Epi 207: Introduction to Using Directed Acyclic Graphs

Maria Glymour
Department of Epidemiology & Biostatistics
University of California, San Francisco

Organization

1. Drawing and using DAGs
2. DAGs for common problems in epidemiology
   - Confounding
   - RCTs
   - Selection biases
   - Measurement error
   - Missing Data
   - Mediation
3. Contrasting IV-based methods and covariate adjustment/propensity scores
4. Limitations and controversies of counterfactuals and DAGs

Counterfactuals or Potential Outcomes

- Represent the outcome value an individual would have had under a particular treatment regime.
- Only observed if the individual actually has that particular treatment regime. All other potential outcomes for that person are defined but not observed.
- X can be a cause of Y even if it is neither necessary nor sufficient to produce Y for some individuals
- Conceptualize individuals as:
  - **Immune**: $Y_{x=0} = Y_{x=1} = 0$
  - **Causative**: $Y_{x=0} = 0$ $Y_{x=1} = 1$
  - **Protective**: $Y_{x=0} = 1$ $Y_{x=1} = 0$
  - **Doomed**: $Y_{x=0} = Y_{x=1} = 1$

Inferring Causation From Association

Statistical *association* between two variables X and Y may be due to:

1. Random fluctuation
2. X caused Y
3. Y caused X
4. X and Y share a common cause
5. The statistical association was induced by conditioning on a common effect of X and Y (as in selection bias).

How we can use this

- Eliminating four of these explanations is usually the goal of an analysis.
- Knowing these five sources of statistical association helps identify the (set of) causal structure(s) that could have generated the observed statistical associations.

Petersen's causal roadmap

1. Specify knowledge about the system to be studied using a causal model.
Represent background knowledge about the system to be studied. A causal model describes the set of possible data-generating processes for this system.

2. Specify the observed data and their link to the causal model: Specify what variables have been or will be measured, and how these variables are generated by the system described by the causal model.

3. Specify a target causal quantity: Translate the scientific question into a formal causal quantity (defined as some parameter of the distribution of counterfactuals)

4. Assess identifiability: Assess whether it is possible to represent the target causal quantity as a parameter of the observed data distribution (estimand), and, if not, what further assumptions would allow one to do so.

5. State the statistical estimation problem: Specify the estimand and statistical model. If knowledge is sufficient to identify the causal effect of interest, commit to the corresponding estimand. If not, choose an estimand that under minimal additional assumptions would equal or approximate the causal effect of interest.

6. Estimate: Estimate the target parameter of the observed data distribution, respecting the statistical model.

7. Interpret: Select among a hierarchy of interpretations, ranging from purely statistical to approximating a hypothetical randomized trial.

Petersen and Van der Laan, 2014

Drawing and using causal DAGs

- Basic rules
- The d-separation criterion
- A tangent on collider bias
- Criteria for testing the null and identifying effects

Causal Directed Acyclic Graphs

Show your assumptions about the causal relationships among X, Y, and possible covariates in a causal diagram:

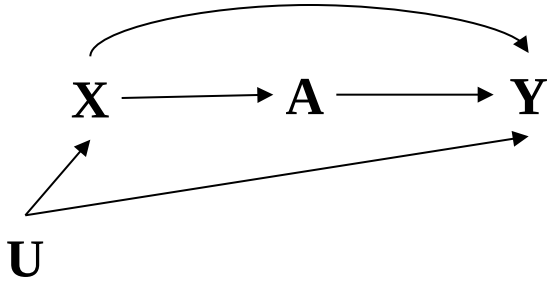

•If two variables shown in the graph have a common cause, you must show the cause in the graph.
•Do not allow causal "loops".

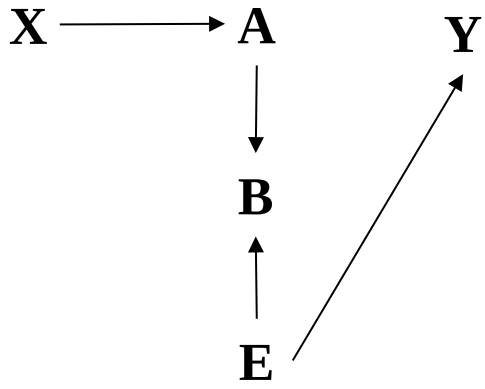

•Useful to think of DAGs as a set of (vague) data generating rules.
•Each variable is generated as a function of its parents and noise/unmeasured variables.

$$X=f(A, \delta_X)$$
$$Y=f(X,U,\delta_Y)$$

•Helpful as a starting point when specifying simulations.

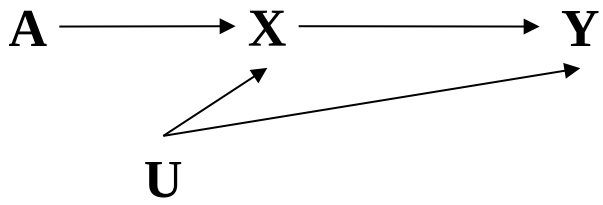

Terminology

- Descendants
  - The direct or indirect effects of a variable
- Paths
  - A sequence of lines (edges) between two variables, regardless of direction of arrows
  - Not retracing any line segments
- Colliders
  - Common effect of two variables in a path: where the arrows 'collide'.
  - The two causes must both be "on the path".
  - Any variable on a path that is not a collider is a "non-collider".
- Conditioning
  - Examining the distribution of one variable within levels of another
  - Regression adjustment, stratification, restriction

Colliders vs Non-Colliders

Colliders: common *effects*

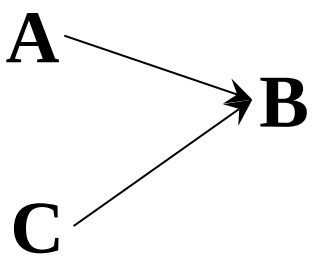

Non-Colliders:
common *causes* (=confounders)

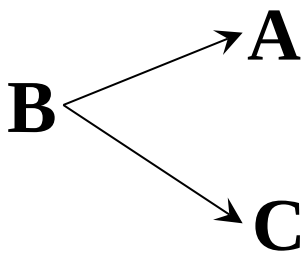

Or mediators

D-separation

- The assumptions shown in a causal diagram imply that a variable X will be independent of a variable Y, after conditioning on a set of variables {Z} if every path between X and Y is blocked by {Z}.

- The set of covariates {Z} blocks a path if and only if either:
  1. The path contains a non-collider that is in {Z} , or
  2. The path contains a collider which is not in {Z} , and no descendant of the collider is in {Z} .

- If there is an unblocked path linking X and Y, then X and Y will typically be statistically dependent (unless there is a perfectly offsetting balance between two paths).

D-separation: intuition

- There may be many reasons that two variables are associated (some confounding, some mediated causation etc).
- Adjusting for a confounder of the two variables blocks that source of association between two variables
- Adjusting for a mediator between the two variables blocks that source of association between two variables
- Adjusting for a common effect of the two variables *creates* an association between the two variables

Example Causal Diagrams

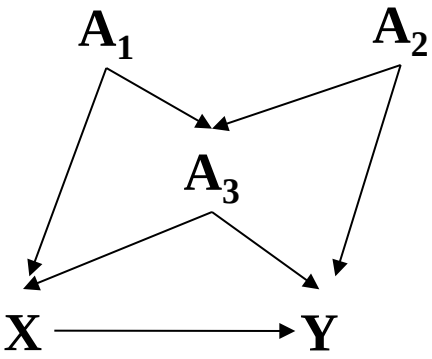

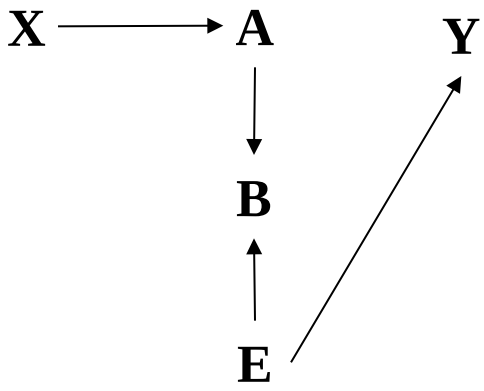

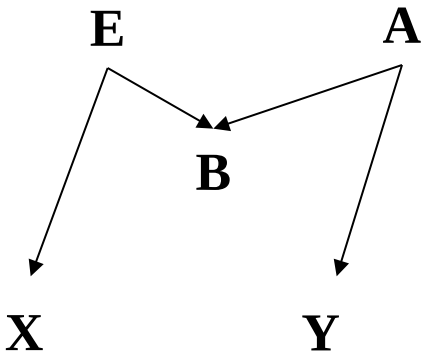

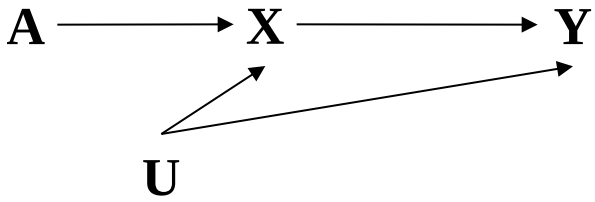

DAGs are data generating rules.

A

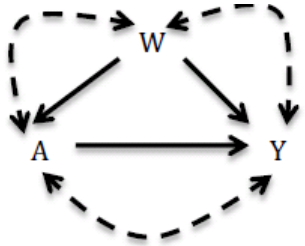

$W = f_W(U_W)$
$A = f_A(W, U_A)$
$Y = f_Y(W, A, U_Y)$
No restrictions on the distribution of $U = (U_W, U_A, U_Y)$

No restrictions on the distribution of $U = (U_W, U_A, U_Y)$

Causal graphs and corresponding structural equations representing alternative data generating processes.
(A), Assumes that W may have affected A, both A and W may have affected Y, and W, A, and Y may all share unmeasured common causes.
The double headed arrow is a convention some use to represent an unmeasured common cause or correlated errors, e.g., the double headed arrow between A and W implies a correlation between U_w and U_A, which is presumably attributable to a common cause of U_w and U_A.

If you wanted to generate data, for example (here ignoring the correlated errors):

gen Uw=rnormal()

gen Ua=rnormal()

gen Uy=rnormal()

gen W=3*Uw+2

gen A=.1*W+Ua or gen A=.7*W+Ua*Ua

gen Y=1*W+4*A+.3*Uy or gen Y=1*W+4*A-2*W*A+Uy

Petersen and Van der Laan, 2014

DAGs are data generating rules.

A

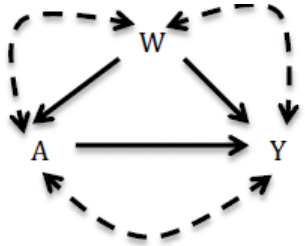

$W = f_W(U_W)$
$A = f_A(W, U_A)$
$Y = f_Y(W, A, U_Y)$
No restrictions on the distribution of $U = (U_W, U_A, U_Y)$

No restrictions on the distribution of $U = (U_W, U_A, U_Y)$

gen W=3*Uw+2

gen A=.1*W+Ua

or

gen A=.7*W+Ua*Ua

gen Y=1*W+4*A+.3*Uy

or

gen Y=1*W+4*A-2*W*A+Uy

The double headed arrow is a convention some use to represent an unmeasured common cause or correlated errors, e.g., the double headed arrow between A and W implies a correlation between U_w and U_A, which is presumably attributable to a common cause of U_w and U_A.

If you wanted to generate data, incorporating the correlated errors):

gen U1=rnormal() /* cause of A and W*/

gen U2=rnormal() /* cause of A and Y*/

gen U3=rnormal() /* cause of W and Y*/

gen Uw=rnormal()+U1+U3

gen Ua=rnormal()+U1+U2

gen Uy=rnormal()+U2+U3

Petersen and Van der Laan, 2014

DAGs are data generating rules.

A

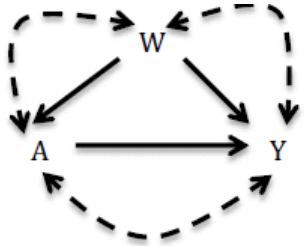

$W = f_W(U_W)$
$A = f_A(W, U_A)$
$Y = f_Y(W, A, U_Y)$
No restrictions on the distribution of $U = (U_W, U_A, U_Y)$

No restrictions on the distribution of $U = (U_W, U_A, U_Y)$

Causal graphs and corresponding structural equations representing alternative data generating processes.

(A), Assumes that W may have affected A, both A and W may have affected Y, and W, A, and Y may all share unmeasured common causes

gen Y=1*W+4*A-2*W*A+Uy

This specifies an interaction between W and A in determining Y. Is that consistent with the DAG?

Petersen and Van der Laan, 2014

DAGs are data generating rules.

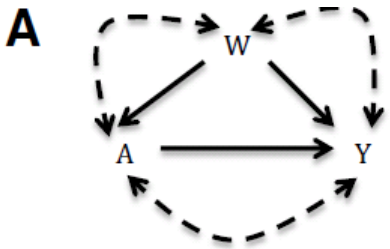

$W = f_W(U_W)$
$A = f_A(W, U_A)$
$Y = f_Y(W, A, U_Y)$

No restrictions on the distribution of $U = (U_W, U_A, U_Y)$

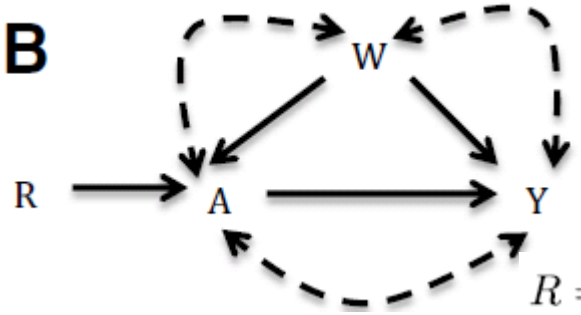

$R = I(U_R < 0.5)$
$W = f_W(U_W)$
$A = f_A(W, R, U_A)$
$Y = f_Y(W, A, U_Y)$
$U_R \sim Unif(0, 1)$
$U_R \perp\!\!\!\perp (U_A, U_W, U_Y)$

Causal graphs and corresponding structural equations representing alternative data generating processes.

(A), Assumes that W may have affected A, both A and W may have affected Y, and W, A, and Y may all share unmeasured common causes. (B), Assumes that subjects were randomly assigned to an exposure arm R with probability 0.5 and that assigned exposure arm R does not affect the outcome Y other than via uptake of the exposure A.

Petersen and Van der Laan, 2014

DAGs are data generating rules.

A

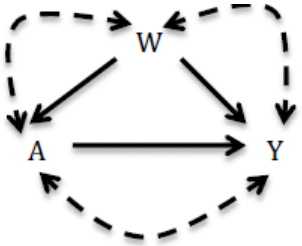

$W = f_W(U_W)$
$A = f_A(W, U_A)$
$Y = f_Y(W, A, U_Y)$
No restrictions on th
distribution of $U = ($

C

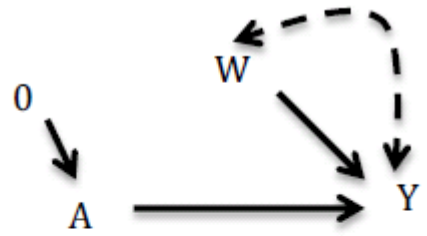

$W = f_W(U_W)$
$A = 0$
$Y = f_Y(W, A, U_Y)$
No restrictions on the
distribution of $U = (U_W, U_Y)$

- C) Figure A after a hypothetical intervention in which all subjects in the target population are assigned exposure level o (however this is defined).

