DAGs and Bad Controls

Francis J. DiTraglia

University of Oxford

Core Empirical Research Methods

Selection on observables again...

Last Time

- Binary treatment D, potential outcomes (Y_0, Y_1).
- Observed outcome $Y = (1 - D)Y_0 + DY_1$.
- Selection on observables: $D \perp\!\!\!\perp (Y_0, Y_1) | \boldsymbol{X}$ for observed covariates $\boldsymbol{X}$.
- Overlap: $0 < p(\boldsymbol{X}) < 1$. (Recall that we can check this.)
- Regression Adjustment, Propensity Score weighting, Matching

Elephant in the Room

We have completely ignored the question of what to include in $\boldsymbol{X}$.

Which covariates should we adjust for?

Identification

- Some choices of $\boldsymbol{X}$ may satisfy $D \perp\!\!\!\perp (Y_0, Y_1) | \boldsymbol{X}$ while others don't.
- If selection on observables assumption fails, we don't get the ATE / TOT.

Efficiency

- There may be multiple choices of $\boldsymbol{X}$ that satisfy selection on observables.
- But some may produce more efficient estimates of the ATE / TOT.

Today: focus on identification; use regression adjustment for simplicity.

The Omitted Variables Bias (OVB) Formula[1]

- To keep things simple, assume a linear model with homogeneous effects:

$$Y = \alpha + \beta D + \gamma X + U, \quad \text{Cov}(D, U) = \text{Cov}(X, U) = \mathbb{E}(U) = 0.$$

- In other words:

$$Y_0 = \alpha + \gamma X + U, \quad Y_1 = Y_0 + \beta, \quad \text{ATE} = \text{TOT} = \beta$$

- What does a regression of Y on D identify?

$$\frac{\text{Cov}(D, Y)}{\text{Var}(D)} = \frac{\text{Cov}(D, \alpha + \beta D + \gamma X + U)}{\text{Var}(D)} = \beta + \gamma \frac{\text{Cov}(D, X)}{\text{Var}(X)}$$

- "Short" regression coefficient only equals β if $\gamma = 0$ or $\text{Cov}(D, X) = 0$.

[1] See, e.g., Section 3.2.2 of *Mostly Harmless Econometrics*.

How *not* to interpret the OVB Formula.

- OVB Formula tells us *when* and *how* the coefficient on D differs depending on whether we include X in the regression.
- "Short" regression includes only D; "Long" regression includes both D and X.
- "Short" and "Long" coefficients for D agree if:
  1. X does not help predict Y in the "Long" regression or
  2. X is uncorrelated with D.
- Only if we **assume** that the long regression is the true causal model does this tell us whether we need to adjust for X.

Bad Advice: "Adjust for any observed variable that is correlated with D and Y."

Example 1: A prototypical bad control.

```r
set.seed(1693)
n <- 5000
d <- rbinom(n, 1, 0.4)
x <- rbinom(n, 1, 0.25 + 0.5 * d)
y <- x + rnorm(n)
mean(y[d == 1]) - mean(y[d == 0])
```

```
## [1] 0.523149
```

```r
library(broom); library(tidyverse)
xtilde <- x - mean(x)
reg <- lm(y ~ d + x + d:xtilde)
tidy(reg) |> filter(term == 'd') |>
  select(estimate, std.error)
```

```
## # A tibble: 1 x 2
##   estimate std.error
##      <dbl>     <dbl>
## 1   0.0113    0.0343
```

$$\begin{aligned}\mathbb{E}[Y_1 - Y_0] &= \mathbb{E}[Y|D=1] - \mathbb{E}[Y|D=0] \\ &= \mathbb{E}[X|D=1] - \mathbb{E}[X|D=0] = 0.5\end{aligned}$$

Why is X a bad control?

Intermediate Outcome

- Example 1: X is *itself* an outcome of D that *goes on* to cause Y.
- Adjusting for an intermediate outcome masks the true causal effect of D.
- E.g. randomized early childhood intervention causes college; college causes wage.
- In the simulation, 100% of the effect of D on Y goes through X.

