Lecture 5: Conditioning Strategies for Identifying Causal Effects

POL-GA 1251
Quantitative Political Analysis II
Prof. Cyrus Samii
NYU Politics

February 11, 2019

Overview

- Conditioning to remove bias with the backdoor criterion
- Conditioning that introduces or increases bias
  - Collider bias
  - Sample selection
  - Post-treatment bias
  - Bias amplification
  - Model dependence and misspecification

Confounding

- Considering arbitrary draws from a population or DGP.
- Treatment, D, outcome, Y, observed covariates, X, and unobserved pre-treatment variables, U.

Confounding

- Considering arbitrary draws from a population or DGP.
- Treatment, D, outcome, Y, observed covariates, X, and unobserved pre-treatment variables, U.
- Interested in effect of D on Y: $p(Y|do(d))$ or, e.g., $\mathrm{E}[Y(d) - Y(d')]$ for $d, d' \in \mathscr{D}_0$.
- Realized bivariate distribution of Y and D does not characterize this effect because of "back door paths."

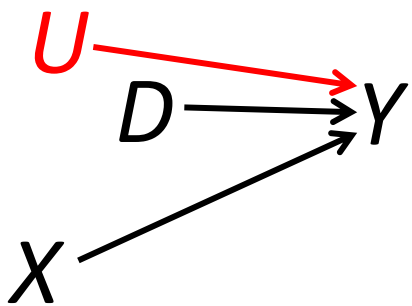

Target

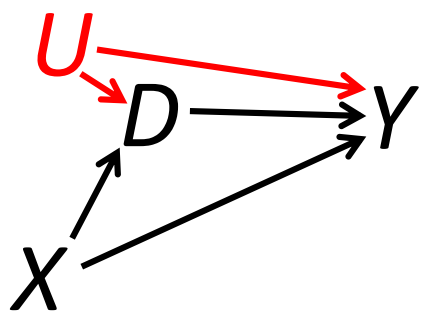

Realized

Confounding

For binary treatment $D = 0, 1$, recall the decomposition:

$$\begin{aligned}\mathrm{E}[\hat{\rho}] &= E[Y|D=1] - E[Y|D=0] = E[Y_1|D=1] - E[Y_0|D=0] \\ &= \underbrace{E[Y_1|D=1] - E[Y_0|D=1]}_{\text{Average treatment effect on the treated (ATT)}} + \underbrace{E[Y_0|D=1] - E[Y_0|D=0]}_{\text{Selection bias wrt } Y_0}.\end{aligned}$$

Biased for ATT when Y_0's of treated differ from those of controls.

Confounding

Similarly:

$$\begin{aligned}\mathrm{E}[\hat{\rho}] &= E[Y|D=1] - E[Y|D=0] = E[Y_1|D=1] - E[Y_0|D=0] \\ &= \underbrace{E[Y_1|D=0] - E[Y_0|D=0]}_{\text{Average treatment effect on the controls (ATC)}} + \underbrace{E[Y_1|D=1] - E[Y_1|D=0]}_{\text{Selection bias wrt } Y_1}.\end{aligned}$$

Biased for ATC when Y_1's of treated differ from those of controls.

Confounding

Finally, letting $\pi = \Pr[D = 1]$ (cf. CCI):

$$\begin{aligned}
\mathrm{E}[\hat{\rho}] &= E[Y|D=1] - E[Y|D=0] \\
&= \underbrace{\rho}_{\text{Average treatment effect}} + \underbrace{E[Y_0|D=1] - E[Y_0|D=0]}_{\text{Selection bias wrt } Y_0} \\
&\quad + \underbrace{(1-\pi)(E[\rho|D=1] - E[\rho|D=0])}_{\text{Selection bias wrt } \rho}
\end{aligned}$$

Biased for ATE when either Y_0's or ρ's of treated differ from those of controls. (Could be rewritten in terms of selection bias wrt to Y_1's and ρ as well.)

Confounding

Recast in terms of regression:

$$Y = \underbrace{E[Y_0]}_{\alpha} + \underbrace{DE[Y_1 - Y_0]}_{+\quad D\rho} + \underbrace{D(Y_1 - E[Y_1]) + (1-D)(Y_0 - E[Y_0])}_{+\quad \varepsilon}$$

or

$$= \underbrace{E[Y_0]}_{\alpha} + \underbrace{DE[Y_1 - Y_0]}_{+\quad D\rho} + \underbrace{(Y_0 - E[Y_0]) + D[(Y_1 - Y_0) - (E[Y_1] - E[Y_0])]}_{+\quad \varepsilon}$$

where correlation between ε and D corresponds to exactly the same forms of selection bias as specified above. OVB formula allows one to characterize the bias.

Confounding

- Can conditioning help? Depends.

Confounding

- Can conditioning help? Depends.
- First, define conditioning as graph operation. Condition on $V=$
  1. If V is a *collider*, link all pairs of parents of V with an undirected edge, connoting induced dependency.
  2. For any *ancestor* of V, if this ancestor is itself a collider, link all pairs of parents of this ancestor with undirected arcs to connote induced dependencies.
  3. Erase V from the graph and all edges connected with V.

Confounding

- Can conditioning help? Depends.
- First, define conditioning as graph operation. Condition on $V=$
  1. If V is a *collider*, link all pairs of parents of V with an undirected edge, connoting induced dependency.
  2. For any *ancestor* of V, if this ancestor is itself a collider, link all pairs of parents of this ancestor with undirected arcs to connote induced dependencies.
  3. Erase V from the graph and all edges connected with V.

  Practice:

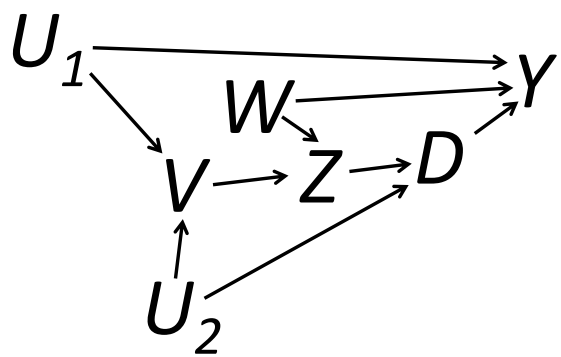

Backdoor paths from D to Y? Consequence on conditioning on (V, W)?

Backdoor criterion

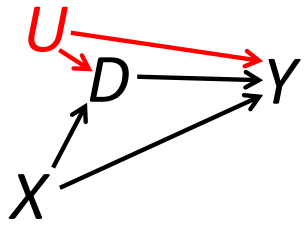

Backdoor criterion

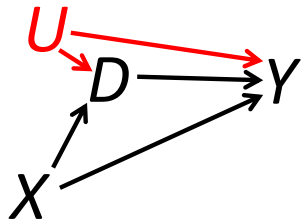

Pearl (2009) and CCI: given backdoor path between D and Y, conditioning on X identifies the effect of D on Y if

Backdoor criterion

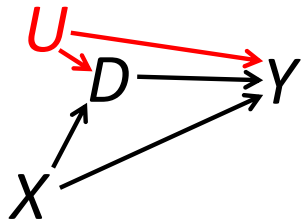

Pearl (2009) and CCI: given backdoor path between D and Y, conditioning on X identifies the effect of D on Y if

1. X blocks (i.e. "d-separates") all backdoor paths, where each backdoor path is
   - a mediation chain of form $A \to C \to B$, where C is in X, or
   - a fork of form $A \leftarrow C \to B$, where C is in X, or
   - an inverted fork where $A \to C \leftarrow B$, and C and all descendants of C are not in X,
2. X does not contain descendants of D lying on directed path to Y.