Petersen and Van der Laan, 2014

Demonstrating Causation

- Block all back door paths (condition on CPCs)

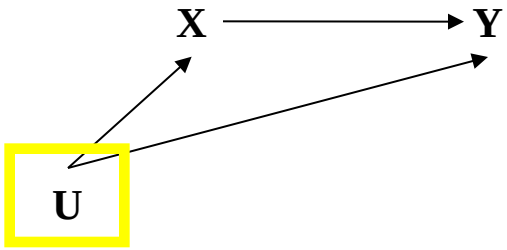

- Measure all front door paths (add up pathways)

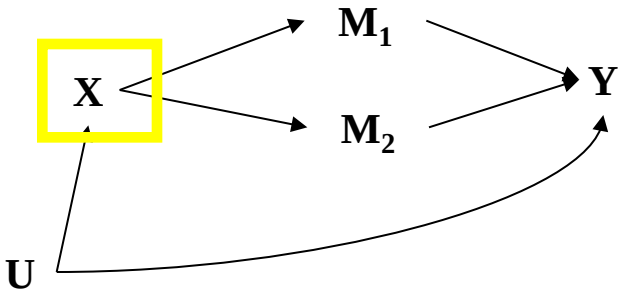

- Use an instrumental variable

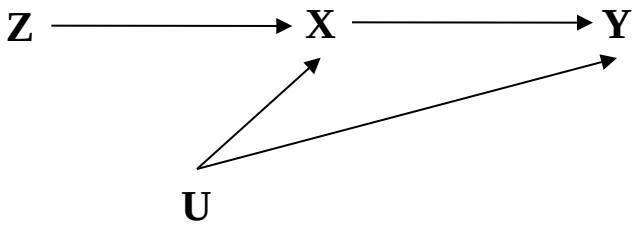

Identifying causal effects: the Back Door Criterion

- A set of variables Z satisfies the back-door criterion relative to X and Y in a DAG if:
  - no variable in Z is a descendant of X
  - Z blocks every path between X and Y that contains an arrow into X

- If a set of variables Z satisfies the back-door criterion relative to X,Y, then the causal effect of X on Y is identifiable and can be calculated by:
  - $P(Y|do(X))=\sum_Z P(Y|x,z)P(z)$

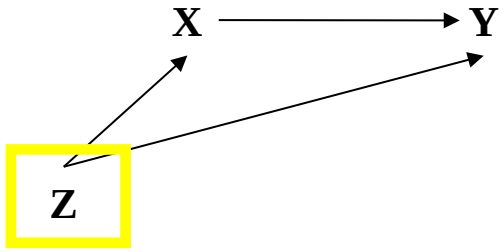

Research Design

- You are interested in estimating the effect of X on Y.
- Imagine all (well, some of the most plausible) the ways the world might work.
- If the world worked like this or like that, would your analysis plan successfully identify the effect of X on Y?
- What assumptions do you have to make to interpret your statistical association as a causal effect?

Quizlet: Can you test which DAG is correct if you have data on these variables?

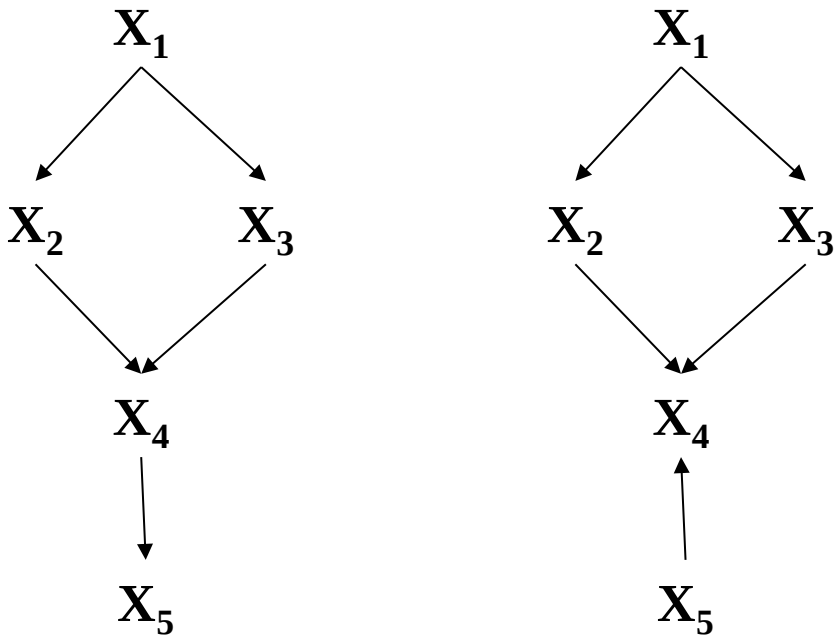

If you know that one of these two causal structures generated the data, and you have perfect measures of all 5 variables, can you tell which one is correct?

Quizlet

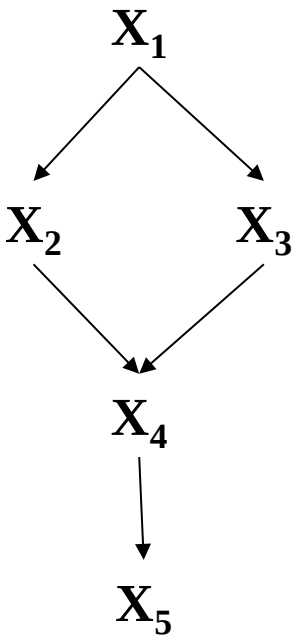

X_1, X_2, and X_3 predict X_5, but are independent of X_5 conditional on X_4.

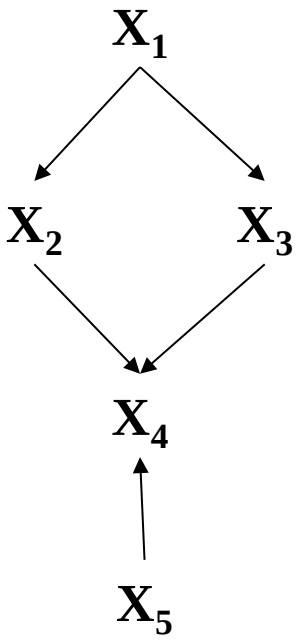

X_1, X_2, and X_3 are independent of X_5, unless you condition on X_4, when they become associated with X_5

Quizlet: Can you test which DAG is correct if you have data on these variables?

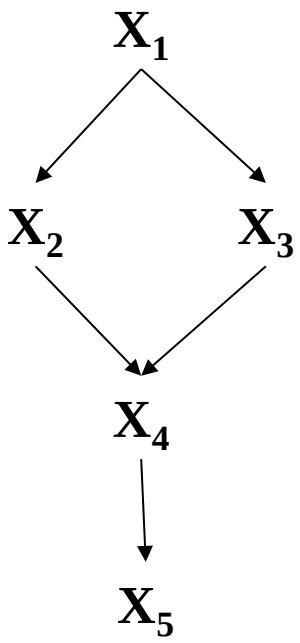

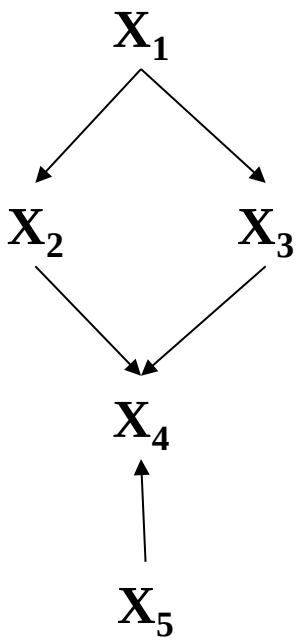

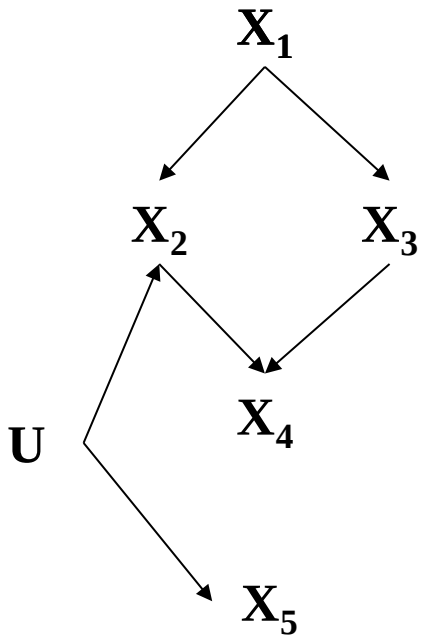

If you know that one of these three causal structures generated the data, and you have perfect measures of all 5 variables, can you tell which one is correct?

Quizlet

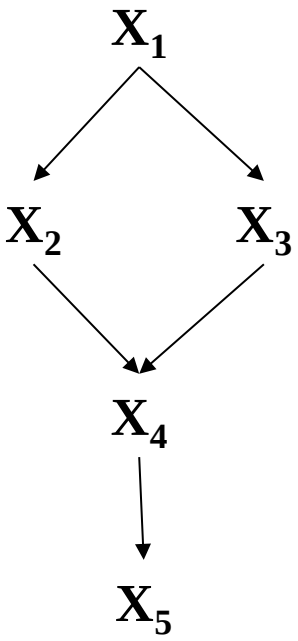

X_1, X_2, and X_3 predict X_5, but are independent of X_5 conditional on X_4.

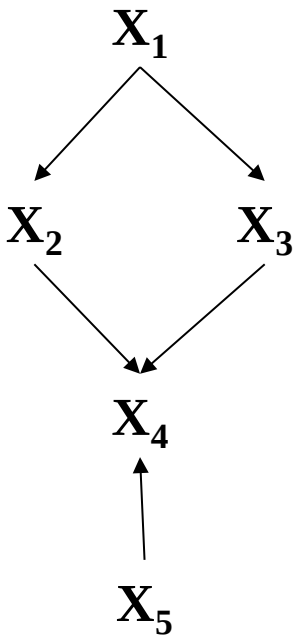

X_1, X_2, and X_3 are independent of X_5, unless you condition on X_4, when they become associated with X_5

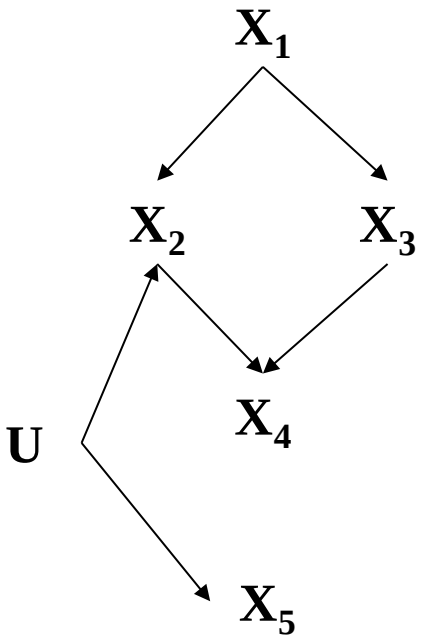

X_1 and X_3 are independent of X_5, unless you condition on X_2 or X_4, when they become associated with X_5

Organization

1. Drawing and using DAGs
2. DAGs for common problems in epidemiology
   - Confounding
   - RCTs
   - Selection biases
   - Measurement error
   - Missing Data
   - Mediation
3. Contrasting IV-based methods and covariate adjustment/propensity scores
4. Limitations and controversies of counterfactuals and DAGs

Confounding

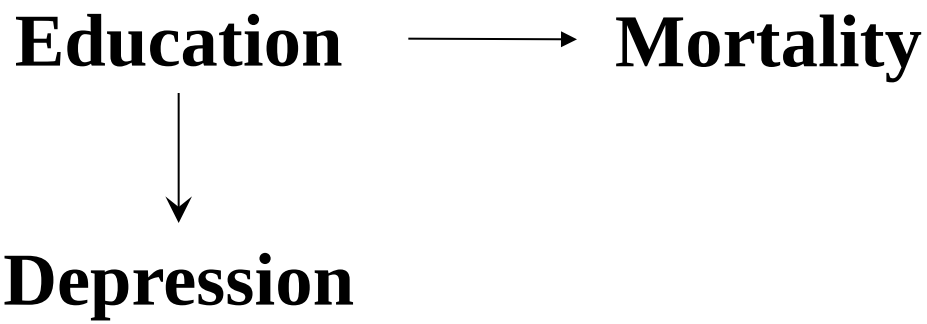

–Education“confounds” the association between Depression and mortality.

–Conditioning on education would be sufficient to identify the effect of depression on mortality

Confounding

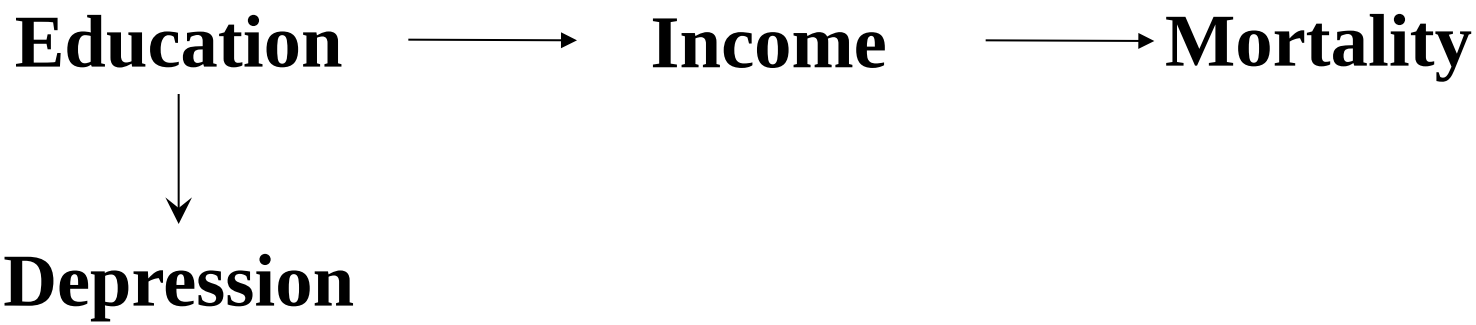

–Education“confounds” the association between Depression and mortality.

