Potential outcomes

Edicson Luna

University of Pennsylvania

Motivation

Rubin causal model

Randomized controlled trials

Motivation

Motivation

So far, we have been cautious about interpreting results as "an increment of 1 unity in X_i generates an increment of β_1 in the outcome variable" (assuming we are in a case where the two variables are in levels). The reason why I have avoided to use words like "generate" is that they are implicitly assigning some causality (from X_i to Y_i).

In the cases we have our main assumptions, we have decided to use the word "identification". Sometimes, researchers use the word "causal identification".

Motivation

In my experience, I've noticed that the only cases that are usually accepted for a "causal" interpretation are when (i) we work with instrumental variables (or better said, you find a source of exogenous variation), or (ii) we can express our analysis in a "potential outcomes" setup (and the conditions are satisfied).

Today, we will study the most basic potential outcomes setup. This is extended to a lot of different scenarios and you will cover some of them in Econometrics (2) and in Advanced Econometrics (if you decide to take it).

Causality

Before getting into the math, let's think about some causal questions.

- What was the effect of the pandemic on employment?
- What is the effect of an increment of the minimum wage on informality?
- What is the effect of studying on the grade of a test?
- What is the effect of having a baby on the number of hours worked by a woman? by a man?

If you think as an economist, you will notice that the majority of our questions are implicitly trying to caught some level of causality.

Causality vs correlation

Let's think about the following question: What is the effect of an increment in the minimum wage (MW) on the informality rate? Two things may happen

- The increment in the minimum wage is generating that some firms are not able to pay the "legal" MW, and are moving their employees to informality. In this case, it is evident that there is causality.
- The government is trying to "benefit" people with low income and has decided to increase the MW. At the same, the government is incrementing the bureaucracy of a firm when registering formally their employees. In this case, there is just a correlation (spurious). The reason is the government and not the increment in the MW.

Causality vs correlation

A famous experiment with Tobacco was developed in 1950 by Richard Doll and Bradford Hill. During these years there were huge debates about the relationship between tobacco and cancer. The mortality rate due to lung cancer had increased by 4 in only 50 years.

Doll and Hill concluded that the probability that the increase in cancer was "random" (spurious correlation) was 1.5 million to 1. A report made by experts said that the relationship was causal.

Tobacco effect

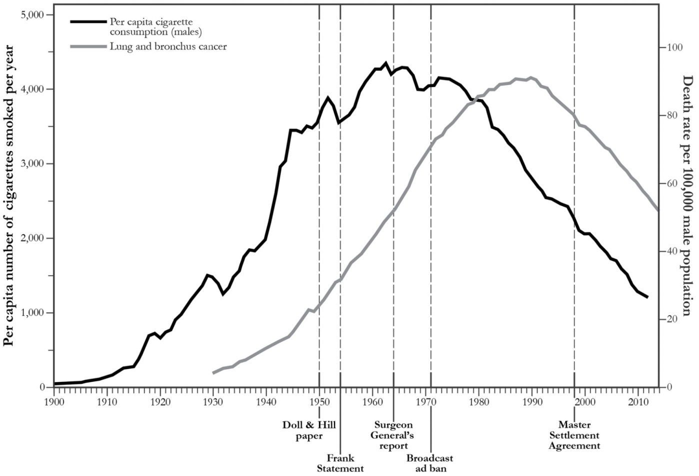

Figure: Tobacco and cancer

Rubin causal model

Counterfactuals

Remember our 4 questions: (i) What was the effect of the pandemic on employment? (ii) What is the effect of an increment of the minimum wage on informality? (iii) What is the effect of studying on the grade of a test? (iv) What is the effect of having a baby on the number of hours worked by a woman? by a man?

