Introduction to causal inference

Session 11

MATH 80667A: Experimental Design and Statistical Methods
HEC Montréal

Outline

Basics of causal inference

Directed acyclic graphs

Causal mediation

Causal inference

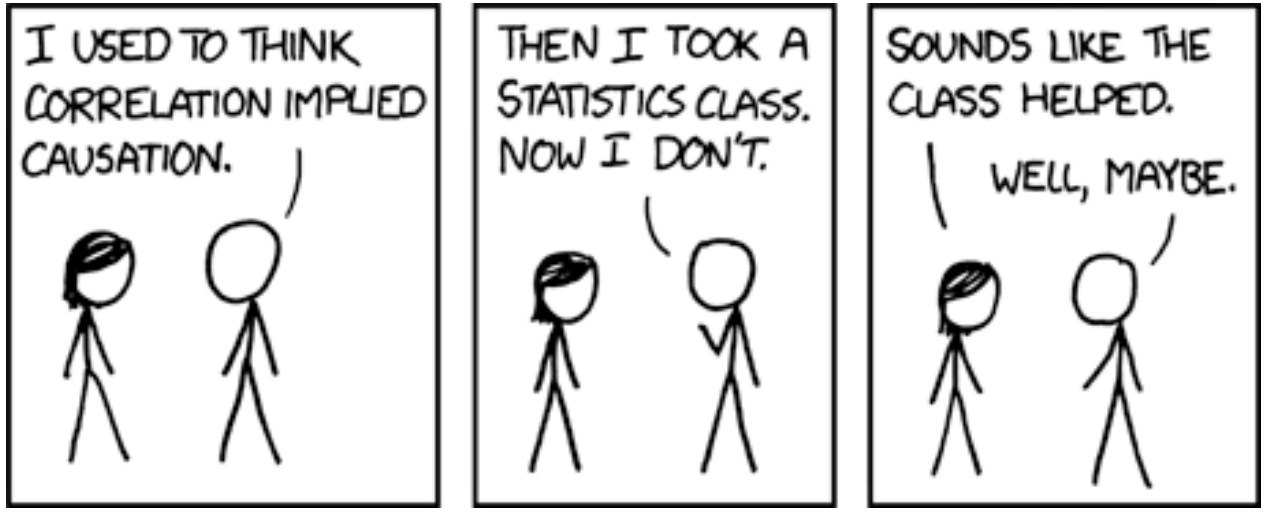

xkcd comic 552 by Randall Munroe, CC BY-NC 2.5 license. Alt text: Correlation doesn't imply causation, but it does waggle its eyebrows suggestively and gesture furtively while mouthing 'look over there'.

Correlation is not causation

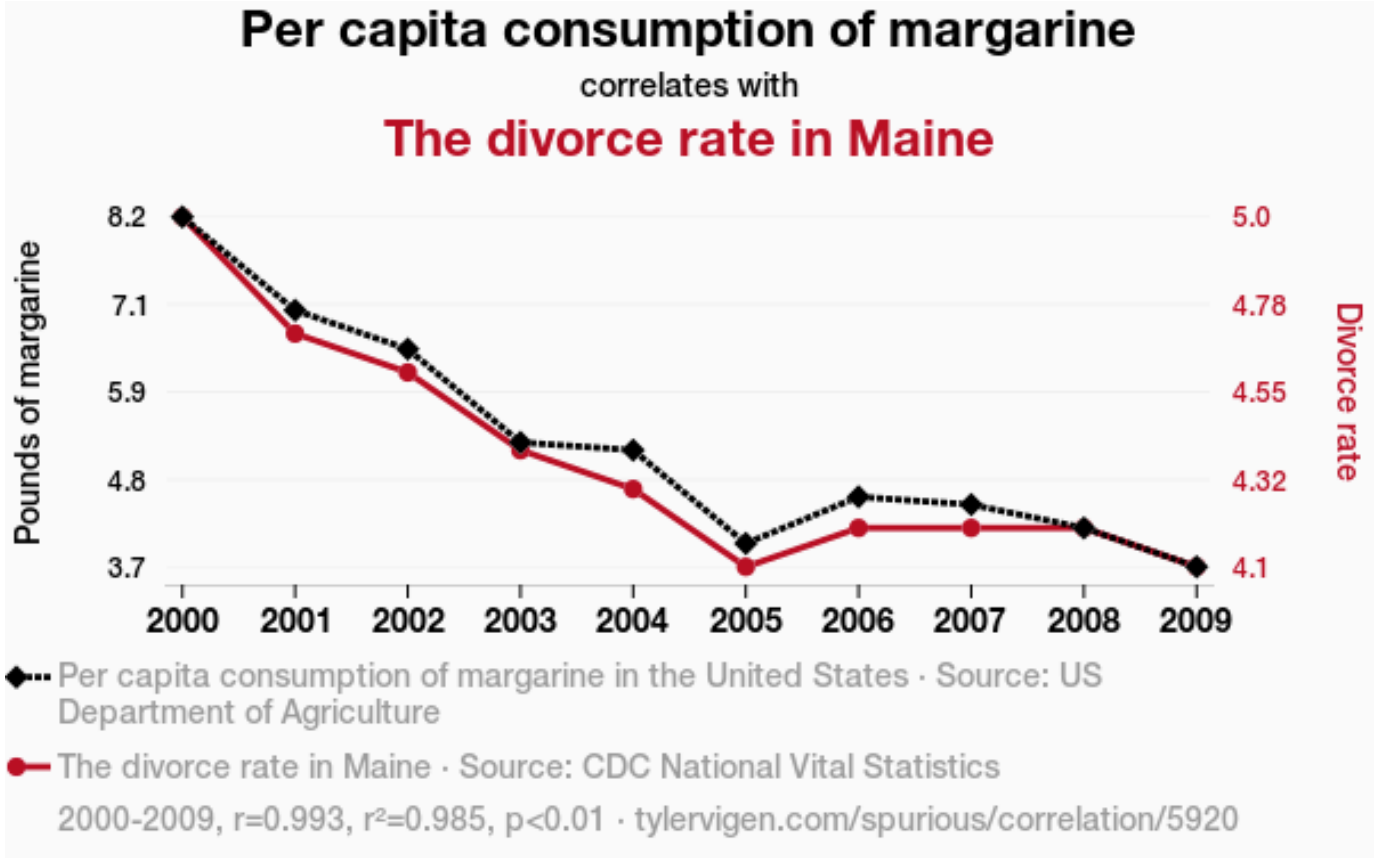

Spurious correlation by Tyler Vigen, licensed under CC BY 4.0

Correlation vs causation

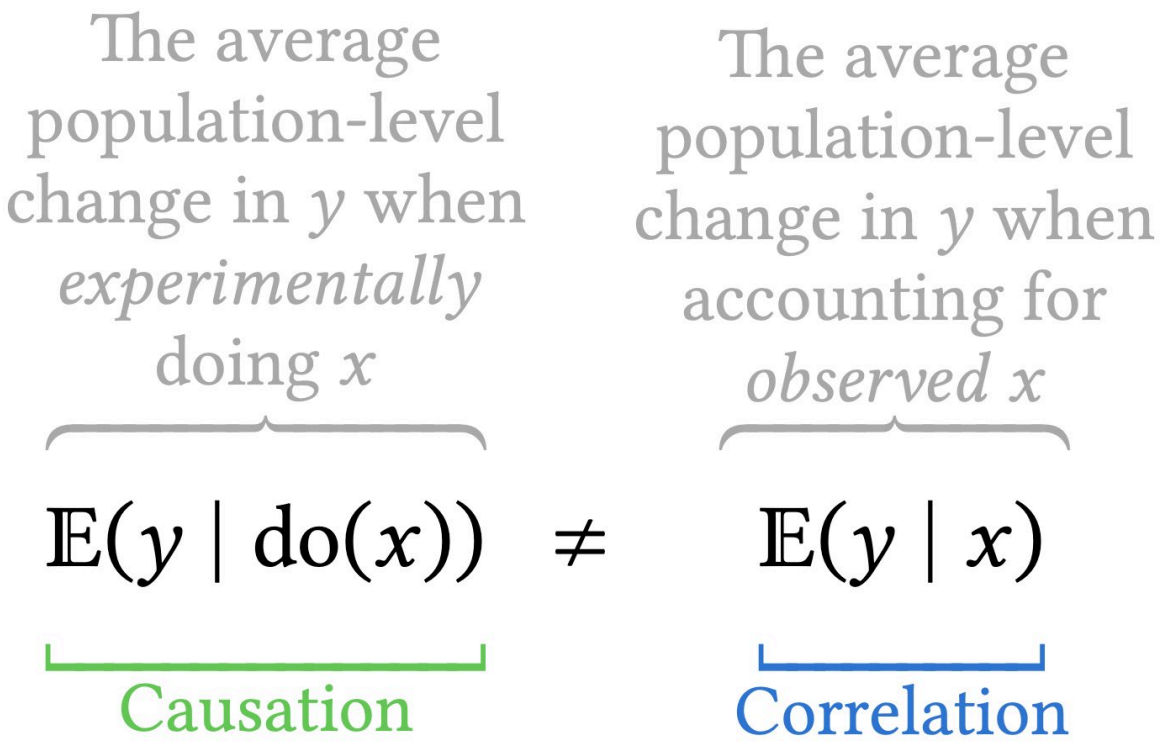

Illustration by Andrew Heiss, licensed under CC BY 4.0

Potential outcomes

For individual i, we postulate the existence of a potential outcomes

- $Y_i(1)$ (response for treatment $X = 1$) and
- $Y_i(0)$ (response for control $X = 0$).

Both are possible, but only one will be realized.

Observe outcome for a single treatment

- Result $Y(X)$ of your test given that you either party ($X = 1$) or study ($X = 0$) the night before your exam.

Fundamental problem of causal inference

i	X	Y(0)	Y(1)	Y(1)-Y(0)
1	1	?	4	?
2	0	3	?	?
3	1	?	6	?
4	0	1	?	?
5	0	5	?	?
6	1	?	7	?

Causal assumptions?

Since we can't estimate individual treatment, we consider the **average** treatment effect (average over population) $\mathsf{E}\{Y(1) - Y(0)\}$.

The latter can be estimated as

$$\underbrace{\mathsf{ATE} =}_{} \quad \underbrace{\mathsf{E}(Y \mid X = 1)}_{\substack{\text{expected response among} \\ \text{treatment group}}} \quad - \quad \underbrace{\mathsf{E}(Y \mid X = 0)}_{\substack{\text{expected response among} \\ \text{control group}}}$$

When is this a valid causal effect?

(Untestable) assumptions

For the ATE to be equivalent to $\mathsf{E}\{Y(1) - Y(0)\}$, we need:

1. conditional *ignorability*, which states that potential outcomes are independent (denoted with the $\perp\!\!\!\perp$ symbol) of assignment to treatment given a set of explanatories $\boldsymbol{Z}$. In notation $\{Y(0), Y(1)\} \perp\!\!\!\perp X \mid \boldsymbol{Z}$
2. lack of interference: the outcome of any participant is unaffected by the treatment assignment of other participants.
3. consistency: given a treatment X taking level j, the observed value for the response $Y \mid X = j$ is equal to the corresponding potential outcome $Y(j)$.

