Applied Microeconometrics I

Lecture 2: Randomized Controlled Trials

Stefano Lombardi

Aalto University

September 7, 2023

Lecture Slides

What did we do last time?

- Economics is increasingly an empirical discipline
- Aim of the course: Core tools for causal inference
- Three types of research questions:
  1. Descriptive
  2. Forecasting
  3. Causal
- What distinguishes causal questions from descriptive/forecasting questions?
- Today's topic: RCT's

Today

- The first problem set is now available at MyCourses. Due on 15/9 at 23.59
- Software tutorials: Stata this Wednesday, R next Wednesday
- Today:
  - Randomized trials
- Recommended reading:
  - Introduction Angrist & Pischke
  - Holland (1986) (parts discussed in class)

How can we estimate a causal relationship?

- Two important (well-known) points:
  1. Correlation does not imply causation
  2. y can cause x even if x takes place before y

Correlation does not imply causation

- $\text{Cor}(x,y) \neq 0$
  1. x implies y
  2. y implies x
  3. z implies x and y
- Some examples:
  - People that sleep less tend to live longer
  - Countries that eat more chocolate receive more Nobel prizes
    - Messerli (2012)

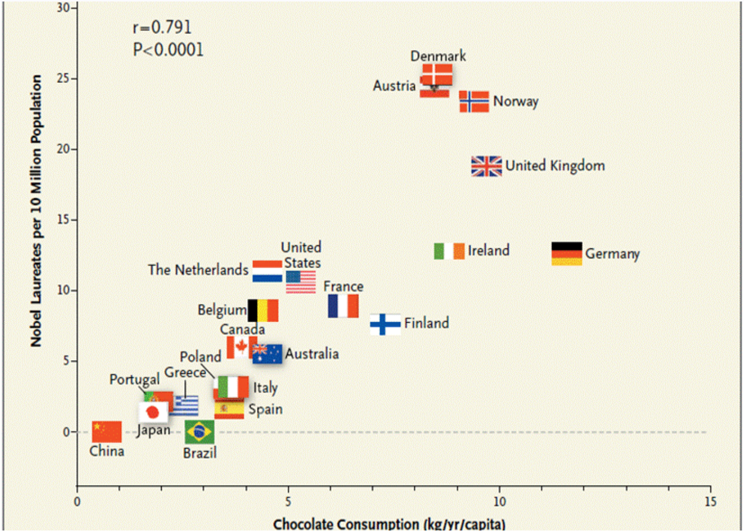
r=0.791
P<0.0001
Nobel Laureates per 10 Million Population
Chocolate Consumption (kg/yr/capita)
Denmark
Austria
Norway
United Kingdom
Ireland
Germany
United States
The Netherlands
France
Belgium
Canada
Finland
Australia
Poland
Greece
Portugal
Italy
Spain
Japan
China
Brazil
30
25
20
15
10
5
0
0
5
10
15

Correlation does not imply causation

- Other examples
  - GDP growth and public debt
  - Students living in households where there are more books tend to perform better in PISA evaluations
  - Women in boards and firm performance
  - SES and health
  - GDP growth and R&D expenditure
  - Marketing/R&D expenditure and firms' profits
- However, often it's just impossible to find meaningful reverse causality/spurious correlation interpretations

Correlation does not imply causation

$$r = 0.789$$

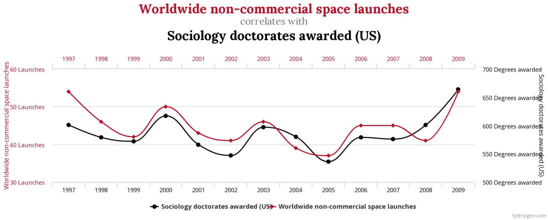

$$r = 0.666$$

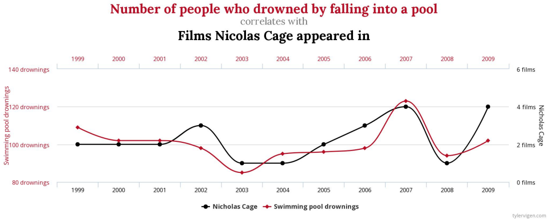

y can cause x even if x takes place before y

The fact that x temporally occurs before y is a necessary but not sufficient condition to talk about $x \longrightarrow y$ causal effects.

Examples of anticipatory effects:

1. **Rain**
   - When many people carry their umbrellas in the morning usually it rains in the afternoon
   - Not a bad idea to use this fact to predict rain...
   - ... However, subsidizing umbrellas is not a great policy to increase rain
2. **Stock exchange**
   - Stock market usually goes down just before an extremist party wins an election

How can we estimate a causal relationship?