–Conditioning on either *education* or *income* would be sufficient to identify the effect of depression on mortality

Randomized Trial: Randomly Assign Therapy to Reduce Depression

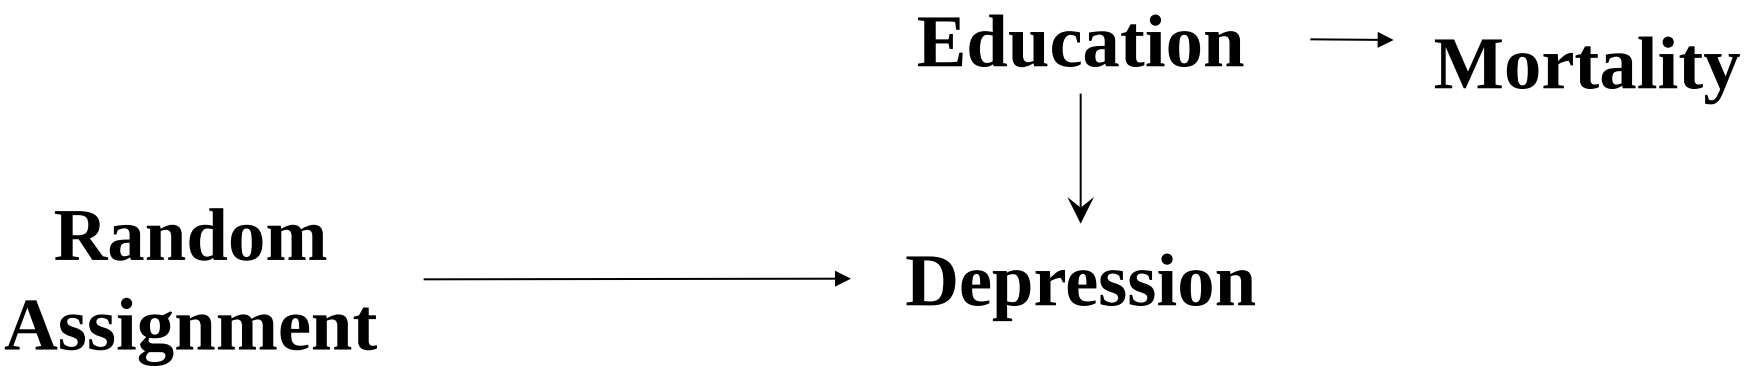

–Random assignment has no arrows into it: there are no "causes" of randomization

–Random assignment influences Depression, but so does the confounder education.

Randomized Trial: Randomly Assign Therapy to Reduce Depression

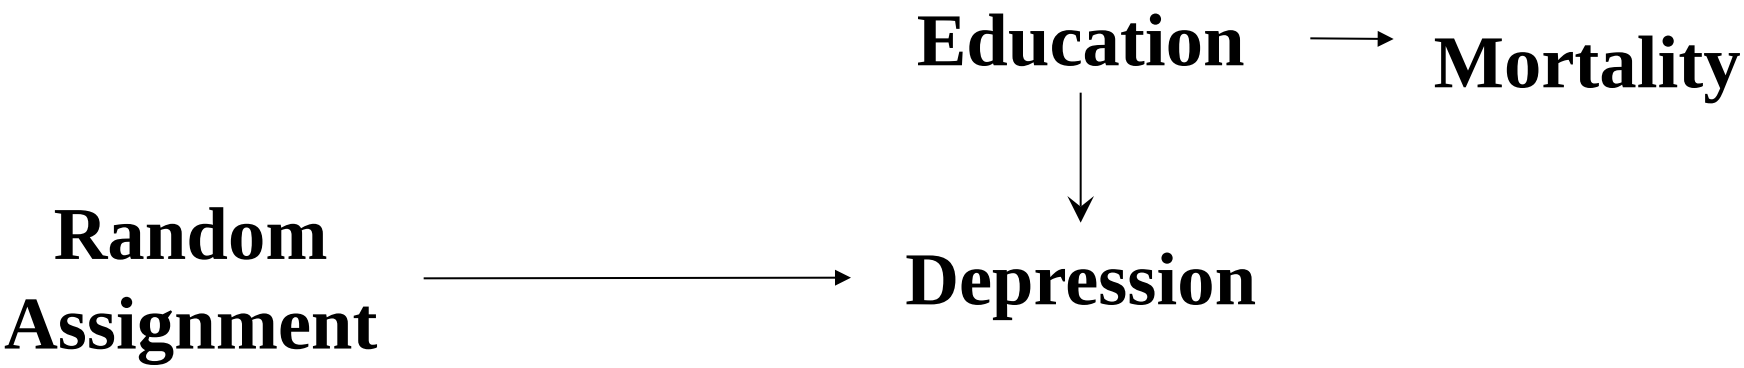

–This is why the analysis must be based on comparing mortality across values of random assignment (ITT) *not* across values of depression.

Things that go wrong with trials

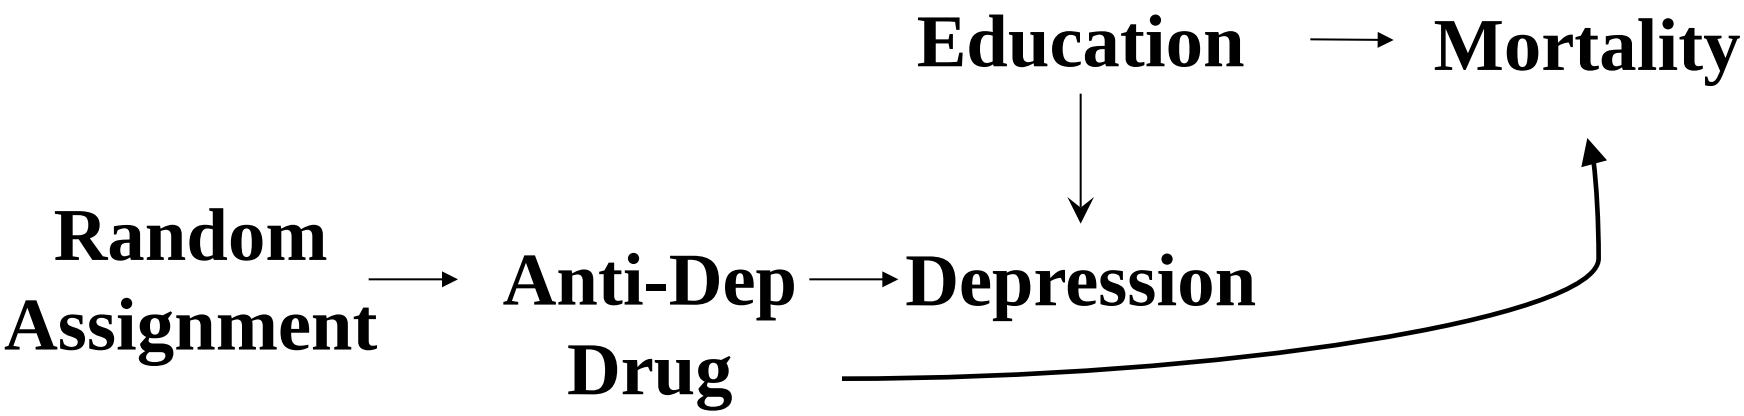

–If there is another pathway linking random assignment to the outcome, then you cannot interpret the ITT analysis as testing whether *depression* affects mortality.

Confounding: Signed DAGs

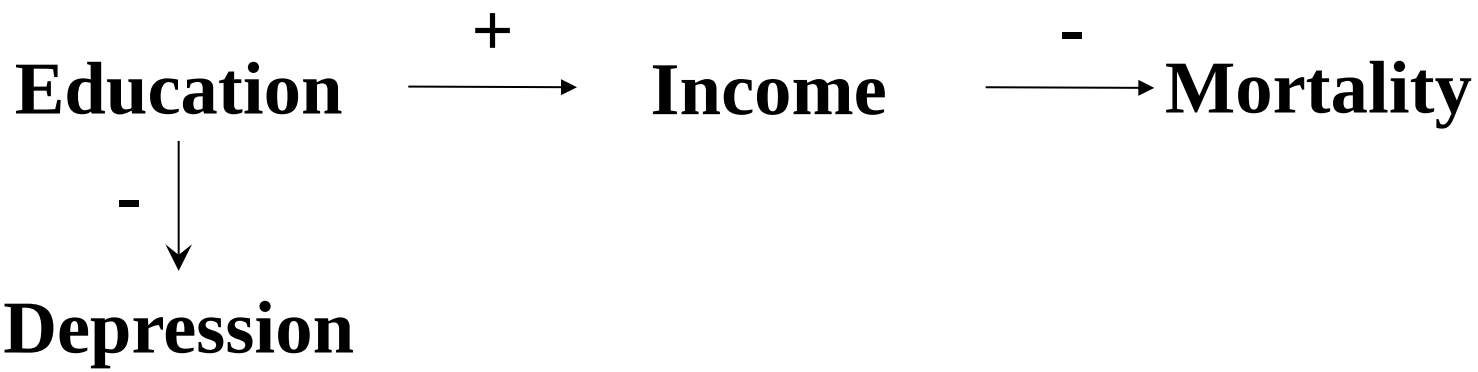

–If you think you can't fully measure education or income but you are confident about the direction of the signs, you *may* still be able to anticipate the sign of bias and potentially infer the sign of a causal effect of depression on mortality.

–With these signs, we expect *bias* to induce a positive association between depression and mortality.

–Caveat "average" vs "distributional" monotonicity.

Confounding: Signed DAGs

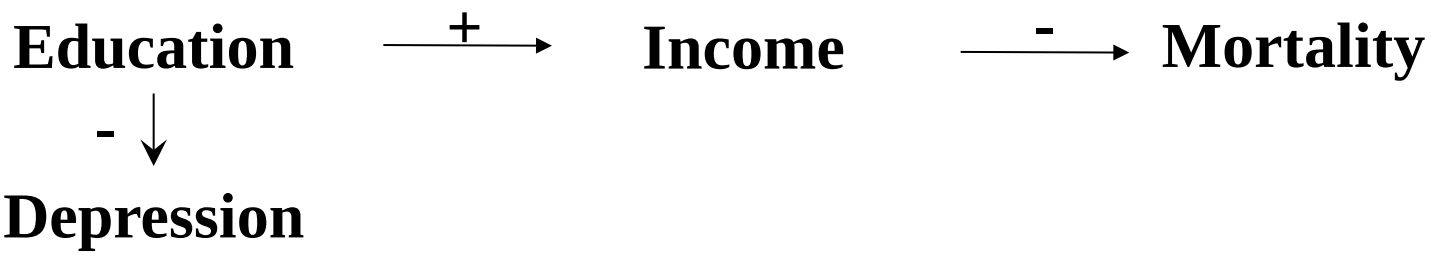

•Average monotonic effect: an increase in the value of the parent variable increases or leaves unchanged the child variable, on average.

• Under average monotonicity, some examples in which V1 is a parent of V2 and V2 is a parent of V3 and in which V1 increases V2 on average and V2 increases V3 on average but V1 *decreases* V3 on average

•That is to say, transitivity of monotonic effects or ‘‘signed edges’’ fails when one uses ‘‘monotonicity on average.’’

• Distributional monotonicity: a parent X has a positive distributional monotonic effect on a child Y if, for each value y, an increase in X renders the event Y > y at least as likely or more likely conditional on the other parents of Y.

•Transitivity of signed edges holds with distributional monotonic effects

Organization

1. Drawing and using DAGs
2. DAGs for common problems in epidemiology
   - Confounding
   - RCTs
   - Selection biases
   - Measurement error
   - Missing Data
   - Mediation
3. Contrasting IV-based methods and covariate adjustment/propensity scores
4. Limitations and controversies of counterfactuals and DAGs

Variations on collider bias: Why would you condition on a collider?

Some "colliders" are not optional:

- Loss to Follow Up
- Survival
- Diagnosis with a disease
- Selection into a study
- Providing complete data

Loss to Follow-Up

– Randomize participants to an intervention to reduce crystal meth use.

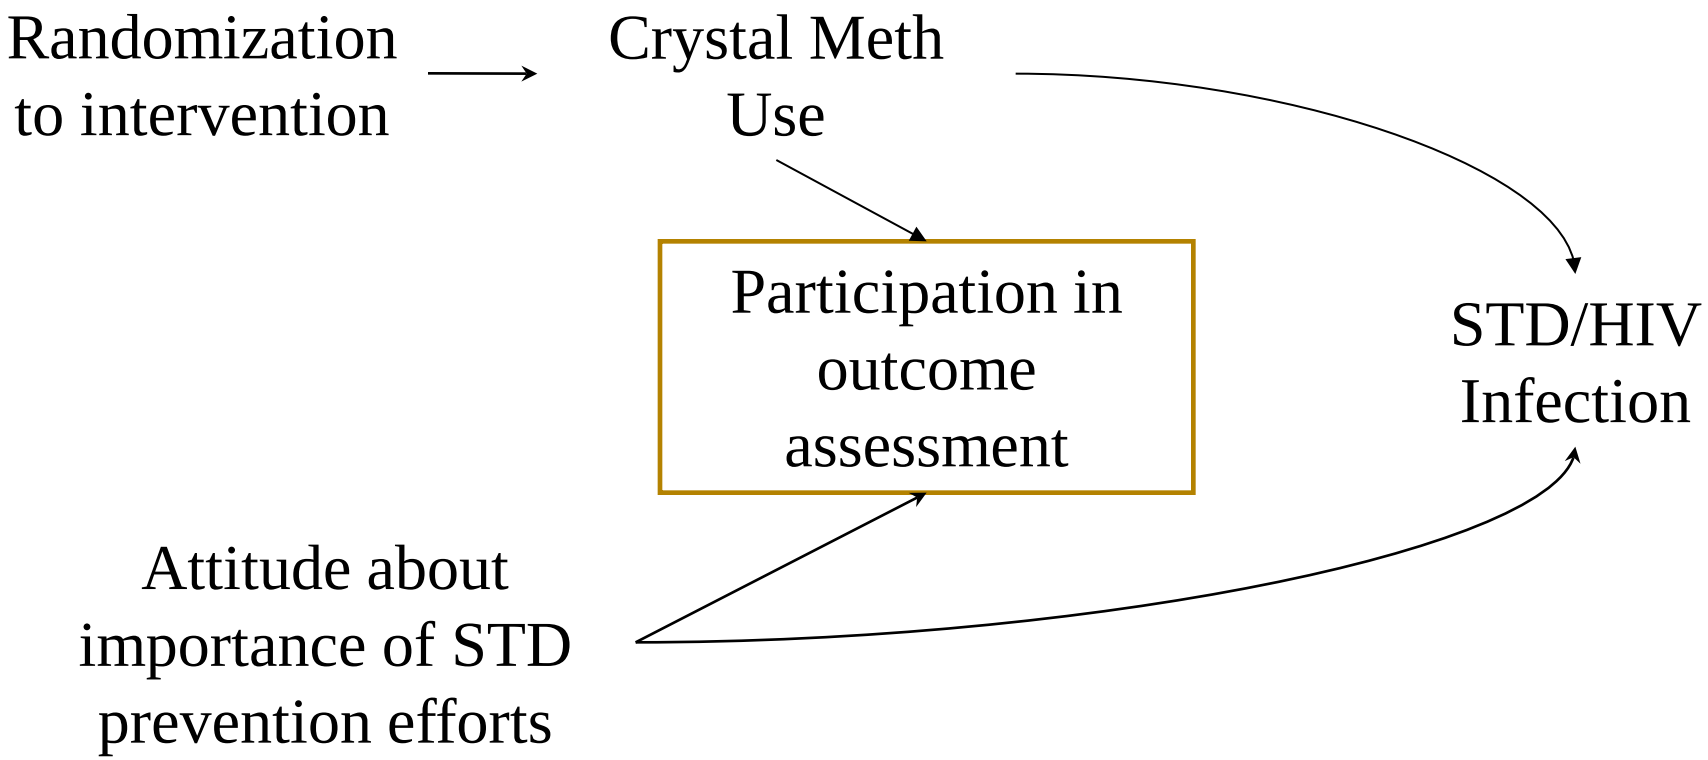

Loss to Follow-Up

– Randomize participants to participate in an Atkins-variety diet.