Backdoor criterion

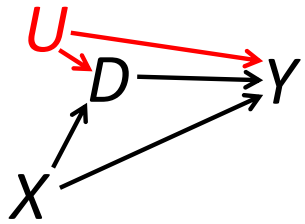

Pearl (2009) and CCI: given backdoor path between D and Y, conditioning on X identifies the effect of D on Y if

1. X blocks (i.e. "d-separates") all backdoor paths, where each backdoor path is
   - a mediation chain of form $A \to C \to B$, where C is in X, or
   - a fork of form $A \leftarrow C \to B$, where C is in X, or
   - an inverted fork where $A \to C \leftarrow B$, and C and all descendants of C are not in X,
2. X does not contain descendants of D lying on directed path to Y.

Can satisfy backdoor criterion *if* graph allows. Graph is the DGP.

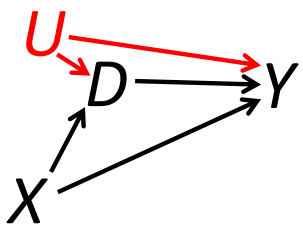

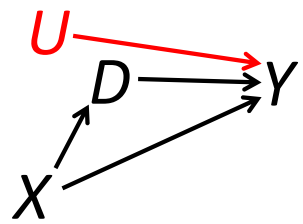

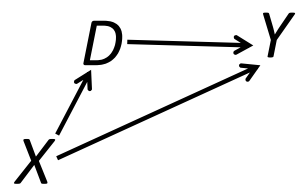

Backdoor criterion and CIA

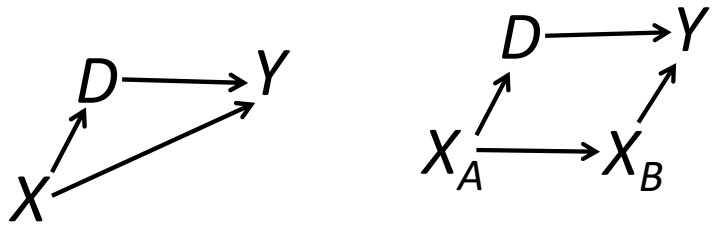

- Satisfaction of backdoor criterion is analogous to CIA:

$$D \perp\!\!\!\perp Y(d)|X \text{ and } 0 < p(d) < 1 \text{ for all } d \in \mathscr{D}$$

Backdoor criterion and CIA

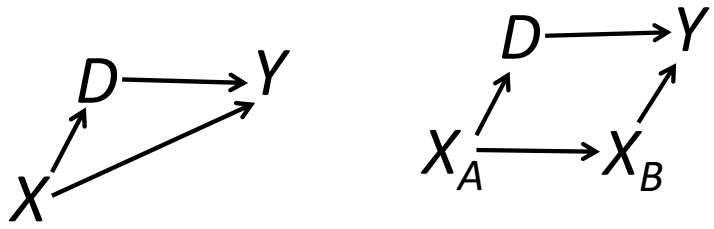

- Satisfaction of backdoor criterion is analogous to CIA:

$$D \perp\!\!\!\perp Y(d)|X \text{ and } 0 < p(d) < 1 \text{ for all } d \in \mathscr{D}$$

- Two ways to think about requisite conditioning (Imbens 2000):
  - Assignment mechanism.
  - Outcome DGP.
- *Either* is *sufficient*.

Backdoor criterion and CIA

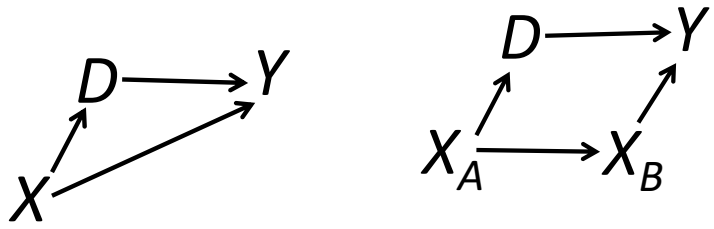

- Satisfaction of backdoor criterion is analogous to CIA:

$$D \perp\!\!\!\perp Y(d)|X \text{ and } 0 < p(d) < 1 \text{ for all } d \in \mathscr{D}$$

- Two ways to think about requisite conditioning (Imbens 2000):
  - Assignment mechanism.
  - Outcome DGP.
- *Either* is *sufficient*.
- Key implication is the ability to rewrite counterfactual distributions as observable distributions, e.g. $\mathrm{E}[Y(d)|D \neq d, X] = \mathrm{E}[Y(d)|D = d, X]$.

Conditioning that increases bias

Conditioning that increases bias

- Collider bias
- Sample selection
- Post-treatment bias
- Bias amplification
- Model dependence and misspecification

Collider bias

Collider:

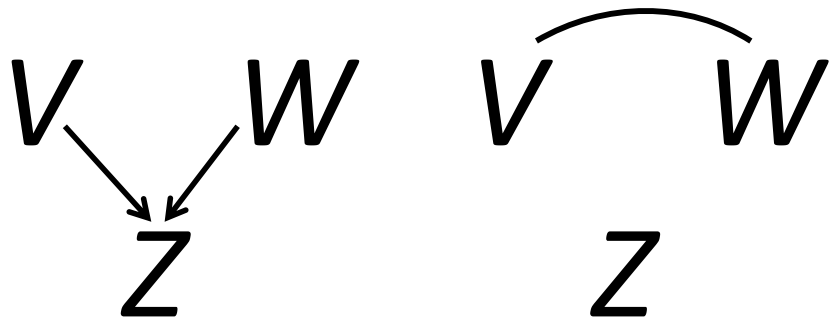

Collider bias

Lagged dependent variable as a collider:

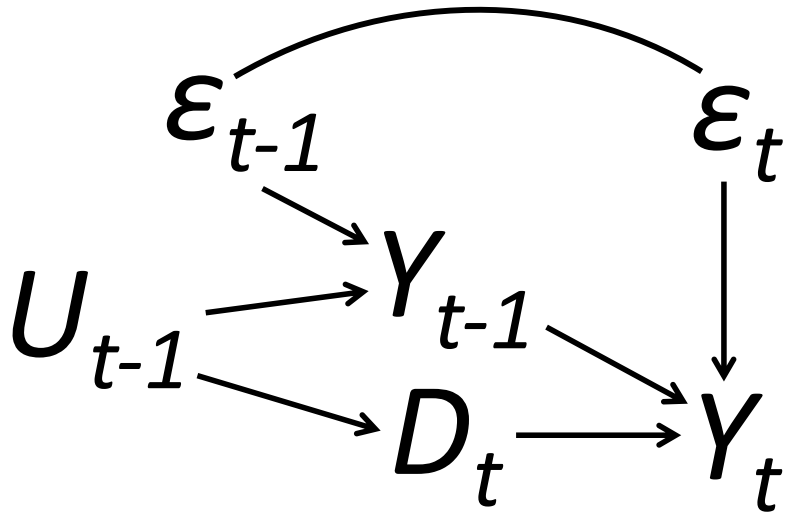

Sample selection

Matthew Hahn
@3rdreviewer

Follow

You can be a professional basketball player, no matter how tall you are!
No correlation between height and scoring success in the NBA:

