


# Econ 5813: Econometrics I

## Lecture 1: The Potential Outcomes Framework



Department of Economics  
University of Colorado, Denver

(Read Chapter 1 in Mastering 'Metrics')



# Introduction to applied econometrics: Questions?

- The tools of econometrics can be used to accomplish many things.  
*(some for good, some for evil)*
- In this class, we will focus on using econometrics to preform causal inference.
- What type of causal inference questions are we interested in?



“What is the effect of a college degree on wages later in life?”



“Does providing free insecticide-treated bed nets to families in Africa reduce childhood malaria?”



“Does increasing the number of police on neighborhood patrol decrease crime?”



These questions boil down to “does X cause Y”, and if so, what is the magnitude of that effect?

# Introduction to applied econometrics

- Causal inference questions start with a cause (education, minimum wage, a new medication, development aid).
- We then ask, “how does the cause affect an outcome of interest (e.g. wages, unemployment, health, civil convict). Does this cause that? **Cause → Effect.**
- What type of questions are not covered in this class?
  - Answer: What caused the effect? **Cause ← Effect.** For example:
    - *What caused World War I?*
    - *Why are poor countries poor?*
    - *What are the determinants of income? We are not forecasting.*
- These are fine questions, but these types of questions are much harder to answer using applied econometrics.
  - Pretty much all events/outcomes have multiple causes.
  - Outcomes have complex proximate and distal (ultimate) causes.

# How to answer questions about a causal relationship?

- We all know the saying, “correlation does not prove causation.”
  - We observe that areas with lots of police on patrol have a lot of crime.
  - Does this mean that police cause crime?
  - Not necessarily. This is likely an example of reverse causation.



I have noticed a strong correlation between my neighbor carrying an umbrella on her way to work in the morning and my lawn being wet when I return home in the evening.

- I am confident that my neighbor’s umbrella does not cause my lawn to become wet (no causation), and that my lawn being wet does not cause my neighbor to carry her umbrella (no reverse causation).
- Instead, a confounder or an omitted variable (rain) creates a correlation between two variables that are not causally related.
- This is very important as a policy matter, because banning umbrellas would be a very misguided policy to prevent my lawn from becoming wet.

# How to answer a question about a causal relationship?

- OK, “correlation does not prove causation.” But is correlation useful at all?
- YES! Correlation is important, and is used in causal inference, just not by itself.
- A more accurate saying would be “causation cannot be inferred from correlation alone,” or more generally, “causation cannot be inferred from sample statistics alone.”
  - An average, a difference in averages, a regression coefficient, etc...
- To perform valid causal inference, we use sample statistics combined with an identification strategy.
  - The identification strategy is what makes us believe the sample statistic is likely revealing the causal effect.
- The gold standard identification strategy is a randomized control trial (RTC).
  - Randomly assign the “cause.”

# Randomized Control Trial (RTC)

- Question: “Does Lipitor lower cholesterol?”
- The sample statistic is the average difference in blood cholesterol levels between patients “treated” with Lipitor and patients “not treated” with Lipitor.
- If all I tell you is that the average blood cholesterol levels among the treated group is higher than the untreated group, then would you conclude that Lipitor causes an increases in cholesterol levels?
  - Answer: NO.
  - People with high blood cholesterol may select into treatment.
  - People with low blood cholesterol may select into control.
  - Treatment and control groups differ systematically along important dimensions.
  - Thus, the difference in blood cholesterol does not reflect a causal effect.