Common Advice

> *Variables measured before the variable of interest [D] was determined are usually good controls. In particular, because these variables were determined before the variable of interest, they cannot themselves be outcomes in the causal nexus.*[2]

[2] From Section 3.2.3 of *Mostly Harmless Econometrics*, but similar statements are common.

Example 2: This bad control *is not* an intermediate outcome.

```r
library(mvtnorm)
R <- matrix(c(1, 0.5, 0.5, 1),
            2, 2)
errors <- rmvnorm(n, sigma = R)
u <- errors[,1]
v <- errors[,2]
x <- rbinom(n, 1, 0.5)
d <- 1 * ((-1 + 2 * x + v) > 0)
y <- -0.3 + d + u

mean(y[d == 1]) - mean(y[d == 0])
```

```
## [1] 1.511498
```

```r
xtilde <- x - mean(x)
lm(y ~ d + x + d:xtilde) |>
  tidy() |>
  filter(term == 'd') |>
  select(estimate, std.error)
```

```
## # A tibble: 1 x 2
##   estimate std.error
##      <dbl>     <dbl>
## 1     1.90    0.0355
```

Why is X a bad control?

Instrumental Variable

- Example 2: X is a valid instrument for the endogenous treatment D.
- But this is the **wrong** way to use an instrumental variable: should run IV.
- Adjusting for X soaks up the **exogenous** part of D, making the bias worse.[3]

```
library(AER)
ivreg(y ~ d | x) |> tidy() |> filter(term == 'd') |>
  select(estimate, std.error)
```

```
## # A tibble: 1 x 2
##   estimate std.error
##      <dbl>     <dbl>
## 1     1.01    0.0430
```

[3] See here for a proof.

"No causes in; no causes out."[4]

Feeling confused?

- How can we tell which variables to adjust for and which are bad controls?
- Is is simply a matter of "I know it when I see it"?

Bad News

- Meaningful causal inference **always** requires assumptions, even in RCTs.
- Causal inference from observational data requires **even more** assumptions.

Good News

- If you make your assumptions explicit, there is a **definitive** solution.
- If it's possible to use selection-on-observables, find the correct $\boldsymbol{X}$; if it's not possible, show why this is so.
- Free bonus: better intuition about bad controls.

[4] Nancy Cartwright

Graph Terminology

- **Graph**: set of **nodes** connected by **edges**.
- Two nodes are **adjacent** if connected by an edge.
- Edges can be **directed** (figure) or **undirected**.
- Directed edge points from **parent** to **child**.
- **Directed graph** has only directed edges.
- **Path**: sequence of connected vertices.
- **Directed Path**: a path that "obeys one-way signs"
- Directed path points from **ancestor** to **descendant**.
- **Cycle**: directed path that returns to starting node.
- **Acyclic Graph**: a graph without any cycles.

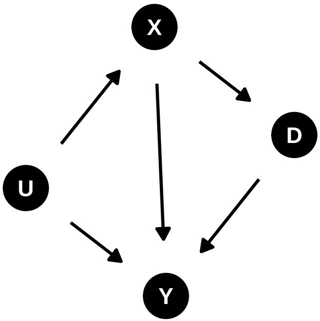

Structural Causal Model

- Structural Equations: $f = (f_d, f_x, f_y)$
- Exogenous "Errors": (U, V, W)
- P and f induce **observational dist** of Endogenous Variables: (D, X, Y)

$$
\begin{aligned}
X &\leftarrow f_x(U, V) \\
D &\leftarrow f_d(X, W) \\
Y &\leftarrow f_y(D, X, U) \\
(U, V, W) &\sim P
\end{aligned}
$$

Directed Acyclic Graph (DAG)[a]