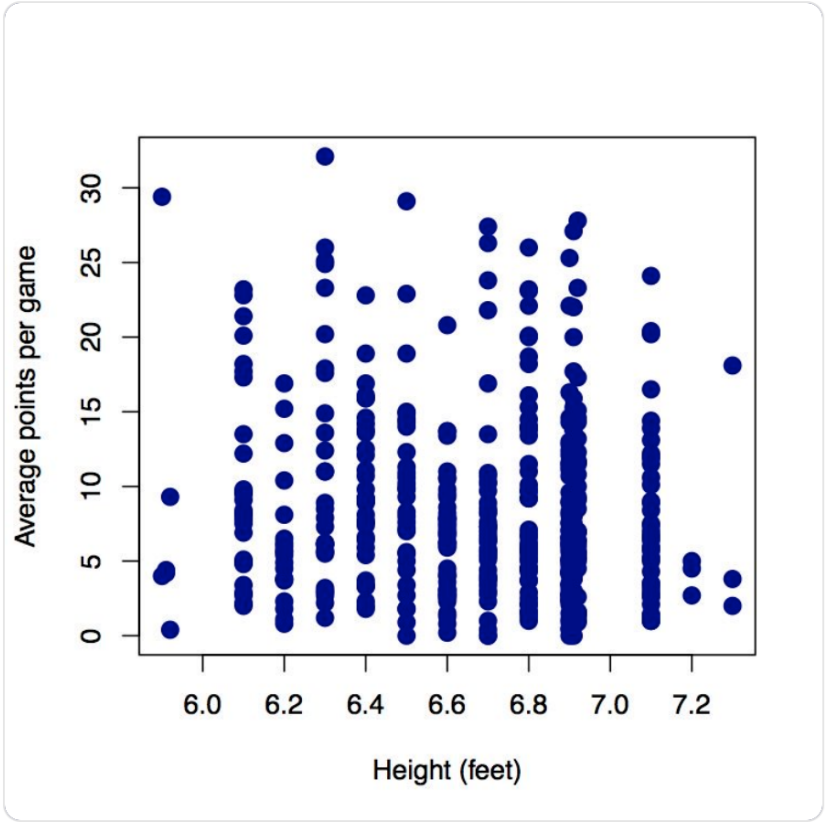

12:22 PM - 10 Mar 2017

Sample selection

Sample selection problem is a collider problem:

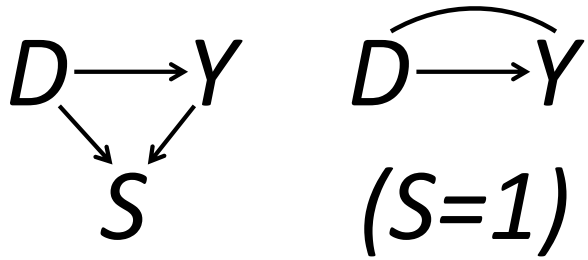

Sample selection

- Let $S_d = 1$ when observed under treatment d and 0 when unobserved. ($D \to S$)
- Consider ATE with $D = 0, 1$.

Sample selection

- Let $S_d = 1$ when observed under treatment d and 0 when unobserved. ($D \to S$)
- Consider ATE with $D = 0, 1$.
- Then by total probability we can decompose ρ as follows,

$$\begin{aligned}
\rho &= \mathrm{E}[Y_1 - Y_0] \\
&= (\mathrm{E}[Y_1|S_1 = 1]\Pr[S_1 = 1] - \mathrm{E}[Y_0|S_0 = 1]\Pr[S_0 = 1]) \\
&\quad + (\underbrace{\mathrm{E}[Y_1|S_1 = 0]}_{A}\Pr[S_1 = 0] - \underbrace{\mathrm{E}[Y_0|S_0 = 0]}_{B}\Pr[S_0 = 0]).
\end{aligned}$$

Sample selection

- Let $S_d = 1$ when observed under treatment d and 0 when unobserved. ($D \to S$)
- Consider ATE with $D = 0, 1$.
- Then by total probability we can decompose ρ as follows,

$$\begin{aligned}\rho &= \mathrm{E}[Y_1 - Y_0] \\ &= (\mathrm{E}[Y_1|S_1 = 1]\Pr[S_1 = 1] - \mathrm{E}[Y_0|S_0 = 1]\Pr[S_0 = 1]) \\ &\quad + (\underbrace{\mathrm{E}[Y_1|S_1 = 0]}_{A}\Pr[S_1 = 0] - \underbrace{\mathrm{E}[Y_0|S_0 = 0]}_{B}\Pr[S_0 = 0]).\end{aligned}$$

- A and B unidentified in data, but needed to estimate ρ.

Sample selection

- Let $S_d = 1$ when observed under treatment d and 0 when unobserved. ($D \to S$)
- Consider ATE with $D = 0, 1$.
- Then by total probability we can decompose ρ as follows,

$$\begin{aligned}\rho &= \mathrm{E}[Y_1 - Y_0] \\ &= (\mathrm{E}[Y_1|S_1 = 1]\Pr[S_1 = 1] - \mathrm{E}[Y_0|S_0 = 1]\Pr[S_0 = 1]) \\ &\quad + (\underbrace{\mathrm{E}[Y_1|S_1 = 0]}_{A}\Pr[S_1 = 0] - \underbrace{\mathrm{E}[Y_0|S_0 = 0]}_{B}\Pr[S_0 = 0]).\end{aligned}$$

- A and B unidentified in data, but needed to estimate ρ.
- Bias arises when:
  - S_0 correlated with Y_0.
  - S_1 correlated with Y_1.
  - i.e., $S \leftarrow Y$, or could also be $S \to Y$ (backdoor criterion condition 2), or (S, Y) descendants variables with such relationships.