We need to think in the counterfactual scenarios

- Unemployment rate w/o pandemic
- Informality rate w/o MW increase
- Grade without studying
- # of hours worked w/o a baby

Potential outcomes

Consider y_i as an observed outcome variable. For simplicity assume y_i is the observed grade in quiz 2.

$$\text{Potential outcomes} = \begin{cases} y_i^1 & \text{if } i \text{ studies} \\ y_i^0 & \text{if } i \text{ doesn't study} \end{cases}$$

The “causal” effect is the difference in potential outcomes

$$\tau_i = y_i^1 - y_i^0$$

What is the fundamental problem here? **We only see one of the states for each person**. The other is the counterfactual (unobserved).

Indicator function

Define D_i as an indicator function

$$D_i = \begin{cases} 1 & \text{if } i \text{ studies} \\ 0 & \text{if } i \text{ doesn't study} \end{cases}$$

What we observe is

$$\begin{aligned} y_i &= y_i^1 D_i + y_i^0 (1 - D_i) \\ &= y_i^0 + (y_i^1 - y_i^0) D_i \\ &= y_i^0 + \tau_i D_i \end{aligned}$$

In the end, if the person does not study we see y_i^0, and if she does we observe $y_i^0 + \tau_i$.

ATE, ATT, ATU

Assume we observe both y_i and D_i. Define

- **Average treatment effect (ATE)**

$$\tau_{ATE} \equiv \mathbb{E}[\tau_i] = \mathbb{E}[y_i^1 - y_i^0]$$

- **Average treatment effect on the treated (ATT)**

$$\begin{aligned} \tau_{ATT} \equiv \mathbb{E}[\tau_i | D_i = 1] &= \mathbb{E}[y_i^1 - y_i^0 | D_i = 1] \\ &= \mathbb{E}[y_i^1 | D_i = 1] - \mathbb{E}[y_i^0 | D_i = 1] \end{aligned}$$

- **Average treatment effect on the untreated (ATU)**

$$\begin{aligned} \tau_{ATU} \equiv \mathbb{E}[\tau_i | D_i = 0] &= \mathbb{E}[y_i^1 - y_i^0 | D_i = 0] \\ &= \mathbb{E}[y_i^1 | D_i = 0] - \mathbb{E}[y_i^0 | D_i = 0] \end{aligned}$$

Selection bias

Let's start by comparing the average among the two groups

$$
\begin{aligned}
\mathbb{E}[y_i|D_i=1]-\mathbb{E}[y_i|D_i=0] &= \mathbb{E}[y_i^1|D_i=1]-\mathbb{E}[y_i^0|D_i=0] \\
&= \mathbb{E}[y_i^1|D_i=1]-\mathbb{E}[y_i^0|D_i=1] \\
&+ \mathbb{E}[y_i^0|D_i=1]-\mathbb{E}[y_i^0|D_i=0] \\
&= \tau_{ATT} + \mathbb{E}[y_i^0|D_i=1]-\mathbb{E}[y_i^0|D_i=0]
\end{aligned}
$$

$\mathbb{E}[y_i^0|D_i=1]-\mathbb{E}[y_i^0|D_i=0]$ is what we call "selection bias". Hence,

$$\text{Mean difference} = \tau_{ATT} + \text{Selection bias}$$

Selection bias

Notice that if we observe the Selection bias, we can obtain the τ_{ATT}. Since we do not observe $\mathbb{E}[y_i^0|D_i = 1]$, this is not possible. This is the main identification challenge.

This implies that we must be careful when we are analyzing data or when we are designing an experiment.

In the end, the problem is solved as long as $\mathbb{E}[y_i^0|D_i = 1] - \mathbb{E}[y_i^0|D_i = 0] = 0$. This is the same as saying that we have a "control" group which acts well as a counterfactual.

Independence

- **Ideal experiment:** the potential outcomes y_i^1 and y_i^0 are independent of D_i. $(y_i^1, y_i^0 \perp D_i)$. In this case, the selection bias is 0.

Think of a counterfactual group for each of the 4 questions made before.

