7. Noncompliance and Instrumental Variables

ISS5096 || ECI
Jaewon ("Jay-one") Yoo
National Tsing Hua University

Outline

Where are we? Where are we going?

- We've covered randomized experiments (no confounding)
- We've covered selection on observables (no unmeasured confounding)
- What if there is unmeasured confounding? What can we do?
- First approach we will explore: **instrumental variables**
    - First: motivate IV through experiments and noncompliance.
    - Next: discuss how it relates to classical econometric methods like 2SLS?

Researcher: You, are in the control group. No need to take the treatment

Defier: But I want it!

Researcher: Just kidding, you are in the treatment group. Here it is

Defier:

Source: Causal Inference for the Brave and True By Matheus Facure Alves

1/ Randomized Experiments with Noncompliance

Noncompliance

- Get-out-to-vote (GOTV) experiment with door-to-door canvassing:

Source: https://www.dreamstime.com/political-meeting-candidates-pre-election-campaign-concept-multinational-people-support-banners-man-woman-voters-flat-image185175984

- Households are randomized so treatment assignment is unconfounded
  - $Z_i = 1$ for assigned to treatment (canvassing attempted),
  - $Z_i = 0$ for assigned to control (no canvassing attempted).
- **Noncompliance**: units don't follow treatment assignment.
  - Units assigned to treatment take control or vice versa.
  - $D_i = 1$ for actually took treatment (heard canvasser message).
  - $D_i = 0$ for actually took control (didn't answer the door).
  - Full compliance means $Z_i = D_i$ for all i.

How to Handle Noncompliance?

- What do we do if we want the effect of *canvassing exposure* (D_i)?
- Two approaches one might consider:
  1. **Intent-to-treat** (ITT) analysis: just use randomization.
     - Use Z_i as the treatment and analyze as a typical experiment.
     - Downside: can't learn about the effect of actually being canvassed!
  2. **As-treated** analysis: just use treatment uptake.
     - As if D_i was randomly assigned.
     - Not valid if uptake is **correlated** with outcome.
     - $\rightsquigarrow$ unmeasured confounding between D_i and **PO**s

The As-Treated Problem, Visualized

$n = 200, \quad n_1 = n_0 = 100, \quad \pi_{co} = 0.7$ (one-sided noncompliance)

$Z = 1$ **group:**

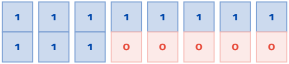

$\cdots$ D_i varies

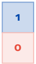

$D_i = 1$

$D_i = 0$

$Z = 0$ **group:**

$\cdots$ all $D_i = 0$

"Treated" ($D = 1$)
70 units
a selected subgroup

vs.

"Control" ($D = 0$)
130 units
everyone else

Apples to oranges!

$D = 1$ are people who *chose* to comply. Systematically different.
Comparing by D reintroduces confounding.

$\rightsquigarrow$ Alternative: leverage latent strata of **compliance types**.

Setup

- Treatment assignment, $Z_i \in \{0, 1\}$; treatment uptake, $D_i \in \{0, 1\}$,
- Treatment uptake is affected by assignment: $D_i(z)$
  - $D_i(1) = 1$ if assigned to canvassing, I **would** open my door.
  - $D_i(1) = 0$ if assigned to canvassing, I **would not** open my door.
  - Noncompliance means $D_i(z) \neq z$ for some i.
- Consistency for the observed treatment as usual:

$$D_i = D_i(Z_i) = Z_i D_i(1) + (1 - Z_i) D_i(0)$$

- Canvassing is an example of **one-sided noncompliance**.
  - People might refuse treatment when offered (i.e., $D_i(1) = 0$)
  - But no one receives treatment if in control (i.e., $D_i(0) = 0, \ \forall\ i$)
  - **Two-sided noncompliance** is when you can refuse to comply with treatment **or** control.