Directed acyclic graphs

Types of data

Experimental

You have control over which units get treatment

Observational

You don't have control over which units get treatment

Causal diagrams

Directed acyclic graphs (DAGs)

Directed: Each node has an arrow that points to another node

Acyclic: You can't cycle back to a node (and arrows only have one direction)

Graph: A set of nodes (variables) and vertices (arrows indicating interdependence)

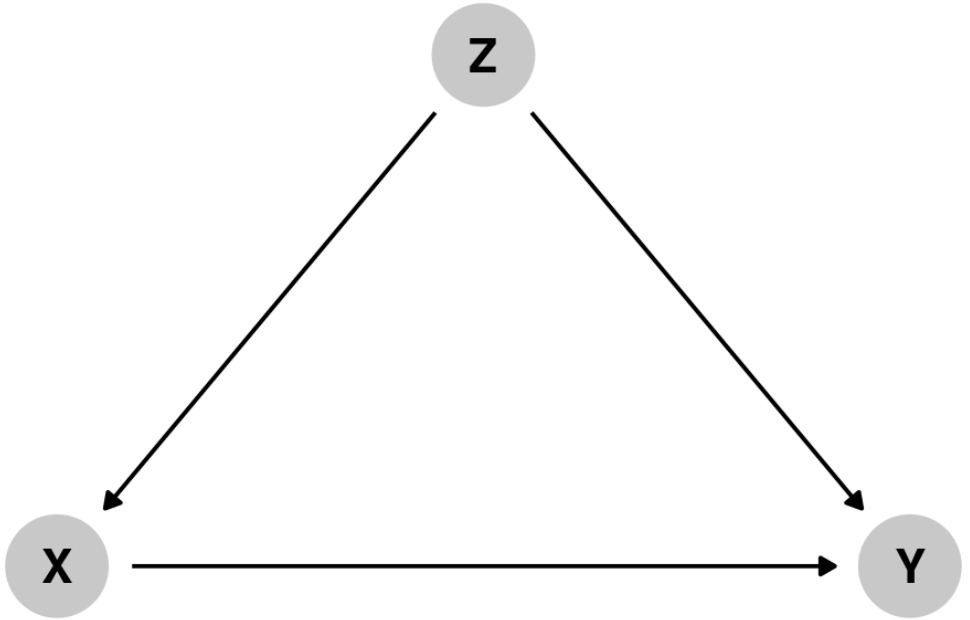

Causal diagrams

Directed acyclic graphs (DAGs)

Graphical model of the process that generates the data

Maps your philosophical model

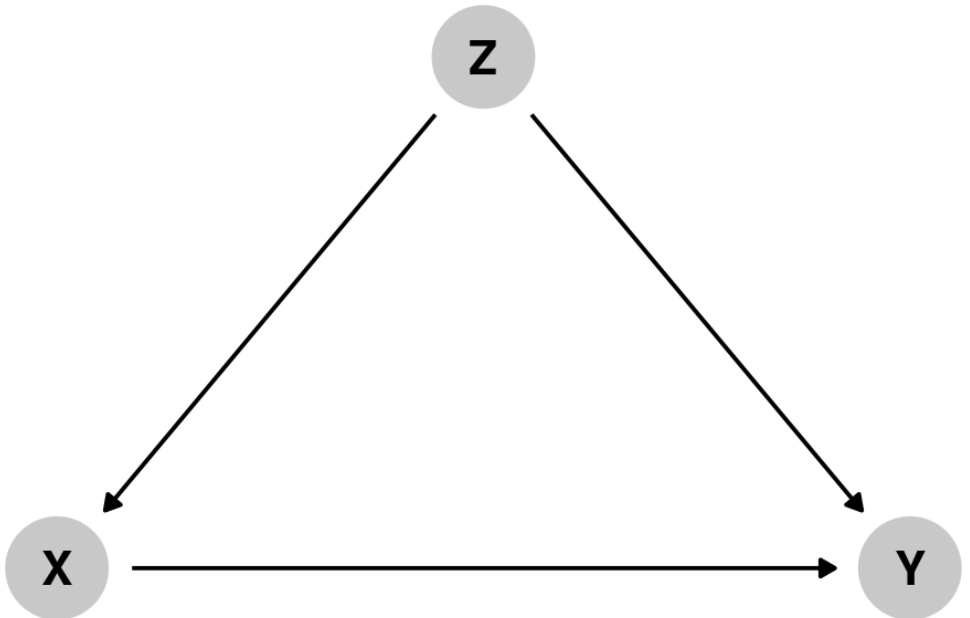

How to draw a DAG

What is the causal effect of an additional year of education on earnings?

Step 1: List variables

Step 2: Simplify

Step 3: Connect arrows

Step 4: Use logic and math to determine which nodes and arrows to measure

1. List variables

Education (treatment) → Earnings (outcome)

Location **Ability** **Demographics**

Socioeconomic status **Year of birth**

Compulsory schooling laws **Job connections**

2. Simplify

Education (treatment) → Earnings (outcome)

Location **Ability** **Demographics**

Socioeonomic status **Year of birth**

Compulsory schooling laws **Job connections**

Background

3. Draw arrows

Education causes earnings

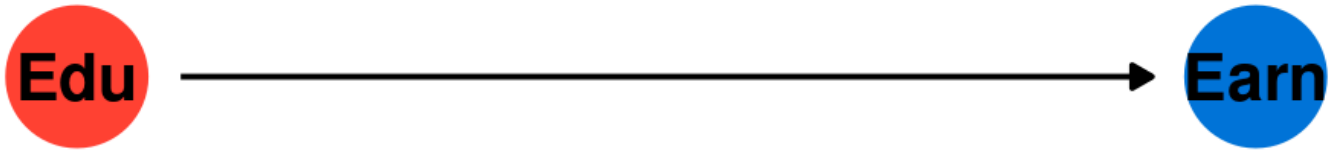

3. Draw arrows

Background, year of birth, location, job connections, and school requirements all cause education