# Randomized Control Trial (RTC)

- Question: “Does Lipitor lower cholesterol?”
- The sample statistic is the average difference in blood cholesterol levels between patients “treated” with Lipitor and patients “not treated” with Lipitor.
- Instead, if I tell you that the treated and control groups were part of a large and carefully designed randomized control trial (RTC) where:
  - Patients were randomly assigned to the treatment and control groups.
  - The RTC is double-blind with perfect compliance and no attrition.
  - Before the treatment, baseline characteristics, including the average blood cholesterol levels, are the same in the treatment and control groups (balanced samples).
  - After the treatment, the average blood cholesterol in the treatment group is significantly lower than in the in the control group.
  - Then you may conclude that Lipitor causes a reduction in cholesterol levels.
- The average difference in blood cholesterol after treatment is the sample statistic.
- Random assignment is the identification strategy.
- The degree of “proof” of causation comes from the strength of the identification strategy, not from the sample statistic (not from the “statistical significance” of a coefficient).

# The Potential Outcomes Framework

What do we mean when we say this caused that?

Every individual  $i$  has two potential outcomes:

- If individual  $i$  is “treated”:  $Y_{1i}$
- If individual  $i$  is “not treated”:  $Y_{0i}$

We only ever observe either  $Y_{1i}$  or  $Y_{0i}$ . Never both.

The causal effect of treatment for individual  $i$  is defined as:

$$Y_{1i} - Y_{0i}$$

The Average Treatment Effect (ATE) in the population is:

$$ATE = E[Y_{1i} - Y_{0i}]$$



Mastering 'Metrics



# The Counterfactual

Each person is either in the treated group ( $D_i = 1$ ), or in the control group ( $D_i = 0$ ), but not both at the same time. Thus, for each person, we can only observe one of their potential outcomes. What we observe ( $Y_i$ ) is:

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

The other (unobserved) outcomes are the **counterfactuals**.

- For individuals in the treated group the counterfactual is:  $(Y_{0i} | D_i = 1)$ .
- For individuals in the control group the counterfactual is:  $(Y_{1i} | D_i = 0)$ .

Therefore, we will never observe the **causal effect of treatment** for individual  $i$  ( $Y_{1i} - Y_{0i}$ ) or the true  $ATE = E[Y_{1i} - Y_{0i}]$ .

# What does this definition of causal effect imply?



Mastering 'Metrics

- The causal effect includes everything along the causal pathway.
  - If “A” causes “B” and “B” causes “C” then “A” causes “C”.
- This definition is not the only possible definition of causal effect (legal definitions include proximate cause, actual cause, cause in fact, substantial factor, etc...).
- We adopt this definition of cause because it is most relevant for policy. I want to know if I enact policy A, what will happen to outcome C?
- Still, in many ways, this definition can seem a bit strange.

Example 1: Will eating chicken four times a week cause individual  $i$  to live longer?

Assume the following stylized facts are true for this example:

- The only protein people eat are red meat, chicken, and fish.
- The cholesterol in red meat “causes” heart attacks.
- The omega-3 fatty acids in fish “prevents” heart attacks.
- Nothing in chicken “*causes*” or “*prevents*” heart attacks.

- Question (1) If eating chicken four times a week causes individual  $i$  to eat less red meat, which in turn causes the individual to live longer, is it correct to say that eating chicken **caused** individual  $i$  to live longer?
- Question (2) If eating chicken four times a week causes individual  $j$  to eat less fish, which in turn causes the individual to die sooner, is it correct to say that eating chicken **caused** individual  $j$  to die sooner?

- The answer to both questions is “**yes,**” at least according to the potential outcome’s definition of causation (which is the definition we will use, ...but why?).
- Think about this from a policy perspective. (1) Would you expect a policy that increases the consumption of chicken to affect the number of heart attacks? ; and (2) would you expect this policy to have the same Average Treatment Effect (ATE) among ranchers in Montana (meat eaters) and the Inuit in Alaska (fish eaters)?

- Answers: (1) Yes, probably. ; (2) NO!

## This raises the question: what is “the” ATE?

- Definition: The Average Treatment Effect (ATE) is the expected effect on the outcome if everyone in the population were treated (or among individuals from the population who were randomly assigned to treatment).
- Notice how the ATE depends on the question we are asking and the population we are asking it about.
- The ATE might not be relevant to policy makers because it may include the effect on persons for whom the treatment was never intended.
  - However, this is partly semantics because we are asking the questions. Most researches define the ATE as what they are trying to estimate among the population that they are interested in.
  - It is OK to be interested in ranchers in Montana.

