Technical Track
Session I:
Causal Inference

Damien de Walque
Amman, Jordan
March 8-12, 2009

Motivation

- The research questions that motivate most studies in the health sciences are causal in nature.
- For example:
  - What is the efficacy of a given drug a given population?
  - What fraction of deaths from a given disease could have been avoided by a given treatment or policy?

Motivation

- The most challenging empirical questions in economics also involve causal-effect relationships:
  - Does school decentralization improve schools quality?
  - Does one more year of education *causes* higher income?
  - Does conditional cash transfers *causes* better health outcomes in children?

Motivation

Interest in these questions is motivated by:

- Policy concerns
  - Do public programs reduce poverty?
- Theoretical considerations
- Problems facing individual decision makers

Intuition of the problem: a hypothetical example

- A conditional cash transfer (CCT) program for pregnant women
  - They will be under periodic medical checkouts
  - The program focuses on better nutrition
  - Periodic meetings about risks of smoking and drinking

- Key question: does the program have any impact on children's birth weight?

Cash transfers (CCTs) and birth weight

- Suppose your data are: women in a town with CCTs and women in a nearby town without CCTs

Y = Average birth weight in grams				
	Women in town with CCT	Women in nearby town without CCT	Difference	Double Difference
Before the CCTs started	?	?		
Once the CCTs started	3,250	3,100	150	?

Cash transfers (CCTs) and birth weight

- Suppose your data are: women in a town with the program and women in a nearby town (without the program)

Y = Average birth weight in grams				
	Women in town with CCTs	Women in nearby town without CCTs	Difference	Double Difference
Before the program started	3,025	2,840	185	
Once the program started	3,250	3,100	150	-35

Cash transfers (CCTs) and birth weight

- Suppose you have data on women who got assigned to the program randomly

Y = average birth weight in grams				
	Women in town with CCTs	Women in nearby town without CCTs	Difference	Double Difference
Before the program started	3,025	2,840	185	
Once the program started	3,250	3,100	150	-35
	Women in town with CCTs	Women in nearby town BY RANDOMIZATION	Difference	Double Difference
Before the program started	3,028	3,028	0	
Once the program started	3,250	3,105	145	145

Standard Statistical Analysis

- **Tools:** likelihood and other estimation techniques
- **Aim**: to infer parameters of a distribution from samples drawn of that distribution.
- **Uses:** With the help of such parameters, one can:
  - Infer association among variables,
  - Estimate the likelihood of past and future events,
  - Update the likelihood of events in light of new evidence or new measurement.
- **Condition** for this to work well: experimental conditions must remain the same.

Causal Analysis

- Causal Analysis goes one step further than Standard Statistical Analysis:
- Its aim is to infer aspects of the data generation process.
- With the help of such aspects, one can deduce
  - the likelihood of events under *static conditions*, (as in Standard Statistical Analysis)
  - and also the dynamics of events under *changing conditions.*

Causal Analysis

- "*dynamics of events under changing conditions*" includes :

1. Predicting the effects of interventions
2. Predicting the effects of spontaneous changes
3. Identifying causes of reported events

Causation versus Correlation

- Standard statistical analysis/ probability theory:
  - The word "*cause*" is not in its vocabulary
  - Allows us to say is that two events are mutually *correlated*, or dependent
    - if we find one, we can expect to find the other
- This is not enough for policy makers
  - They look rationales for policy decisions: if we do XXX, then will we get YYY?
  - Hence we must supplement the language of probability with a vocabulary for *causality*.

The Rubin Causal Model

- Define the population by U.
  Each unit in U is denoted by u.

- For each $u \in U$, there is an associated value $Y(u)$ of the variable of interest Y, which we call: the response variable.

- Rubin **takes** the position that **causes** are only those things that could be treatments in hypothetical experiments.

- For simplicity, we assume that there are just two causes or level of treatment.
- Let *D* be a variable that indicates the cause to which each unit in *U* is exposed:

$$D = \begin{cases} 1 & \text{if unit u is exposed to treatment} \\ 0 & \text{if unit u is exposed to control} \end{cases}$$

- In a controlled study, *D* is constructed by the experimenter.
- In an uncontrolled study, *D* is determined by factors beyond the experimenter's control.

- The response Y is potentially affected by whether u receives treatment or not.

- Thus, we need two response variables:

  $Y_1(u)$ is the outcome if unit u is exposed to treatment

  $Y_0(u)$ is the outcome if unit u is exposed to control

$$D = \begin{cases} 1 & \text{if unit u is exposed to treatment} \\ 0 & \text{if unit u is exposed to control} \end{cases}$$

$Y_1(u)$ is the outcome if unit u is exposed to treatment

$Y_0(u)$ is the outcome if unit u is exposed to control

↓

Then, the outcome of each unit u can be written as:

$$Y(u) = DY_1(u) + (1 - D)Y_0(u)$$

Note: This definition assumes that the treatment status of one unit does not affect the potential outcomes of other units.

Definition: For every unit u, treatment causes the effect

$$\delta_u = Y_1(u) - Y_0(u)$$

Fundamental Problem of Causal Inference:

For a given u, we observe either $Y_1(u)$ OR $Y_0(u)$

We cannot observe the value of $Y_1(u)$ and $Y_0(u)$ on the same unit u

$\Rightarrow$ it is impossible to observe the effect of treatment on u by itself.

Issue: We do not have the counterfactual evidence for u

i.e. what would have happened to u in the absence of treatment.