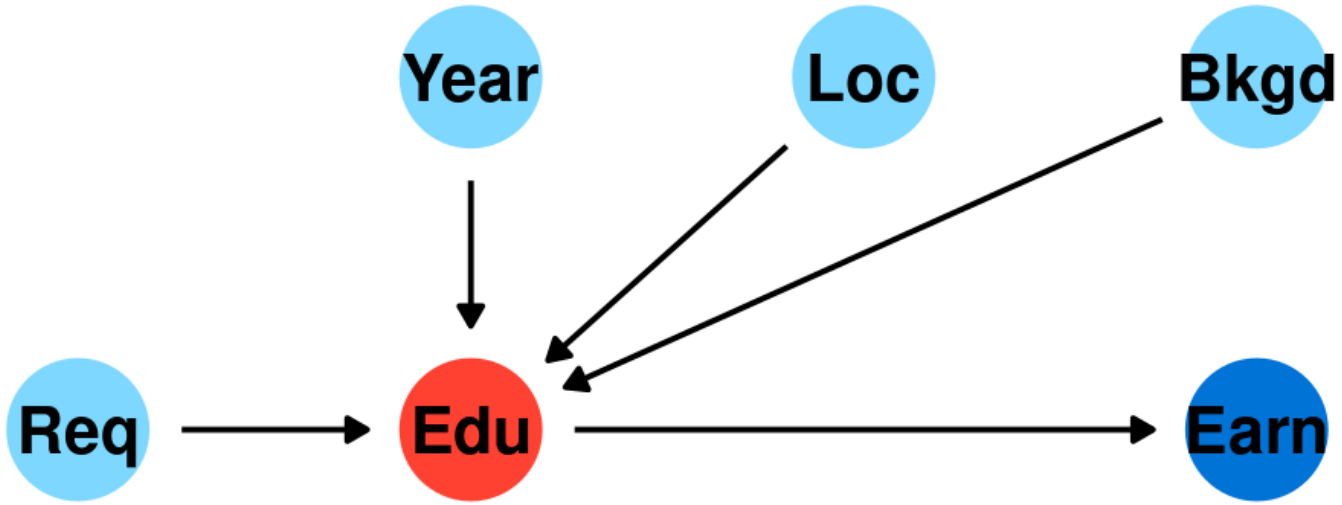

3. Draw arrows

Background, year of birth, and location all cause earnings too

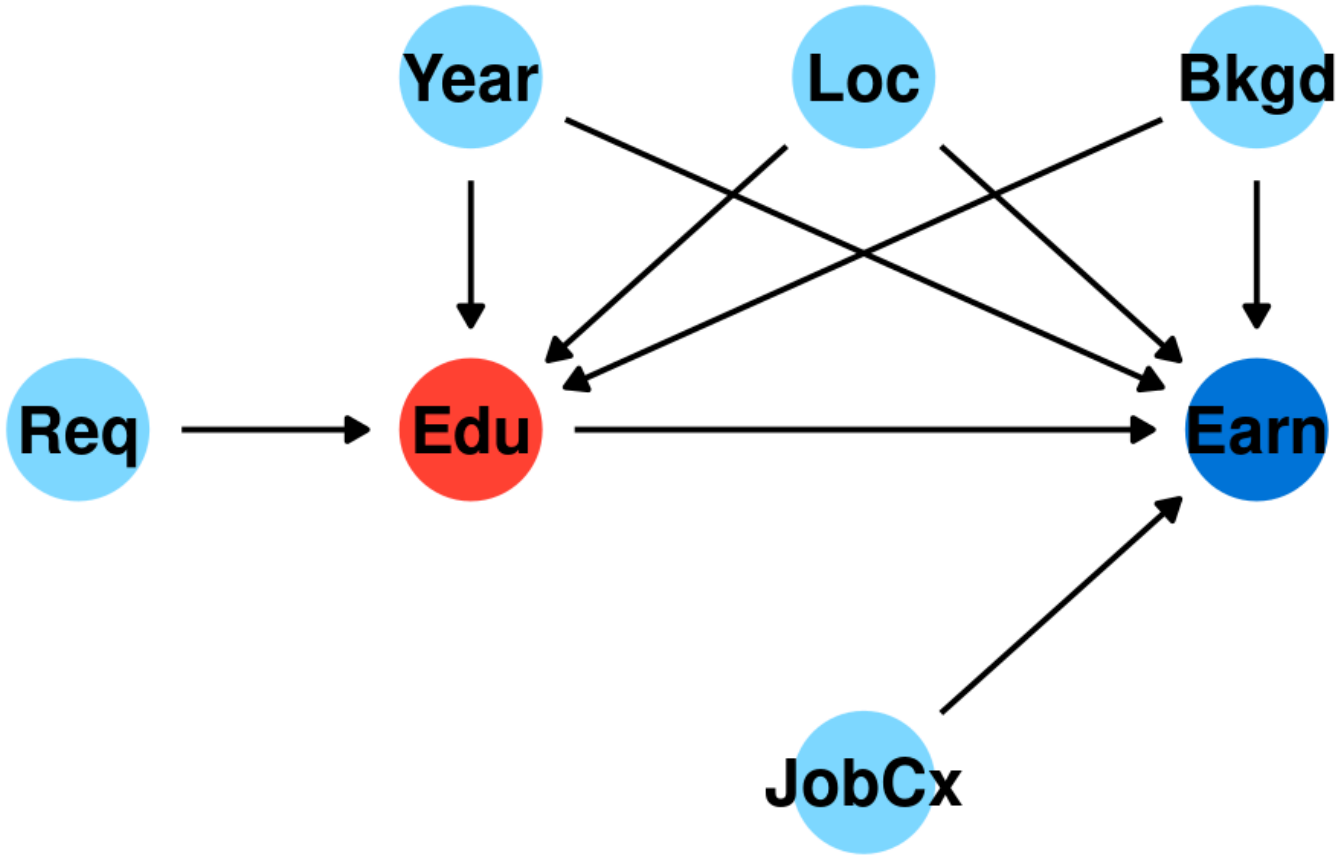

3. Draw arrows

Education causes job earnings

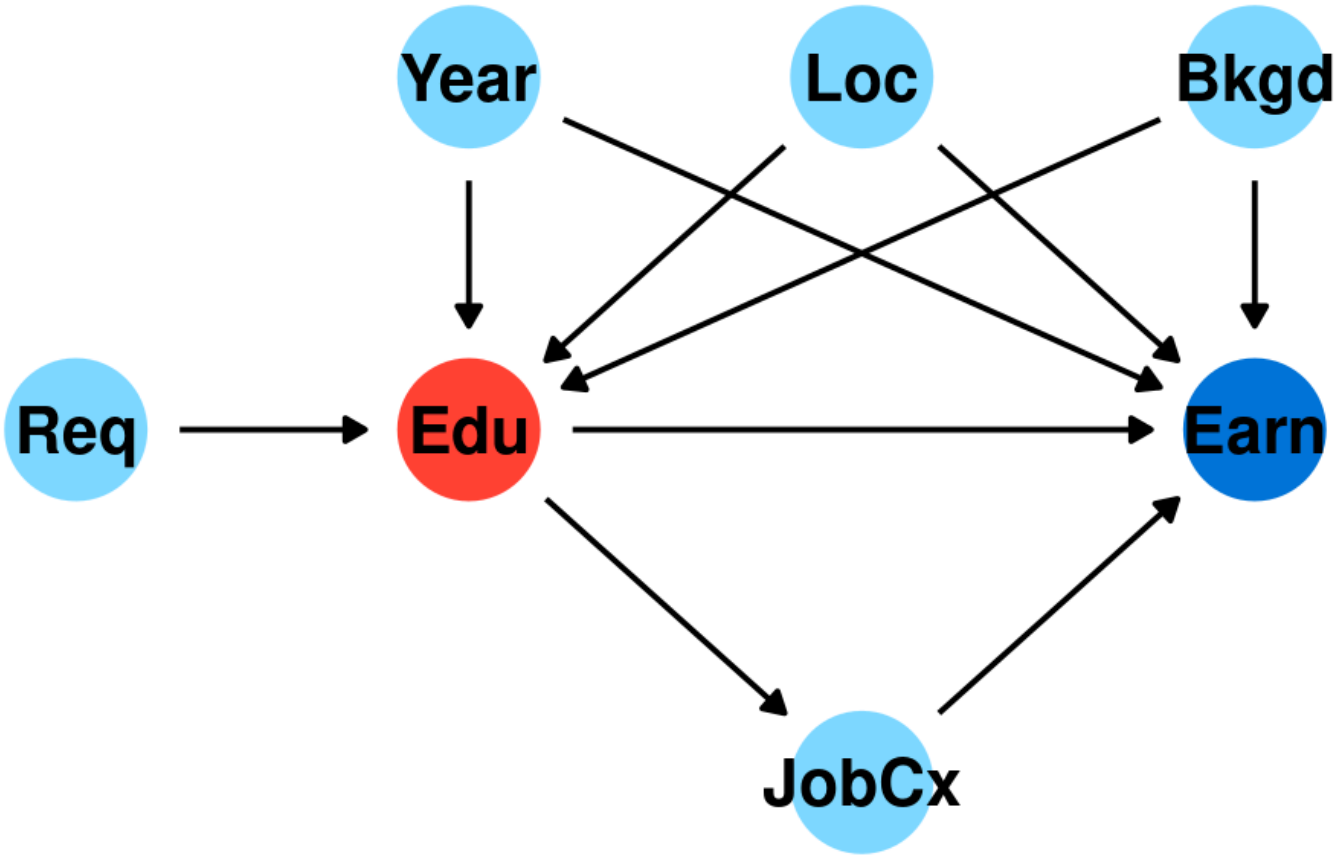

3. Draw arrows

Location and background are probably related, but neither causes the other. Something unobservable (U1) does that.

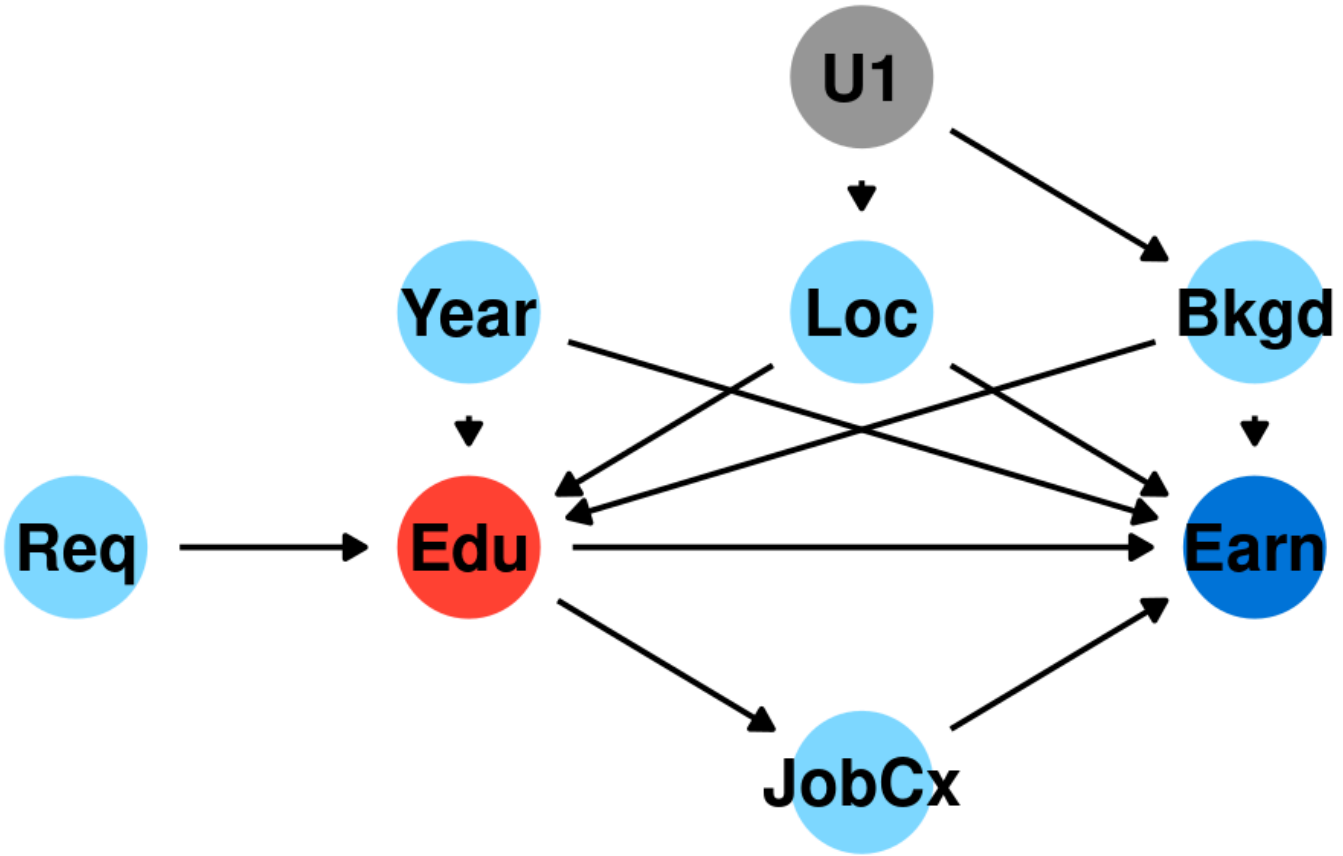

Causal identification

A causal effect is *identified* if the association between treatment and outcome is propertly stripped and isolated

Paths and associations

Arrows in a DAG transmit associations

You can redirect and control those paths by "adjusting" or "conditioning"