Sample selection

Heckman's framework:

- Suppose the decision to work is a function of whether expected wage, Y^*, which is a linear function of X (observed) and ν (unobserved), is greater than "reservation wage", R,

$$\text{work if } Y^* = X\gamma + \nu > R.$$

Sample selection

Heckman's framework:

- Suppose the decision to work is a function of whether expected wage, Y^*, which is a linear function of X (observed) and ν (unobserved), is greater than "reservation wage", R,

$$\text{work if } Y^* = X\gamma + \nu > R.$$

- Given that the person works, suppose actual wages, Y, are determined by X (observed) and ε,

$$Y = X\beta + \varepsilon.$$

Sample selection

Heckman's framework:

- Suppose the decision to work is a function of whether expected wage, Y^*, which is a linear function of X (observed) and ν (unobserved), is greater than "reservation wage", R,

$$\text{work if } Y^* = X\gamma + \nu > R.$$

- Given that the person works, suppose actual wages, Y, are determined by X (observed) and ε,

$$Y = X\beta + \varepsilon.$$

- If we just look at people working, we have,

$$\mathrm{E}[Y|X] = X\beta + \mathrm{E}[\varepsilon|X, X\gamma + \nu > R].$$

- A working person with small X likely had large ν to clear R.
- If ε and ν are positively correlated, then $\mathrm{E}[\varepsilon|X, X\gamma + \nu > R]$ is large when X is small.
- Thus, X and ε are correlated in the "selected" sample.

Sample selection

Heckman's selection bias as a collider problem:

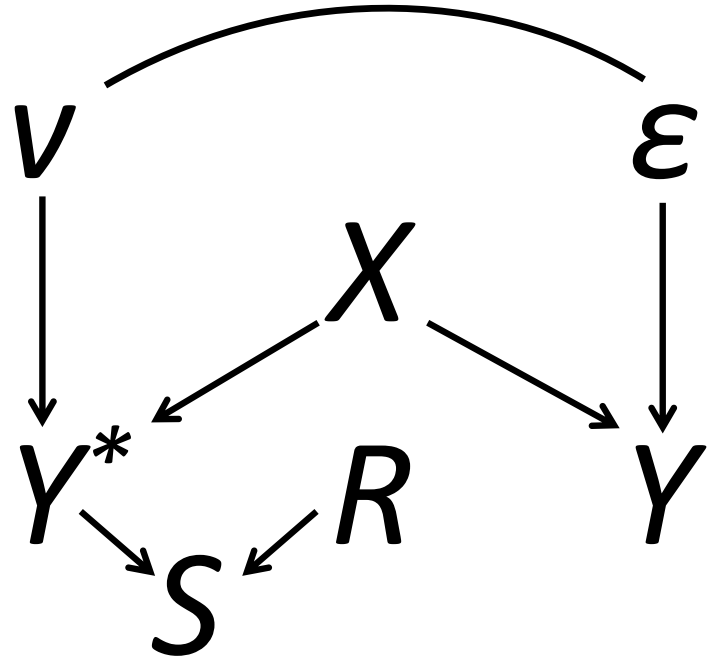

Sample selection

Other examples:

- GRE scores are a poor predictor of grad student performance. Why?
- A country's military strength is a poor predictor likelihood of winning a war. Why?

Sample selection

What are the solutions?

Sample selection

What are the solutions?

- Design-based solutions:
  - Seek research opportunities that are not plagued by these problems—natural or controlled experiments.
  - Be clear about the target population for which you are identified.

Sample selection

What are the solutions?

- Design-based solutions:
  - Seek research opportunities that are not plagued by these problems—natural or controlled experiments.
  - Be clear about the target population for which you are identified.
- Analytical solutions
  - Bounds and modeling (more on that later in the semester).

Post-treatment bias

- Recall condition 2 of the back door criterion: "X does not contain descendants of D lying on directed path to Y"

Post-treatment bias

- Recall condition 2 of the back door criterion: "X does not contain descendants of D lying on directed path to Y"
- Conditioning on such variables "steals away" from the treatment effect.

Post-treatment bias

- Recall condition 2 of the back door criterion: "X does not contain descendants of D lying on directed path to Y"
- Conditioning on such variables "steals away" from the treatment effect.
- Moreover, new dependencies can be induced (collider bias).

Post-treatment bias

- Recall condition 2 of the back door criterion: "X does not contain descendants of D lying on directed path to Y"
- Conditioning on such variables "steals away" from the treatment effect.
- Moreover, new dependencies can be induced (collider bias).
- Various ways to characterize this.

Post-treatment bias

- Suppose following DAG.
- Can we use S to remove bias due to U? Red arrows are key.

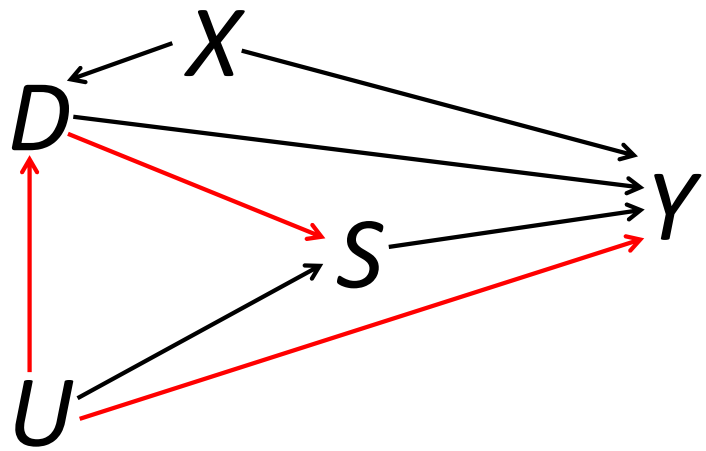

Post-treatment bias

- Suppose following DAG.
- Can we use S to remove bias due to U? Red arrows are key.

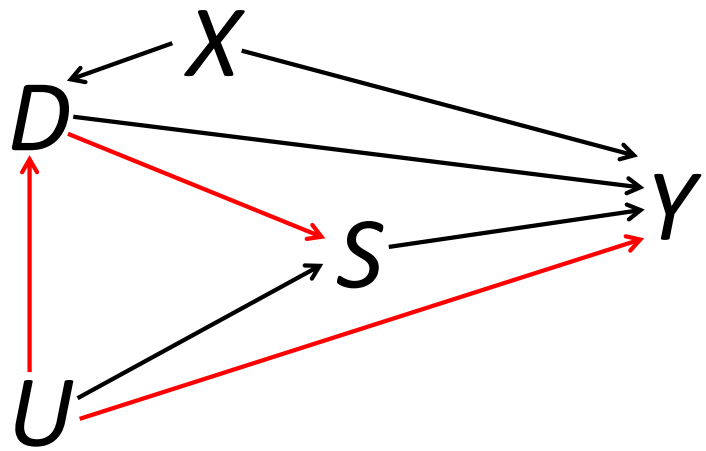

- If D affects S, conditioning on S *induces* new dependencies between U and D, even if there were none ex ante.
- If U affects Y in a manner that is not channeled exclusively through S, then conditioning on S does not block the confounding that U introduces to the effect of D on Y.

Post-treatment bias

Rosenbaum (1984) is a canonical analysis. First, suppose,

- as usual, $\mathscr{D} = \{0, 1\}$, SUTVA, (Y_1, Y_0), X.
- target estimand is $\rho = \mathrm{E}[Y_1 - Y_0]$.
- possible identifying assumptions are random assignment or CIA.