Potential Outcomes

- Outcomes now dependant on assignment and uptake: $Y_i(z, d)$.
  - $Y_i(1, 1)$: would I vote if I were assigned to canvassing and received it?
- Can only observe two potential outcomes: $Y_i(1, D_i(1))$ and $Y_i(0, D_i(0))$
  - $Y_i(1, D_i(1))$: potential outcome when assigned canvassing and whether uptake occurs for unit i when assigned to canvassing.
  - $Y_i(1, 1 - D_i(1))$: not possible to ever observe! (cross-world counterfactual)
    - Contradicts the unit's own compliance type.
    - Not needed under the IV framework, but enters the definition of causal mediation parameters $\rightsquigarrow$ makes their identification harder (more later).
- Consistency assumption: $Y_i = Y_i(Z_i, D_i(Z_i))$

Some Notations

- Let's use 0/1 subscripts for assignment, and t/c subscripts for uptake:

$$n_1 = \sum_{i=1}^{n} Z_i \quad n_0 = \sum_{i=1}^{n} (1 - Z_i) \quad n_t = \sum_{i=1}^{n} D_i \quad n_c = \sum_{i=1}^{n} (1 - D_i)$$

- Average outcomes and uptake in each assignment group:

$$\overline{Y}_1 = \frac{1}{n_1} \sum_{i=1}^{n} Z_i Y_i \quad \overline{Y}_0 = \frac{1}{n_0} \sum_{i=1}^{n} (1 - Z_i) Y_i$$

$$\overline{D}_1 = \frac{1}{n_t} \sum_{i=1}^{n} Z_i D_i \quad \overline{D}_0 = \frac{1}{n_c} \sum_{i=1}^{n} (1 - Z_i) D_i$$

- Assumption 1: **randomization** $[\{Y_i(d, z), \forall d, z\}, D_i(1), D_i(0)] \perp\!\!\!\perp Z_i$
  - For observational uses of IV, might condition on some $\mathbf{X}_i$.

Defining ITT Effects

- **Intent-to-treat** (ITT) effects are just the ATEs of Z_i:

$$\text{ITT}_D = \frac{1}{n}\sum_{i=1}^{n} D_i(1) - D_i(0) \qquad \text{ITT}_Y = \frac{1}{n}\sum_{i=1}^{n} Y_i(1, D_i(1)) - Y_i(0, D_i(0))$$

  - SATE of assignment on treatment uptake & the outcome.
  - If noncompliance is one-sided, then $\text{ITT}_D \geq 0$

- Standard estimators for these quantities:

$$\widehat{\text{ITT}}_D = \overline{D}_1 - \overline{D}_0 \qquad \widehat{\text{ITT}}_Y = \overline{Y}_1 - \overline{Y}_0$$

- Under randomization of Z_i, everything just like Neyman approach.
  - Variances, tests, CIs all standard.

- Problem: ITT_Y is a combination of true effect of D_i and noncompliance
  - Effect of D_i may be more externally valid than Z_i.

2/ Compliance Types

Compliance Status

- We can stratify units by their **compliance type**.
    - Compliance type = how units would respond to treatment assignment.
    - Basically, it is the value of $(D_i(0), D_i(1))$ for any unit.
- Under one-sided noncompliance, there are two types:
    - **Compliers** with $D_i(1) = 1$ and **noncompliers** with $D_i(1) = 0$.
    - Compliers answer the door when assigned to canvassing.
    - Noncompliers don't answer the door when assigned to canvassing.
    - Everyone has $D_i(0) = 0$, so no noncomliance there.
- Compliance is a function of potential outcomes so it is **pretreatment**!
    - $\rightsquigarrow$ treatment assignment independent of C_i

ITTs Among the Compliance Groups

- Compliance type indicator, $C_i \in \{\text{co}, \ \text{nc}\}$:
  - Number of compliers: $n_{\text{co}} = \sum_{i=1} \mathbb{1}(C_i = \text{co})$.
  - Proportion of compliers: $\pi_{\text{co}} = n_{\text{co}}/n$.
  - Same for noncompliers: n_{nc} and π_{nc}.
- ITT on uptake is directly related to compliance type:

$$\text{ITT}_{D,\text{nc}} = \frac{1}{n_{\text{nc}}} \sum_{i:C_i=\text{nc}} D_i(1) - D_i(0) = 0$$

$$\text{ITT}_{D,\text{co}} = \frac{1}{n_{\text{co}}} \sum_{i:C_i=\text{co}} D_i(1) - D_i(0) = 1$$

