Front-Door Adjustment

Ethan Fosse
Princeton University

Fall 2016

Preliminaries

Two Problems of Causal Inference

1. Given the causal structure, what are the estimated effects?
2. Given data about a system, what is the causal structure?

- Social scientists focus almost entirely on #1, ignoring #2
- Both are vulnerable to unobserved common confounders (i.e., violation of causal sufficiency)
- Bias formulas allow us to expand the causal structure to include unobserved common confounders and adjust estimates accordingly

Basic Notation and Definitions

- Let A denote the treatment received and Y denote the observed post-treatment outcome
- X is an observed confounder, U is an unobserved confounder, and M is a mediator between A and Y
- Let Y_a denote the potential outcome for an individual if treatment A is fixed to a
- We will compare potential outcomes for any two treatment levels of A, a_1 and a_0, where a_0 is the reference level
- The corresponding potential outcomes are Y_{a_1} and Y_{a_0}

Assumptions

- **SUTVA**: for each individual the potential outcomes Y_{a_1} and Y_{a_1} do not depend on treatments received by other individuals (i.e., there is no interference)
- **Positivity**: all conditional probabilities of treatment are greater than zero
- These are standard assumptions in the potential outcomes (or counterfactual) framework

DAG Focusing on Back-Door Adjustment

- We can specify a basic causal structure as a directed acyclic graph (DAG):

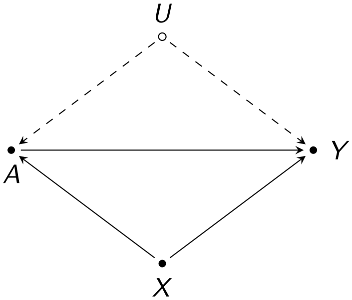

- The true causal effect of A on Y requires conditioning on X and U (that is, $Y_a \perp\!\!\!\perp A | X, U$)

Incorporating a Mechanism

- We can specify a DAG with a mechanism:

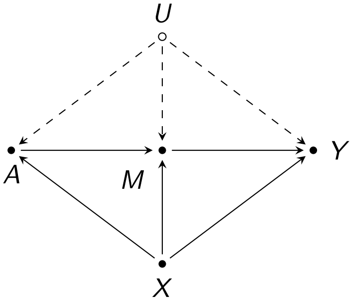

- Now we can estimate causal effects using back-door *or* front-door adjustment, but either way we have bias due to U!

Examples of Mechanisms in Sociology

Mechanisms in Sociology

- Mechanism-based models are common in sociology
- Popularized by path-breaking (pun) work by Hubert Blalock, Herbert Simon, and Otis Dudley Duncan
- Commonly modeled as linear structural equation models (LSEM)
- Approaches did not use a formal causal calculus and identification assumptions have not always been enunciated clearly

Example: Weberian Theory

Figure 1: Weber's Theory of Capitalism

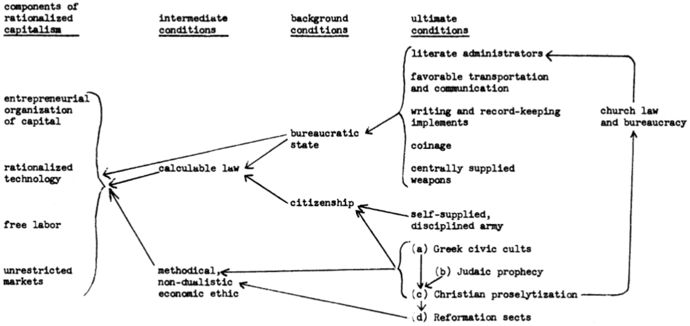

Figure 1. The Weberian Causal Chain

Example: Extension of Adorno's Theory

Figure 2: Right-Wing Authoritarianism

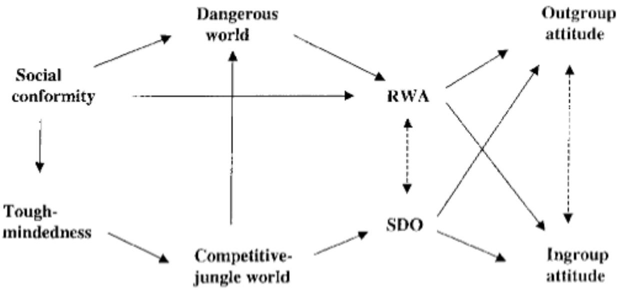

Figure 1. A causal model of the impact of personality and social world view on the two ideological attitude dimensions of right-wing authoritarianism (RWA) and social dominance orientation (SDO) and their impact on intergroup attitudes. The dashed lines indicate context-dependent causal paths.