Post-treatment bias

Rosenbaum (1984) is a canonical analysis. First, suppose,

- as usual, $\mathscr{D} = \{0,1\}$, SUTVA, (Y_1, Y_0), X.
- target estimand is $\rho = \mathrm{E}[Y_1 - Y_0]$.
- possible identifying assumptions are random assignment or CIA.
- Add to the mix (S_1, S_0), post-treatment variables, where

$$S = DS_1 + (1-D)S_0$$

- Suppose we control for both S and X, yielding in expectation:

$$\begin{aligned}\tilde{\rho} &= \mathrm{E}_{S,X}\{\mathrm{E}[Y_i|D=1, S_i=s, X=x] - \mathrm{E}[Y_i|D=0, S_i=s, X=x]\} \\ &= \mathrm{E}_{S,X}\{\mathrm{E}[Y_1|D=1, S_1=s, X=x] - \mathrm{E}[Y_0|D=0, S_0=s, X=x]\}\end{aligned}$$

- Alarm bells should be ringing. Do you see why?

Post-treatment bias

To characterize the bias, let's first define,

$$\tilde{\nu} = \mathrm{E}_{S,X}\{\mathrm{E}[Y_1|S_1 = s, X = x] - \mathrm{E}[Y_0|S_0 = s, X = x]\},$$

the "net treatment difference." It is not a proper causal quantity.

Post-treatment bias

An expression for post-treatment bias:

$$\tilde{\rho} - \rho = \underbrace{(\tilde{\rho} - \tilde{\nu})}_{A} + \underbrace{(\tilde{\nu} - \rho)}_{B}$$

Post-treatment bias

An expression for post-treatment bias:

$$\tilde{\rho} - \rho = \underbrace{(\tilde{\rho} - \tilde{\nu})}_{A} + \underbrace{(\tilde{\nu} - \rho)}_{B}$$

- A assured to be zero only if random assignment or CIA holds for both (Y_1, Y_0) and (S_1, S_0):

$$(Y_1, S_1, Y_0, S_0) \perp\!\!\!\perp D|X \Rightarrow Y_{di} \perp\!\!\!\perp D|X, S_{di}$$

$$\begin{aligned}\tilde{\rho} - \tilde{\nu} = {} & \mathrm{E}_{S,X}\{\mathrm{E}[Y_1|D=1, S_1 = s, X = x] - \mathrm{E}[Y_0|D=0, S_0 = s, X = x]\} \\ & - \mathrm{E}_{S,X}\{\mathrm{E}[Y_1|S_1 = s, X = x] - \mathrm{E}[Y_0|S_0 = s, X = x]\}.\end{aligned}$$

Post-treatment bias

An expression for post-treatment bias:

$$\tilde{\rho} - \rho = \underbrace{(\tilde{\rho} - \tilde{\nu})}_{A} + \underbrace{(\tilde{\nu} - \rho)}_{B}$$

- A assured to be zero only if random assignment or CIA holds for both (Y_1, Y_0) and (S_1, S_0):

$$(Y_1, S_1, Y_0, S_0) \perp\!\!\!\perp D|X \Rightarrow Y_{di} \perp\!\!\!\perp D|X, S_{di}$$

$$\begin{aligned}\tilde{\rho} - \tilde{\nu} = \mathrm{E}_{S,X}\{\mathrm{E}[Y_1|D=1, S_1 = s, X = x] - \mathrm{E}[Y_0|D=0, S_0 = s, X = x]\} \\ - \mathrm{E}_{S,X}\{\mathrm{E}[Y_1|S_1 = s, X = x] - \mathrm{E}[Y_0|S_0 = s, X = x]\}.\end{aligned}$$

- But B assured to be zero only if $S_1 = S_0$ (no effect) for all i:

$$\tilde{\nu} - \rho = \mathrm{E}_{S,X}\{\mathrm{E}[Y_1|S_1 = s, X = x] - \mathrm{E}[Y_0|S_0 = s, X = x]\} - \mathrm{E}[Y_1 - Y_0]$$

- So, even if CIA or random assignment holds with respect to X, controlling for S may induce bias that could have been avoided by leaving S out.

Post-treatment bias

Intuitions about post-treatment bias:

- Suppose $\mathrm{E}[Y_{di}|S_{di}=s,X=x]=\alpha_d+x\beta+s\gamma$.
- Then the true treatment effect is given by,

$$\begin{aligned}\rho &= \mathrm{E}[(\alpha_1+X\beta+S_1\gamma)-(\alpha_0+X\beta+S_0\gamma)] \\ &= (\alpha_1-\alpha_0)+\mathrm{E}[S_1-S_0]\gamma.\end{aligned}$$

- The estimate that conditions on S is given by,

$$\begin{aligned}\tilde{\rho} &= \mathrm{E}[(\alpha_1+X\beta+S\gamma)-(\alpha_0+X\beta+S\gamma)] \\ &= (\alpha_1-\alpha_0).\end{aligned}$$

- The bias, $\tilde{\rho}-\rho=-\mathrm{E}[S_1-S_0]\gamma$ amounts to the portion of ρ that has been "stolen away" by conditioning on S.

Post-treatment bias

Sometimes we feel that we need to control for S because it proxies for some pre-treatment variable, U, which was not measured.

- Suppose that CIA holds when conditioning on X and U.
- We call S a "surrogate" for U when,

$$Y_{di} \perp\!\!\!\perp U | (S_{di}, X) \text{ for } d = 0, 1$$

  Then, controlling for S_{di} breaks dependency between Y_{di} and U.
- If so, bias component $A = 0$ (D's drop out of conditioning sets).
- Then, if $S_1 = S_0$ for all i as well, post-treatment bias is zero.
- These three conditions imply that controlling for S is justified for identifying ρ.

Post-treatment bias

- Conditions are the same as what we noted on the DAG:

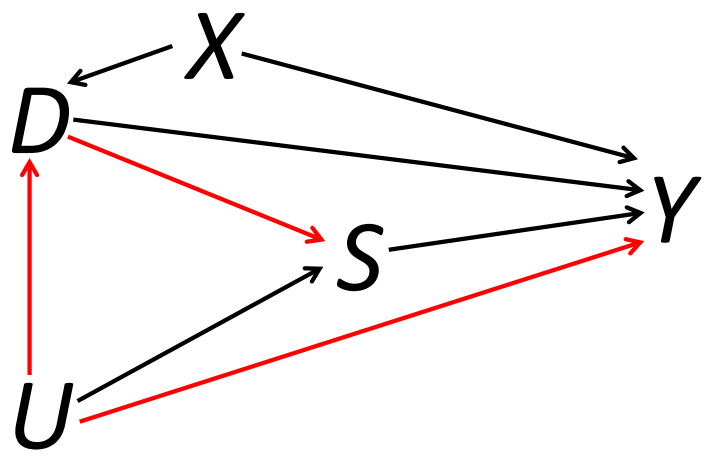

Post-treatment bias

Generally speaking,

- Controlling for post-treatment variables is usually a *bad idea* for estimating ATEs, SATEs, etc. in experiments or well-identified natural experiments.