Three types of associations

Confounding

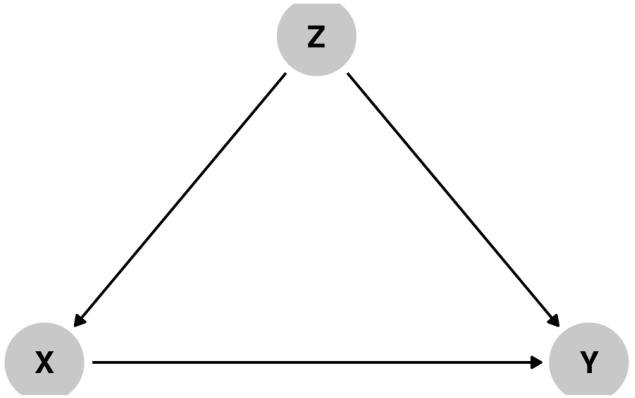

Common cause

Causation

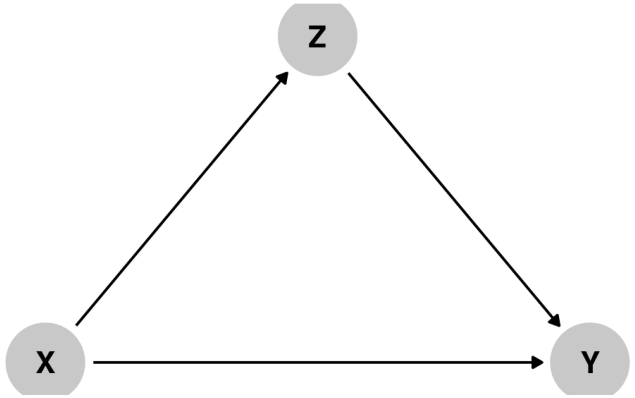

Mediation

Collision

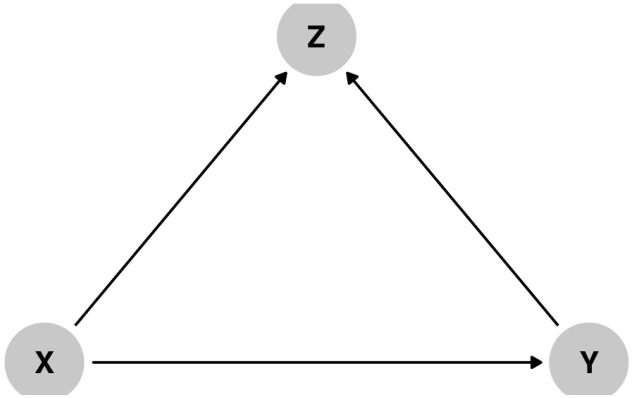

Selection / endogeneity

Confounding

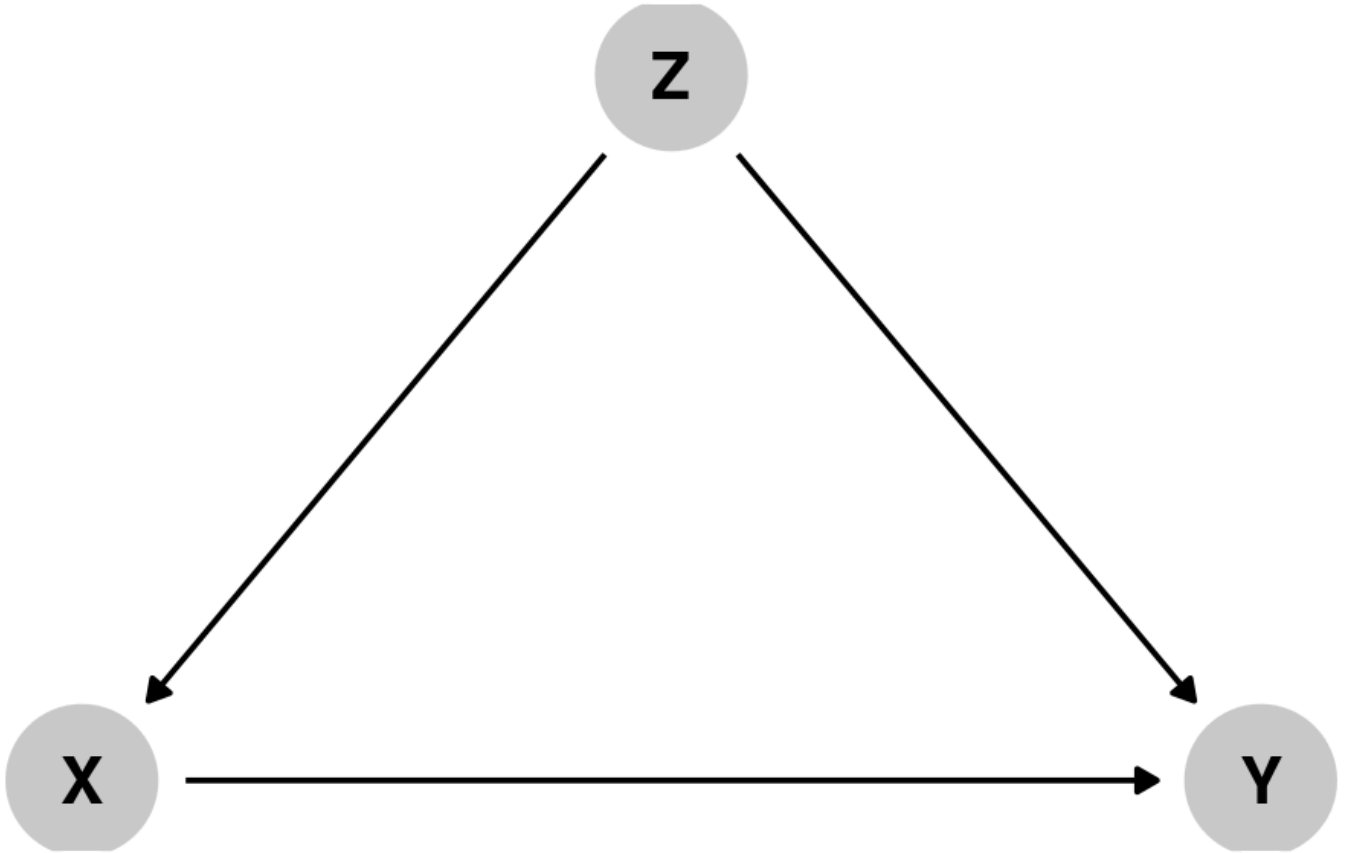

X causes Y

But Z causes both X and Y

Z *confounds* the X → Y association

Paths

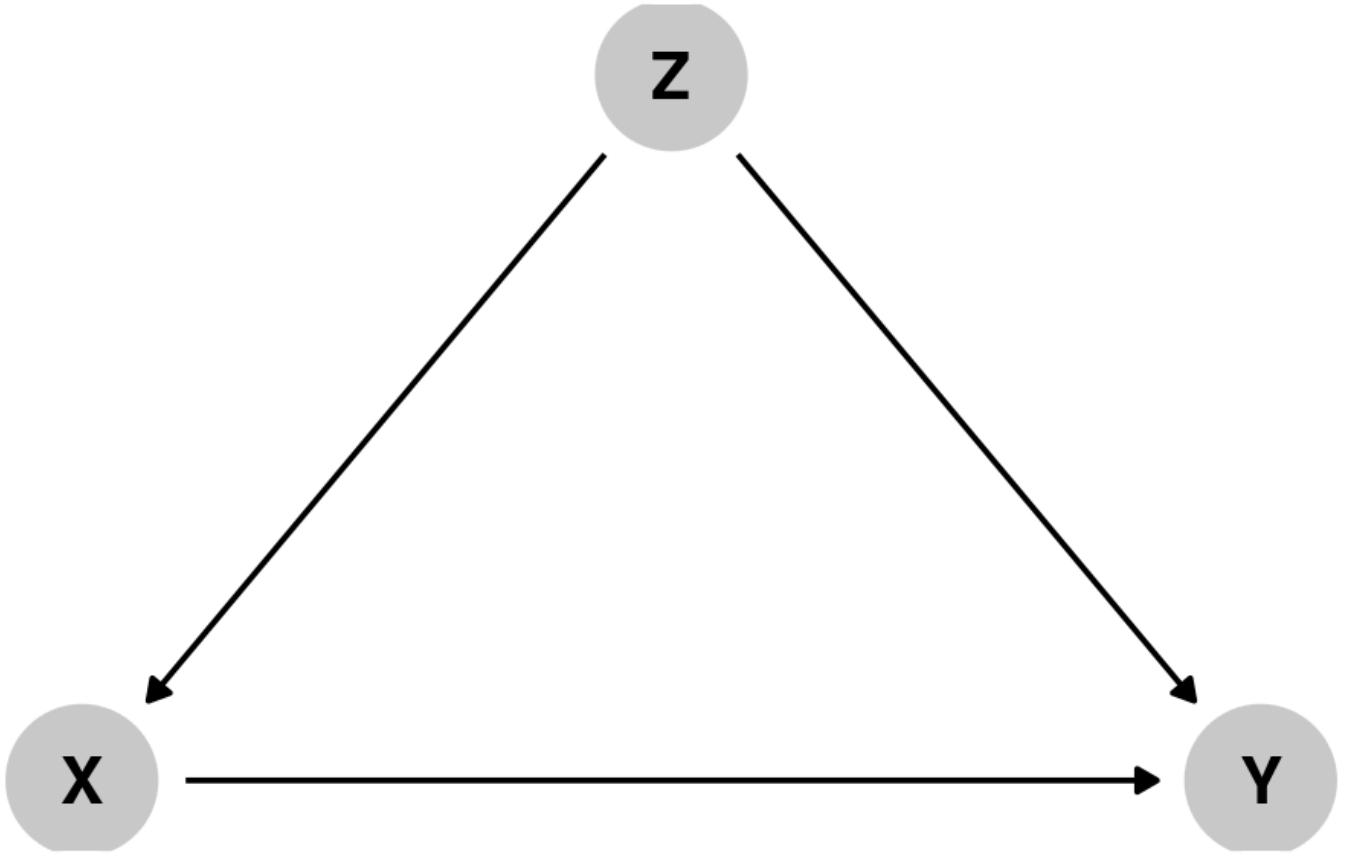

Paths between X and Y?

X → Y

X ← Z → Y

Z is a *backdoor*

d-connection

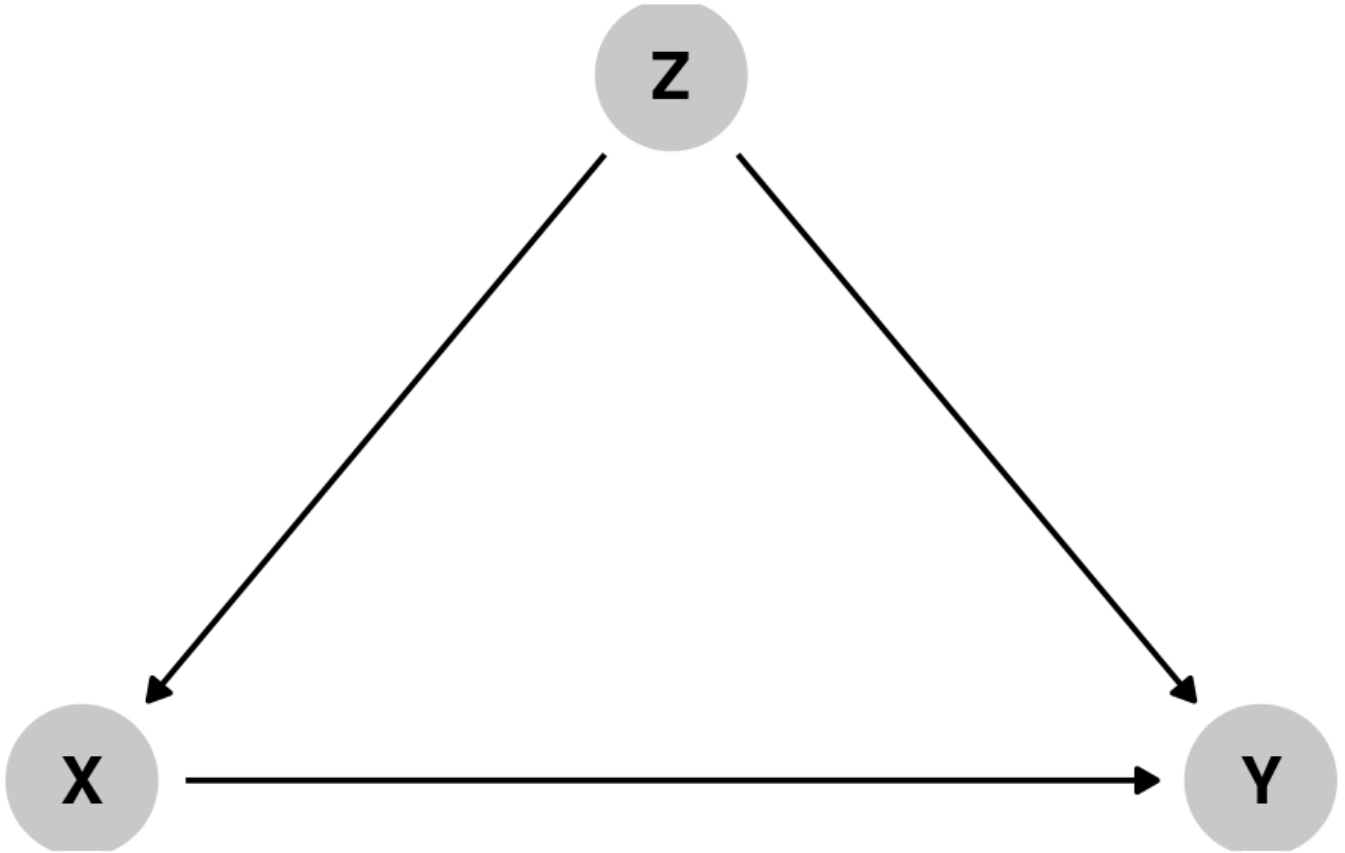

X and Y are "*d*-connected" because associations can pass through Z

The relationship between X and Y is not identified / isolated

Effect of money on elections

What are the paths between money and win margin?

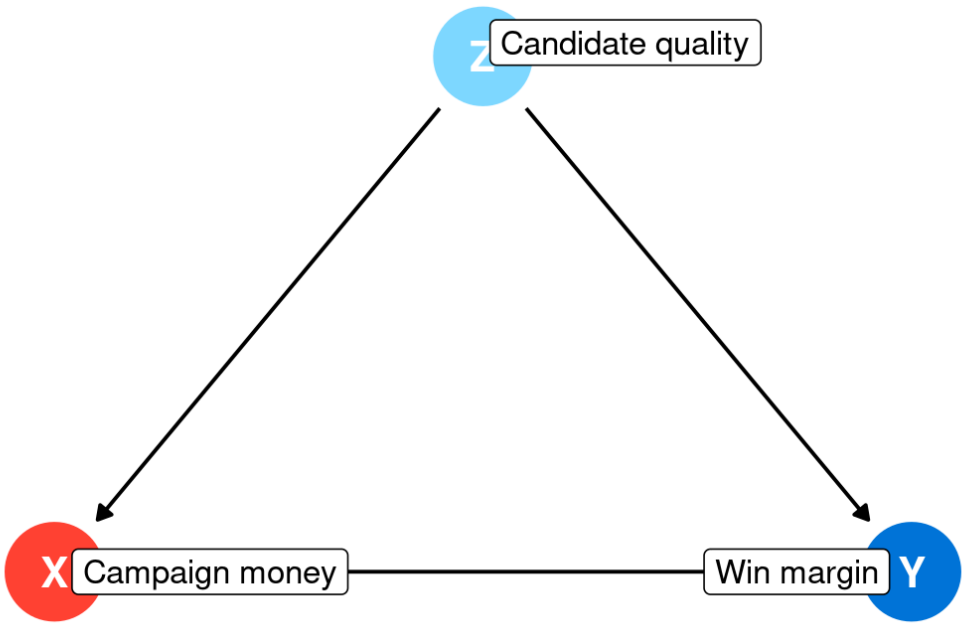

Money → Margin

Money ← Quality → Margin

Quality is a *backdoor*

Closing doors

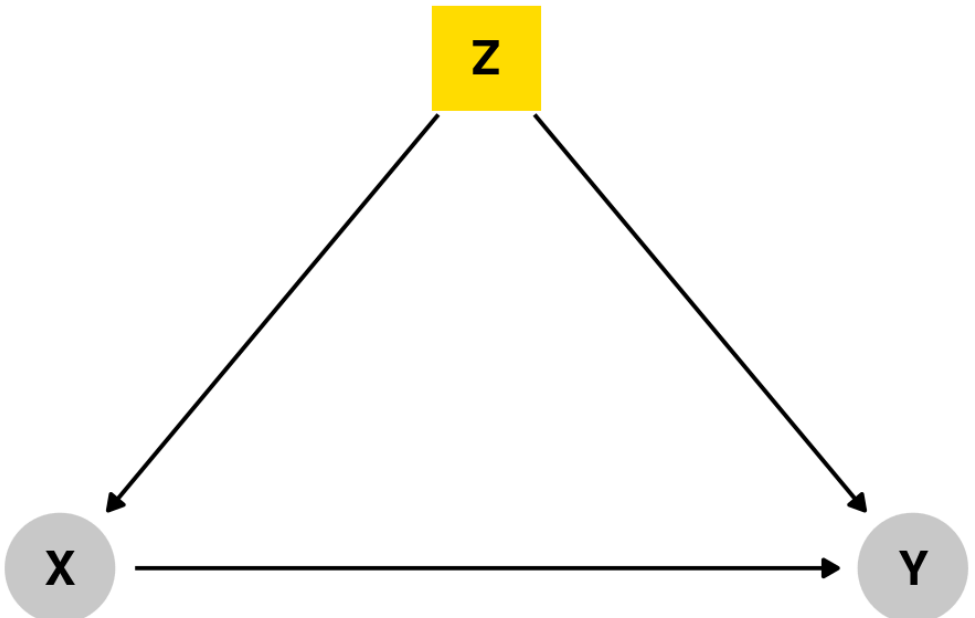

Close the backdoor by adjusting for Z

Closing doors

Find the part of campaign money that is explained by quality, remove it.
This is the residual part of money.