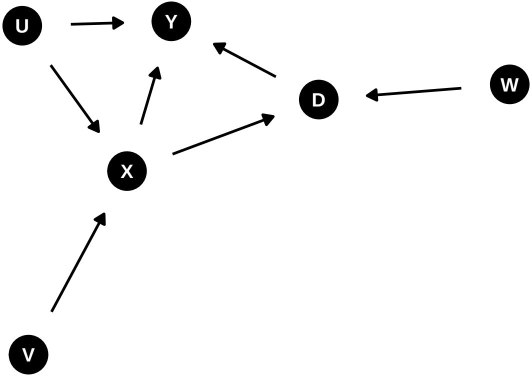

[a] We'd usually suppress W and V.

- Graphical causal model (DAG): "stylized" structural model (non-parametric)
- Directed edges encode **assumptions** about flow of causal influence or lack thereof.
- (Parent $\rightarrow$ Child $=$ **Direct Cause**); (Ancestor $\rightarrow$ Descendant $=$ **Potential Cause**)

How Graphical Causal Models (DAGs) Can Help

Key Point

DAGs encode **conditional dependence structure** under our modelling assumptions.

Back Door Criterion

- Is it possible to learn the causal effect of D on Y via selection on observables?
- If so, which observed variables should we adjust for?

Placebo Tests

- What conditional **independencies** between endog. variables does our model imply?
- Since these variables are observed, we can **partially test** our model.

Do-Calculus

Determine if causal effect of interest is non-parametrically identified **full stop**, whether by selection on observables, IV, or something else.

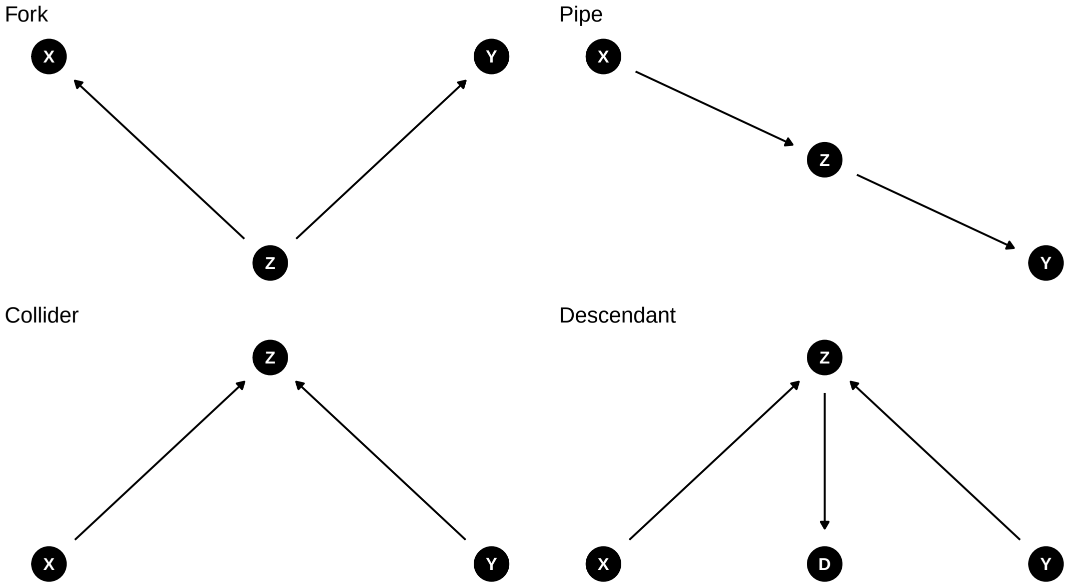
Fork
X
Y
Z
Pipe
X
Z
Y
Collider
Z
X
Y
Descendant
Z
X
D
Y

The Fork: A Common Cause

Fork Lemma

If Z is a common cause of X and Y and there is only one path between X and Y, then $X \perp\!\!\!\perp Y | Z$.

So what?

- X and Y are correlated even if neither causes the other.
- Adjusting for Z solves the problem: **blocks the path**.
- Z = health status, X = hospital, Y = mortality.

