The LATE Theorem

Hunan University

Heterogeneous Effects

- We have implicitly assumed that treatment X has the same effect (β) on everyone.
- In real life, treatment effects are usually *heterogeneous*.
- Eg: it's unlikely that each person would get the same benefit from an **additional year of schooling**;
  - for some people, the effect can even be negative.

- Since treatment effects can be *heterogeneous*, we economists always say that we are estimating the **average** treatment effect (ATE)
  - **BAD:** when treatment effect is heterogeneous, the instrumental variable approach is not valid.
  - **GOOD:** All may not be lost! We just need a different interpretation of the estimate.
- The remedy is simple: Just interpret our estimate as the average effect for **a subset of the population**.
  - The corresponding *estimand* is called the Local Average Treatment effect.
- BTW, do you know the diff between *estimand, estimate* and *estimator*?

A brief aside: estimands, estimators and estimates

- **Estimand**: the quantity to be estimated
- **Estimate**: the approximation of the estimand using a finite data sample
- **Estimator**: the method or formula for arriving at the estimate for an estimand

A brief aside: estimands, estimators and estimates

- **Estimand**: the quantity to be estimated
- **Estimate**: the approximation of the estimand using a finite data sample
- **Estimator**: the method or formula for arriving at the estimate for an estimand

For example, we specify the data generation process $y = 1 + \beta x + \epsilon$ with $\beta = 2$. We use R to simulate the data $\{x_i, y_i\}_{i=1}^{N}$ and `lm(y~x)` yields $\hat{\beta} = 2.1$. What are estimands, estimators and estimates?

Heterogeneous Effects: Model

$$Y_i = \alpha + \beta_i X_i + \epsilon_i$$

- Intuitively, different "research designs" (e.g. instruments) may capture different effects of the same treatment — even when all are valid
  - Different IVs may yield very different estimates!
- This idea is formalized in the (Nobel-winning) Imbens and Angrist 1994 LATE theorem.
  - It uses a general potential outcomes framework.

Potential Outcome Setup

Let $Y_i(0)$ and $Y_i(1)$ be an individual i's potential outcomes given a binary treatment $D_i \in \{0, 1\}$.

- Observed outcomes: $Y_i = (1 - D_i)Y_i(0) + D_iY_i(1)$

Potential Outcome Setup

Let $Y_i(0)$ and $Y_i(1)$ be an individual i's potential outcomes given a binary treatment $D_i \in \{0, 1\}$.

- Observed outcomes: $Y_i = (1 - D_i)Y_i(0) + D_i Y_i(1) = \alpha_i + \beta_i D_i$
- (Individaul-level) treatment effects $\beta_i = Y_i(1) - Y_i(0)$

Imbens-Angrist: we can also do this for an IV first stage:

- Let $D_i(0)$ and $D_i(1)$ denote individual i's potential treatment given a binary treatment $Z_i \in \{0, 1\}$.
- **What is** `ivreg(Y ~ D|Z)`?

Imbens and Angrist (1994) Assumptions

1. **As-good-as-random assignment:**
    - The assignment of Z_i is independent of Y_i and D_i
2. **Exclusion:** Z_i only affects Y_i through D_i.
    - Implicit in our potential outcomes notation: $Y_i(D)$ not indexed by Z_i
3. **Relevance:** Z_i is correlated with D_i.
    - Equivalently, $E[D_i(1) - D_i(0)] \neq 0$.

Imbens and Angrist (1994) Assumptions

1. **As-good-as-random assignment:**
   - The assignment of Z_i is independent of Y_i and D_i
2. **Exclusion:** Z_i only affects Y_i through D_i.
3. **Relevance:** Z_i is correlated with D_i.
4. **Monotonicity:** $D_i(1) \geq D_i(0)$ for all i.
   - The instrument can only shift the treatment in one direction.
   - Egs: College distance, rainfall, birth-quarter,...

Local Average Treatment Effect (LATE) Identification

Imbens and Angrist showed that under these assumptions:

$$\beta^{IV} = E[Y_i(1) - Y_i(0) \mid D_i(1) > D_i(0)]$$

The IV estimator β^{IV} identifies a **LATE**:

- the **average treatment effect** $Y_i(1) - Y_i(0)$ among *compliers:* those with $1 = D_i(1) > D_i(0) = 0$.

Local Average Treatment Effect (LATE) Identification

Imbens and Angrist showed that under these assumptions:

$$\beta^{IV} = E[Y_i(1) - Y_i(0) \mid D_i(1) > D_i(0)]$$

The IV estimator β^{IV} identifies a **LATE**:

- the **average treatment effect** $Y_i(1) - Y_i(0)$ among *compliers:* those with $1 = D_i(1) > D_i(0) = 0$.

Intuitively, IV can't tell us anything about the treatment effects of *never-takers* and *always-takers*.

Complier, Always Taker, Never Taker, Defier

1. *Complier:* $\Pr(D=1|Z=1,C)=\Pr(D=0|Z=0,C)=1$
2. *Always Taker:* $\Pr(D=1|Z=1,A)=\Pr(D=1|Z=0,A)=1$
3. *Never Taker:* $\Pr(D=0|Z=1,N)=\Pr(D=0|Z=0,N)=1$
4. *Defier:* $\Pr(D=0|Z=1,De)=\Pr(D=1|Z=0,De)=1$

Complier, Always Taker, Never Taker, Defier

1. *Complier:* $\Pr(D = 1 | Z = 1, C) = \Pr(D = 0 | Z = 0, C) = 1$
2. *Always Taker:* $\Pr(D = 1 | Z = 1, A) = \Pr(D = 1 | Z = 0, A) = 1$
3. *Never Taker:* $\Pr(D = 0 | Z = 1, N) = \Pr(D = 0 | Z = 0, N) = 1$
4. *Defier:* $\Pr(D = 0 | Z = 1, De) = \Pr(D = 1 | Z = 0, De) = 1$

LATE Theorem: IV identifies the average treatment effect among *Compliers*.