Find the part of win margin that is explained by quality, remove it. This is the residual part of win margin.

Find the relationship between the residual part of money and residual part of win margin.
This is the causal effect.

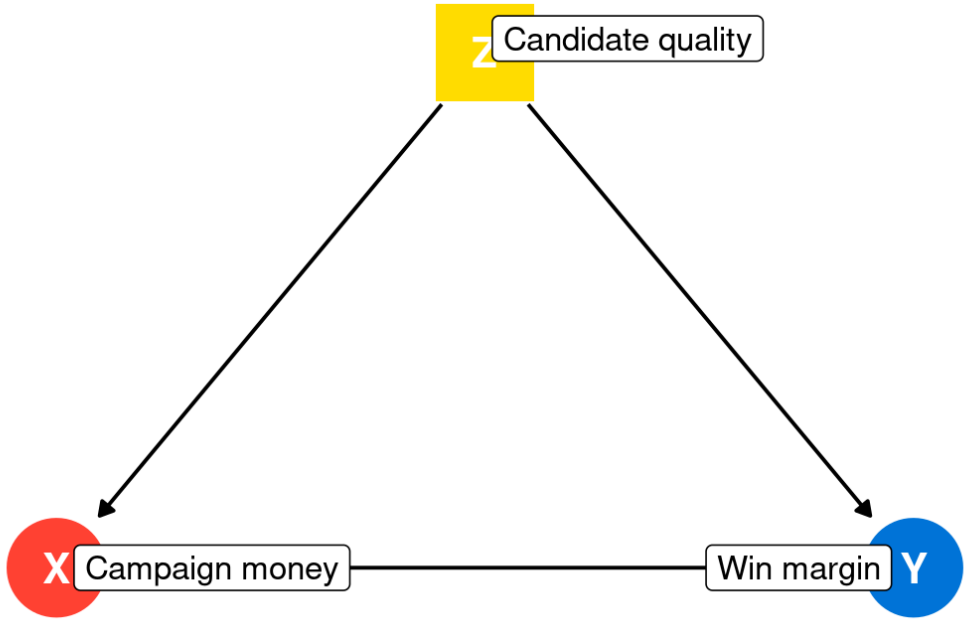

Closing doors

Compare candidates as if they had the same quality

Remove differences that are predicted by quality

Hold quality constant

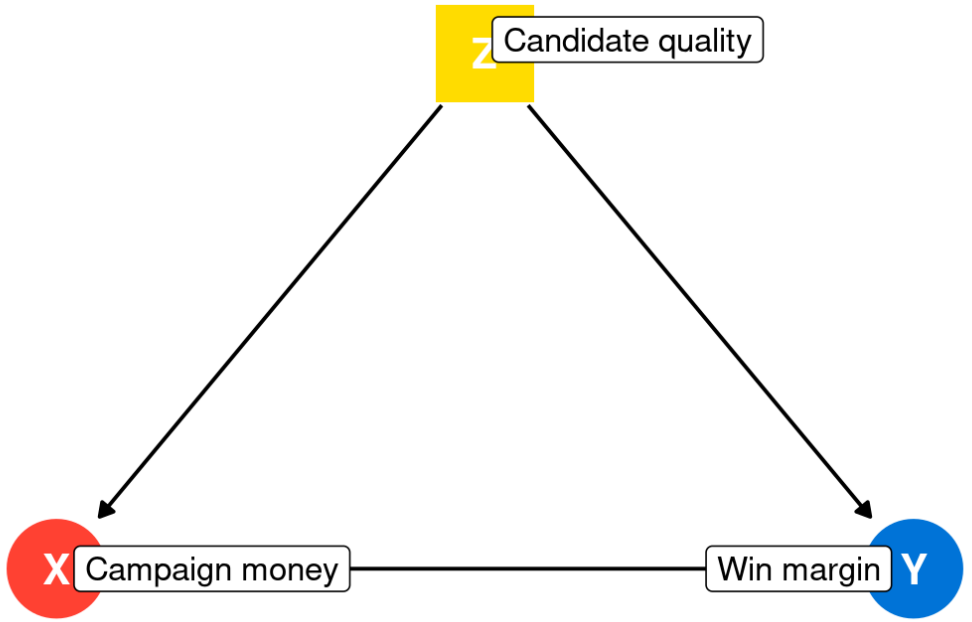

How to adjust

Include covariate in regression

Matching | **Stratifying** | **Inverse probability weighting**

d-separation

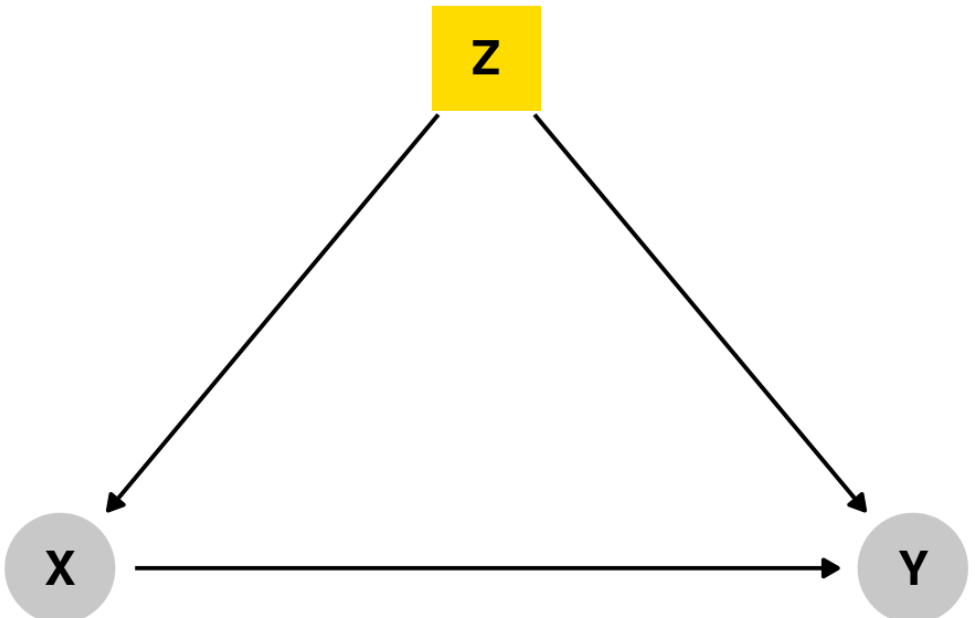

If we control for Z, X and Y are now "*d*-separated" and the association is isolated!

Closing backdoors

Block all backdoor paths to identify the main pathway you care about

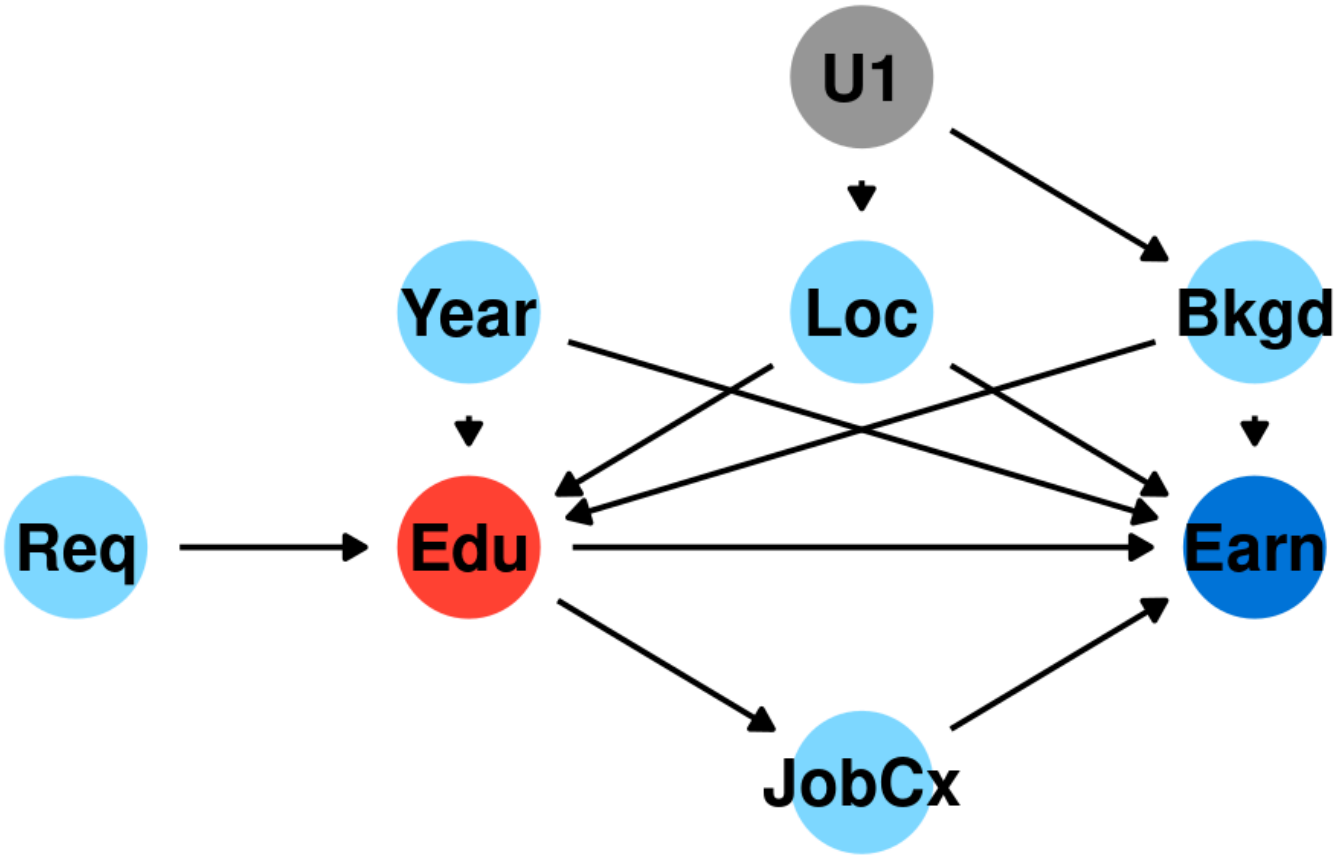

All paths

Education → Earnings

Education → Job connections → Earnings

Education ← Background → Earnings

Education ← Background ← U1 → Location → Earnings

Education ← Location → Earnings

Education ← Location ← U1 → Background → Earnings

Education ← Year → Earnings

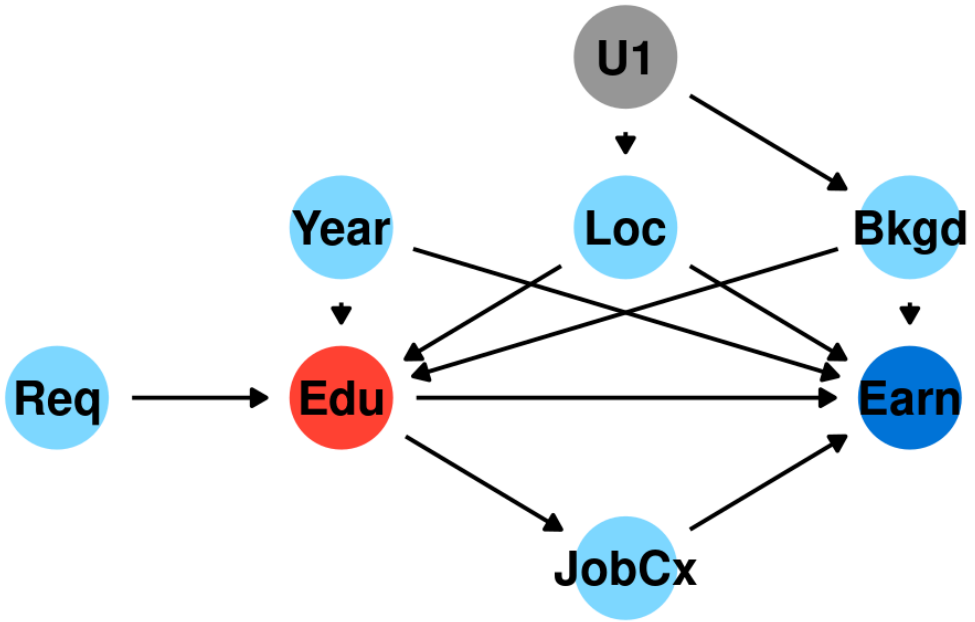

All paths

Adjust for **Location**, **Background** and **Year** to isolate the **Education → Earnings** causal effect