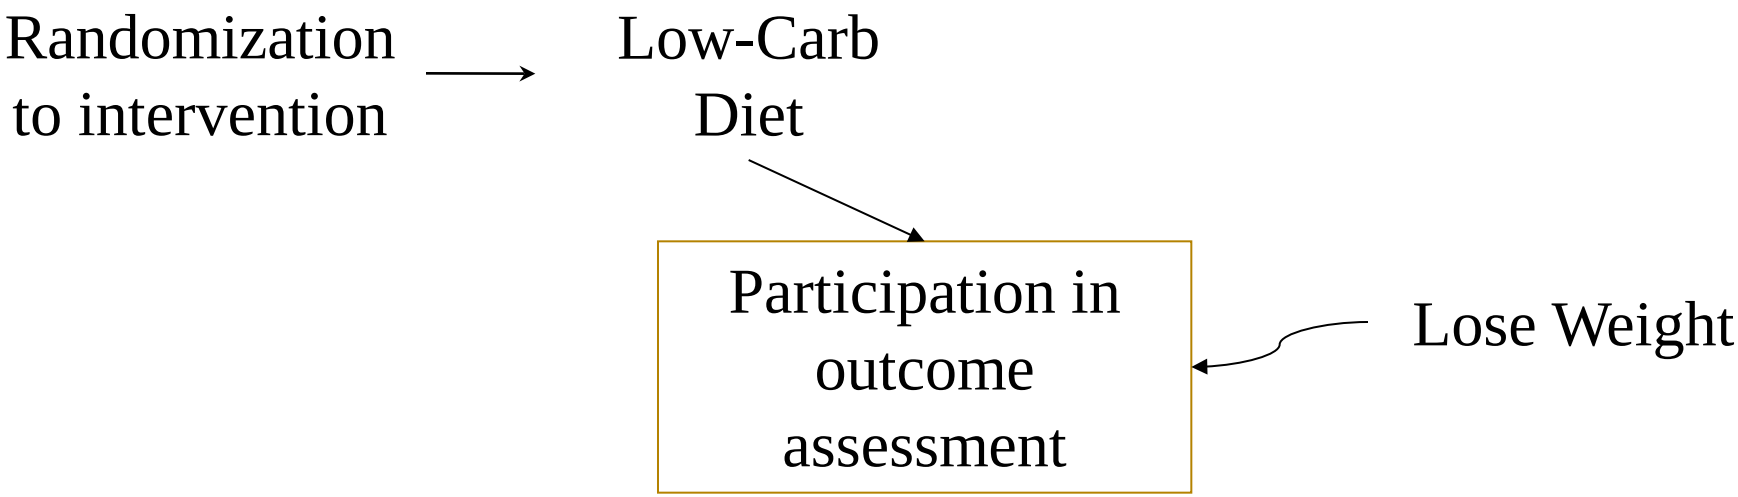

Selection/Survival Bias as a Threat to Internal Validity

- Imagine studying education and dementia in EPESE data.
- Education, completed ~age 25, affects survival to age 65.
- EPESE enrollment ~age 65

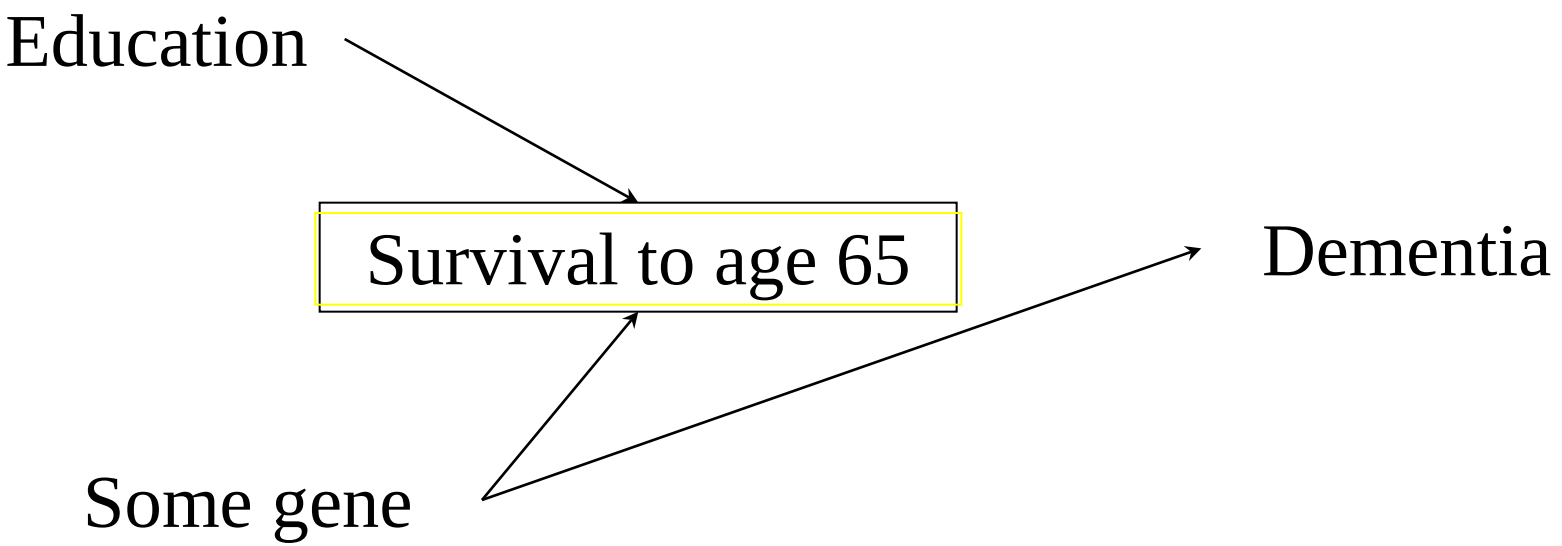

Selection Bias

Here, we *assume* education has no effect on dementia.

Would it be statistically associated with dementia among EPESE enrollees?

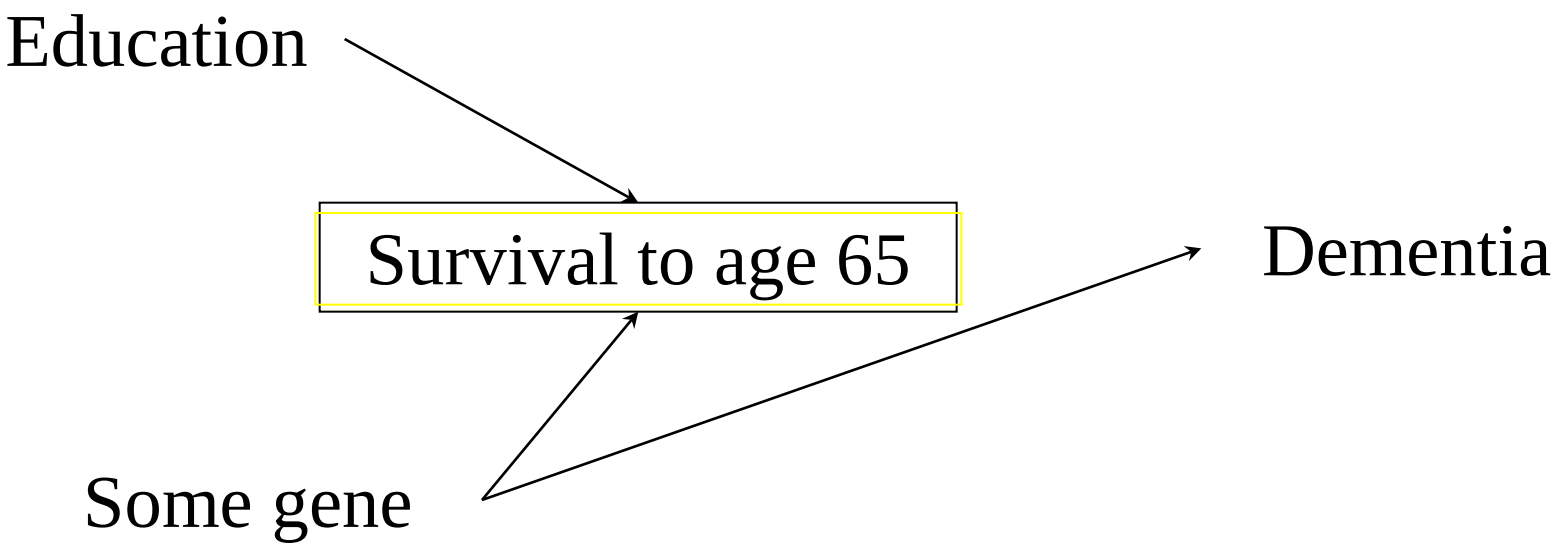

A DAG for Selection Bias

Yes.

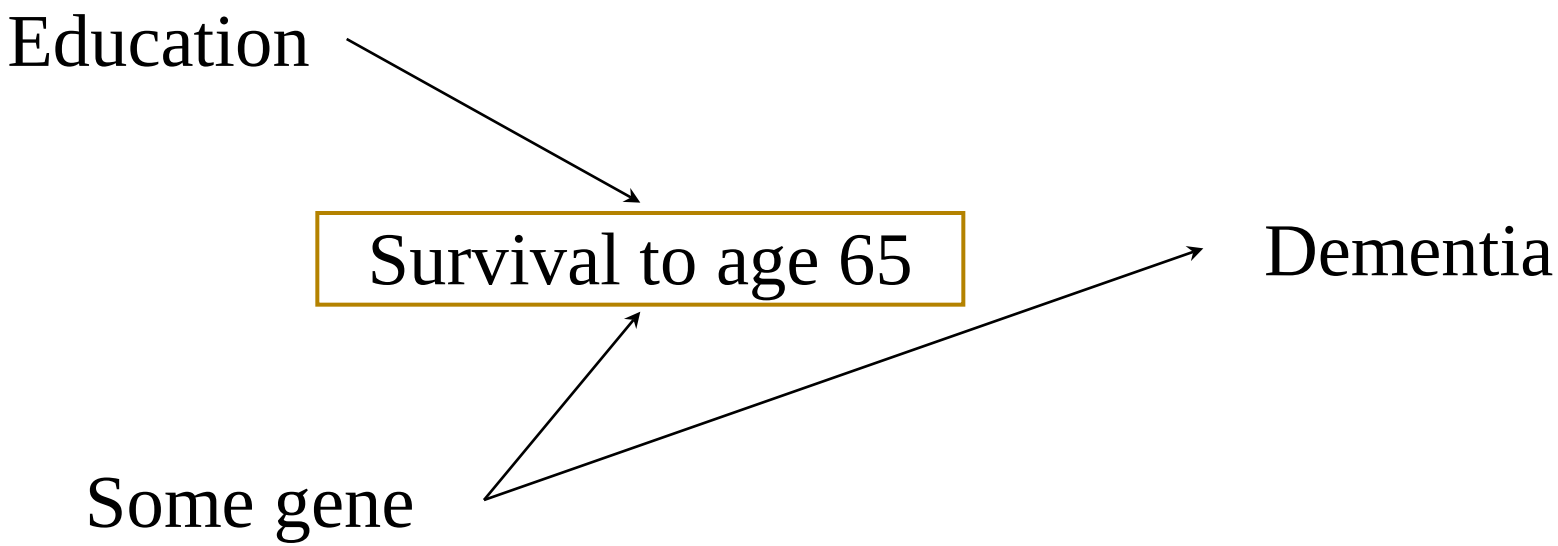

Selection Bias Compromises *Internal* Validity

- This is relevant in any study of patient samples (ie, almost any clinical epidemiology
- Does obesity affect mortality of heart failure patients?

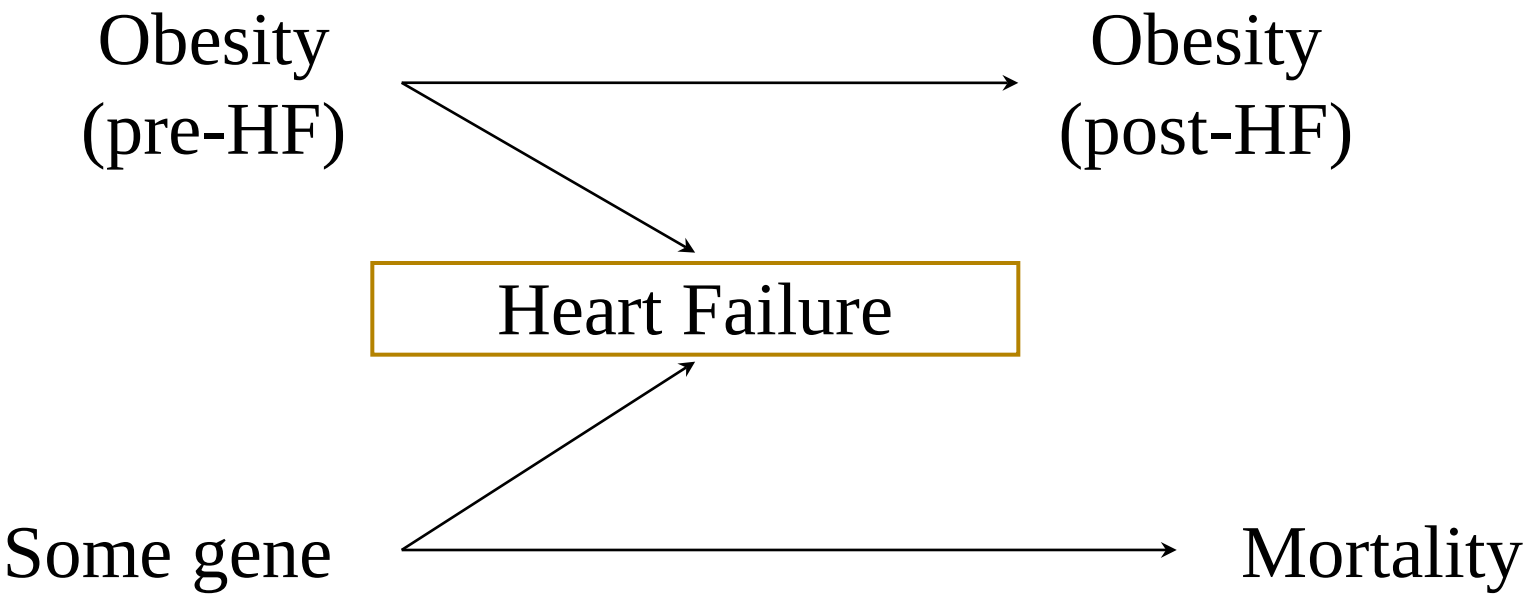

Prevalent Users Bias

- Analyses are restricted to eligible individuals who initiate a treatment strategy and remain under follow-up through time T_0, and only time after T_0 is included in the analysis.
- Arises in studies that include individuals who initiated one of the treatment strategies of interest before the start of follow-up and who continue to follow the same strategy during the follow-up, i.e., prevalent, current, or persistent users.
- This creates a potential collider bias because exposed individuals who do not have the event prior to T_0 may have other characteristics (U) that protected them from an early event and also bias associations with later event rates.