Example: Classic Status Attainment Model

Figure 3: Blau and Duncan's Status Attainment Model

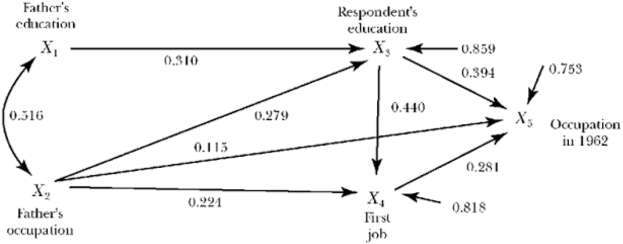

Refinement of Status Attainment Model

Figure 4: Wisconsin Model of Status Attainment

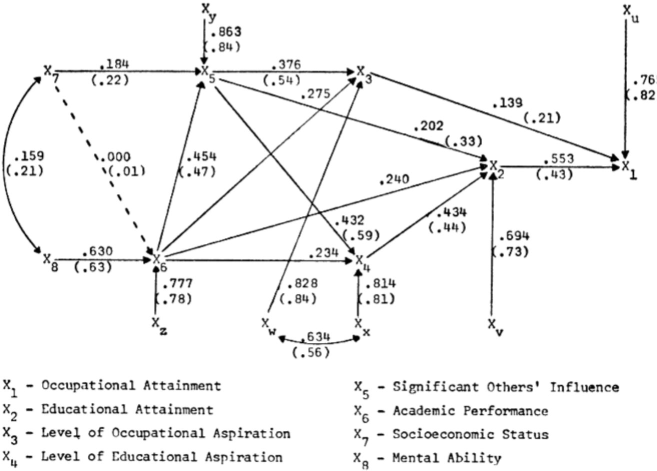

Example: Age-Period-Cohort Models

Figure 5: Young Generation (YG) Outcome

Which statement do you agree more with?

The younger generation should be taught by their elders to do what is right.

The younger generation should be taught to think for themselves even though they may do something their elders disapprove of.

Example: Age-Period-Cohort Models (cont.)

Figure 6: Duncan's Age-Period-Cohort Model

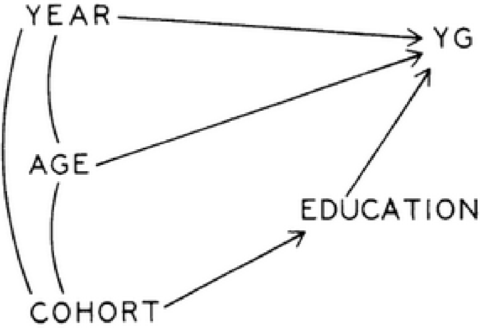

Big Problem

- Initial work on mechanisms in sociology did not formalize provide a formal mathematical calculus for identifying causal relationships
- Judea Pearl and colleagues developed a formal approach starting in the 1980s
- Two main adjustment strategies: back-door criterion and front-door criterion
- Both are prone to failure in the presence of unobserved common-cause confounders

Bias Formulas and the Front-Door Criterion

Three Main Kinds of Average Causal Effects

- The **Average Treatment Effect** (ATE):

$$E[Y_{a_1}] - E[Y_{a_0}] \tag{1}$$

- The **Average Treatment Effect for the Treated** (ATT):

$$E[Y_{a_1}|a_1] - E[Y_{a_0}|a_1] \tag{2}$$

- The **Average Treatment Effect for the Untreated** (ATU):

$$E[Y_{a_1}|a_0] - E[Y_{a_0}|a_0] \tag{3}$$

- In general, sensitivity analyses for the ATT (and ATU) require fewer assumptions about U than for the ATE
- Glynn and Kashin focus on the ATT