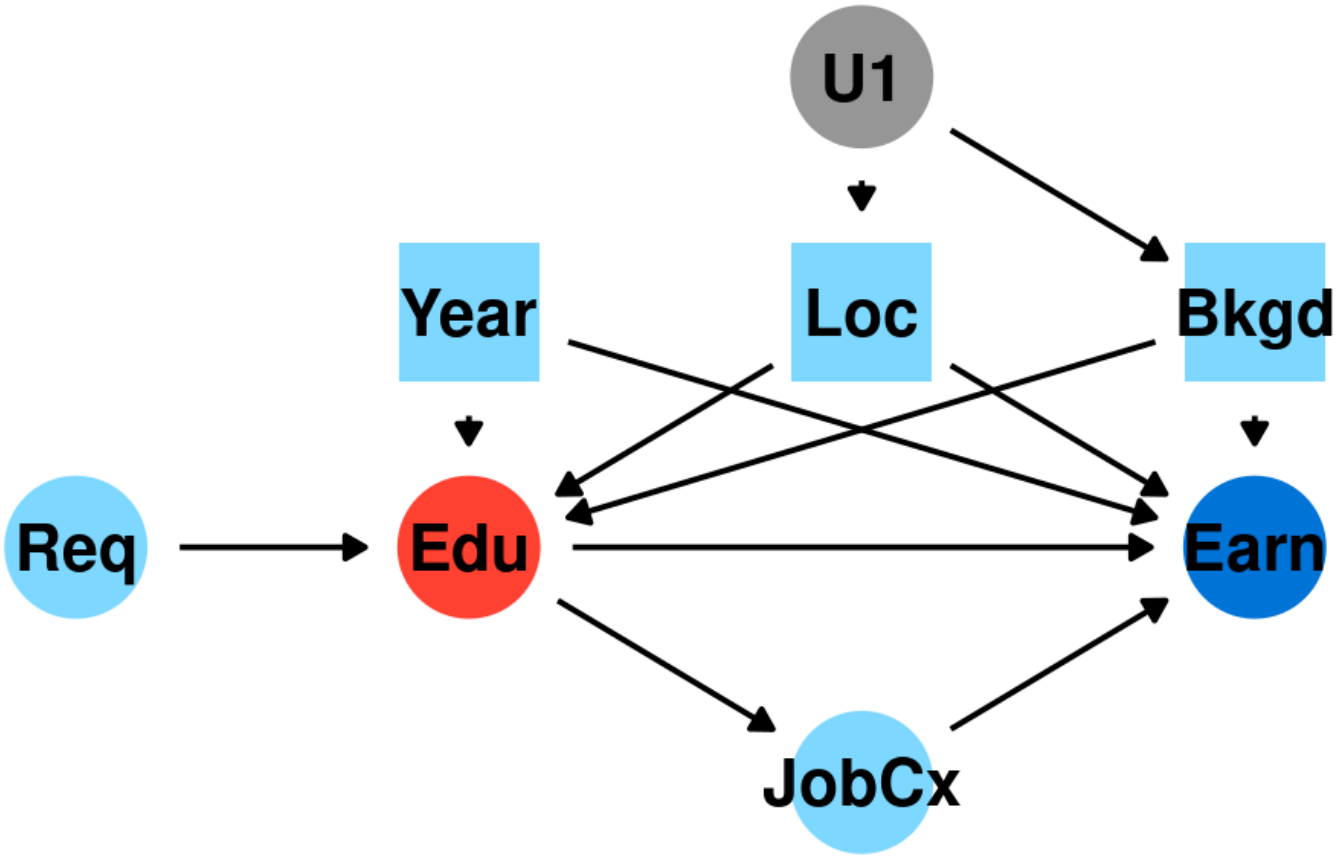

How do you know if this is right?

You can test the implications of the model to see if they're right in your data

$$X \perp\!\!\!\perp Y \mid Z$$

X is independent of Y, given Z

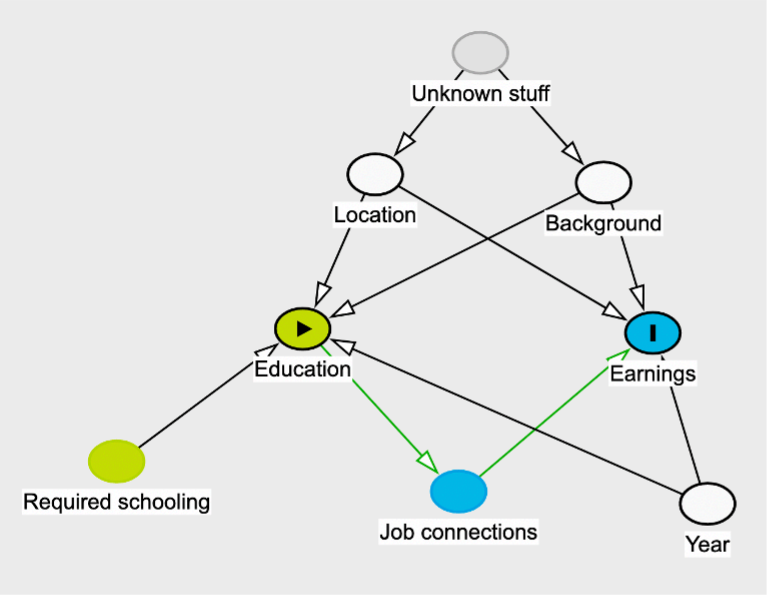

Testable implications

The model implies the following conditional independences:

- Education ⊥ Earnings | Background, Job connections, Location, Year
- Required schooling ⊥ Job connections | Education
- Required schooling ⊥ Year
- Required schooling ⊥ Earnings | Background, Job connections, Location, Year
- Required schooling ⊥ Earnings | Background, Education, Location, Year
- Required schooling ⊥ Background
- Required schooling ⊥ Location
- Job connections ⊥ Year | Education
- Job connections ⊥ Background | Education
- Job connections ⊥ Location | Education
- Year ⊥ Background
- Year ⊥ Location