- Note: The monotonicity assumption rules out *Defiers*.

What Does This Mean *Practically*?

Two conceptually distinct considerations: **internal vs. external validity**

- **Context of an IV**, and who the **compliers** are, matter.
- Usual "over-identification" test logic fails: two valid IVs may have different estimands. (Kitagawa, 2015)

What Does This Mean *Practically*?

Two conceptually distinct considerations: **internal vs. external validity**

- **Context of an IV**, and who the **compliers** are, matter.
- Usual "over-identification" test logic fails: two valid IVs may have different estimands. (Kitagawa, 2015)

In addition to as-good-as-random assignment and exclusion, we may need to worry about **monotonicity** when we do IV.

- Monotonicity holds in almost all IV examples
- But it may fail. Eg, judge IV.

Judge (or examiner) IV design

Imbens and Angrist (1994) provide an example in which Monotonicity fails.

EXAMPLE 2 (Administrative Screening):[5] Suppose applicants for a social program are screened by two officials. The two officials are likely to have different admission rates, even if the stated admission criteria are identical. Since the identity of the official is probably immaterial to the response, it seems plausible that Condition 1 is satisfied. The instrument is binary so Condition 3 is trivially satisfied. However, Condition 2 requires that if official A accepts applicants with probability $P(0)$, and official B accepts people with probability $P(1) > P(0)$, official B must accept *any* applicant who would have been accepted by official A. This is unlikely to hold if admission is based on a number of criteria. Therefore, in this example we *cannot* use Theorem 1 to identify a local average treatment effect nonparametrically despite the presence of an instrument satisfying Condition 1.

[5] This example was suggested to us by Geert Ridder.

Judge (or examiner) IV design

Imbens and Angrist (1994) provide an example in which Monotonicity fails.

A judge (or examiner) IV design leverages the *idiosyncratic assignment* of individuals to a set of judges. A more tolerant judge is likely to be more merciful in deciding the prison sentence.

- Kling (2006): sentencing judges
- Doyle (2007): foster care investigators
- Maestas et al. (2013): SSDI benefit examiners
- Doyle et al. (2015): ambulance companies

Judge (or examiner) IV design

Imbens and Angrist (1994) provide an example in which Monotonicity fails.

A judge (or examiner) IV design leverages the *idiosyncratic assignment* of individuals to a set of judges. A more tolerant judge is likely to be more merciful in deciding the prison sentence.

All these works need to take special care of the monotonicity restriction.

- Eg: Frandsen et al. (2019) formalize a weaker **"average monotonicity"** condition.

Extensions of Angrist and Imbens (1994)

Angrist and Imbens worked out the original LATE theroem for binary D_i, discrete Z_i, and no included controls.

- Angrist/Imbens '95: multi-valued (ordered) D_i , saturated covariates
- Angrist/Graddy/Imbens ’00: continuous D_i (supply/demand setup)
- Heckman/Vytlicil ’05: continuous Z_i
- Multiple unordered treatments is harder (e.g. Behaghel et al. 2013)

Extensions of Angrist and Imbens (1994)

Angrist and Imbens worked out the original LATE theroem for binary D_i, discrete Z_i, and no included controls.

- Angrist/Imbens '95: multi-valued (ordered) D_i , saturated covariates
- Angrist/Graddy/Imbens ’00: continuous D_i (supply/demand setup)
- Heckman/Vytlicil ’05: continuous Z_i
- Multiple unordered treatments is harder (e.g. Behaghel et al. 2013)

Recent discussions highlight importance of including flexible controls

- Sloczynski ’20, Borusyak and Hull ’21, Mogstad et al. ’22

Better LATE than nothing

- LATE is not a new *estimator*. We are still using 2SLS as before.
- LATE is about the *estimand*. When we are doing IV, we are really just estimating the average treatment effect on the *Compliers*.
    - Besides **exclusion** and **as-good-as-random assignment**, the **monotonicity restriction** is also needed.
- LATE implies that we should always beaware of *Context of an IV* and who the *Compliers* are.
    - In other words, LATE makes it **harder to misuse the IV research design**.
- For these reasons, Imbens (2010) argues that we should always think in the LATE framework --- better LATE than nothing.

Further readings on IV

- Lots of stuff about IV are still not covered.
  - Specifically, how to **characterize compliers**? How to use IV with **panel data** or in a **DiD** setting? **GMM version** of IV?

- I hope that starting from this course, you can continue your study of IV and use it in your own research.
  - Mixtape IV
  - Pinkham's Applied Empirical Methods at Yale

Final words on IV

- It's getting more popular to view IV (and DiD, RD...) as experimental methods, and use "causal inference" (instead of regression-based analysis) as a unifying framework.
- We do not really talk too much about the potential outcome framework, the underlying framework used for causal inference so far.
  - If you are interested, there is an online course that can be helpful: Crash course in causality.