  - Intuition: no effect of assignment on uptake for noncompliers!
- Implies that overall ITT on uptake equals the **proportion of compliers**:

$$\text{ITT}_D = \pi_{\text{co}} \cdot \text{ITT}_{D,\text{co}} + \pi_{\text{nc}} \cdot \text{ITT}_{D,\text{nc}} = \pi_{\text{co}}$$

3/ Instrumental Variables

Identification Assumptions

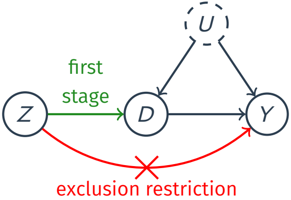

- Assumption 2. **First stage** (relevance) $\text{ITT}_D = \pi_{\text{co}} \neq 0$
  - At least one person complies with treatment.
- Assumption 3. **exclusion restriction** Z_i only affects Y_i through D_i
  - $Y_i(z, d) = Y_i(z', d)$ for all z, z', and d.
  - Assignment to canvassing only affects turnout through actual canvassing
  - Not a testable assumption and cannot be guaranteed by design!
- Implies that potential outcomes only a function of D_i:

$$Y_i(1) = Y_i(D_i = 1) = Y_i(Z_i = 1, D_i = 1)$$
$$Y_i(0) = Y_i(D_i = 0) = Y_i(Z_i = 1, D_i = 0)$$

When Does the Exclusion Restriction Fail?

- ER says: Z_i affects Y_i **only through** D_i.
  - Not testable. Not guaranteed by design. Must argue substantively.
- **GOTV example**: does assignment to canvassing affect turnout *only* through hearing the canvasser's message?
  - Canvassers often leave a *flyer on the door* when no one answers. Noncompliers still receive election information. ($Z \longrightarrow Y$ directly)
  - A knock itself might *remind* you about the election, even if you don't open.
- **Draft lottery** (Angrist, 1990): does being drafted affect earnings *only* through military service?
  - Low draft number could cause stress, change college plans, or alter career expectations, even before serving.
- ER is often the **most controversial** IV assumption.
  - Think carefully: is there *any* channel from Z to Y besides D?

Outcome ITTs and Compliance Types

- We can define the ITTs on the outcome by compliance type as well.
  - $\text{ITT}_{Y,\text{co}} \rightsquigarrow$ effect of assignment on outcome among compliers.
  - $\text{ITT}_{Y,\text{nc}} \rightsquigarrow$ effect of assignment on outcome among noncompliers.
- Exclusion restriction has implications for these:
  - Implies that $\text{ITT}_{Y,\text{nc}} = 0$: if D_i does not change, Y_i cannot change.
  - Implies that $\text{ITT}_{Y,\text{co}}$ is due entirely to treatment uptake.
- Allows us to connect the ITT_Y on the outcome to compliance groups:

$$\begin{aligned}\text{ITT}_Y &= \pi_{\text{co}} \cdot \text{ITT}_{Y,\text{co}} + \pi_{\text{nc}} \cdot \text{ITT}_{Y,\text{nc}} \\ &= \text{ITT}_D \cdot \text{ITT}_{Y,\text{co}}\end{aligned}$$

LATE

- Under the exclusion restriction, $\text{ITT}_{Y,\text{co}}$ is the effect of treatment receipt/uptake:

$$\begin{aligned}\text{ITT}_{Y,\text{co}} &= \frac{1}{n_{\text{co}}} \sum_{i:C_i=\text{co}} Y_i(1, D_i(1)) - Y_i(0, D_i(0)) \\ &= \frac{1}{n_{\text{co}}} \sum_{i:C_i=\text{co}} Y_i(D_i = 1) - Y_i(D_i = 0) = \tau_{\text{LATE}}\end{aligned}$$

- This quantity is known as the **local ATE** (LATE), local to compliers.
  - It's a conditional ATE, where we condition on units being a complier.
  - Also referred to as **complier average causal effect** (CACE).