Bias for the Front-Door Approach for ATT

- The ATT is: $E[Y_{a_1}|a_1] - E[Y_{a_0}|a_1] = \mu_{1|a_1} - \mu_{0|a_1}$
- They assume consistency such that: $E[Y_{a_1}|a_1] = E[Y|a_1]$
- Thus, their quantity of interest is:

$$\tau_{ATT} = E[Y|a_1] - E[Y_{a_0}|a_1] = \mu_{1|a_1} - \mu_{0|a_1} \tag{4}$$

The True ATT

- They assume that $\mu_{0|a_1}$ is identifiable by conditioning on a set of observed covariates X and unobserved covariates U
- For simplicity they assume discrete variables such that

$$\mu_{0|a_1} = \sum_x \sum_u E[Y|a_0, x, u] P(u|a_1, x) P(x|a_1) \tag{5}$$

- This formula is just saying that, after adjusting for U and X, we can obtain an estimate for $\mu_{0|a_1}$
- Of course, we don't know the true ATT since we don't observe U

Front-Door Adjustment

- However, by using information on the mechanism M we can obtain an estimate using the front-door criterion
- Front-door adjustment for a set of measured post-treatment variables M is:

$$\mu^{fd}_{0|a_1} = \sum_x \sum_m P(m|a_0, x)E[Y|a_1, m, x]P(x|a_1) \tag{6}$$

- Thus the front-door estimator is:

$$\tau^{fd}_{ATT} = \mu_{1|a_1} - \mu^{fd}_{0|a_1} \tag{7}$$

Bias in the Front-Door Estimator

- The front-door estimator is biased because it doesn't adjust for U
- The bias in the front-door estimator can be expressed as:

$$B^{fd}_{ATT} =$$
$$\sum_x P(x|a_1) \sum_m \sum_u P(m|a_0, x, u) E[Y, a_0, m, x, u] P(u|a_1, x) -$$
$$\sum_x P(x|a_1) \sum_m \sum_u P(m|a_0, x) E[Y, a_1, m, x, u] P(u|a_1, m, x)$$

- Zero front-door bias when: (1) $U \perp\!\!\!\perp M|(a_0, x)$, (2) $U \perp\!\!\!\perp M|(a_1, x)$, (3) Y is mean independent of A conditional on U,M,X (a weaker assumption that conditional independence itself)
- If U is affecting M or A has a direct effect on Y, then our front-door estimator is biased

A Few Problems

- This formula is very general, and requires a lot of information about U
- We need to make simplifying assumptions
- Besides focusing on the ATT, Glynn and Kashin do two things:

1. Assume one-side compliance
2. Use simplifying assumptions of VanderWeele and Arah (2011): (a) relationships don't vary across strata of X and (b) U is binary

Simplifying the Sensitivity Analysis

Nonrandomized Program Evaluations with One-Sided Compliance

- Consider an ATT where A is treatment assigned (e.g., "Take this pill!") and M is treatment received (e.g., "I actually take the pill")
- One-sided compliance assumption: $P(M = 0|a_0, x) = P(M = 0, a_0, x, u) = 1$ for all values of X and U
- This implies that only people assigned to treatment (i.e., $A = 1$) can receive treatment (i.e., $M = 1$)
- This helps us simplify the front-door estimator

Front- and Back-Door Estimators Under One-Sided Noncompliance

- A is program sign-up, M is program participation
- Front-door adjustment:

$$E[Y|a_1] - \sum_x E[Y|a_1, m_0, x]P(x|a_1) \tag{8}$$

- Back-door adjustment:

$$E[Y|a_1] - \sum_x E[Y|a_0, x]P(x|a_1) \tag{9}$$

- Treated non-compliers (rebels): $E[Y|a_1, m_0, x]$
- Controls: $E[Y|a_0, x]$

Front- and Back-Door Estimators Under One-Sided Noncompliance

- Back-door estimates match individuals assigned to treatment ($A = 1$) to individuals assigned control ($A = 0$)
- Front-door estimates match individuals assigned and who received treatment ($A = 1$ and $M = 1$) to individuals assigned treatment but who did not receive treatment ($A = 1$ but $M = 0$)
- Assumes, conditional on covariates, that non-compliance was assigned as "if it were random"
- But aren't $A = 1, M = 0$ people fundamentally different (rebels without a cause)?