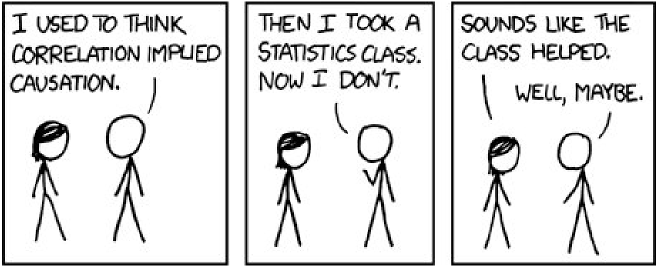

How can we estimate a causal relationship?

Suppose we are interested in the causal effect of getting university degree (treatment) on lifetime income (outcome).

- Let us define the *individual-level treatment effect* as the difference in the individual-level outcome when taking the course vs. not taking the course.
- **Fundamental problem of causal inference**: For a person, we never observe the outcome in both states of the world.
- Statisticians have been very skeptical about studying causal effects and instead (mostly) focus on *associational inference*.

Two approaches for *causal inference* (Holland, 1986): the **scientific solution** and **statistical solution**

How can we estimate a causal relationship?

Scientific solution:

- In the physical sciences, one can answer causal questions by running a 'scientific' experiment in a controlled environment.
- Galileo's Leaning Tower of Pisa (thought?) experiment:
  - Aristotle's theory of gravity: objects fall at speed relative to their mass
  - Galileo's hypothesis: time of descent is independent of object mass (instead it depends on the volume)
  - Throw two balls of same size but different weight from tower

Leaning Tower of Pisa

The scientific solution

- Assumptions:
  1. **Temporal stability**: the response does not change if the time when a treatment is applied is varied slightly (constancy of response over t).
  2. **Causal transience**: the response to one treatment is not affected by prior exposure of i to another treatment status
  3. **Unit homogeneity**: units are homogeneous with respect to the assigned treatment (it doesn't matter *which* unit is treated).
- Under 1. and 2., take one unit i and:
  - expose i to "control" status and measure outcome
  - expose same i to "treatment" status and measure outcome
- Under 3., the outcome after a given exposure status is the same regardless of whether i or i' are exposed to it
  - expose i to treatment or control, and i' to the other status.
- In Social Sciences, 1–3 are typically not plausible.

The potential outcomes framework

- Counterfactual outcomes are the language/building blocks of the **statistical solution** used in economics.
- Think of a treatment as a binary random variable $D_i = \{0, 1\}$
- And potential outcomes (counterfactuals):

$$Y_i = \begin{cases} Y_{1i} & \text{if } D_i = 1 \\ Y_{0i} & \text{if } D_i = 0 \end{cases}$$

- The causal effect of treatment for an individual i is $Y_{1i} - Y_{0i}$
- Unfortunately, for each i, we only observe:

$$Y_i = D_i Y_{1i} + (1 - D_i) Y_{0i} = Y_{0i} + (Y_{1i} - Y_{0i}) D_i$$

The Fundamental Problem of Causal Inference

It is impossible to observe for the same individual i the values $D_i = 1$ and $D_i = 0$ as well as the values Y_{i1} and Y_{i0} and, therefore, it is impossible to observe $Y_{1i} - Y_{0i}$

(Holland, 1986)

Observed difference in outcomes vs. treatment effect

- We would want to know the **average treatment effect on the treated** (ATET), $E[Y_{1i} - Y_{0i} | D_i = 1]$
- A naive comparison of averages does not tell us this:

$$\begin{aligned} E[Y_i|D_i = 1] - E[Y_i|D_i = 0] &= E[Y_{i1}|D_i = 1] - E[Y_{i0}|D_i = 0] \\ &= \underbrace{E[Y_{1i} - Y_{0i}|D_i = 1]}_{\text{Average treatment effect on the treated}} \\ &+ \underbrace{E[Y_{0i}|D_i = 1] - E[Y_{0i}|D_i = 0]}_{\text{Selection bias}} \end{aligned}$$

- Comparison of observable averages is contaminated by selection bias
- What is the direction of the bias?

Observed difference in outcomes vs. treatment effect

- The difference in the observed outcomes (between the treated and the non treated) is equal to the average treatment effect on the treated if there is no selection bias:
- If $E[Y_{0i}|D_i = 1] = E[Y_{0i}|D_i = 0]$
  then $E[Y_i|D_i = 1] - E[Y_i|D_i = 0] = E[Y_{1i} - Y_{0i}|D_i = 1]$
- What does this mean? Treated and the non-treated do not differ in their counterfactual outcomes (no differential selection into treatment based on counterfactual outcomes)
  - this is much to ask in observational data. Why?