- **LATE theorem** under one-sided noncompliance, exclusion restriction, first-stage, and randomization:

$$\tau_{\text{LATE}} = \text{ITT}_{Y,\text{co}} = \frac{\text{ITT}_Y}{\text{ITT}_D}$$

Wald Estimator

- **Wald** or **instrumental variable estimator** for the LATE:

$$\widehat{\tau}_{\text{iv}} = \frac{\widehat{\text{ITT}}_Y}{\widehat{\text{ITT}}_D}$$

  - Ratio of the two unbiased ITT estimators.
  - Not 'unbiased', but it is **consistent** for τ_{LATE}.
    - $\mathbb{E}[A/B] \neq \mathbb{E}[A] \,/\, \mathbb{E}[B]$.

- Equivalent to the **two-stage least squares** estimator:
  - Regress D_i on Z_i to get fitted values, $\widehat{D}_i$.
  - Regress Y_i on $\widehat{D}_i$.

- We can use the delta method to find the (superpopulation) variance:

$$\mathbb{V}[\widehat{\tau}_{\text{iv}}] = \frac{1}{\text{ITT}_D^2}\mathbb{V}\left[\widehat{\text{ITT}}_Y\right] + \frac{\text{ITT}_Y^2}{\text{ITT}_D^4}\mathbb{V}\left[\widehat{\text{ITT}}_D\right] - 2 \cdot \frac{\text{ITT}_Y}{\text{ITT}_D^3} \cdot \text{cov}\left[\widehat{\text{ITT}}_Y, \widehat{\text{ITT}}_D\right]$$

4/ Two-Sided Noncompliance

Two-Sided Noncompliance

- Two-sided noncompliance: those in control can select into treatment.
- **Encouragement design**: randomly assign an encouragement of some treatment.
  - Some may refuse encouragement and opt not to take treatment.
  - Some may take treatment even without encouragement.
- Z_i is the encouragement, and D_i is the treatment.
- No change in estimation, just different identification assumptions.

Compliance Types

- Four compliance types (or **principal strata**) in this setting:
  - Complier: $D_i(1) = 1$ and $D_i(0) = 0$
  - Always-taker: $D_i(1) = D_i(0) = 1$
  - Never-taker: $D_i(1) = D_i(0) = 0$
  - Defier: $D_i(1) = 0$ and $D_i(0) = 1$
- Connections between observed data and compliance types:

	$Z_i = 0$	$Z_i = 1$
$D_i = 0$	Never-taker or Complier	Never-taker or Defier
$D_i = 1$	Always-taker or Defier	Always-taker or Complier

- Let π_{co}, π_{at}, π_{nt}, and π_{df} be the proportions of each compliance type.
- ITT effects on D_i are more murky: $\text{ITT}_D = \pi_{\text{co}} - \pi_{\text{df}}$
  - Defiers really makes things messy!

Instrumental Variables Assumptions

- Canonical IV assumptions for Z_i to be a valid instrument:
  1. Randomization of Z_i
  2. Presence of some compliers $\pi_{\text{co}} \neq 0$ (first-stage)
  3. Exclusion restriction $Y_i(z, d) = Y_i(z', d)$
  4. **Monotonicity**: $D_i(1) \geq D_i(0)$ for all i (no defiers)

- Implies ITT effect on treatment equals proportion of compliers: $\text{ITT}_D = \pi_{\text{co}}$
- Implies that ITT for the outcome has the same interpretation:

$$\begin{aligned}\text{ITT}_Y &= \pi_{\text{co}} \cdot \text{ITT}_{Y,\text{co}} + \pi_{\text{at}} \cdot \underbrace{\text{ITT}_{Y,\text{at}}}_{=0 \text{ (ER)}} + \pi_{\text{nt}} \cdot \underbrace{\text{ITT}_{Y,\text{nt}}}_{=0 \text{ (ER)}} + \text{ITT}_{Y,\text{df}} \cdot \underbrace{\pi_{\text{df}}}_{=0 \text{ (mono)}} \\ &= \text{ITT}_{\text{co}} \pi_{\text{co}}\end{aligned}$$

- $\rightsquigarrow$ Same identification result: $\tau_{\text{LATE}} = \text{ITT}_Y / \text{ITT}_D$

Is the LATE Useful?