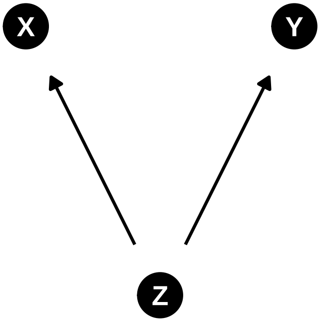

The Pipe: A Mediator[5]

Pipe Lemma

If there is only one directed path from X to Y and Z intercepts that path, then $X \perp\!\!\!\perp Y | Z$.

So what?

- Adjusting for Z **blocks the path**, just like in the pipe.
- But if we want the causal effect of X on Y, this is **bad**.
- This is *precisely* example 1 from above.
- $X =$ childhood intervention, $Z =$ college, $Y =$ wage.

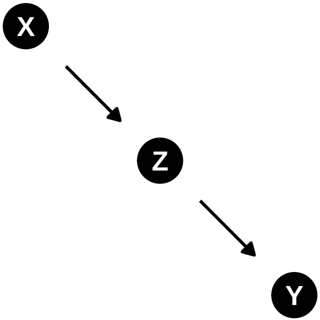

[5]Also known as a "chain".

The Collider: A Common Effect

Collider Lemma

If there is only one path between X and Y and Z is their common effect, then $X \perp\!\!\!\perp Y$ but $X \not\perp\!\!\!\perp Y | Z$.

So what?

- This is the weird/interesting/confusing one.
- Adjust for $Z \Rightarrow$ **spurious association** between X and Y.
- X, Y indep. coins; $Z =$ bell rings if at least one HEADS.
- $X =$ beauty, $Y =$ talent, $Z =$ movie star.

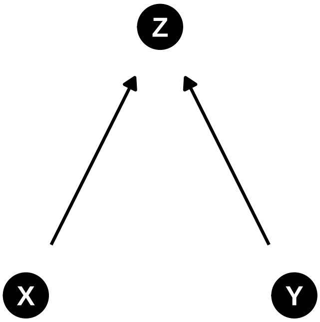

The Descendant

Descendant Lemma

Conditioning on a descendant D of Z has the effect of *partially conditioning* on Z itself.

Collider Corollary

In the figure, $X \perp\!\!\!\perp Y$ but $X \not\!\perp\!\!\!\perp Y | D$.

Discussion

- What this means depends on the situation.
- In the figure Z is a collider.
- Could also have Z as the middle node in pipe/fork.
- Pipe/fork: adjust for $D \Rightarrow$ **partially block** X, Y path.

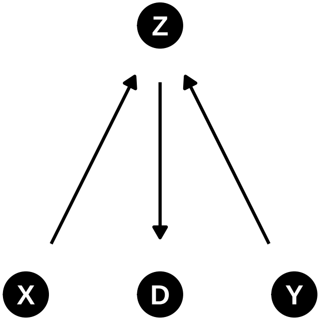

Blocked Path

A set of nodes Z **blocks a path** p iff p contains:

1. a pipe/fork whose middle node is in Z or
2. a collider that is not in Z and has no descendants in Z.

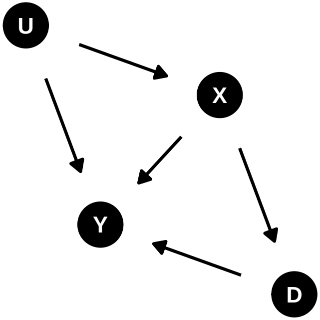

d-Separation

If Z blocks all paths between X and Y, then $X \perp\!\!\!\perp Y | Z$. We say that X and Y are *d*-**separated** conditional on Z.