- Given that the treatment effect for a single unit *u* cannot be observed, we aim to identify the *average treatment effect* for the population *U* (or for sub-populations).
- The *average treatment effect ATE* over *U* (or sub-populations of *U*):

$$TE_u = \delta_u = Y_1(u) - Y_0(u)$$

$$\Downarrow$$

$$\begin{aligned} ATE_U &= E_U\ [Y_1(u) - Y_0(u)] \\ &= E_U[Y_1(u)] - E_U[Y_0(u)] \\ &= \overline{Y}_1 - \overline{Y}_0 \\ &= \overline{\delta} \qquad (1) \end{aligned}$$

- The statistical solution replaces the impossible-to-observe treatment effect of *t* on a specific unit *u* with the possible-to-estimate *average* treatment effect of *t* over a population *U* of such units.

- Although $E_U(Y_1)$ and $E_U(Y_0)$ cannot both be calculated, they can be estimated.

- Most econometrics methods attempt to construct from observational data consistent estimates of

$$E_U(Y_1) = \bar{Y}_1 \qquad \text{and} \qquad E_U(Y_0) = \bar{Y}_0$$

So we are trying to estimate:

$$ATE_U = E_U[Y_1(u)] - E_U[Y_0(u)]$$
$$= \bar{Y}_1 - \bar{Y}_0 \qquad (1)$$

Consider the following simple estimator of ATE_U:

$$\hat{\bar{\delta}} = [\hat{\bar{Y}}_1 \mid D = 1] - [\hat{\bar{Y}}_0 \mid D = 0] \qquad (2)$$

- equation (1) is defined for the whole population,
- equation (2) is an estimator to be evaluated on a sample drawn from that population

Lemma: If we assume that

$$[\bar{Y}_1 \mid D = 1] = [\bar{Y}_1 \mid D = 0 \quad]$$

and $[\bar{Y}_0 \mid D = 1] = [\bar{Y}_0 \mid D = 0$

then

$$\hat{\bar{\delta}} = [\hat{\bar{Y}}_1 \mid D = 1] - [\hat{\bar{Y}}_0 \mid D = 0]$$

is a consistent estimator of

$$\bar{\delta} = \bar{Y}_1 - \bar{Y}_0$$

- Thus, a sufficient condition for the simple estimator to consistently estimate the true *ATE* is that:

$$[\overline{Y}_1 \mid D = 1] = [\overline{Y}_1 \mid D = 0 \quad]$$

and

$$[\overline{Y}_0 \mid D = 1] = [\overline{Y}_0 \mid D = 0$$

- The average outcome under treatment $\overline{Y}_1$ is the same for the treatment (D=1) and the control (D=0) groups

- The average outcome under control $\overline{Y}_0$ is the same for the treatment (D=1) and the control (D=0) groups

When will those conditions be satisfied?

- It is sufficient that treatment assignment D be uncorrelated with the potential outcome distributions of Y_0 and Y_1.
  - Intuitively: there can be no correlation between
    - Whether someone gets the treatment
    - How much that person potentially benefits from the treatment

- The easiest way to achieve this uncorrelatedness is through random assignment of treatment.

Another way of looking at it

- After a bit of algebra, it can be shown that:

$$\underbrace{\hat{\bar{\delta}}}_{\substack{\textit{simple}\\ \textit{estimator}}} = \underbrace{\bar{\delta}}_{\substack{\textit{true}\\ \textit{impact}}} + \underbrace{\left([\bar{Y}_0 \mid D=1] - [\bar{Y}_0 \mid D=0\right)}_{\textit{Baseline Difference}}$$

$$+ (1-\pi)\underbrace{\left(\bar{\delta}_{\{D=1\}} - \bar{\delta}_{\{D=0\}}\right)}_{\textit{Treatment Heterogeneity}}$$

Another way of looking at it (in words)

- There are two sources of biases that need to be eliminated from estimates of causal effects from observational studies.
  - Baseline difference. (selection bias)
  - Treatment Heterogeneity.
- Most of the methods available only deal with selection bias

Treatment on the Treated

- *ATE* is not always the parameter of interest.

- Often, it is the average treatment effect *for the treated* that is of substantive interest:

$$TOT = E\left[Y_1(u) - Y_0(u) \mid D = 1\right]$$
$$= E\left[Y_1(u) \mid D = 1\right] - E\left[Y_0(u) \mid D = 1\right]$$

Treatment on the Treated

If we need to estimate TOT

$$TOT = E[Y_1(u) \mid D = 1] - E[Y_0(u) \mid D = 1]$$

Then the simple estimator (2)

$$\hat{\delta} = [\hat{\bar{Y}}_1 \mid D = 1] - [\hat{\bar{Y}}_0 \mid D = 0]$$

consistently estimates *TOT* if:

$$[\bar{Y}_0 \mid D = 1] = [\bar{Y}_0 \mid D = 0$$

"No baseline difference between the treated and control groups"

References

- Judea Pearl (2000): Causality: Models, Reasoning and Inference, Cambridge University press. (Book) Chapters 1, 5 and 7.
- Trygve Haavelmo (1944): “The probability approach in econometrics,” *Econometrica* 12, pp. iii-vi+1-115.
- Arthur Goldberger (1972): “Structural Equations Methods in the Social Sciences,” *Econometrica 40*, pp. 979-1002.
- Donald B. Rubin (1974): “Estimating causal effects of treatments in randomized and nonrandomized experiments,” *Journal of Educational Psychology 66*, pp. 688-701.
- Paul W. Holland (1986): “Statistics and Causal Inference,” *Journal of the American Statistical Association 81*, pp. 945-70, with discussion.