# Why is the causal effect not the same for everyone?

- Does going to college have the same effect on everyone? If not, why not?

Two possible reasons:

- **Heterogeneous effects.** Individuals are different, so you might think that going to college could affect each person in a *systematically* different way. By *systematically*, we mean that:  $E[Y_{1i} - Y_{0i}] \neq E[Y_{1j} - Y_{0j}]$ .

Put another way, the expected treatment effect varies from individual to individual. This is important because, unaccounted for, this can lead to **selection bias** and to **selection on the returns to treatment** (more on these concepts later).

- **Random events.** Going to college puts a person on an entirely different path in life. This can expose a person to a different environment filled with different random events. However, this alone does not mean that the expected treatment effect varies among individuals.

## More on random events

Example: Suppose that if Jane goes to college, she will be killed in a random car accident while driving to class one day. If she does not go to college, she will not get in the car accident and will instead die of a heart attack at the age of 95. These are Jane's potential outcomes, even if the probability of her getting into a fatal car accident by age 25 is the same if she goes to college or not. Her potential outcomes are:

$$Y_{1,Jane} = 25 ; Y_{0,Jane} = 95 \quad ; \quad (Y_{1,Jane} - Y_{0,Jane}) = -75.$$

Question: Would you say that going to college causes Jane to die young?

- The answer is: yes. Again, this definition can seem a bit strange, but it is the truth for Jane (i.e., it is the difference in her potential outcomes).

Bottom line: We define the causal effect of treatment for an individual  $i$  to be the difference in the individual's potential outcomes ( $Y_{1i} - Y_{0i}$ ). While *random* events can affect the potential outcomes of any specific individual  $i$ , if they are truly independent, they average out of the ATE for the population, and do not lead to the **selection bias problem**.

## This leads to more types of questions not covered in this class

- Generally, we avoid case study questions (the effect of a unique one-time event).
- We want to estimate  $E[Y_{1i} - Y_{0i}]$  (i.e., the ATE), not  $(Y_{1i} - Y_{0i})$  for one observation. There may be many observations, but how many random draws?
  - We do not ask: “Did going to college cause Jane to die in a car accident?”
  - We do ask: “Does going to college increase a person’s probability of dying in a car accident?”
  - We ask: “What is the effect of minimum wage laws on youth labor force participation?”
  - We avoid: “What was the effect of Seattle’s minimum wage increase to \$11 an hour in April of 2015 on Seattle’s youth labor force participation?”
- In this class, and in this program, we are most interested in using econometrics to generate generalizable knowledge. All scientific questions/conclusions should be subject to replication by other experiments in the future.

# The Causal Pathway (Mechanism/Channel)

Suppose that, for most people, the causal pathway from college to life expectancy looks something like:

College → increased wealth → better health → longer life.

- If this case, we might think that going to college will increase the life expectancy of everyone, which we would write as:  $Y_{1i} - Y_{0i} > 0 \forall i$ .
- Even in this case, we may or may not expect that going to college would increase everybody's life expectancy by the same amount (due to Heterogeneous effects or random events).



# Constant treatment effects

Most of the time in this course, and in the Mastering 'Metrics book, it simplifies things to think in terms of **constant treatment effects**. ...But what exactly does this mean? The simplest way to think about constant treatment effects is:

$$Y_{1i} - Y_{0i} = \kappa \quad \forall i.$$

- This suggests that the causal effect ( $\kappa$ ) is the same for everyone. This assumption is stronger than is necessary, but it is often made for simplicity.
- There are situations where this assumption would be plausible, and others, when taken *literally*, would be unrealistic. For example, consider the example of a binary treatment and a binary outcome.