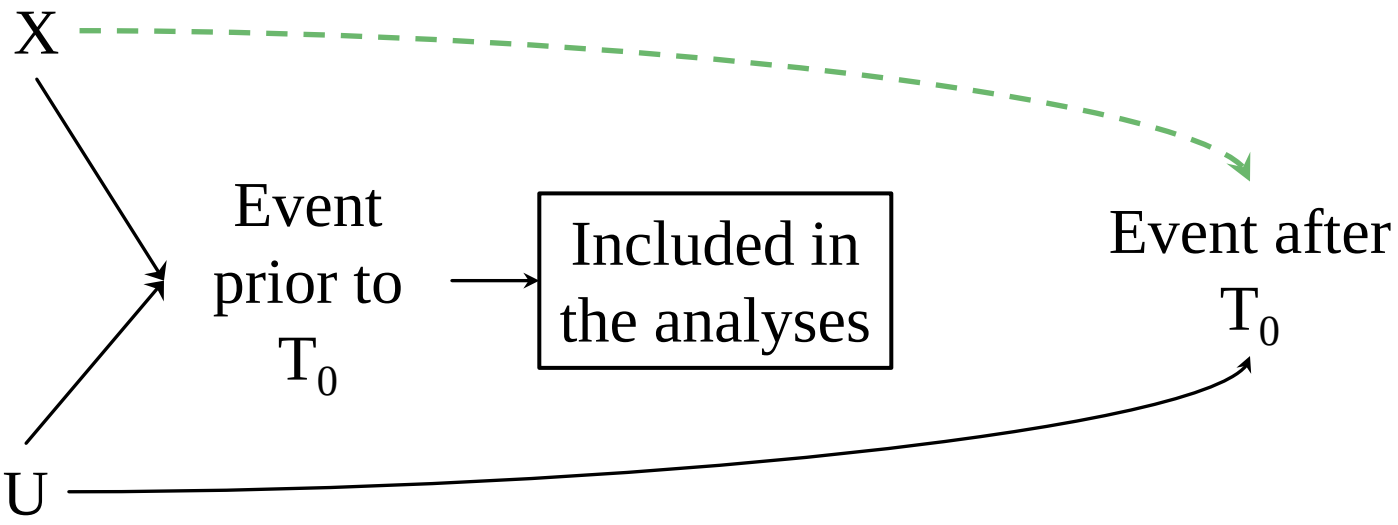

Prevalent Users Bias Variant

- Analyses are restricted to eligible individuals who initiate a treatment strategy and remain under follow-up through time T_0, and only time after T_0 is included in the analysis.
- Arises in studies that include individuals who initiated one of the treatment strategies of interest before the start of follow-up and who continue to follow the same strategy during the follow-up, i.e., prevalent, current, or persistent users.
- This may also create a bias if people who experience side effects alter treatment status before T_0.

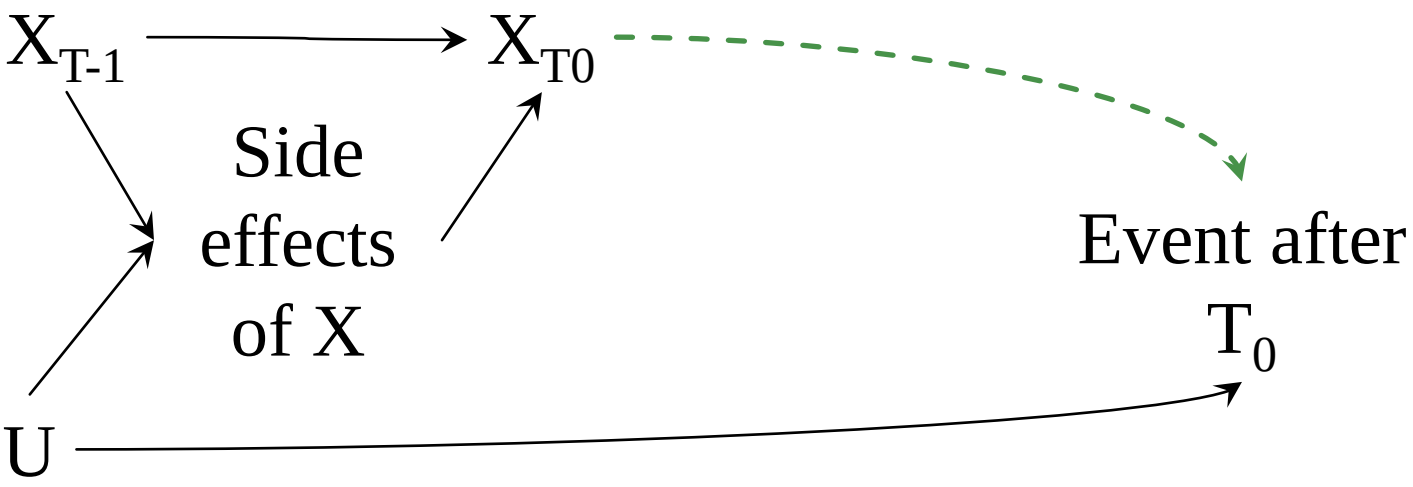

Immortal time bias

- "A span of time in the observation or follow-up period of a cohort during which the outcome under study could not have occurred."
- Example: comparing effect of heart transplant receipt on survival time, beginning follow up at the time patient is placed on the wait list. The time from wait list placement to receipt of the transplant is "immortal time" for transplant recipients in that, had they died during this period, they would not have received the transplant (organ donation rules being limited to living recipients).
- Sometimes called "survivor treatment selection bias" but that is a very confusing label.

DAGs for the most innocuous variant of immortal time bias

- The immortal time is equal for the exposed and unexposed, i.e., neither exposed nor unexposed who had the event during the immortal period are included.
- This attenuates all rate estimates (because there's too much time in the denominator). It will therefore typically bias effect estimates based on the ratios of event rates.
- If there is any unmeasured cause of the event which persists from the immortal time to the follow-up time (e.g., U in the DAG), it will also induce collider bias and affect any comparison (e.g., risk differences).

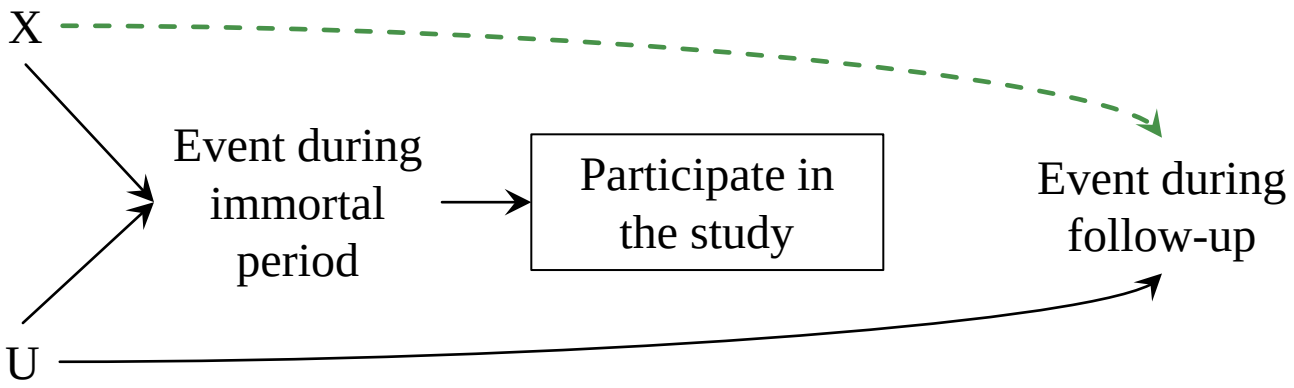

DAGs for the most innocuous variant of immortal time bias

- Under the null hypothesis that exposure doesn't influence the outcome, this type of immortal time shouldn't create a bias though, because there is no collider.
- Note that you would still be calculating event rates incorrectly, but the effect estimate for X should be null regardless.

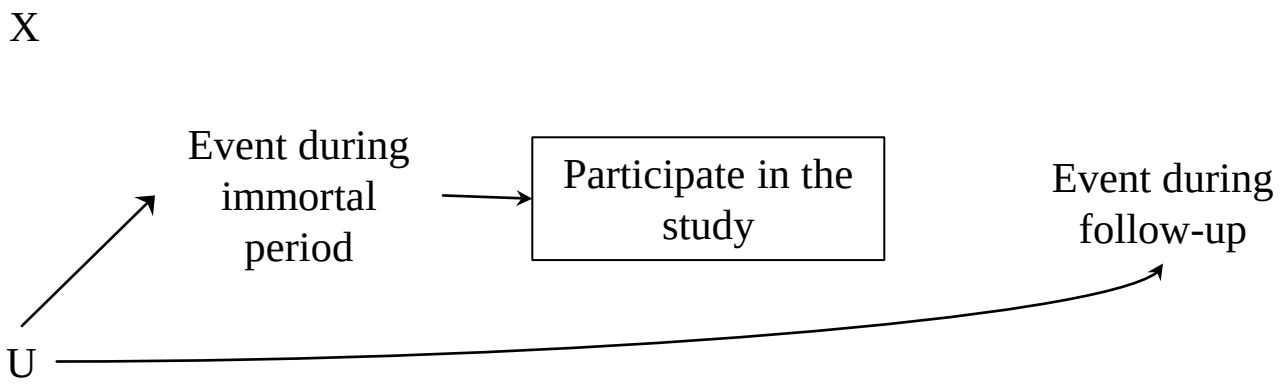

DAGs for the worst variant of immortal time bias

- The immortal time applies differentially to the exposed versus the unexposed.
- For example, if exposure was only characterized or occurred at the conclusion of the immortal time (the transplant example) any cases that occurred during that time would be classed as unexposed.
- This creates an extreme bias under the null or any alternative hypothesis, because having an event during the immortal period effectively confounds exposure and the event rate.

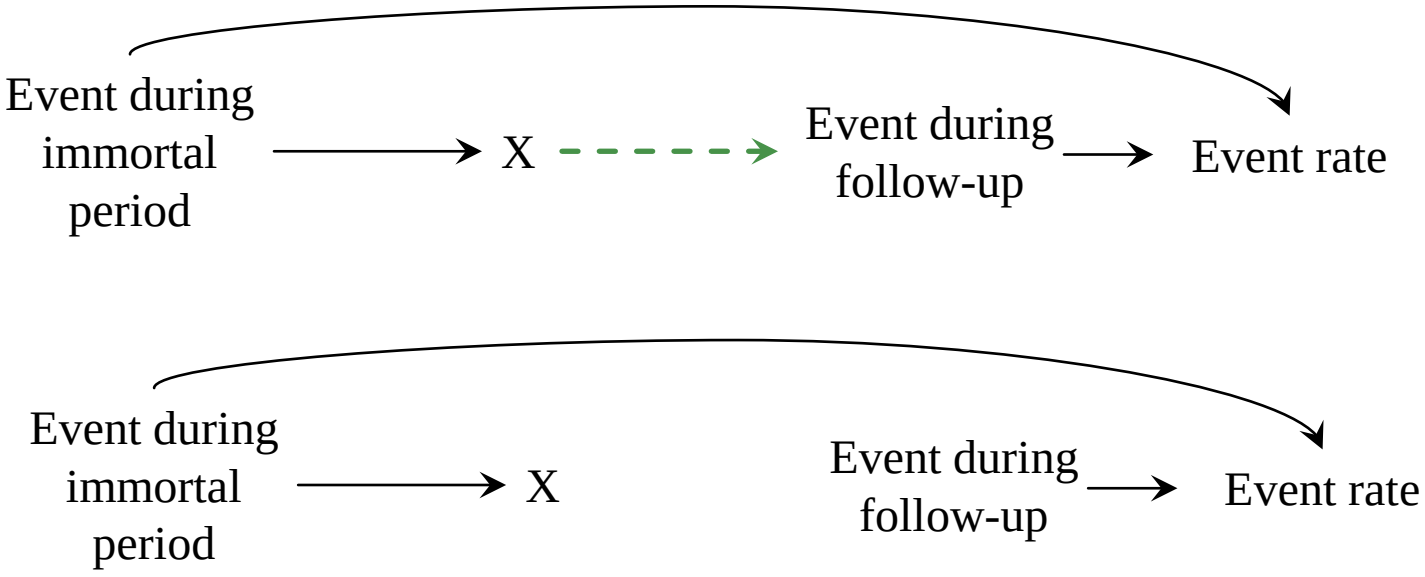

DAGs for an especially seductive variant of immortal time bias

- The independent variable is a count of episodes of exposures, which includes episodes occurring throughout the follow-up time. For ex, estimate the effect of number of waves of medication use (0 to 3) on stroke risk, counting medication use from wave 1 to wave 3, and starting follow-up time for stroke onset at wave 1.
- Anybody with 3 waves of medication use must not have had a stroke before wave 3.
- Time is not strictly immortal, but exposure status is influenced by event rate.
- This creates large bias under the null or any alternative hypothesis.
- Same figure applies if the "cumulative X" is "ever vs never exposed"

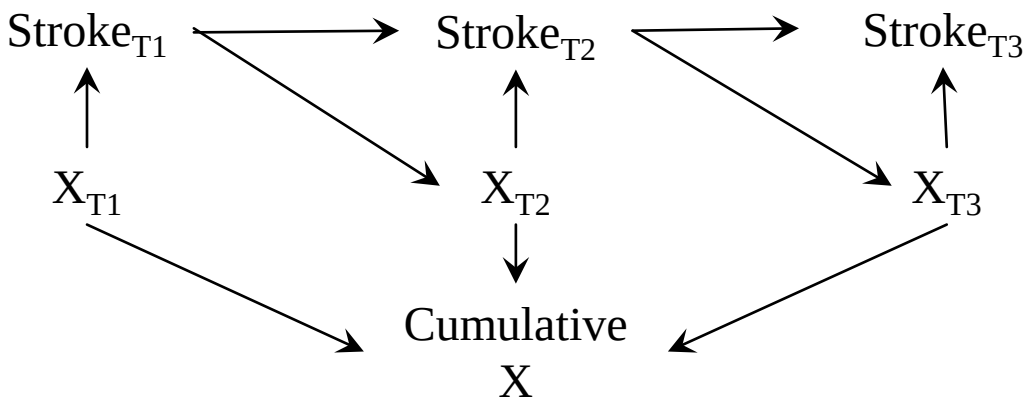

Can you quantify the bias from collider stratification?

- Must make assumptions about the magnitude and direction of each causal association
- This is not specified in the DAG, the DAG only tells you conditional dependence/independence.
- Often the bias is small, but not always.
- Often the bias is *negative*.