Post-treatment bias

Generally speaking,

- Controlling for post-treatment variables is usually a *bad idea* for estimating ATEs, SATEs, etc. in experiments or well-identified natural experiments.
- In observational studies where lots of controls are needed for CIA, it may be justified, but *only if the necessary surrogacy and zero-effect assumptions* are plausible.

Post-treatment bias

Generally speaking,

- Controlling for post-treatment variables is usually a *bad idea* for estimating ATEs, SATEs, etc. in experiments or well-identified natural experiments.
- In observational studies where lots of controls are needed for CIA, it may be justified, but *only if the necessary surrogacy and zero-effect assumptions* are plausible.
- At the very least, if you feel that you have to control for a post-treatment variable, conduct a sensitivity analysis.

Post-treatment bias

Generally speaking,

- Controlling for post-treatment variables is usually a *bad idea* for estimating ATEs, SATEs, etc. in experiments or well-identified natural experiments.
- In observational studies where lots of controls are needed for CIA, it may be justified, but *only if the necessary surrogacy and zero-effect assumptions* are plausible.
- At the very least, if you feel that you have to control for a post-treatment variable, conduct a sensitivity analysis.
- We will look into questions of mediation and "direct" vs. "indirect" effects later in the semester.

Bias amplification

Bias amplification

- Another scenario where conditioning increases bias is the "Z-bias" situation (Clarke 2005; Pearl 2010).

Bias amplification

- Another scenario where conditioning increases bias is the "Z-bias" situation (Clarke 2005; Pearl 2010).
- Imagine that we have conditioned on X, but this is insufficient for CIA. We have another variable Z:

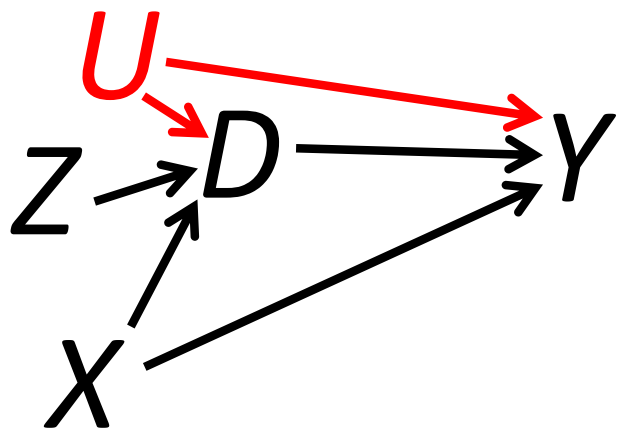

Bias amplification

- Another scenario where conditioning increases bias is the "Z-bias" situation (Clarke 2005; Pearl 2010).
- Imagine that we have conditioned on X, but this is insufficient for CIA. We have another variable Z:

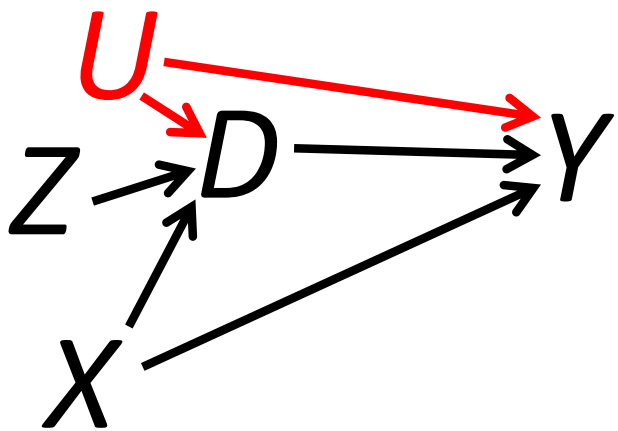

- Claim: the bias from conditioning on Z can be worse than leaving it out.

Bias amplification

- Regression illustration—suppose the specification

$$Y = \beta_1 + \beta_2 D + \beta_3 X + U,$$

but U still correlated with D even when controlling for X.

Bias amplification

- Regression illustration—suppose the specification

$$Y = \beta_1 + \beta_2 D + \beta_3 X + U,$$

but U still correlated with D even when controlling for X.

- Suppose Z correlated with D but not U or X, and that we try to fit the following specification with OLS,

$$Y = \beta_{1,Z} + \beta_{2,Z} D + \beta_{3,Z} X + \beta_4 Z + U.$$

Bias amplification

- Regression illustration—suppose the specification

$$Y = \beta_1 + \beta_2 D + \beta_3 X + U,$$

but U still correlated with D even when controlling for X.

- Suppose Z correlated with D but not U or X, and that we try to fit the following specification with OLS,

$$Y = \beta_{1,Z} + \beta_{2,Z} D + \beta_{3,Z} X + \beta_4 Z + U.$$

We can rewrite partialing out Z:

$$\tilde{Y} = \beta_{1,Z} + \beta_{2,Z} \tilde{D} + \beta_{3,Z} X + U.$$

Bias amplification

- Regression illustration—suppose the specification

$$Y = \beta_1 + \beta_2 D + \beta_3 X + U,$$

but U still correlated with D even when controlling for X.

- Suppose Z correlated with D but not U or X, and that we try to fit the following specification with OLS,

$$Y = \beta_{1,Z} + \beta_{2,Z} D + \beta_{3,Z} X + \beta_4 Z + U.$$

We can rewrite partialing out Z:

$$\tilde{Y} = \beta_{1,Z} + \beta_{2,Z} \tilde{D} + \beta_{3,Z} X + U.$$

- $\tilde{D}$ is purged of variation due to Z, *which was uncorrelated with* U, and so the correlation with U has been "concentrated"!

Bias amplification

- Regression illustration—suppose the specification

$$Y = \beta_1 + \beta_2 D + \beta_3 X + U,$$

but U still correlated with D even when controlling for X.

- Suppose Z correlated with D but not U or X, and that we try to fit the following specification with OLS,

$$Y = \beta_{1,Z} + \beta_{2,Z} D + \beta_{3,Z} X + \beta_4 Z + U.$$

We can rewrite partialing out Z:

$$\tilde{Y} = \beta_{1,Z} + \beta_{2,Z} \tilde{D} + \beta_{3,Z} X + U.$$

- $\tilde{D}$ is purged of variation due to Z, *which was uncorrelated with* U, and so the correlation with U has been "concentrated"!
- Bias from controlling for Z may be worse than omitting it.

Bias amplification

- Generalizing beyond "instruments" (Clarke 2005, 2009), suppose:

$$Y = \beta_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4 + \varepsilon,$$

with ε independent.