Paths between D and U

$\underbrace{(D \rightarrow Y \leftarrow U)}_{\text{Path I}}$ $\underbrace{(D \rightarrow Y \leftarrow X \leftarrow U)}_{\text{Path II}}$ $\underbrace{(D \leftarrow X \rightarrow Y \leftarrow U)}_{\text{Path III}}$ $\underbrace{(D \leftarrow X \leftarrow U)}_{\text{Path IV}}$

- Paths I–III already blocked: Y is a collider. Conditioning on X blocks Path IV.
- Therefore, D and U are *d*-separated conditional on X.

Building and Plotting a DAG in R

Using the ggdag package.

```
library(ggdag)
mydag <- dagify(
  Y ~ X + U + D,
  X ~ U,
  D ~ X
)

mydag |>
  ggdag(node_size = 8, text_size = 3) +
  theme_dag()
```

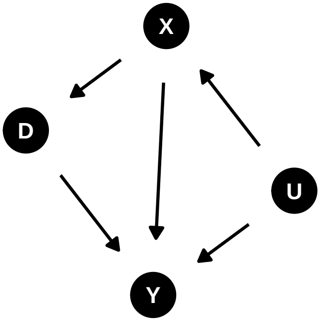

Finding Paths: Open and Closed

Using the dagitty package

```
library(dagitty)
paths(mydag, from = 'D', to = 'U')
```

```
## $paths
## [1] "D -> Y <- U"      "D -> Y <- X <- U" "D <- X -> Y <- U" "D <- X <- U"
##
## $open
## [1] FALSE FALSE FALSE  TRUE
```

Graph Surgery

Observational Distribution: $\mathbb{P}(Y|X = x)$

Actual distribution of Y among people observed to have $X = x$.

Interventional Distribution: $\mathbb{P}(Y|do(X = x))$

Distribution of Y that that we *would obtain* if we *intervened* and set $X = x$ for everyone.

- DAG encodes observational distribution.
- Delete edges **pointing into** D to obtain $do(D)$ interventional distribution.
- Causal effect of interest is the path from D to Y in this "modified" graph.
- This is what an experiment does!
- $\text{ATE} = \mathbb{E}(Y_1 - Y_0) = \mathbb{E}(Y|do(D = 1)) - \mathbb{E}(Y|do(D = 0))$

Graph Surgery: Delete Edges Pointing Into *D*

Observational Distribution

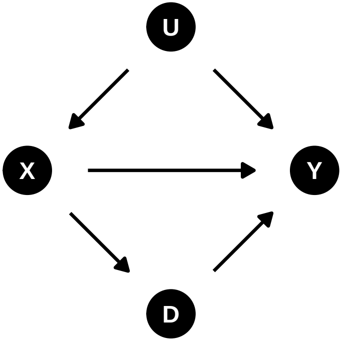

Interventional Distribution

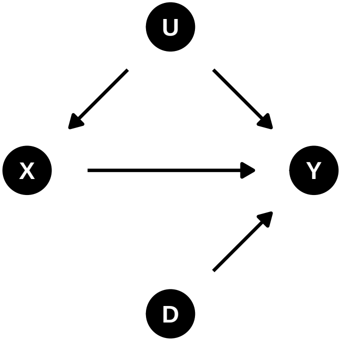

- Absent an experiment, all we have is the observational distribution.
- How can we use it to learn about the (unobserved) interventional distribution?

Shutting Back-doors

Back-door Paths

- Path between treatment and outcome starting with edge pointing *into* treatment.
- Noncausal: only edges pointing *out* from treatment represent causal effects.

Intuition

- Experiments *delete* back-door paths: remove edges pointing into treatment.
- If we can't delete the back-door paths, maybe we can **block** them instead.

Back-door Criterion

A set of nodes Z satisfies the back-door criterion relative to (X, Y) if no node in Z is a descendant of X and Z blocks every back-door path between X and Y.

A Less Formal Statement of the Back-door Criterion

1. List all the paths that connect treatment and outcome.
2. Check which of them *open*. A path is *open* unless it contains a collider.
3. Check which of them are *back-door paths*: contain an arrow pointing at D.
4. If there are no open back-door paths, you're done. If not, look for nodes you can condition on to **block** remaining open back-door paths without opening new ones.