Causation

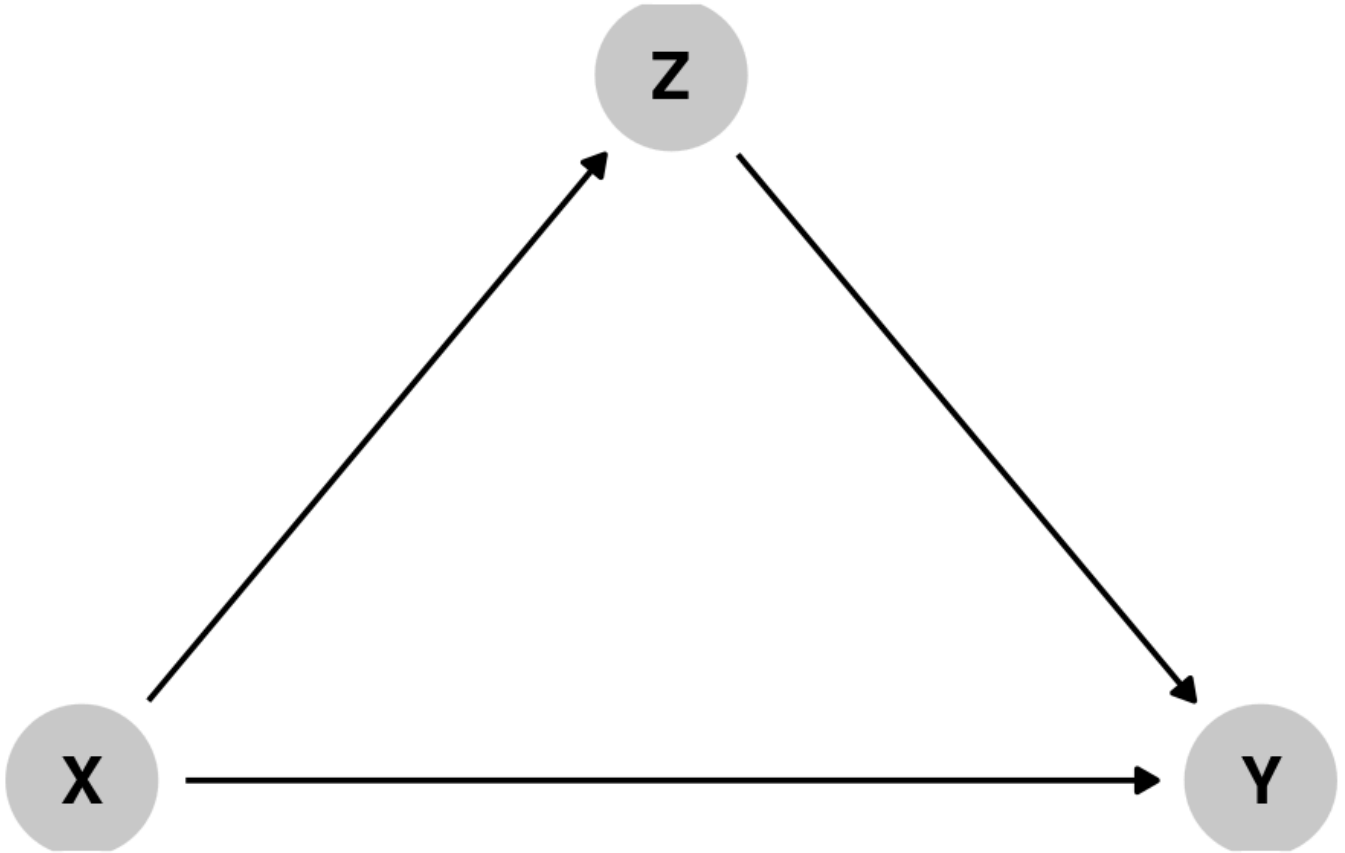

X causes Y

X causes Z which causes Y

Z is a mediator

Colliders

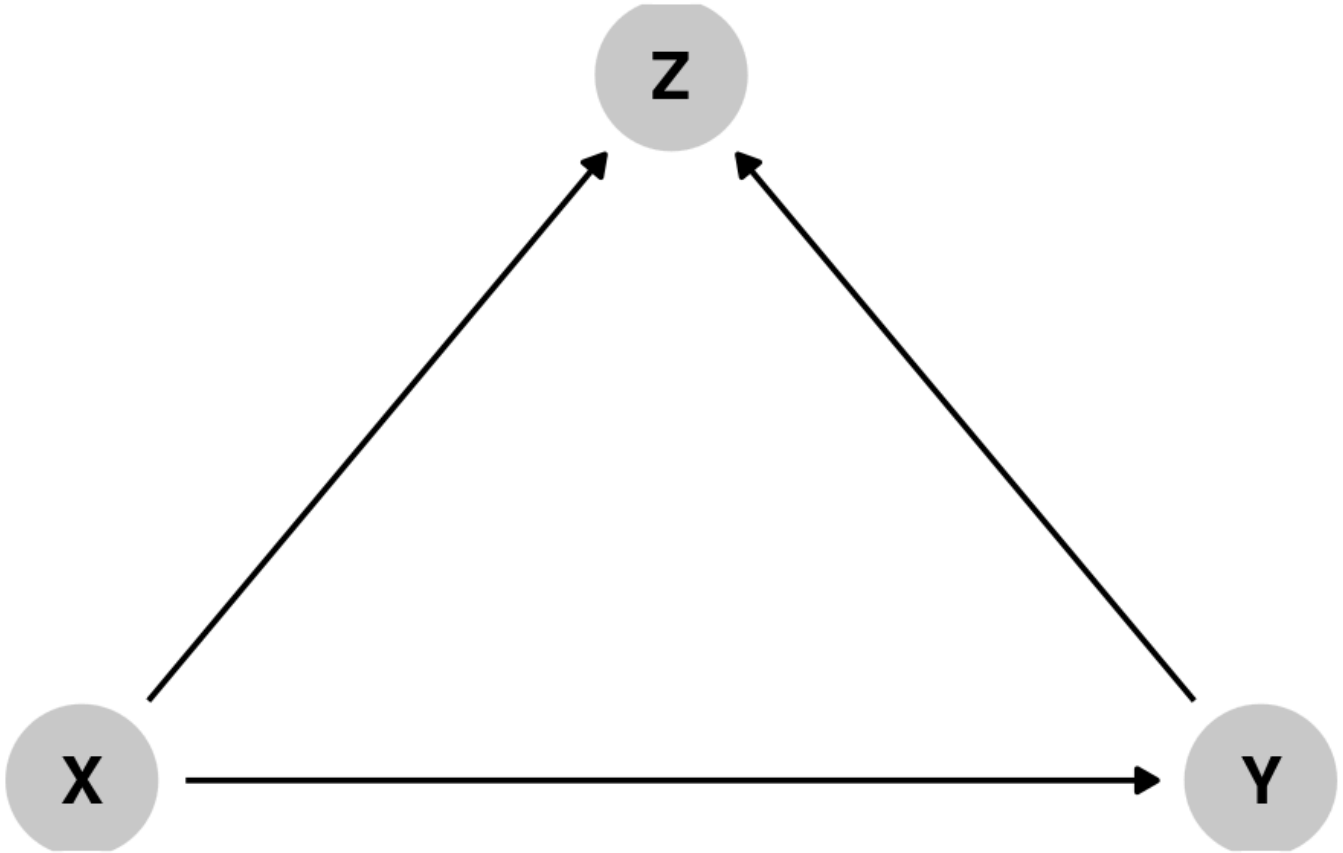

X causes Z

Y causes Z

Should you control for Z?

Programming and social skills

Do programming skills reduce social skills?

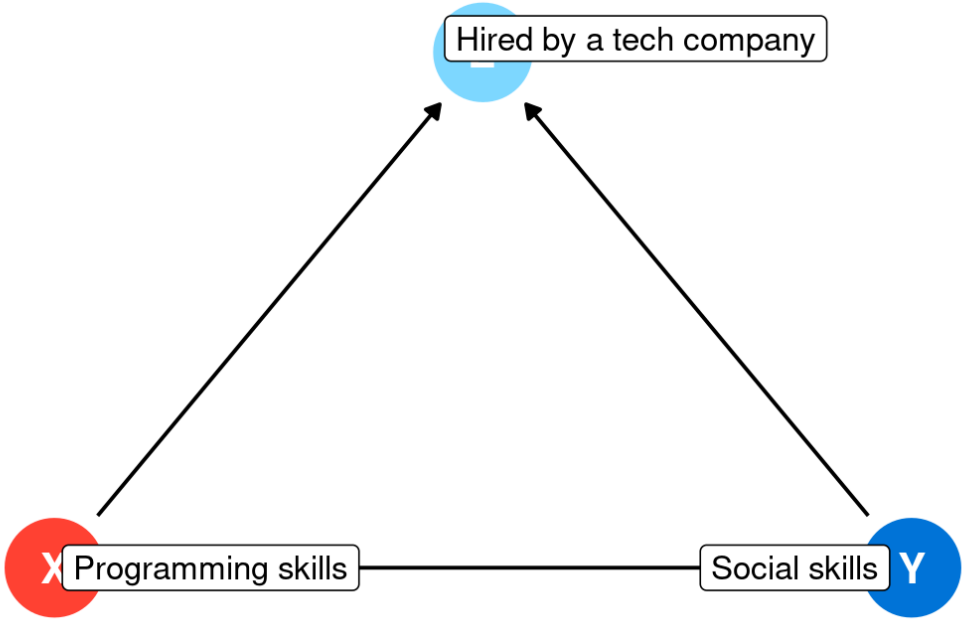

You go to a tech company and conduct a survey. You find a negative relationship!
Is it real?

Programming and social skills

Do programming skills reduce social skills?

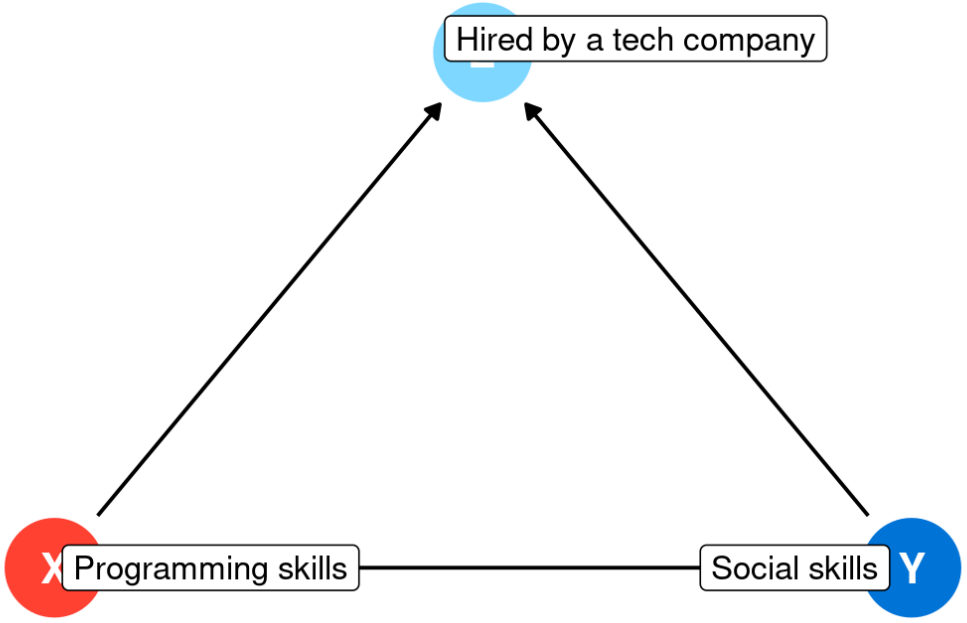

No! Hired by a tech company is a collider and we controlled for it.

This inadvertently connected the two.

Colliders can create fake causal effects

Colliders can hide real causal effects

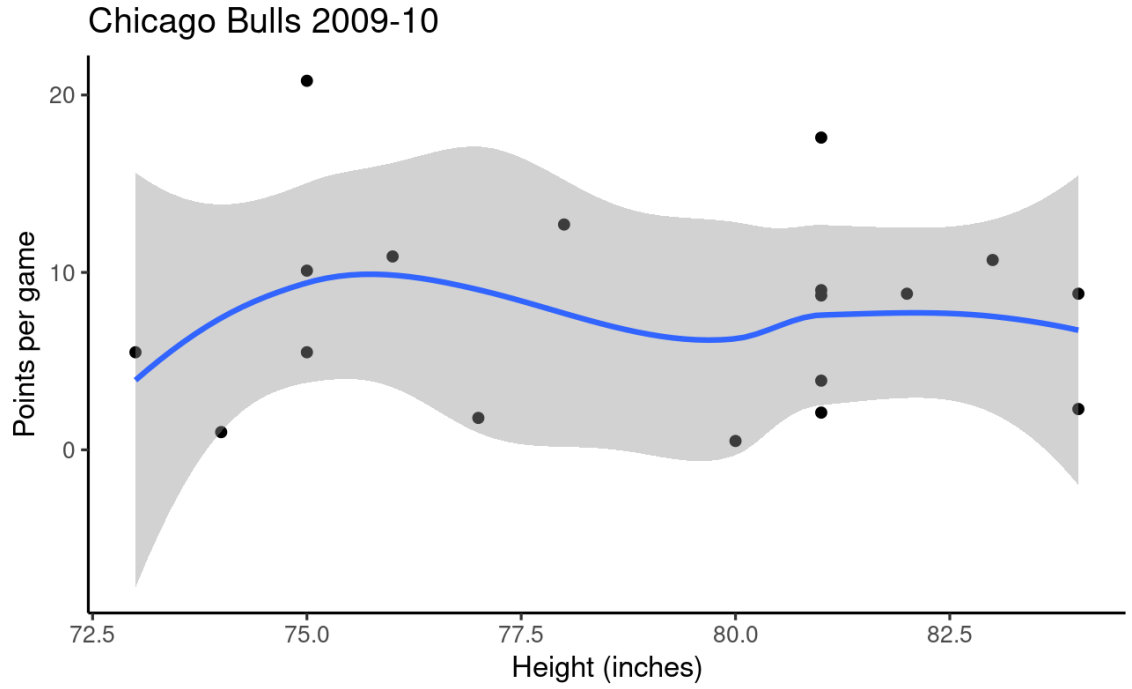

Height is unrelated to basketball skill... among NBA players

Colliders and selection bias

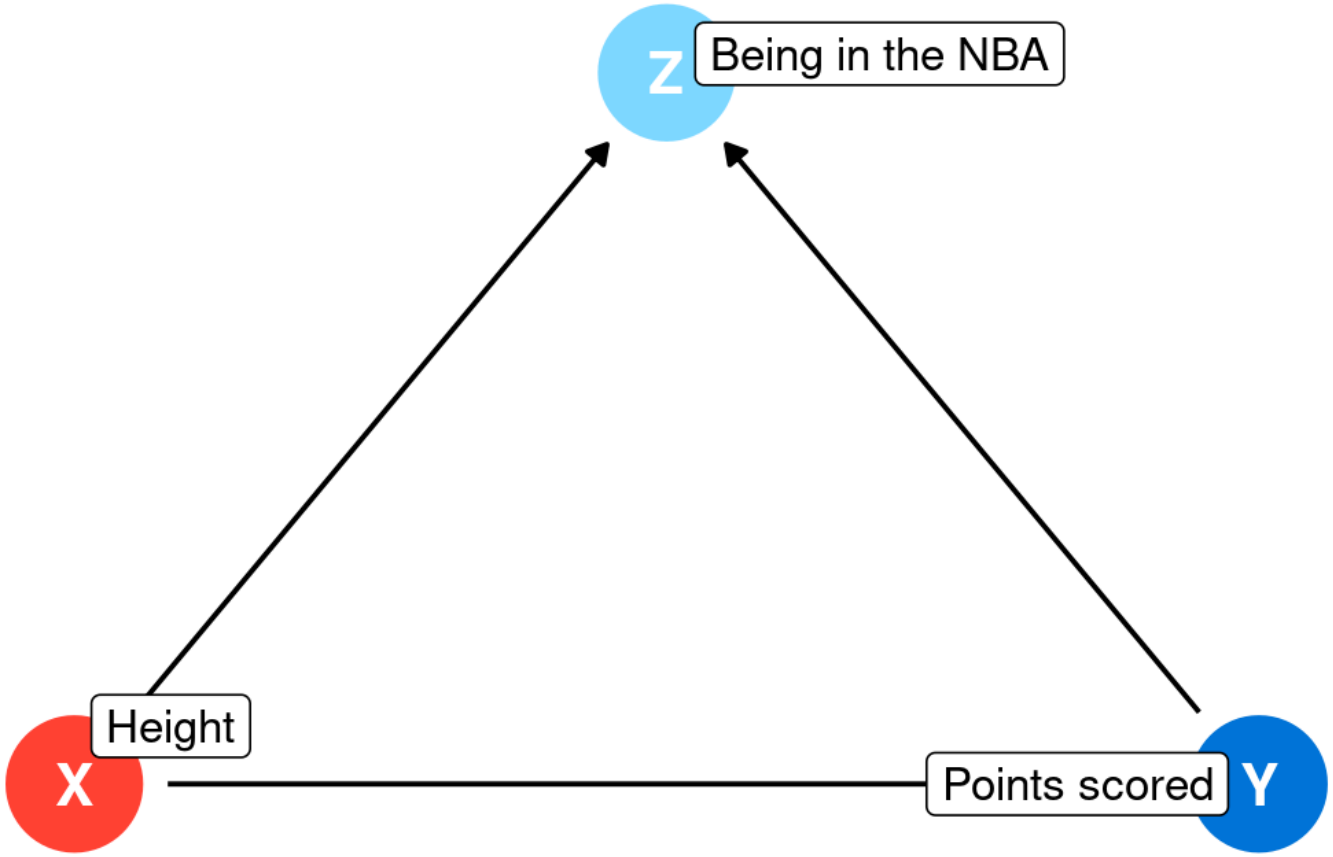

Conditioning on colliders

- Omnipresent in the literature
- Example: When and how does the number of children affect marital satisfaction? An international survey
- Example: The Predictive Validity of the GRE Across Graduate Outcomes