- The LATE is an effect on an unknown subset of the data.
  - Treated units are a mix of always takers and compliers.
  - Control units are a mix of never takers and compliers.
- Without further assumptions, $\tau_{\text{LATE}} \neq \tau$.
- Complier group depends on the instrument $\rightsquigarrow$ different IVs will lead to different identified estimands.
  - But, we cannot do any better in terms of point estimation without more assumptions.

Pseudo-Code: Wald Estimator

```
# *Recall what we did in Neyman's approach*
my_data # data includes Z, D, and Y

# Proportion of compliers (using ITT_D)
pi_co <- mean(my_data$D[my_data$Z == 1]) - mean(my_data$D[my_data$Z == 0])

# Compute ITT's
ITT_Y <- mean(my_data$Y[my_data$Z == 1]) - mean(my_data$Y[my_data$Z == 0])
ITT_D <- mean(my_data$D[my_data$Z == 1]) - mean(my_data$D[my_data$Z == 0])
# (ITT_D = pi_co)

# Wald estimator
Wald_est <- ITT_Y / ITT_D

# Variance (delta method)
Var_ITT_Y <- var(my_data$Y[my_data$Z==1])/sum(my_data$Z==1) +
             var(my_data$Y[my_data$Z==0])/sum(my_data$Z==0)
Var_ITT_D <- var(my_data$D[my_data$Z==1])/sum(my_data$Z==1) +
             var(my_data$D[my_data$Z==0])/sum(my_data$Z==0)

# Covariance term (within treated assignment group)
demeaned_y <- my_data$Y[my_data$Z == 1] - mean(my_data$Y[my_data$Z == 1])
demeaned_d <- my_data$D[my_data$Z == 1] - mean(my_data$D[my_data$Z == 1])
denom <- sum(my_data$Z == 1) * (sum(my_data$Z == 1) - 1)
Covar_est <- sum(demeaned_y * demeaned_d) / denom

# Delta method variance
Var_Wald <- (1/ITT_D^2)*Var_ITT_Y + (ITT_Y^2/ITT_D^4)*Var_ITT_D -
            2*(ITT_Y/ITT_D^3)*Covar_est

# Or use TSLS with AER::ivreg
ivmodel <- AER::ivreg(Y ~ D | Z, data = my_data)
ivpack::robust.se(ivmodel)
```

5/ In-Class Exercise

In-Class Exercise: Noncompliance in Practice

Setting: A pharmaceutical company runs a randomized trial for a new drug. $n = 200$ patients are randomized: 100 assigned to treatment ($Z_i = 1$), 100 to placebo ($Z_i = 0$). However, some patients in the treatment group refuse to take the drug.

	n	Took drug ($\overline{D}$)	Health score ($\overline{Y}$)
Assigned treatment ($Z = 1$)	100	0.70	75
Assigned placebo ($Z = 0$)	100	0.00	70

Exercise Questions

1. Is this one-sided or two-sided noncompliance? Why?
2. Compute $\widehat{\text{ITT}}_Y$, $\widehat{\text{ITT}}_D$, and $\widehat{\tau}_{\text{iv}}$. Interpret each in words.
3. What is $\widehat{\pi}_{\text{co}}$? What proportion of the treatment group are noncompliers?
4. A colleague suggests: "Just compare patients who actually took the drug vs. those who didn't." What is wrong with this?
5. Think of a plausible violation of the exclusion restriction in this setting. Is there a channel from Z to Y that bypasses D?

Onto the presentation!