## Example: Helmet laws and motorcycle deaths

Individuals are treated or not ( $D_i = \{0,1\}$ ), they live or die ( $Y_i = \{0,1\}$ )

There are only four possible combinations of potential outcome

---

	Potential outcomes ( $n = 1,000$ )			
	Type A ( $n = 100$ )	Type B ( $n = 795$ )	Type C ( $n = 100$ )	Type D ( $n = 5$ )
If not treated ( $Y_{0i}$ )	1 (die)	0 (survive)	1 (die)	0 (survive)
If treated ( $Y_{1i}$ )	0 (survive)	0 (survive)	1 (die)	1 (die)
Causal Effect of treatment ( $Y_{1i} - Y_{0i}$ )	-1	0	0	1

---

Individuals are:

- Type A: Helmets causes 100 lives to be saved.
- Type B & C: Helmets has no effect on 895 individuals.
- Type D: Helmets caused 5 deaths (*“possible side effect include: death”*).

Conceptual question: How were individual assigned their potential outcomes?  
Ask yourself: Why are there 100 people who are Type A, etc...

- Potential outcomes are deterministic concepts.
  - $Y_{1i}$  and  $Y_{0i}$  are deterministic (fixed).
  - Individuals are just Type A, B, C or D, and that's all there is to it. It is just in their DNA.
  - Constant treatment effects means:  $Y_{1i} - Y_{0i} = \kappa$ . In this setting, everyone would have to be the same type (A, B/C, or D). This does not make sense (no treatment saves everyone's life).
- Potential outcomes are probabilistic concepts.
  - Each individual  $i$  has a different probability of death.
  - Constant treatment effects means:  $E[Y_{1i} - Y_{0i}] = \kappa$ . (not  $\kappa_i$ )
  - Treatment lowers everyone's probability of death by 5%.

# The Selection Problem

The expected difference in outcomes between the treatment and control groups is:




$$E[Y_{1i}|D_i = 1] - E[Y_{0i}|D_i = 0] \quad (\text{Expected difference in group means})$$

If we add and subtract  $E[Y_{0i}|D_i = 1]$ :  $\longrightarrow$

What the average outcome for the treated group *would have been* if they were not treated.

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



$$\text{Expected difference in group means} = E[Y_{1i}|D_i = 1] - E[Y_{0i}|D_i = 1] + E[Y_{0i}|D_i = 1] - E[Y_{0i}|D_i = 0]$$

*Average effect of Treatment on the Treated (ATT)*

*Selection Bias (SB)*

Summary: What we want to estimate is the Average Treatment Effect (ATE) in the population that we are interested in. However, what we observe in our data is the difference in the groups means between the control and treatment group:  $(\bar{Y}_{D=1} - \bar{Y}_{D=0})$ . This difference is equal to:

$$(\bar{Y}_{D=1} - \bar{Y}_{D=0}) = ATT + SB + \textit{Sampling Error (SE)}$$

- For now, we will assume that we have a large enough sample so that sampling error is close to zero. So, we will simply write:

$$(\bar{Y}_{D=1} - \bar{Y}_{D=0}) = ATT + SB$$

- To compare what we observe to the true ATE, write (add and subtract ATE):

$$\underbrace{(\bar{Y}_{D=1} - \bar{Y}_{D=0})}_{\substack{\textit{Difference in} \\ \textit{group means}}} = \underbrace{ATE}_{\substack{\textit{Average} \\ \textit{Treatment} \\ \textit{Effect}}} + \underbrace{(ATT - ATE)}_{\substack{\textit{Selection on} \\ \textit{the returns to} \\ \textit{treatment}}} + \underbrace{SB}_{\substack{\textit{Selection} \\ \textit{Bias}}}$$

## What are these pieces?

$$(\bar{Y}_{D=1} - \bar{Y}_{D=0}) = ATE + (ATT - ATE) + SB$$

$ATE = E[Y_{1i} - Y_{0i}]$ . The thing we want to estimate in the population we are interested in.