Three types of associations

Confounding

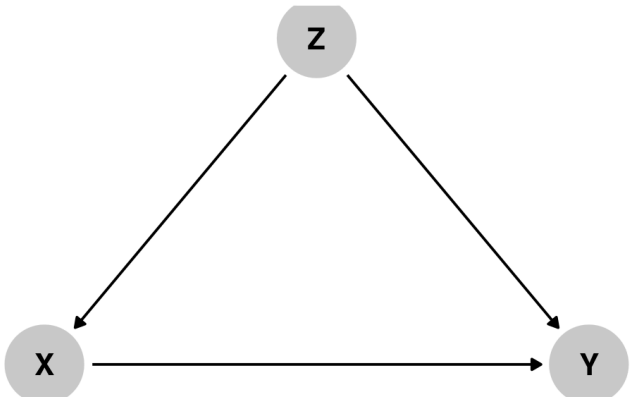

Common cause

Causal forks $X \leftarrow Z \rightarrow Y$

Causation

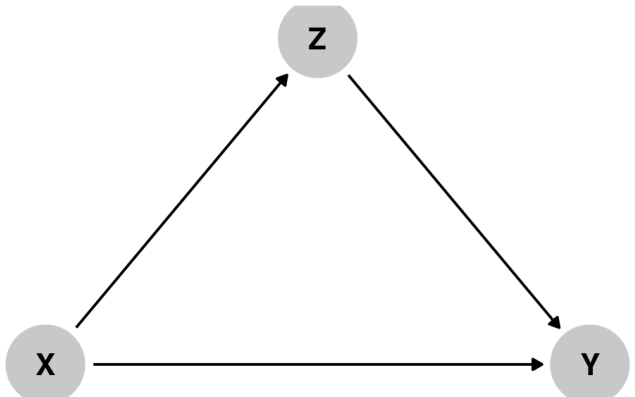

Mediation

Causal chain $X \rightarrow Z \rightarrow Y$

Collision

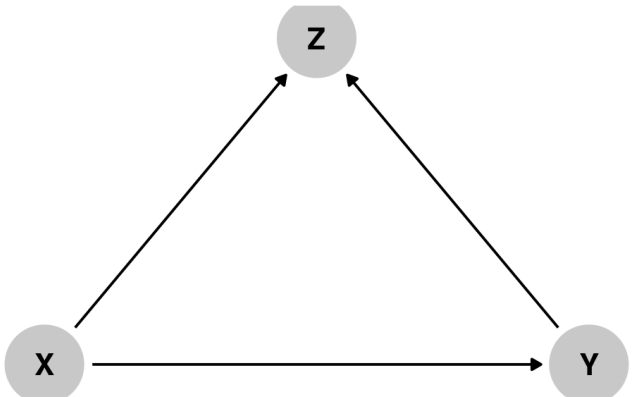

Selection / endogeneity

inverted fork $X \rightarrow Z \leftarrow Y$

Life is inherently complex

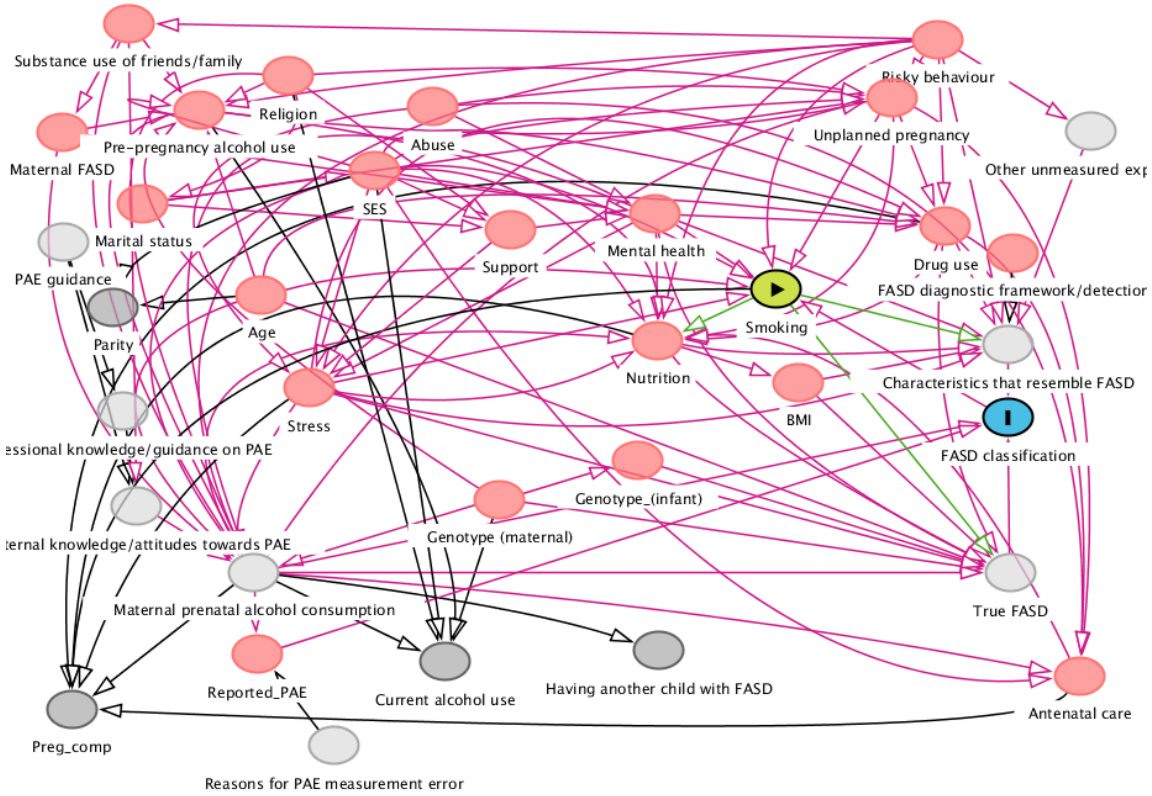

Postulated DAG for the effect of smoking on fetal alcohol spectrum disorders (FASD)

Causal mediation

Key references

- Imai, Keele and Tingley (2010), A General Approach to Causal Mediation Analysis, *Psychological Methods*.
- Pearl (2014), Interpretation and Identification of Causal Mediation, *Psychological Methods*.
- Baron and Kenny (1986), The Moderator-Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic, and Statistical Considerations, *Journal of Personality and Social Psychology*

Limitations:

- Bullock, Green, and Ha (2010), Yes, but what's the mechanism? (don't expect an easy answer)
- Uri Simonsohn (2022) Mediation Analysis is Counterintuitively Invalid

Sequential ignorability assumption

- potential mediation given treatment x as $M_i(x)$ and
- potential outcome for treatment x and mediator m as $Y_i(x, m)$.

Given pre-treatment covariates $\boldsymbol{Z}$, potential outcomes for mediation and treatment are conditionally independent of treatment assignment.

$$Y_i(x', m), M_i(x) \perp\!\!\!\perp X_i \mid \boldsymbol{Z}_i = \boldsymbol{z}$$

Given pre-treatment covariates and observed treatment, potential outcomes are independent of mediation.

$$Y_i(x', m) \perp\!\!\!\perp M_i(x) \mid X_i = x, \boldsymbol{Z}_i = \boldsymbol{z}$$

Total effect

Total effect: overall impact of X (both through M and directly)

$$\mathsf{TE}(x, x^*) = \mathsf{E}[Y \mid \mathrm{do}(X = x)] - \mathsf{E}[Y \mid \mathrm{do}(X = x^*)]$$

X → M → Y
plus
X → Y

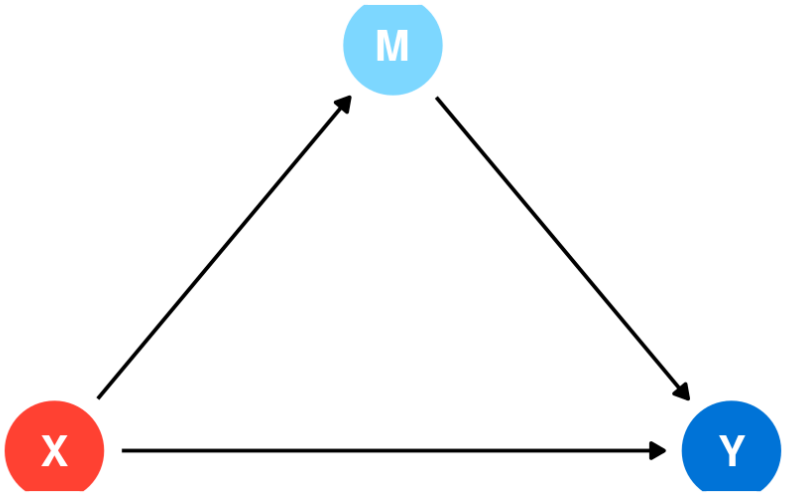

Average controlled direct effect

$$\begin{aligned}\mathsf{CDE}(m, x, x^*) &= \mathsf{E}[Y \mid \mathrm{do}(X = x, m = m)] - \mathsf{E}[Y \mid \mathrm{do}(X = x^*, m \\ &= \mathsf{E}\{Y(x, m) - Y(x^*, m)\}\end{aligned}$$

Expected population change in response when the experimental factor changes from x to x^* and the mediator is set to a fixed value m.

Problem: this forces manipulation of the mediator, and only gives outcome for a fixed value m.

Direct and indirect effects

Natural direct effect:

$$\mathsf{NDE}(x, x^*) = \mathsf{E}[Y\{x, M(x^*)\} - Y\{x^*, M(x^*)\}]$$

- expected change in Y under treatment x if M is set to whatever value it would take under control x^*

Natural indirect effect:

$$\mathsf{NIE}(x, x^*) = \mathsf{E}[Y\{x^*, M(x)\} - Y\{x^*, M(x^*)\}]$$

- expected change in Y if we set X to its control value and change the mediator value which it would attain under x

Counterfactual conditioning reflects a physical intervention, not mere (probabilistic) conditioning.

Necessary and sufficiency of mediation

From Pearl (2014):

> The difference $\mathsf{TE} - \mathsf{NDE}$ quantifies the extent to which the response of Y is owed to mediation, while NIE quantifies the extent to which it is explained by mediation. These two components of mediation, the necessary and the sufficient, coincide into one in models void of interactions (e.g., linear) but differ substantially under moderation

- In linear systems, changing the order of arguments amounts to flipping signs
- This definition works under temporal reversal and gives the correct answer (the regression-slope approach of the linear structural equation model does not).