Further Simplifying Back-Door Bias

- If we assume (a) relationships don't vary across strata of X and (b) U is binary, then the back-door bias formula becomes:
- (Back-Door Bias) = (Direct "effect" of U) $\times$ (Back-door imbalance)

$$B^{bd}_{ATT} = (E[Y|U=1, a_0, x] - E[Y, U=0, a_0, x])\times$$
$$(P(U=1|a_1, x) - P(U=1|a_0, x))$$

Further Simplifying Front-Door Bias

- If we assume (a) relationships don't vary across strata of X and (b) U is binary, then the front-door bias formula becomes:
- (Front-Door Bias) = (Direct "effect" of U) $\times$ (Front-door imbalance) $-$ (Direct "effect" of A)

$$
\begin{aligned}
B^{fd}_{ATT} = (E[Y|U = 1, a_0, x] - E[Y, U = 0, a_0, x])\times \\
(P(U = 1|a_1, x) - P(U = 1|a_1, x, m_0)) \\
\sum_u P(u|a_1, m_0, x)(E[Y|u, a_1, m_0, x] - E[Y|u, a_0, m_0, x])
\end{aligned}
$$

Interesting Cases for Sensitivity Analysis

- (Back-Door Bias) $=$ (Direct "effect" of U) $\times$ (Back-door imbalance)
- (Front-Door Bias) $=$ (Direct "effect" of U) $\times$ (Front-door imbalance) $-$ (Direct "effect" of A)
- What if the direct effect of A on Y is zero, and the absolute value of the front-door imbalance is smaller than the absolute value of the back-door imbalance?
- This implies we would prefer the front-door approach

Example: National JTPA Study

National JTPA Study

- Job training evaluation program with both experimental data and nonexperimental comparison group
- Nonexperimental group different from experimental controls, particularly on labor force participation and earnings histories (Heckman et al., 1997, 1998; Heckman and Smith, 1999)
- Measure program sign-up impact as ATT on earnings post-randomization
- One-sided noncompliance: people who didn't sign-up not allowed to receive JTPA services and some sign-ups drop out

Results of JTPA

Figure 7: Results for Males

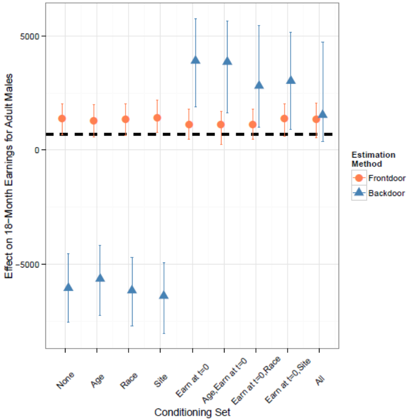

Sensitivity Analyses

- Suppose diligence is an unobservable U, with a positive relationship on showing up (M) and earnings (Y)
- Assumption is that the treated non-compliers ($A = 1$ but $M = 0$) are more diligent and have higher incomes than those with just $A = 0$
- Direct effect of U on Y will dominate the direct effect of A on Y
- Results show front-door estimates are preferable!

Conclusions

Conclusions and Future Directions

- Mechanisms are common in sociology and have a long history in the field
- Front-door estimators can provide causal identification with observational data, but require causal sufficiency
- Sensitivity analyses are vital for these estimators
- Need more research comparing back-door with front-door estimators
- Unclear how strong the assumption of a binary U really is, but we need simplifications on the unobservables for tractability

Front-Door Adjustment with Unmeasured Confounding

Thank you