Causality: Approaches and Pitfalls

Emily Oster

February 26, 2016

The Problem

- Consider an outcome Y, treatment T, covariate vector X
  - Example: job training on wages, low fat diet on heart disease, etc.
  - y_{i1}: outcome for person i if they get $T = 1$
  - y_{i0} : outcome for person i if they get $T = 0$
- What is the impact of T on Y?
- Object of Interest

$$E(y_{i1} - y_{i0})$$

The Problem

- Observe y_{i1} or y_{i0} for a given person
  - You cannot see the same person at the same moment with and without the treatment
- What can you see?

$$\bar{y}_1 - \bar{y}_0$$

- Wht is the problem?
  - People with $T = 1$ may be systematically different from those with $T = 0$
- How can we figure out the causal impact of T?

Example and Outline

- Does eating a low fat diet lower your risk of mortality?
  - Y: Did you die?
  - T: Did you eat a low fat diet?
  - X: Other stuff about you
- Gold Standard
- Options with no special data requirements
- Options with special data requirements

Gold Standard: RCT

- Randomized Controlled Trial
- How to do?
  - Recruit people
  - Randomly tell half of them to eat a low fat diet
  - Follow over time, observe mortality
  - Compare mortality in treatment, control groups
- Why does this work?
- Downsides: expensive, hard to do.

Options with Common Data

- Example: NHANES Data
  - Dietary information
  - Good demographic data (education, income, etc)
  - Other health data
  - Can link to later mortality
- Selection on observables
- Selection on unobservables

Selection on Observables

- Consider:

$$E(y_{i1}|T=1,X_i) - E(y_{i0}|T=0,X_i)$$

- i.e. a regression of Y on T which controls for observed covariates X
- Uncovers causal effect if T is random conditional on X
- In example: T diet, X education, income, race, age, gender.
- Pitfalls?

Selection on Unobservables

- There are controls we are important in determining treatment which we do not observe
  - Without further assumptions, can say nothing
- Assume relationship between X and T is informative about relationship between T and unobserved variables W
- May be able to bound effect size.

Comment

- These techniques can work with limited data requirements
- It is easy to look for heterogeneity across people
- Very transparent
- May not (probably does not) generate causal effect.
- Very difficult to figure out how far you are from causal.

“Special” Data Techniques

- Imagine some variable Z which (randomly) pushes some people into treatment.
- Imagine isolating variation in treatment which is due to variation in Z
- Use that part of the variation to estimate the impact of T
- Examples:
    - Regression discontinuity
    - Instrumental Variables
    - Propensity Score

Regression Discontinuity

- (Fake) Example.
- Doctors adopt new rule: BMI over 25, told to go on low-fat diet. BMI under 25, not told.
- Two individuals
  - BMI 24.9: Not told to go on low-fat diet
  - BMI 25.1: Told to go on low fat diet
- Otherwise similar (before/after breakfast)
- “Random” determination of diet advice
- If anyone listens to diet advice, can use this
  - Graph

Regression Discontinuity: Minuses

- Effects are specific to the characteristics off of which you estimate them
  - What if low fat diet matters more if you are at a BMI of 40?
- Very strong data requirements
- Must be sure other things do not vary across the threshold
  - Why did they put the rule there in the first place?

Instrumental Variables

- (Fake) Example.
- Religion Z requires being a vegetarian; on average, vegetarians eat less fat
  - Key assumption ("exclusion restriction"): religion does not otherwise affect mortality
- Relate Z to T , calculate $\hat{T}$
- Relate Y to $\hat{T}$
  - Effect is driven only by the variation in T that is driven by Z.
  - Z only impacts Y through T .

Instrumental Variables: Minuses

- Effects local to people impacted by the instrument
- Stringent data requirements
    - Exclusion restriction often implausible.

Propensity Score

- I know full set of variables that determine T. Actual T contains some random-ness.
- I do not know how variables enter to determine T
  - polynomials, interactions, etc
- Entering all variables in regression over fits
- Generate $Pr(T)$ by regressing T on all variables and combinations and predicting
  - Note on LASSO
- Control for $Pr(T)$, T.
- Pitfalls

Final Thoughts

- Causality is hard to show.
- Tension between strength of causality and locality of effect.
- I am in some ways more bullish than most on the first set of things,
    - Especially in the era of big data.