Organization

1. Drawing and using DAGs
2. DAGs for common problems in epidemiology
   - Confounding
   - RCTs
   - Selection biases
   - Measurement error
   - Missing Data
   - Mediation
3. Contrasting IV-based methods and covariate adjustment/propensity scores
4. Limitations and controversies of counterfactuals and DAGs

Unreliable Measures

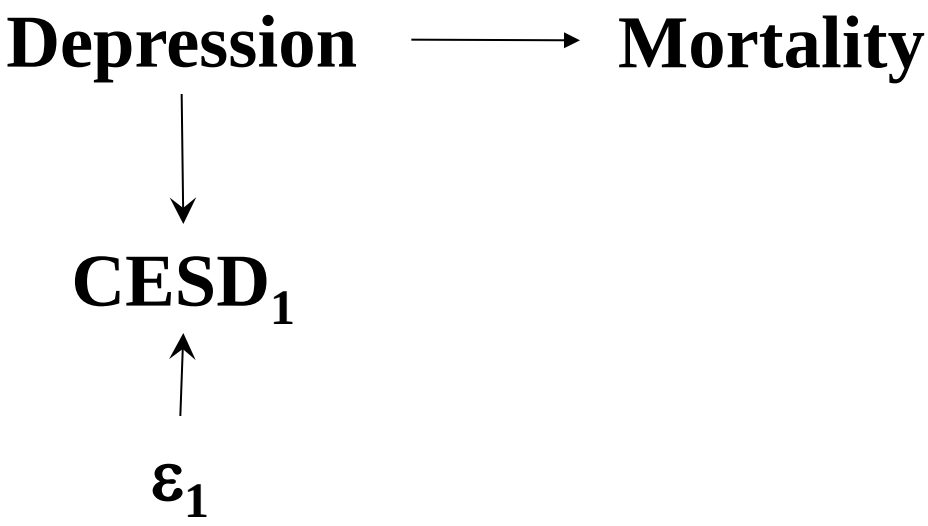

Unreliable Measures

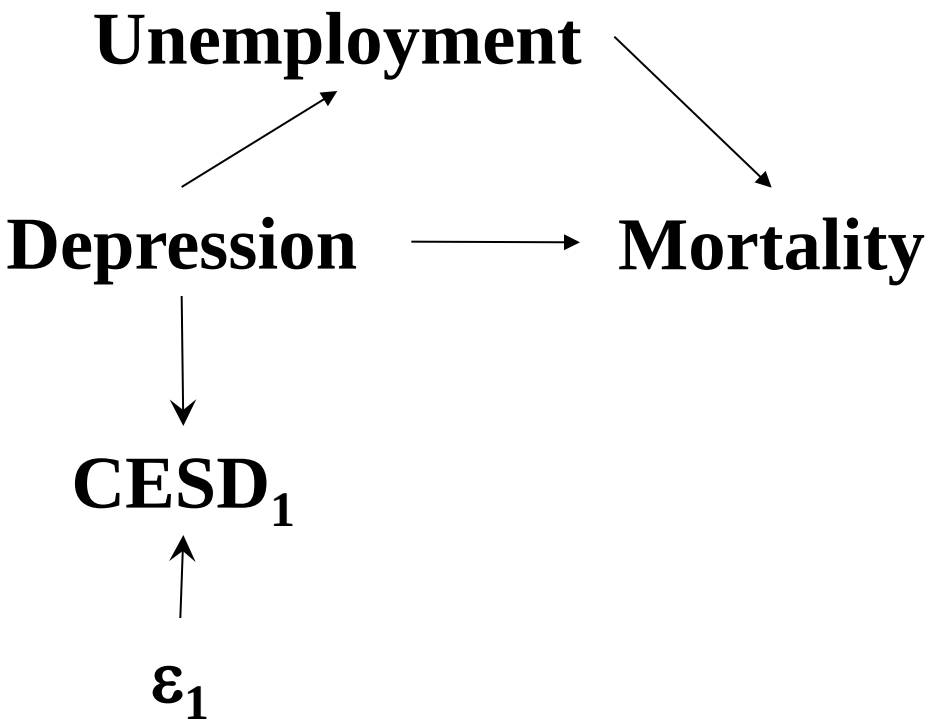

Unreliable Measures in Analyses of Change

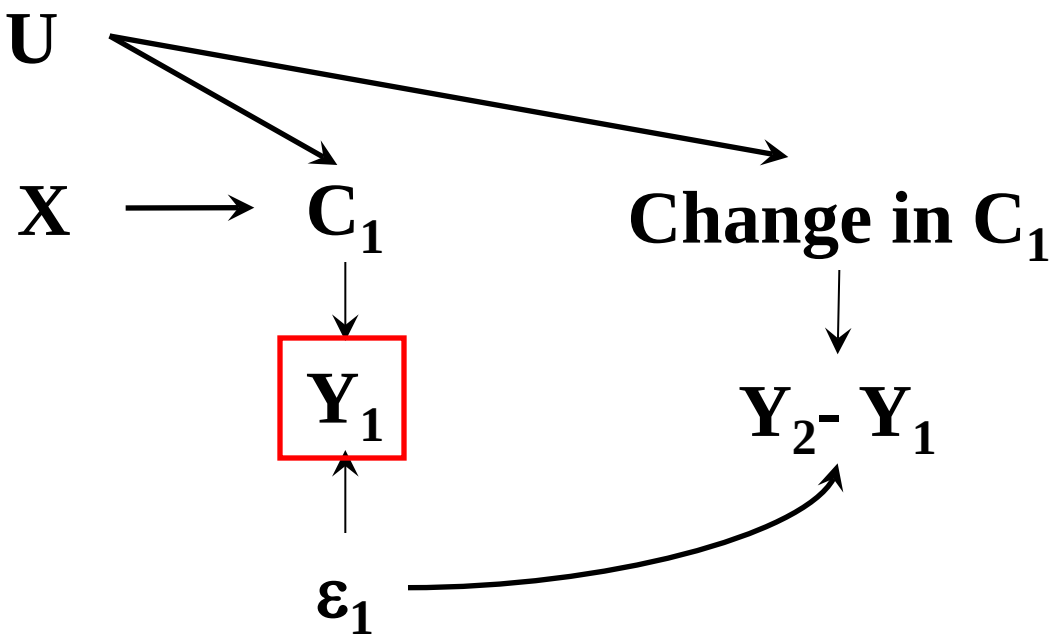

DAGs for missing data

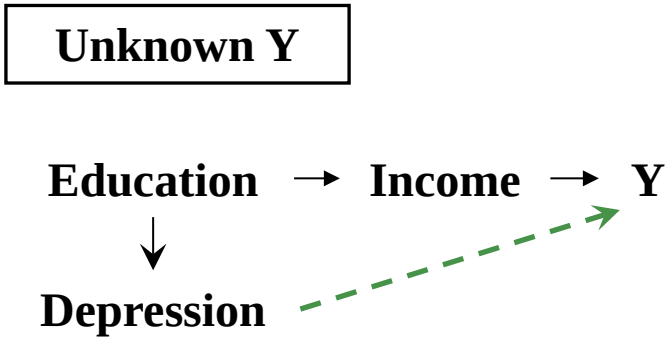

Missing completely at random (Yay!)

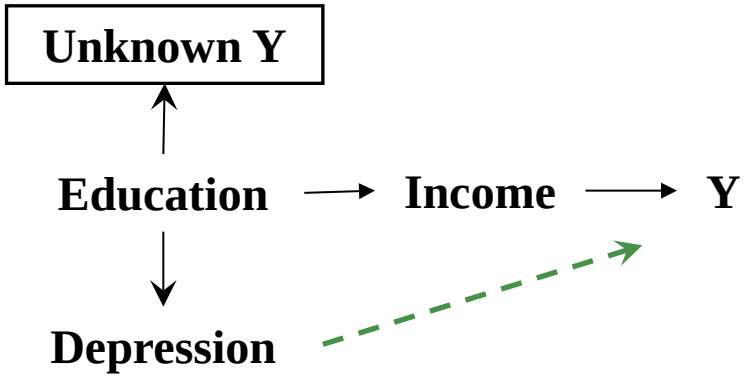

Missing at random (condition on Education, which you needed to condition on anyway… problem solved)

DAGs for missing data

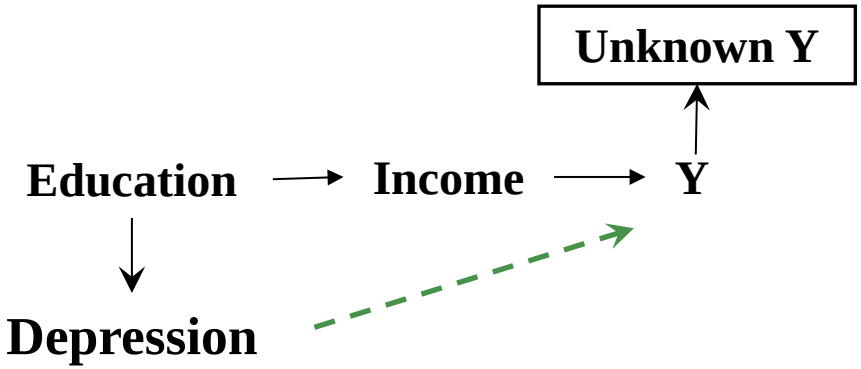

Missing Not at Random (Yikes)

Mediation, Effect Decomposition, and Estimating Direct or Indirect Effects

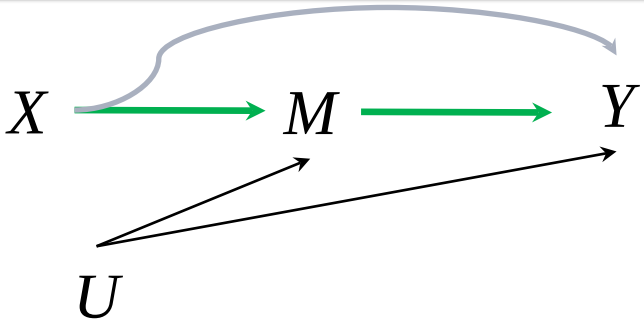

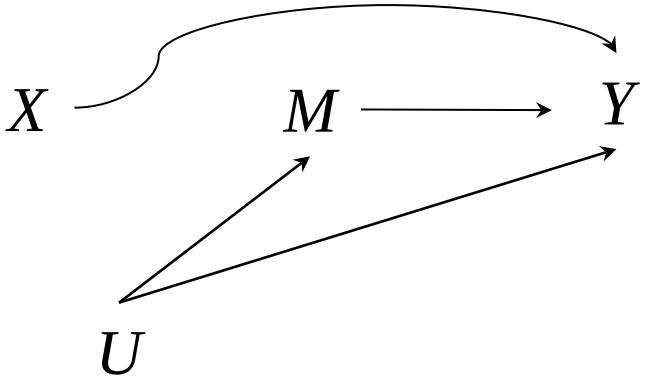

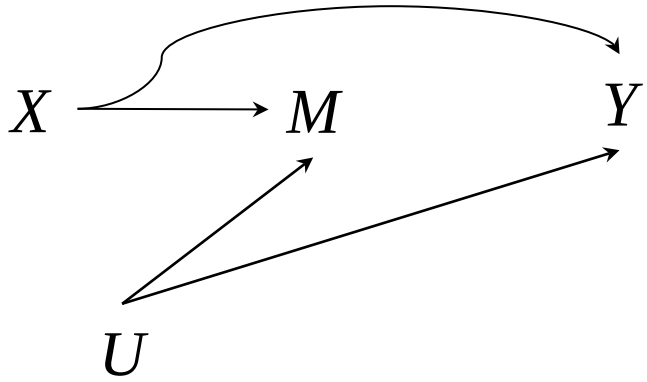

- If X affects M and M affects Y, then M mediates (part of) the effect of X on Y.
  - There are some weird cases where this is not true, but these are exceptions.
- M *cannot* mediate the effect of X on Y if either:
  - X does not affect M,

  or
  - M does not affect Y.

Mediation: lots of possible structures

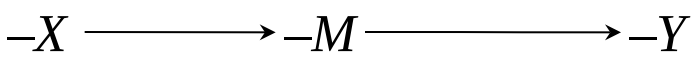

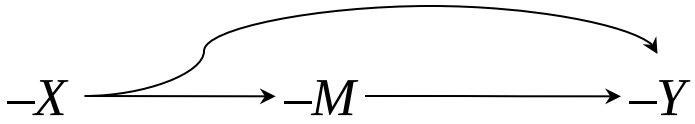

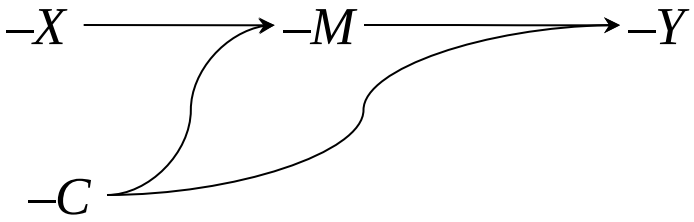

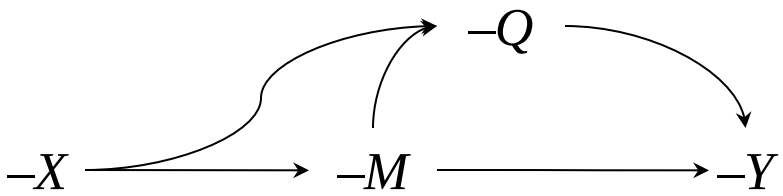

Mediation: Classic Decomposition Approach

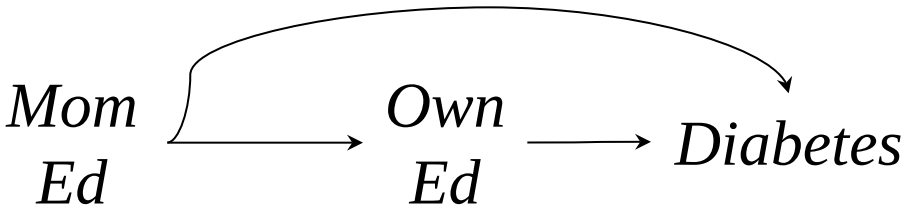

- Total effect= direct effect + indirect effect
- Indirect effect= total effect-direct effect
- Total effect estimate= E(Db | MomEd=1)-E(Db|MomEd=0)
- Direct effect estimate=E(Db|MomEd=1, OwnEd)-E(Db|MomEd=0, OwnEd)
- Sometimes called “Barron-Kenny” decomposition

Mediation: Classic Decomposition Example

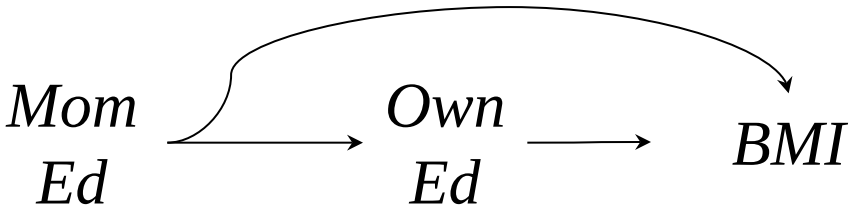

E(Own Ed)=12+0.5*(Mom Ed-12)

E(BMI)= 25-0.5*(Own Ed-12) – 0.25*(Mom Ed-12)

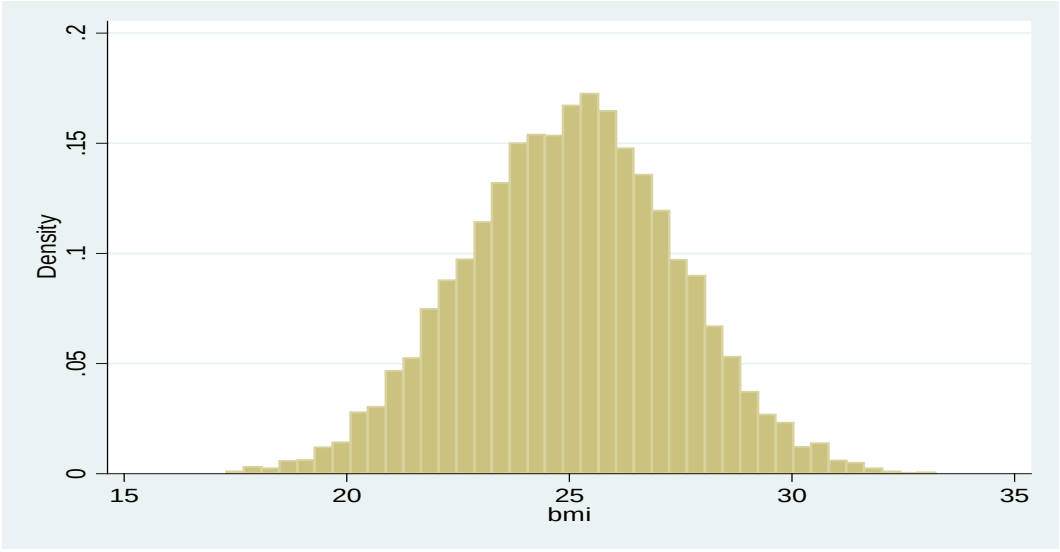

Mediation: Classic Decomposition Example

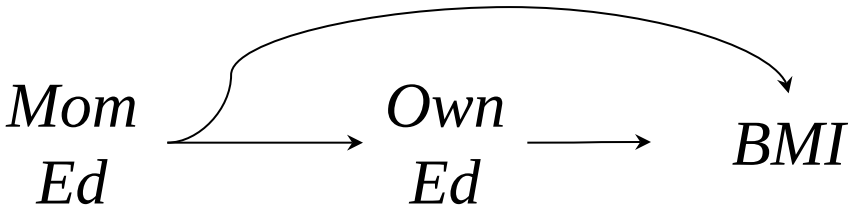

E(Own Ed)=12+0.5*(Mom Ed-12)

E(BMI)= 25-0.5*(Own Ed-12) – 0.25*(Mom Ed-12)

What is the direct effect of Mom Ed on BMI?