Bias amplification

- Generalizing beyond "instruments" (Clarke 2005, 2009), suppose:

$$Y = \beta_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4 + \varepsilon,$$

with ε independent.Want β_2, but have two misspecified models:

$$(A) \qquad Y = \beta_1^a + \beta_2^a X_2 + \varepsilon^a,$$

$$(B) \qquad Y = \beta_1^b + \beta_2^b X_2 + \beta_3^b X_3 + \varepsilon^b$$

Bias amplification

- Generalizing beyond "instruments" (Clarke 2005, 2009), suppose:

$$Y = \beta_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4 + \varepsilon,$$

with ε independent.Want β_2, but have two misspecified models:

$$(A) \qquad Y = \beta_1^a + \beta_2^a X_2 + \varepsilon^a,$$
$$(B) \qquad Y = \beta_1^b + \beta_2^b X_2 + \beta_3^b X_3 + \varepsilon^b$$

By FWL and OVB formula,

$$\mathrm{E}[\hat{\beta}_2^a] = \beta_2 + \beta_3 \delta_{3,2} + \beta_4 \delta_{4,2}$$
$$\mathrm{E}[\hat{\beta}_2^b] = \beta_2 + \beta_4 \tilde{\delta}_{4,2(3)}$$
$$\left|\mathrm{E}[\hat{\beta}_2^a] - \mathrm{E}[\hat{\beta}_2^b]\right| = \left|\beta_3 \delta_{3,2} + \beta_4 (\delta_{4,2} - \tilde{\delta}_{4,2(3)})\right|$$

Bias amplification

- Generalizing beyond "instruments" (Clarke 2005, 2009), suppose:

$$Y = \beta_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4 + \varepsilon,$$

with ε independent.Want β_2, but have two misspecified models:

$$(A) \qquad Y = \beta_1^a + \beta_2^a X_2 + \varepsilon^a,$$
$$(B) \qquad Y = \beta_1^b + \beta_2^b X_2 + \beta_3^b X_3 + \varepsilon^b$$

By FWL and OVB formula,

$$\mathrm{E}[\hat{\beta}_2^a] = \beta_2 + \beta_3 \delta_{3,2} + \beta_4 \delta_{4,2}$$
$$\mathrm{E}[\hat{\beta}_2^b] = \beta_2 + \beta_4 \tilde{\delta}_{4,2(3)}$$
$$\left|\mathrm{E}[\hat{\beta}_2^a] - \mathrm{E}[\hat{\beta}_2^b]\right| = \left|\beta_3 \delta_{3,2} + \beta_4 (\delta_{4,2} - \tilde{\delta}_{4,2(3)})\right|$$

- Possible that (B) is more biased than (A). E.g., sps all positive biases, but β_3 near zero, β_4 large, and $\delta_{4,2} << \tilde{\delta}_{4,2(3)}$. Note last condition implies substantial correlation between X_{i2} and X_{i3}.

Bias amplification more generally

- Pearl (2010):

  [B]eing a good predictor of treatment assignment [may compromise] the bias-reducing potential of a covariate, for it tends to amplify bias due to other, uncontrolled confounders. One would do better therefore to rank order covariates based on their importance with respect to the outcome variable.

Bias amplification more generally

- Pearl (2010):

  [B]eing a good predictor of treatment assignment [may compromise] the bias-reducing potential of a covariate, for it tends to amplify bias due to other, uncontrolled confounders. One would do better therefore to rank order covariates based on their importance with respect to the outcome variable.

- This may be overstating it, but the general recommendation to pay attention to relations to outcomes as well as assignment is warranted when trying to meet the CIA .

Bias amplification more generally

- Pearl (2010):

  [B]eing a good predictor of treatment assignment [may compromise] the bias-reducing potential of a covariate, for it tends to amplify bias due to other, uncontrolled confounders. One would do better therefore to rank order covariates based on their importance with respect to the outcome variable.

- This may be overstating it, but the general recommendation to pay attention to relations to outcomes as well as assignment is warranted when trying to meet the CIA .
- Bound and Solon (1999) and Middleton et al. (2016) show how fixed effects can be especially troublesome in terms of bias amplification and "bias unmasking."

Model dependence and misspecification’

Model dependence and misspecification'

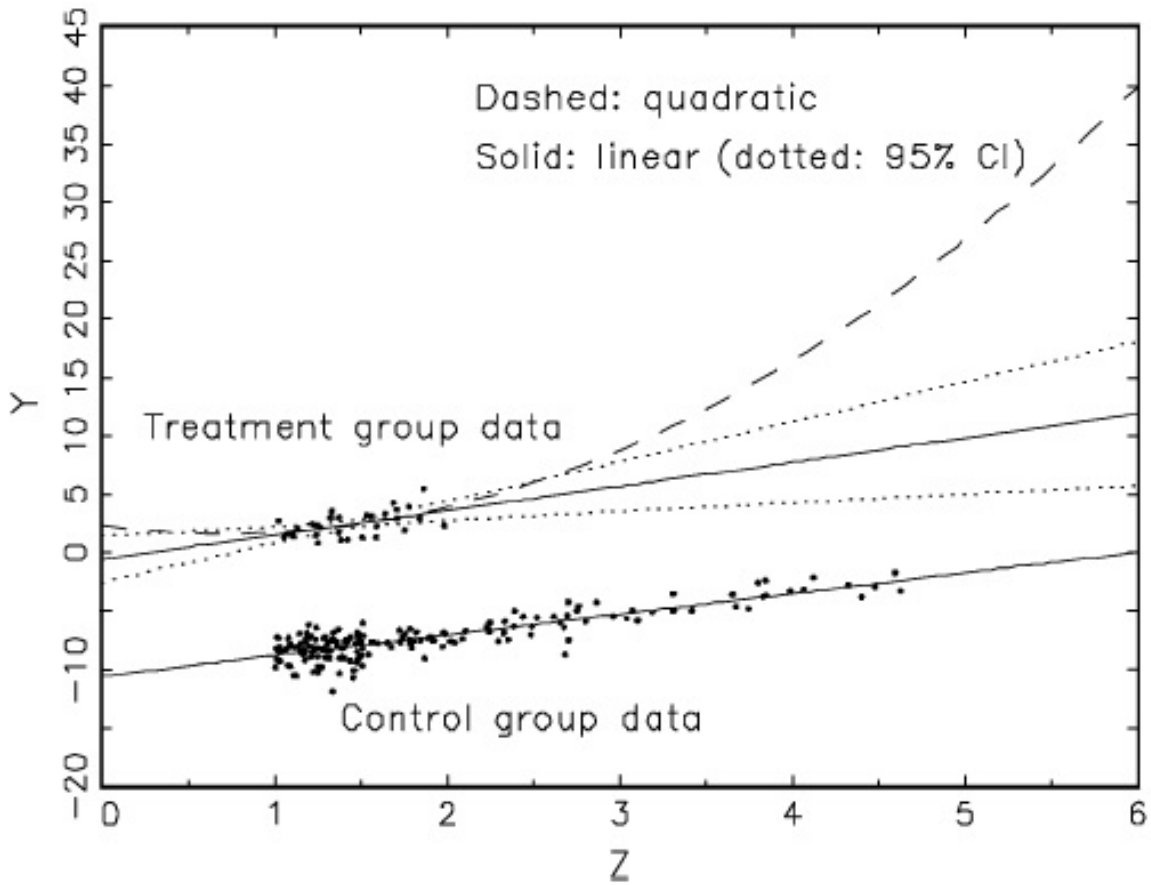

Fig. 4 An illustration of how the degree of extrapolation bias is more severe (and model dependent) than interpolation bias.