$SB = E[Y_{0i}|D_i = 1] - E[Y_{0i}|D_i = 0]$ . Selection bias: the outcomes of those treated and not treated would have been different even in the absence of the treatment (i.e., no SB means no difference in  $Y_{0i}$  between treated and control groups). For example:

- Suppose  $E[Y_{0i}|D_i = 1] > E[Y_{0i}|D_i = 0]$ . Then individuals in the treated group would have had better outcomes than those in the control group even if they were not treated. This is a positive selection bias.
- Suppose  $\{Y_{0i}, Y_{1i}\} \perp D_i$ . Treatment is statistically independent of all potential outcomes. Then  $ATT = E[Y_{1i} - Y_{0i}|D_i = 1] = E[Y_{1i} - Y_{0i}] = ATE$ , and  $SB = E[Y_{1i}] - E[Y_{0i}] = 0$ . (With random assignment,  $ATE = ATT$ , and there is no selection bias).

$ATT = E[Y_{1i}|D_i = 1] - E[Y_{0i}|D_i = 1]$ . Can be written as  $E[Y_{1i} - Y_{0i}|D_i = 1]$ .

ATT is the average treatment effect among the population that gets treated.

# Selection on Returns –vs– Selection Bias

- Selection on returns (i.e.,  $ATT \neq ATE$ ) means that the group that was treated responds differently to the treatment than the general population. Wearing a helmet may have a different impact on the group of motorcycle riders who wear helmets. Going to college may have a different effect on those who went to college than those who did not.
- Selection on returns is not an issue of bias. The ATT is the correct Average Treatment Effect among the population that gets treated. For example, it is the causal effect of earning a college degree for those who earned a college degree. This would be important to know.
- However, the ATT may not be the average treatment effect (ATE) of requiring everyone to go to college, or of requiring all motorcycle riders to wear helmets.
- Selection Bias means that the treatment group would have been different than the control group even if they did not receive the treatment. Therefore, with SB it would be incorrect to conclude that the differences between the treatment and control groups were caused by the treatment. In fact, with SB, the estimate is not the correct estimate for any group.

# Selection on returns is a question internal/external validity

- Suppose we want to estimate the impact of some policy on the U.S. population (i.e., the ATE among the U.S. population). However, we conduct our study in Colorado. Suppose there is no selection bias, so our estimate is internally valid for Colorado. We would call our estimate the ATT. This ATT is the average treatment effect (ATE) in Colorado.
- If the policy has the same effect in Colorado as it does in the U.S. population, then  $ATT = ATE$ , and the estimate is externally valid for the population we are interested in.
- The difference between ATT and ATE is a matter of perspective. If we change the research question to be “what is the impact of the policy in Colorado?” then what we called ATT would now be called the ATE.
- Although the ATT/ATE (internal/external validity) distinction is a bit subjective, it is important. Research is often done using subjects that are not randomly sampled from the population of interest, but are randomly sampled from some other population. Hopefully, this other population is representative of the population of interest, but this is the question.



# Conclusion: Association vs Causation

$$\underbrace{(\bar{Y}_{D=1} - \bar{Y}_{D=0})}_{\text{Difference in group means}} = \underbrace{ATE}_{\text{Average Treatment Effect}} + \underbrace{(ATT - ATE)}_{\text{Selection on the returns to treatment}} + \underbrace{SB}_{\text{Selection Bias}} + \underbrace{SE}_{\text{Sampling Error}}$$

In other words:

*Association = Causation + External validity + Selection bias + Sampling Error*

If we can eliminate *selection bias*, have *external validity*, and a *sufficiently large sample*, then:

*Association = Causation*

- It is the identification strategy's job to eliminate selection bias, and tell us if we are estimating the **ATE or ATT** (and later on, the **LATE**). For example, if treatment is randomly assigned in the population of interest, compliance is perfect, and there is no selective attrition, then  $ATT = ATE$  and  $SB = 0$ . In this case, we say that “we are estimating the causal ATE.”