-0.25

What is the indirect effect of Mom Ed on BMI?

0.5*-0.5=-0.25

What is the total effect of Mom Ed on BMI?

-0.25 + 0.5*(-.5)=-0.5

Mediation: Classic Decomposition Example

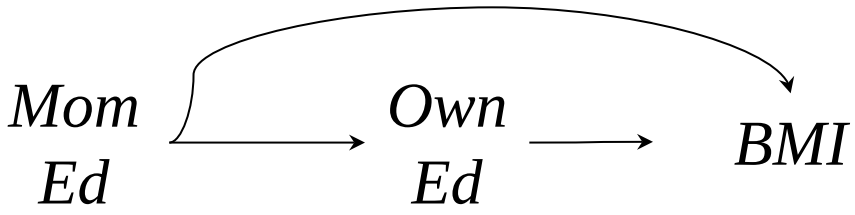

E(Own Ed)=12+0.5*(Mom Ed-12)

E(BMI)= 25-0.5*(Own Ed-12) – 0.25*(Mom Ed-12)

What is the total effect of Mom Ed on BMI?

-0.25 + 0.5*(-.5)=-0.5

```
. reg bmi momed
```

Source	SS	df	MS
Model	9760.68979	1	9760.68979
Residual	46410.3396	9998	4.64196235
Total	56171.0294	9999	5.61766471

```
Number of obs =   10000
F(  1,  9998) = 2102.71
Prob > F      =  0.0000
R-squared     =  0.1738
Adj R-squared =  0.1737
Root MSE      =  2.1545
```

bmi	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
momed	-.4936806	.0107661	-45.86	0.000	-.5147842	-.472577
_cons	30.95059	.1309641	236.33	0.000	30.69387	31.20731

Mediation: Classic Decomposition Example

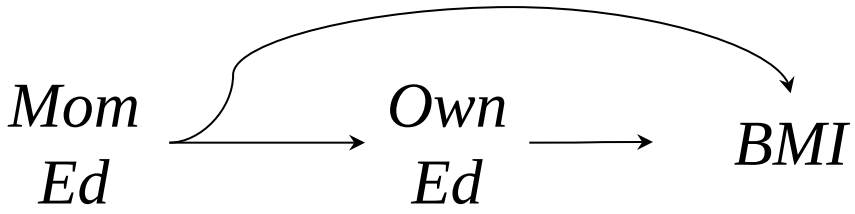

E(Own Ed)=12+0.5*(Mom Ed-12)

E(BMI)= 25-0.5*(Own Ed-12) – 0.25*(Mom Ed-12)

What is the direct effect of Mom Ed on BMI?

-0.25

```
. reg bmi momed own_ed
```

Source	SS	df	MS
Model	15300.174	2	7650.08698
Residual	40870.8555	9997	4.08831204
Total	56171.0294	9999	5.61766471

```
Number of obs =   10000
F(  2,  9997) = 1871.21
Prob > F      =  0.0000
R-squared     =  0.2724
Adj R-squared =  0.2722
Root MSE      =   2.022
```

bmi	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
momed	-.245425	.0121478	-20.20	0.000	-.2692371	-.2216128
own_ed	-.4988862	.0135531	-36.81	0.000	-.525453	-.4723194
_cons	33.95703	.1475695	230.11	0.000	33.66777	34.2463

Estimating Direct Effects When There is a Mediator-Outcome Confounder

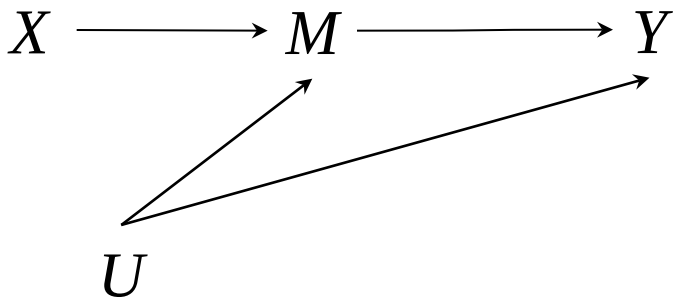

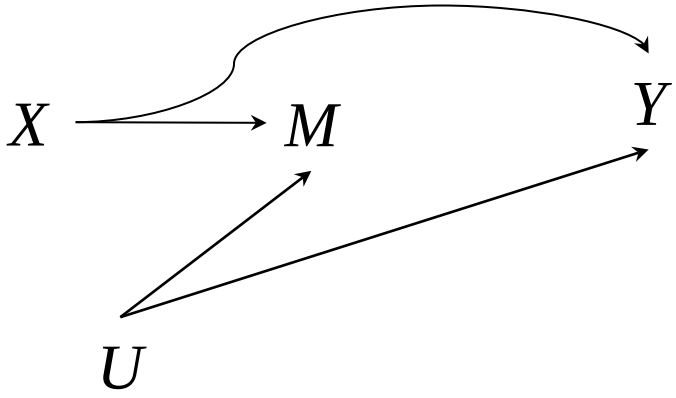

- Conventional approaches to estimating direct effects condition/adjust for the mediator (M), but if there is a confounder of the M-Y association, then *X* and Y would be associated conditional on M even if there was no indirect effect.

Organization

1. Drawing and using DAGs
2. DAGs for common problems in epidemiology
   - Confounding
   - RCTs
   - Selection biases
   - Measurement error
   - Missing Data
   - Mediation
3. Contrasting IV-based methods and covariate adjustment/propensity scores
4. Limitations and controversies of counterfactuals and DAGs

Some objections to DAGs

- DAGs do not display effect modification
  - If Z modifies the effect of X on Y, then Z must itself affect Y.
  - Effect modification is scale dependent, so showing an effect modifier implies a scale and DAGs do not represent any scale
  - Some progress on this: e.g., VanderWeele's sufficient cause framework
  - If Z and X both affect Y, then there is *some* scale on which Z modifies the effect of X and vice versa.

Some objections to DAGs

- DAGs do not represent the magnitude of bias
  - Major limitation
  - Lots of work now on (a) signing the bias, and (b) estimating the magnitude
  - Simulations and bias bounding formulas

Some objections to DAGs

- DAGs do not match the complex, dynamic feedback loops in the real world
  - DAGs can represent feedback between two variables
  - DAGs can represent complex interventions (e.g., deliver this set of 4 treatments vs only 2 of these)
  - DAGs match (one of the) things you want to know: will intervening in a specific way elicit a change in health?
  - To say "it's all very complicated and impossible to specify which variables come first" is to abdicate any hope of guiding clinical or policy decisions.
  - DAGs do not tell you everything, and complex systems simulations are often useful as a next step to predict spillovers and interactions with variables you do not have in your DAG. The simulation will presumably rely on causal parameter inputs.

Some objections to DAGs

- I'm not sure about which causal structure is right
  - Do you know enough to rule some DAGs out?
  - What are the commonalities across all plausible DAGs?
  - Useful to state: under these assumptions, we can draw these inferences; need evidence on the assumptions from other sources
- I don't know anything about the causal structure
  - This is not an objection to DAGs, this is an objection to ever drawing causal inferences

Can I show counterfactuals on DAGs?

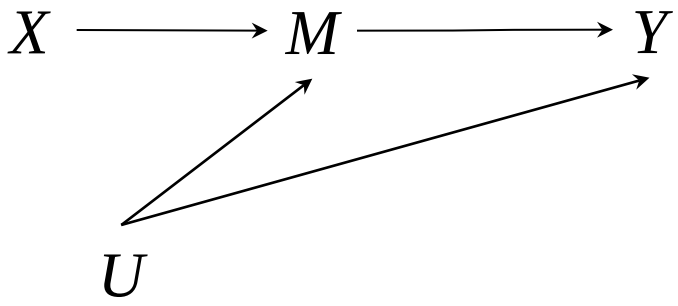

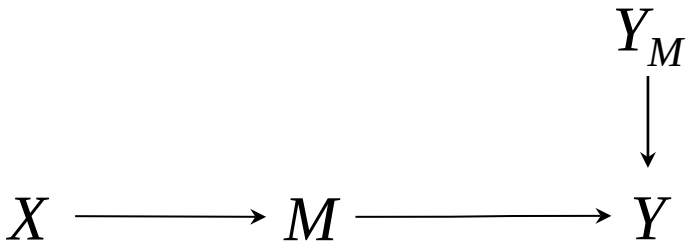

If you show the counterfactual for Y_M, there are no other arrows pointing into Y except for the counterfactual and M, because the counterfactual is the value that Y will take if M were set to any particular value.

Counterfactuals do not always satisfy intuitions about "causation"

- Can non-modifiable characteristics (sex, race) be causes?
  - What's the counterfactual for race and sex? Can you imagine the same individual but of a different race or sex?
  - More challenging for race than sex, because race generally identified based on race of parents, so moment of intervention to change someone's race is unclear.
  - Some say "race is not a cause, racism is". This is a red herring. Saying that race is a cause makes no claims about the mechanism one way or the other, may well be fully mediated by racism.
  - Issue remains controversial. Most people's research seems to treat these as causal but VanderWeele and Robinson offer some language to help you work around the intellectual incoherence if you are one of the people who studies race but doesn't want ot discuss it as a cause.
  - Hernan sometimes seems to take this a bit farther and argue that only "actions" not "states" can be causes.
- Generally useful to ask: what RCT would deliver the parameter I am interested in?

Counterfactuals do not always satisfy intuitions about "causation"

- What are the causes of over-determined outcomes?
  - MaryAnne and Wanda are traveling through the desert with Earl. They both hate him. One night, MaryAnne poisons Earl's water canteen and sneaks away. Wanda, not realizing MaryAnne has poisoned the water, pokes a hole in Earl's canteen so all the water leaks out, and she sneaks away. In the morning, Earl wakes up alone, with no water. He dies of dehydration in the desert shortly thereafter. Who killed Earl?
  - Rex the gang leader orders his subordinate Joe to shoot rival gang member, Ace. Rex and Joe shoot Ace at the same time. Either shot would have been sufficient to kill Ace. Is Joe responsible for Ace's death?
- Important for judicial decisions attributing responsibility.
- Also important for thinking about public health implications of proposed interventions. Intervening on something mechanistically related to an outcome may not improve the outcome if it is over-determined.

Counterfactuals do not always satisfy intuitions about "causation"

- Survival bias: what if some values of X render Y undefined?
  - If Earl smokes, he will die at age 70 of CVD; if Earl doe not smoke, he will develop dementia at age 71
  - What is the effect of smoking on Earl's dementia?

Often unclear how to draw the DAG: Difference in difference

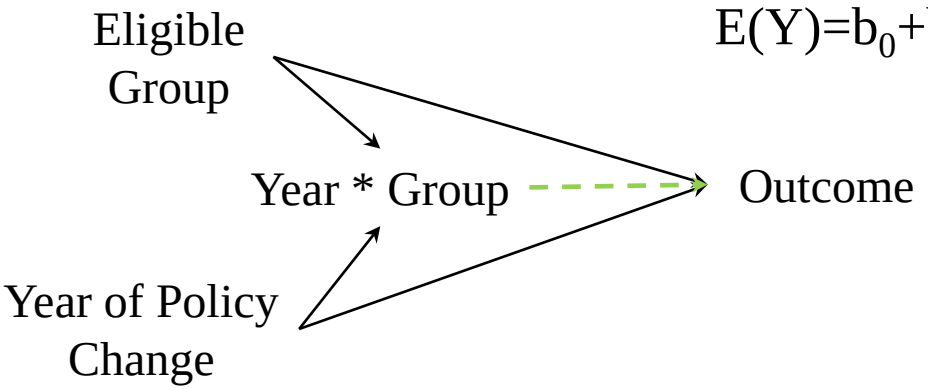

$$E(Y)=b_0+b_1*Group+b_2*Year+b_3*Group*Year$$

Conditional on group and year of change, there are no confounders of the association between the interaction and the outcome.