Contact Information:
jaewon.yoo@iss.nthu.edu.tw
https://j1yoo.github.io/

Appendix

Who Are the Compliers? (Abadie, 2003)

- We cannot identify *which* units are compliers from observed data.
  - Compliance type depends on $(D_i(0), D_i(1))$, both never observed together.
- But we can learn about compliers' **characteristics**.
  - For any pretreatment covariate X_i:

$$\mathbb{E}[X_i \mid \text{complier}] = \frac{\mathbb{E}[X_i \cdot D_i \mid Z_i = 1] - \mathbb{E}[X_i \cdot D_i \mid Z_i = 0]}{\mathbb{E}[D_i \mid Z_i = 1] - \mathbb{E}[D_i \mid Z_i = 0]}$$

  - Same Wald ratio structure, but with $X_i \cdot D_i$ in place of Y_i.
- Why does this matter?
  - **External validity**: are compliers representative of the population of interest?
  - **Policy relevance**: if a program expands, will new participants look like current compliers?
  - Compare complier means to population means to assess.

Noncompliance Dilutes the ITT

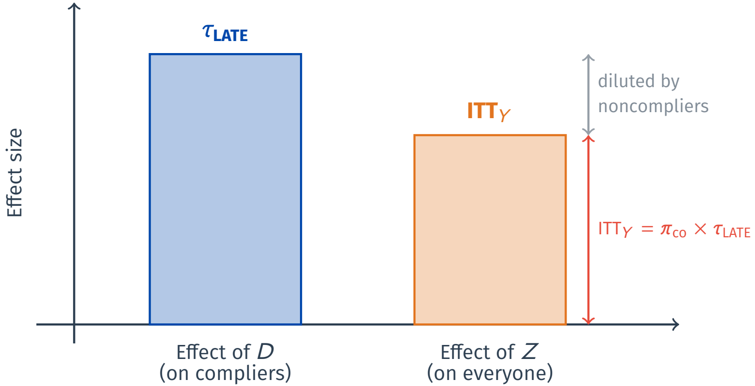

$\tau_{\text{LATE}} = \text{ITT}_Y / \pi_{\text{co}}$: the Wald estimator "undoes" the dilution by dividing by π_{co}

- $\tau_{\text{LATE}} \geq \text{ITT}_Y$ always. Equal only when $\pi_{\text{co}} = 1$ (full compliance).

Numerical Example: God's-Eye View

Unit	$D_i(0)$	$D_i(1)$	Type	$Y_i(0,0)$	$Y_i(1, D_i(1))$	τ_i
1	0	1	Complier	4	7	3
2	0	1	Complier	3	7	4
3	0	0	Noncompl.	5	5	–
4	0	1	Complier	2	7	5
5	0	0	Noncompl.	6	6	–
6	0	1	Complier	1	5	4
7	0	0	Noncompl.	5	5	–
8	0	1	Complier	2	6	4

- One-sided: $D_i(0) = 0$ for all. $\pi_{\text{co}} = 5/8$, $\pi_{\text{nc}} = 3/8$.
- **Noncompliers**: Z has no effect on D or Y. τ_i not defined via D.
- **Compliers**: treatment effect is *only* defined for these units.
  - $\tau_{\text{LATE}} = (3 + 4 + 5 + 4 + 4)/5 = 4.0$

Numerical Example: What We Actually Observe

Suppose Z is randomly assigned: units 1–4 get $Z = 1$, units 5–8 get $Z = 0$.

Unit	Z_i	D_i	Y_i	Type	(Hidden)
1	1	1	7	Complier	
2	1	1	7	Complier	
3	1	0	5	Noncompl.	
4	1	1	7	Complier	
5	0	0	6	Noncompl.	
6	0	0	1	Complier	
7	0	0	5	Noncompl.	
8	0	0	2	Complier	

- $\widehat{\text{ITT}}_Y = \overline{Y}_1 - \overline{Y}_0 = 6.5 - 3.5 = 3.0$
- $\widehat{\text{ITT}}_D = \overline{D}_1 - \overline{D}_0 = 0.75 - 0 = 0.75$
- $\widehat{\tau}_{\text{iv}} = 3.0/0.75 = 4.0 = \tau_{\text{LATE}}$
    - Compliance type is **never observed**, but Wald recovers the LATE!
    - Exact match here is specific to this example. With small n, Wald generally $\neq$ LATE; it converges asymptotically.