Of course we can only condition on *observed* variables!

The Payoff

Back-door Theorem

If Z satisfies the back-door criterion relative to (X, Y), then

$$\mathbb{P}(Y = y | do(X = x)) = \sum_z \mathbb{P}(Y = y | X = x, Z = z)\mathbb{P}(Z = z)$$

Counterfactual Interpretation

If Z satisfies the back-door criterion relative to (X, Y), then for all x $Y_x \perp\!\!\!\perp X | Z$.

Translating to Potential Outcomes

- The "counterfactuals" Y_x are our potential outcomes from last lecture.
- Back-door criterion implies selection on observables assumption for X given Z.
- The formula above is nothing more than **regression adjustment**.

Finding Adjustment Sets with daggity

```
mydag <- dagify(
  Y ~ X + U + D,
  X ~ U,
  D ~ X
)
adjustmentSets(mydag,
               exposure = 'D',
               outcome = 'Y')
```

```
## { X }
```

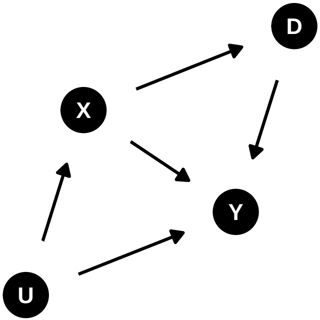

Example: Causal effect of exercise on cancer.[6]

Observed

- $D =$ physical activity
- $Y =$ cervical cancer (Yes/No)
- $X =$ positive pap smear test result (Yes/No)

Unobserved

- $U =$ pre-cancer lesion (Yes/No)
- $V =$ health-consciousness

Story

Health-conscious $\Rightarrow$ more physically active; more visits to doctor. Should we adjust for X?

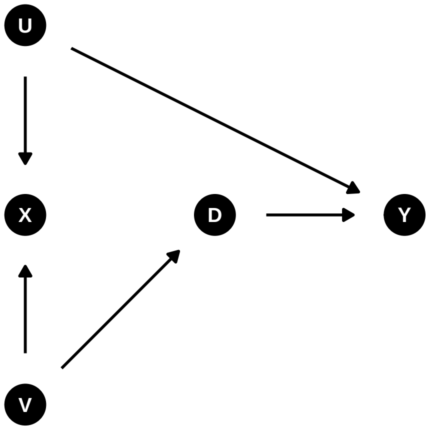

[6]Adapted from Hernan & Robins (2020) Section 7.4.

Exercise / Cancer Example Continued

- X is a **collider**: it *blocks* the back-door path between D and Y through (U, V).
- Adjusting for X *opens* this blocked path, so X is a **bad control**.
- Back door criterion is satisfied with $Z = \emptyset$: don't condition on *anything*!

```
library(dagitty)
library(ggdag)
dagify(Y ~ D + U, D ~ V,  X ~ U + V) |>
  paths(from = 'D', to = 'Y')
```

```
## $paths
## [1] "D -> Y"                    "D <- V -> X <- U -> Y"
##
## $open
## [1]  TRUE FALSE
```

Futher Reading on Graphical Causal Models

- Shalizi (2021) - Advanced Data Analysis from an Elementary Point of View Chapters 19–22.
- Cinelli, Forney & Pearl (2022) - A Crash Course in Good and Bad Controls
- Pearl, Glymour & Jewell (2016) - Causal Inference in Statistics: A Primer
- Pearl (2009) - Causality: Models, Reasoning & Inference (2nd ed)
- Huntington-Klein (2022) - The Effect
- Pearl & Mackenzie (2018) - The Book of Why
- Huntington-Klein (2022) - Pearl before Economists
- Traag & Waltman (2022) - Causal Foundations of Bias, Disparity and Fairness
- Heiss - Causal Inference and this blog post on the *do*-calculus.