(King and Zeng, 2006)

Model dependence and misspecification'

This is the problem of extreme counterfactuals—predictions, what-if questions, and causal inferences that are so far from the data that inferences wind up being drawn on the basis of minor model specification choices no one would like to defend, rather than empirical evidence....Our confidence interval for counterfactuals farther from the data are wider, but the inference may be considerably more uncertain than the confidence interval indicates.
(King and Zeng, 2006, p. 132)

Model dependence and misspecification

- Recall under CIA, we get ATE as,

$$\rho = \mathrm{E}_X \{\mathrm{E}[Y_{1i}|D_i = 1, X_i = x] - \mathrm{E}[Y_{0i}|D_i = 0, X_i = x]\}$$

- Potential for bias when we cannot implement this directly – that is, when we cannot really line up treated and control observations for all values of x.
- *To make causal inferences in situations with nonoverlapping densities, we must therefore either eliminate the region outside of common support or attempt to extrapolate to the needed data, such as by using a parametric model.* (p. 149)
- Extrapolation to areas with no overlap can be perilous.
- Motivates strategy of "changing the goal posts" by focusing on effects in areas of common support and non-parametric methods like matching and weighting.

Model dependence and misspecification

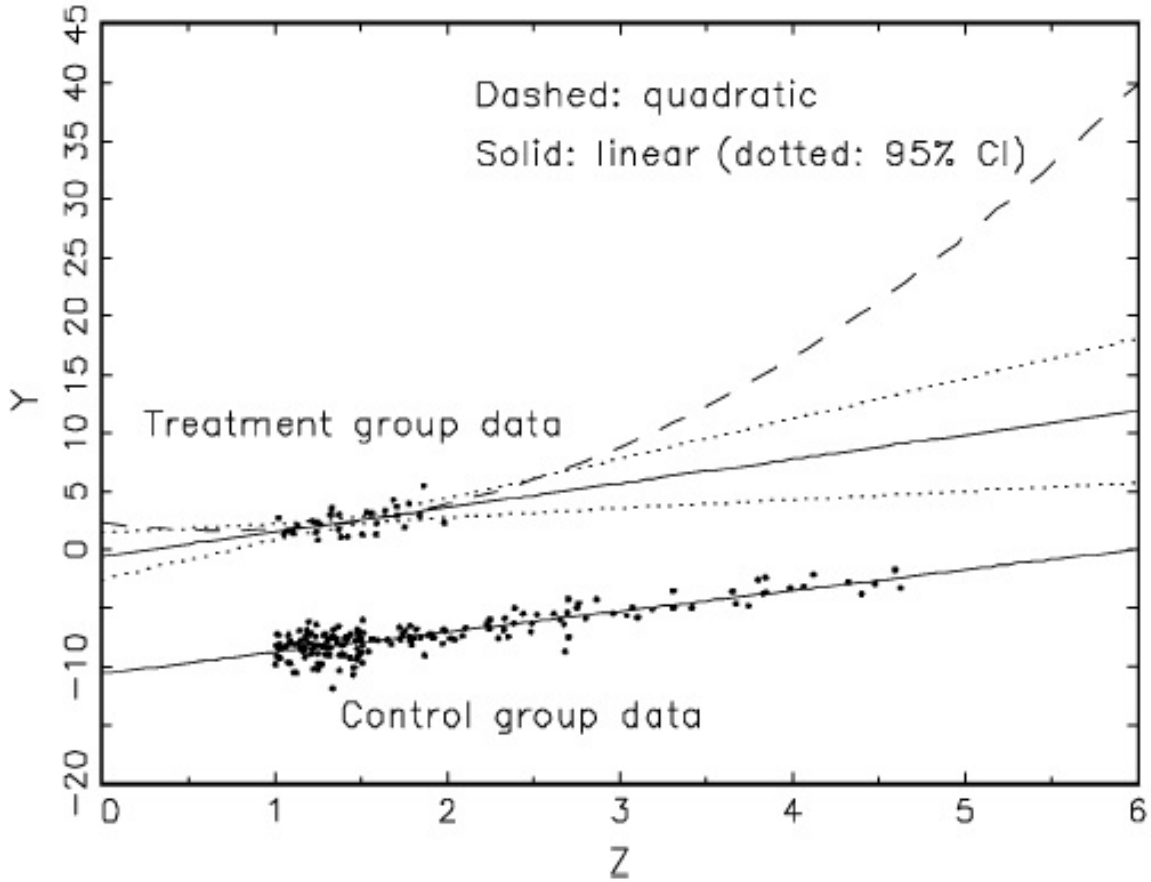

Fig. 4 An illustration of how the degree of extrapolation bias is more severe (and model dependent) than interpolation bias.

Remarks on conditioning strategies for identification

- These biases raise concerns about CIA studies.
- E.g., suppose in a regression study someone adds a control variable to the specification and the effect estimate changes. What is the explanation?

Remarks on conditioning strategies for identification

- These biases raise concerns about CIA studies.
- E.g., suppose in a regression study someone adds a control variable to the specification and the effect estimate changes. What is the explanation?
  1. Is it because you are actually reducing confounding?

Remarks on conditioning strategies for identification

- These biases raise concerns about CIA studies.
- E.g., suppose in a regression study someone adds a control variable to the specification and the effect estimate changes. What is the explanation?
  1. Is it because you are actually reducing confounding?
  2. Or is it because you are shifting the effective sample (remember last lecture)?

Remarks on conditioning strategies for identification

- These biases raise concerns about CIA studies.
- E.g., suppose in a regression study someone adds a control variable to the specification and the effect estimate changes. What is the explanation?
  1. Is it because you are actually reducing confounding?
  2. Or is it because you are shifting the effective sample (remember last lecture)?
  3. Or is it because you are introducing bias amplification?

Remarks on conditioning strategies for identification

- These biases raise concerns about CIA studies.
- E.g., suppose in a regression study someone adds a control variable to the specification and the effect estimate changes. What is the explanation?
  1. Is it because you are actually reducing confounding?
  2. Or is it because you are shifting the effective sample (remember last lecture)?
  3. Or is it because you are introducing bias amplification?
  4. Or is it because you are compounding misspecification?

Remarks on conditioning strategies for identification

- These biases raise concerns about CIA studies.
- E.g., suppose in a regression study someone adds a control variable to the specification and the effect estimate changes. What is the explanation?
  1. Is it because you are actually reducing confounding?
  2. Or is it because you are shifting the effective sample (remember last lecture)?
  3. Or is it because you are introducing bias amplification?
  4. Or is it because you are compounding misspecification?
  5. Or is it another form of bad control, like post-treatment control?
- You may have no good way to tell.
- Not just regression either, forms of each apply to, e.g., matching and weighting.

Remarks on conditioning strategies for identification

- If one can minimize reliance on CIA, control strategies, and modeling, then one can reduce the potential for these "errors of commission."
- When control strategies are *unavoidable*, the general recommendation is sensitivity analysis, to check for sensitivity to violations of CIA, assumptions about post-treatment bias, bias amplification, etc.