **missing** arrows carry the content.
- **Association flows along all paths; causation flows only along directed paths.**
- Three structures, and the whole calculus follows from them:
  - **Chain** $D \to M \to Y$: open; conditioning closes it. Do not control for mediators.
  - **Fork** $D \leftarrow X \to Y$: open; conditioning closes it. Do control for confounders.
  - **Collider** $D \to C \leftarrow Y$: closed; conditioning **opens** it. Never control for colliders.
- **Backdoor criterion**: block every backdoor path, and condition on no descendant of D.

Summary [Part 2]

- **Collider bias is real and large.** Independent variables became correlated at -0.64 purely through sample selection; the return to skill fell from 1.00 to 0.36.
- **Selection is conditioning.** Dropping observations, or using a dataset assembled on a criterion, conditions just as surely as a regression control does.
- **Randomization does not immunize you**: a randomized treatment with one post-treatment control produced -0.5 when the truth was $+1$.
- Whether a control is good is a property of its **position in the graph**, not of the variable.
- Identify first (bias), then choose among valid sets (variance). Never the reverse.
- Next lecture: potential outcomes and DAGs both assumed we could condition our way to identification. What if we cannot? — **Unconfoundedness and matching**.

References

- Cunningham, S. *Causal Inference: The Remix*, Ch. 3.
- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge University Press.
- Pearl, J. and D. Mackenzie (2018). *The Book of Why*. Basic Books.
- Cinelli, C., A. Forney, and J. Pearl (2024). A Crash Course in Good and Bad Controls. *Sociological Methods & Research* 53, 1071–1104.
- Angrist, J. and J.-S. Pischke (2009). *Mostly Harmless Econometrics*, §3.2.3. Princeton University Press.
- Elwert, F. and C. Winship (2014). Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable. *Annual Review of Sociology* 40, 31–53.
- Kisgen, D. and P. Strahan (2010). Do Regulations Based on Credit Ratings Affect a Firm's Cost of Capital? *Review of Financial Studies* 23, 4324–4347.