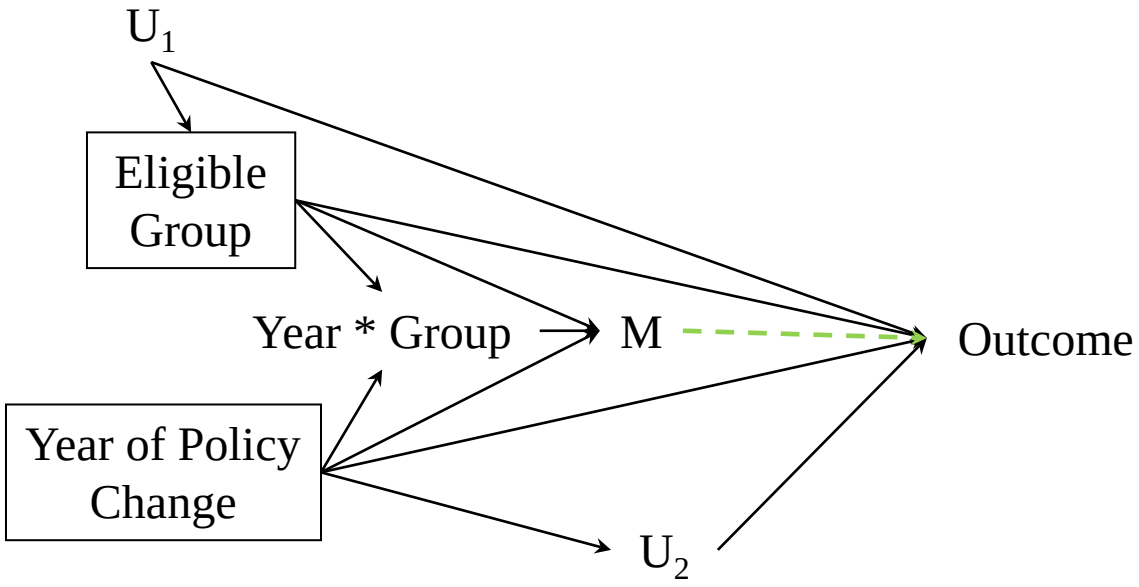

Often unclear how to draw the DAG: Change scores vs change

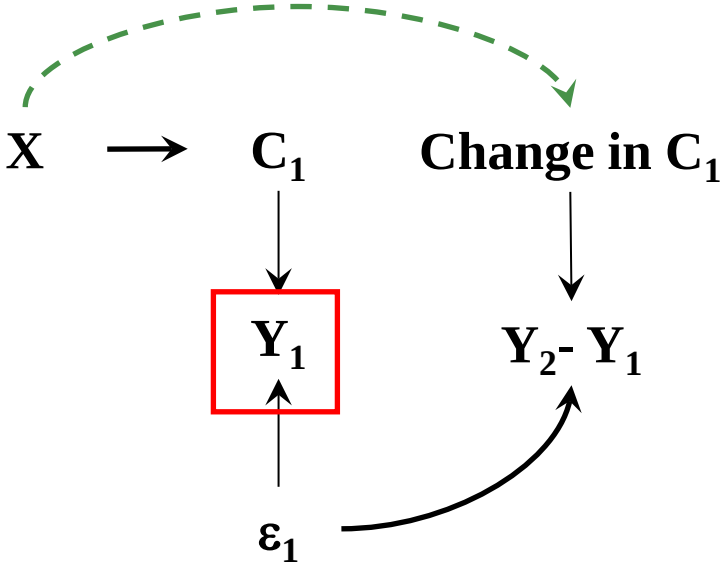

If you draw the DAG as on the left, adjusting for Y1 induces collider bias and a spurious association between X and change score.

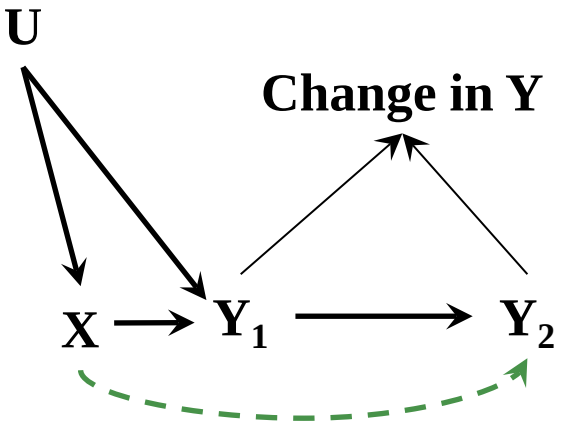

If you draw the DAG as on the right, adjusting for Y1 is necessary to block the direct effect of Y1 on Y2.

I chose the DAG on the left because I consider change in C a biological phenomenon, indirectly measured with test scores. Changing the test scores would not change your brain, but changing your brain would change your test scores.

Causal discovery

- We always infer causal structures from statistical associations, but we typically start with a very narrow range of causal structures under consideration: X→Y or X Y
- With the d-separation rule, you can see that we can often rule in or out a much larger set of causal structures.
- Simple algorithm (SGS in TETRAD- many other/better available):
  - Start with a complete undirected graph on all variables, with edges between all variables.
  - For each pair of variables X and Y, and each set of other variables S, see if X is independent of Y given S; if so, remove the edge between X and Y.
  - Find colliders by checking for conditional dependence (X→S←Y implies X and Y are marginally independent but conditional on S become dependent); orient the edges of colliders.
  - Try to orient undirected edges by consistency with already-oriented edges; do this recursively until no more edges can be oriented.
- Controversial!

Conclusions

- Since the early 1980s, there has been a revolution in how we think about the link between statistical evidence and causal inferences
- This revolution has affected many disciplines and elicited many debates
- The tool set is in many ways different from typical statistical analysis, promoting clear articulation of questions and clear thinking about sources of bias
- These tools foster efficient communication, better reasoning about study design and analysis, and improved interpretation

Backup slides

Conditioning on a Collider

If two variables are statistically independent, but have a common effect, then, within levels of this effect, they will be statistically dependent.

Really.

Usually.

78

A collider anecdote

Some tall people are fast, and some are slow.

Some short people are fast, and some are slow.

Knowing that somebody in the general population is short does not give you information about whether they are fast or slow.

NBA ball players must be either very tall, or very fast.

If you know an NBA ball player is short… what do you know about his speed?

A collider anecdote

Some tall people are fast, and some are slow.

Some short people are fast, and some are slow.

Knowing that somebody in the general population is short does not give you information about whether they are fast or slow.

NBA ball players must be either very tall, or very fast.

If you know an NBA ball player is short… what do you know about his speed?

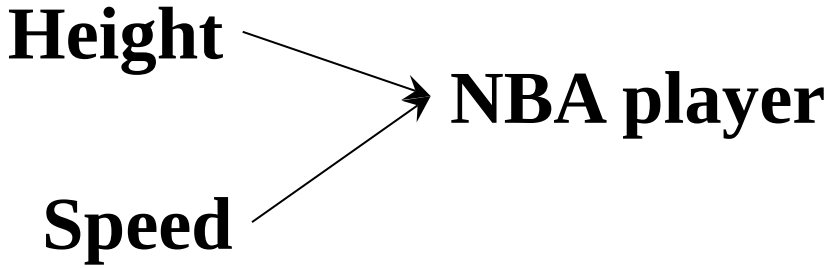

80

A collider anecdote

I throw a party, and I only invite people who are either very rich or very funny.

You come to my party (you are very funny) and get stuck talking to the most **boring** person you have ever met.

Is he rich?

81

A collider anecdote

I throw a party, and I only invite people who are either very rich or very funny.

You come to my party (you are very funny) and get stuck talking to the most **boring** person you have ever met.

Is he rich?

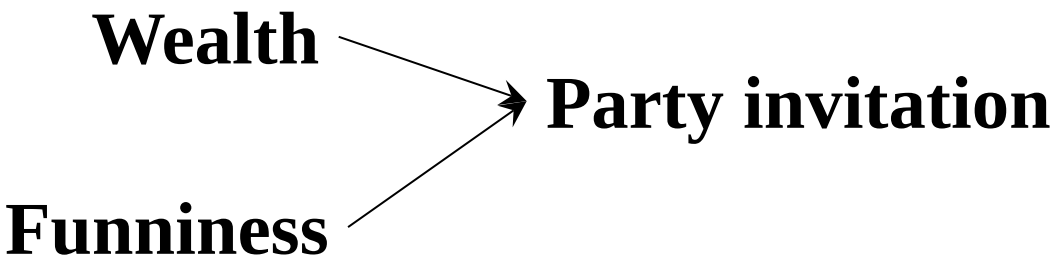

82

A Collider Simulation: Please try this at home.

Make up some rules following the structure of the DAG, such that X and Y are independent causes of Z, then assess association between X and Y after conditioning on Z:

- X~N(0,1)
- Y~N(0,1)
- e~N(0,1)
- Z=X+Y+e
- n=100

83

X and Y are independent; X and Z are positively associated

Scatter X , Y

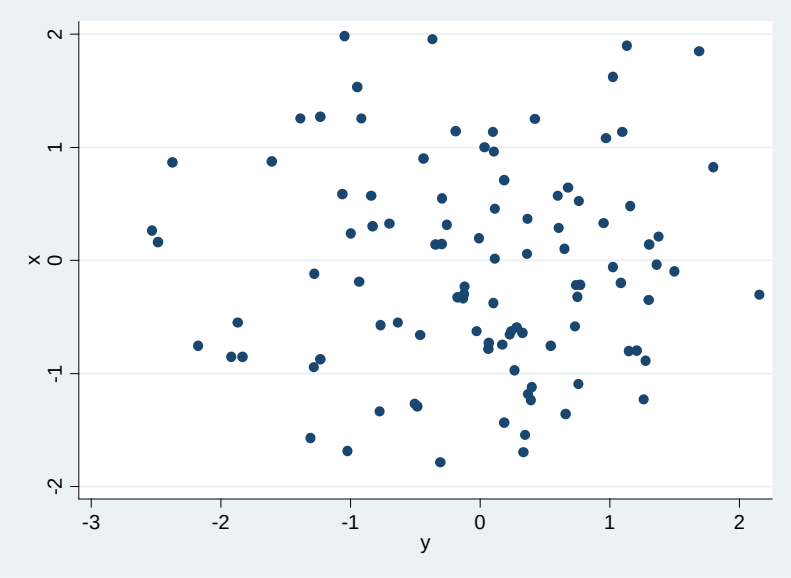

Scatter X , Z

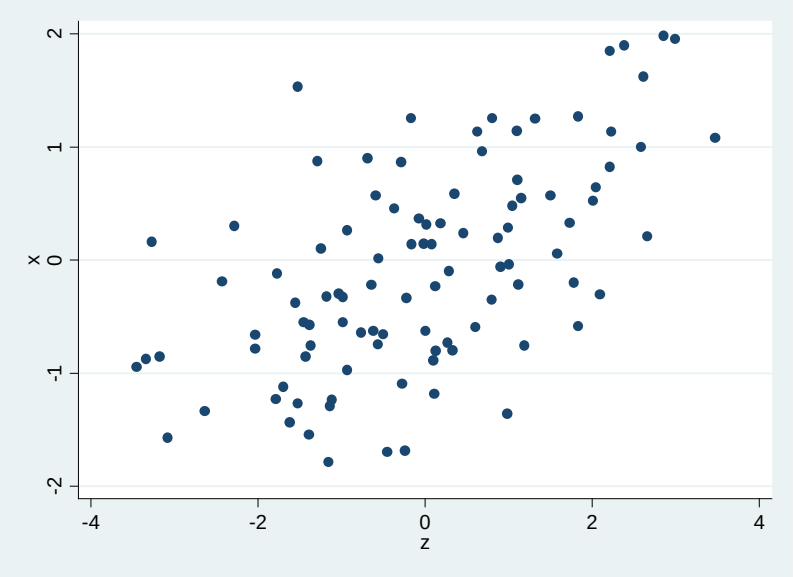

Independent

Positively Associated

84

X is inversely associated with Y, conditional on Z

Scatter X , residual (Y|Z)

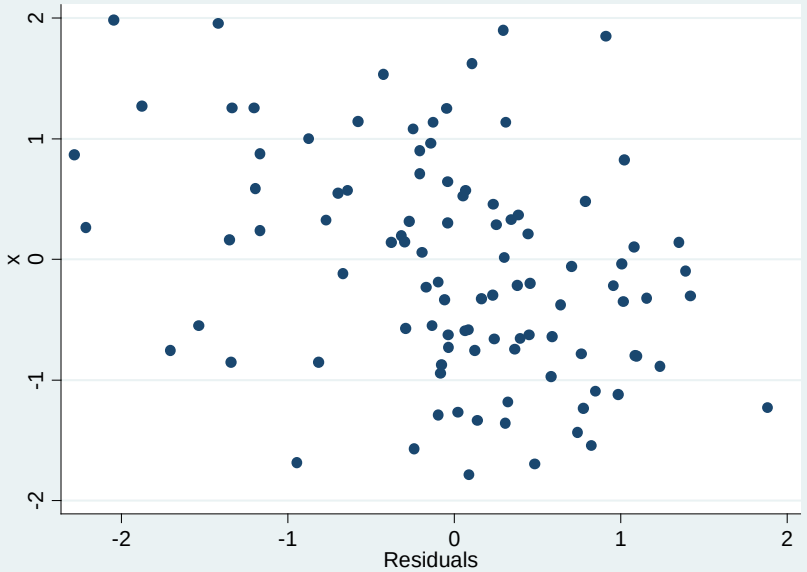

Inversely Associated

85

Getting Rich on Collider Bias

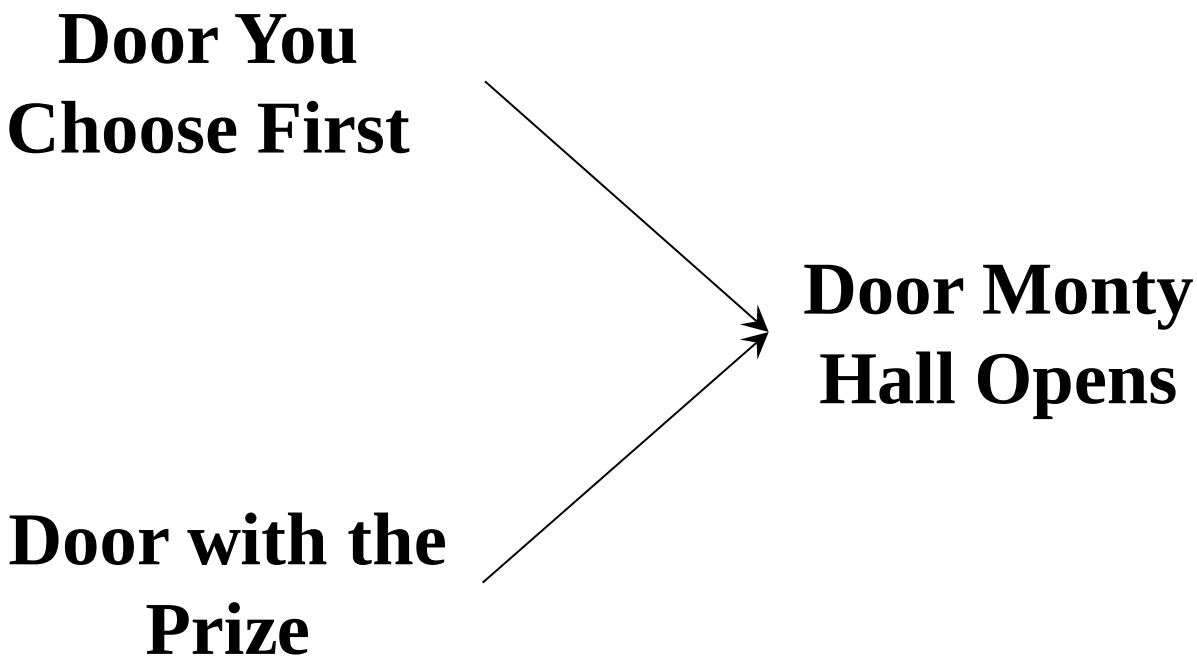

86

Getting Rich on Collider Bias

- The “Monty Hall Problem” / Let’s make a deal
- You can choose among 3 doors:
  - Behind one is a fabulous prize
  - Behind the other two is nothing (or perhaps, something crummy).
- You make a preliminary choice, and Monty Hall will open one of the other doors, always choosing one with nothing behind it.
- You may now revise your choice: do you keep the door you first chose or should you switch?
- Marilyn vos Savant had a famous fight about the Monty Hall problem… but the effect is so large that you can figure it out empirically in about 10 minutes.

87

Limitations and controversies

88