IV and Mechanisms

- Notice: in the IV setup, D is an **intermediate outcome** between Z and Y.
    - Same DAG as mediation analysis: $Z \xrightarrow{\beta} D \xrightarrow{\delta} Y$, direct: $Z \xrightarrow{\gamma} Y$.
- Can we use IV to study mechanisms (direct vs. indirect effects)?
    - IV identifies δ only when $\gamma = 0$ (exclusion restriction).
    - But $\gamma = 0$ means *no direct effect* $\rightsquigarrow$ decomposition is trivial.
- When decomposition matters ($\gamma \neq 0$): IV breaks down.
    - Need a separate framework: **causal mediation analysis**.
    - Requires **sequential ignorability** (see Week 5 Notes; covered in detail later in the course).
    - Strong assumption $\rightsquigarrow$ sensitivity analysis is essential (Imai et al., 2011).
- Takeaway: IV tells us *whether* D affects Y, not *how*.

Numerical Example: GOTV Canvassing

- To tie everything together, let's simulate the GOTV canvassing experiment we've been discussing, where we **know the true data-generating process**.
- Setting: $n = 500$ households, randomized to canvassing (Z_i).
  - 60% are compliers: they open the door if canvassed ($\pi_{co} = 0.6$).
  - 40% are noncompliers: they never open the door.
  - One-sided noncompliance: $D_i(0) = 0$ for all.
- Key design choice: compliers are **more civically engaged** at baseline.
  - People who open the door $\neq$ random sample of the population.
  - This is exactly why as-treated analysis fails!
- True LATE $= 5$ percentage points (effect of hearing the message on voter turnout, among compliers).

Simulating the DGP

The code generates data consistent with the GOTV setup:

```r
set.seed(02138); n <- 500
complier <- rbinom(n, 1, 0.6)
Z <- rbinom(n, 1, 0.5)
D <- Z * complier   # one-sided noncompliance

# Compliers have higher baseline turnout (+3 pp)
Y0 <- 10 + 3 * complier + rnorm(n)
Y1 <- Y0 + 5 * complier  # true LATE = 5
Y  <- D * Y1 + (1 - D) * Y0
```

- $Y_i(0)$: baseline turnout. Compliers average 13, noncompliers 10.
- $Y_i(1) - Y_i(0) = 5$ for compliers, 0 for noncompliers.
- After assignment, we observe only (Z_i, D_i, Y_i), not compliance type.

Three Estimators, Three Answers

With the simulated data, we can compute and compare all three approaches:

```
ITT_Y <- mean(Y[Z==1]) - mean(Y[Z==0])   # 2.96
ITT_D <- mean(D[Z==1]) - mean(D[Z==0])   # 0.60
Wald  <- ITT_Y / ITT_D                     # 4.91
as_tr <- mean(Y[D==1]) - mean(Y[D==0])    # 7.82
```

Estimator	**Estimate**	**Truth**	**Interpretation**
$\widehat{\text{ITT}}_Y$	2.96	3.0	Effect of *assignment* on turnout
$\widehat{\text{ITT}}_D$ $(= \widehat{\pi}_{\text{co}})$	0.60	0.6	Proportion of compliers
$\widehat{\tau}_{\text{iv}}$ (Wald)	4.91	5.0	Effect of *canvassing* on compliers
As-treated	7.82	–	Confounded! (+3 from selection)

- As-treated is ≈ 8 instead of 5: the +3 baseline gap inflates the estimate.

Confirming with `ivreg`

The AER package implements 2SLS. With binary Z and D, this is exactly the Wald estimator:

```
library(AER)
iv_model  <- ivreg(Y ~ D | Z, data = my_data)
ols_model <- lm(Y ~ D, data = my_data)
```

	IV (`ivreg`)	**OLS** (`lm`)
Coefficient on D	4.91	7.82
Interpretation	$\widehat{\tau}_{\text{LATE}}$	As-treated (biased)

- `ivreg(Y ~ D | Z)`: instrument Z for endogenous D. Consistent for LATE.
- `lm(Y ~ D)`: treats D as exogenous. Biased when D is confounded.
- Because we know the DGP, we can verify: IV recovers the truth, OLS does not.