Example

EXAMPLE: The effect of taking a course on the final thesis grade

		Osku	Mia
Potential grade without the course	Y_{oi}	**3**	**5**
Potential grade with the course	Y_{1i}	**4**	**4**
Treatment (took the course)	D_i	**1**	**0**
Realized thesis grade	Y_i	**4**	**5**
Treatment effect	$Y_{1i} - Y_{0i}$	**1**	**-1**

What is the observed difference in the realized grades of the treated and non-treated?

What is the effect of treatment on the treated?

What is the average treatment effect?

Note that

- Average treatment effect on the treated: $E[Y_{1i} - Y_{0i} | D_i = 1]$
- Different from the population average treatment effect: $E[Y_{1i} - Y_{0i}]$
- In general, they are not the same, and the distinction is often important (also for policy makers!)
- They will be the same only if the average treatment effect on the treated and on the non treated are the same:

  $E[Y_{1i}|D_i = 1] - E[Y_{0i}|D_i = 1] = E[Y_{1i}|D_i = 0] - E[Y_{0i}|D_i = 0]$
  (i.e.: the effect of treatment is the same for everybody)

Random assignment

- We want to understand what would have happened to the treated in the absence of treatment and thus overcome the selection problem...
- Solution $\rightarrow$ Random assignment
- Randomization is the **statistical solution** to the fundamental problem of causal inference
- Let us consider randomization in the case of an RCT

Random assignment

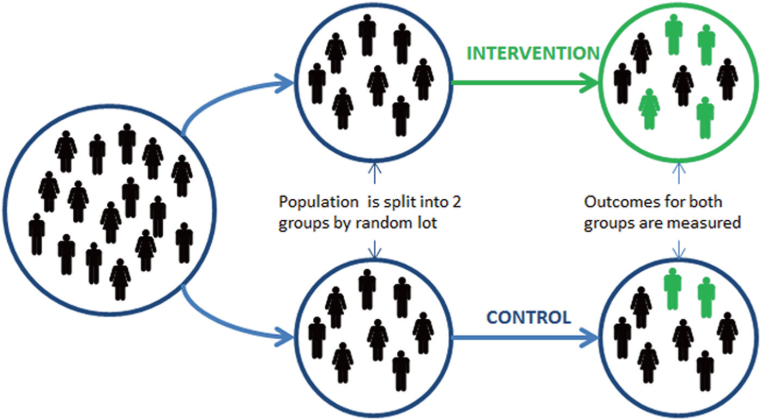

Random assignment

- Random assignment solves the selection problem since it makes D_i independent of potential outcomes, hence:

$$E[Y_{1i}|D_i = 1] = E[Y_{1i}|D_i = 0] = E[Y_{1i}]$$
$$E[Y_{0i}|D_i = 1] = E[Y_{0i}|D_i = 0] = E[Y_{0i}]$$

- Therefore, a comparison of the treatment and control group outcomes provides information about the causal impact of the average treatment effect on the treated (ATET):

$$\begin{aligned} E[Y_i|D_i = 1] - E[Y_i|D_i = 0] &= E[Y_{i1}|D_i = 1] - E[Y_{i0}|D_i = 0] \\ &= E[Y_{1i}|D_i = 1] - E[Y_{0i}|D_i = 1] \\ &= E[Y_{1i} - Y_{0i}|D_i = 1] \end{aligned}$$

Random assignment

- and the average treatment effect on the non-treated:

$$
\begin{aligned}
E[Y_{1i}|D_i = 1] - E[Y_{0i}|D_i = 0] &= E[Y_{1i}|D_i = 0] - E[Y_{0i}|D_i = 0] \\
&= E[Y_{1i} - Y_{0i}|D_i = 0]
\end{aligned}
$$

- and the population average treatment effect (ATE):
$E[Y_{1i} - Y_{0i}|D_i = 1] = E[Y_{1i} - Y_{0i}|D_i = 0] = E[Y_{1i} - Y_{0i}]$

- **Note**: With other empirical methods very often the ATET will not be equal to the ATE.

Let us see some of what we just discussed in a STATA example.

Difference between means and selection bias

$$
\begin{aligned}
E[Y_i|D_i=1]-E[Y_i|D_i=0] &= E[Y_{i1}|D_i=1]-E[Y_{i0}|D_i=0] \\
&= E[Y_{i1}|D_i=1]-E[Y_{i0}|D_i=1] \\
&+ E[Y_{i0}|D_i=1]-E[Y_{i0}|D_i=0] \\
&= \underbrace{E[Y_{1i}-Y_{0i}|D_i=1]}_{\text{Average effect of the treatment on treated}} \\
&+ \underbrace{E[Y_{0i}|D_i=1]-E[Y_{0i}|D_i=0]}_{\text{Selection bias}}
\end{aligned}
$$