Randomized controlled trials

Definition

- A **randomized controlled trial (RCT)** is a scientific study design in which participants are randomly allocated to either an experimental group receiving the intervention being tested $(D_i = 1)$ or a control group receiving standard treatment or placebo $(D_i = 0)$. This randomization helps eliminate bias, allowing for a more reliable comparison of outcomes between the groups to determine the effectiveness of the intervention.
- For interpretation, assume that D_i is the indicator variable of receiving a program. y_i would be the variable the program is aiming to affect. The math done so far applies the same.

Linear regression

Remember that previously we found $y_i = y_i^0 + \tau_i D_i$. We can rewrite this as

$$\begin{aligned} y_i &= y_i^0 + \tau_i D_i \\ &= E[y_i^0] + \tau_i D_i + y_i^0 - E[y_i^0] \\ &= \beta_0 + \beta_1 D_i + \epsilon_i \end{aligned}$$

Where $E[y_i^0] = \beta_0$, $\tau_i = \beta_1$, and $y_i^0 - E[y_i^0] = \epsilon_i$. We call ϵ_i the idiosyncratic component of each i. This is nothing different from the error term we are used to. We are also assuming the τ_i is similar across i's. In reality, we are just trying to capture an average effect (ATT).

Identification challenge

Notice $\mathbb{E}[y_i|D_i = 0] = \beta_0 + \mathbb{E}[\epsilon_i|D_i = 0]$ and $\mathbb{E}[y_i|D_i = 1] = \beta_0 + \beta_1 + \mathbb{E}[\epsilon_i|D_i = 1]$. Hence,

$$\mathbb{E}[y_i|D_i = 1] - \mathbb{E}[y_i|D_i = 0] = \beta_1 + \mathbb{E}[\epsilon_i|D_i = 1] - \mathbb{E}[\epsilon_i|D_i = 0]$$

Given $y_i^0 - E[y_i^0] = \epsilon_i$, we see

- $\mathbb{E}[\epsilon_i|D_i = 0] = \mathbb{E}[y_i^0|D_i = 0] - \beta_0$
- $\mathbb{E}[\epsilon_i|D_i = 1] = \mathbb{E}[y_i^0|D_i = 1] - \beta_0$

Therefore,

$$\begin{aligned}\mathbb{E}[\epsilon_i|D_i = 1] - \mathbb{E}[\epsilon_i|D_i = 0] &= \mathbb{E}[y_i^0|D_i = 1] - \mathbb{E}[y_i^0|D_i = 0] \\ &= \text{Selection bias}\end{aligned}$$

Identification challenge

In conclusion

$$\mathbb{E}[y_i|D_i = 1] - \mathbb{E}[y_i|D_i = 0] = \beta_1 + \text{Selection bias}$$

The selection bias is the same as having a correlation different than 0 between ϵ_i and D_i. This means there are systematic differences between the group of people who received the program and those who did not.

How can we solve this?

Randomizing

Selection bias is immediately solved if we just randomized the treated and untreated groups. If we have randomization, we get

$$y_i^1, y_i^0 \perp D_i$$

To be very explicit, this means the treatment is the only difference between the two groups.

We also need the Stable Unit Treatment Value Assumption (SUTVA). This ensures no interference among i's and that the outcomes are the expected ones (y_i^0 or y_i^1 depending on the treatment of i).

Estimation

To get an estimator for our parameter of interest τ_{ATT} or β_1, it is enough to

- Run a linear regression
- Take the means difference

As long as we have **randomization** in the group assignment and **SUTVA**, the estimation is trivial.

Limitations

- Not all the types of questions economists want to study can be solved by an RCT
- Internal vs external validity
- RCTs are usually extremely costly
- There may be attrition

If you are interested in these kinds of experiments you may see the work of Abhijit Banerjee and Esther Duflo. If you find boring reading articles, “Poor Economics” is a good book to start with.