# Lecture 9
# Non-compliance in randomized experiments, instrumental variables
# Part I

# Outline

- Assignment, receipt, and the vitamin A experiment
- Principal strata and the IV assumptions
- Complier average effects and the Wald estimator
- IV in observational studies
- Uncertainty: linearization, covariance, and confidence intervals
- Textbook Chapters: Imbens and Rubin Chapters 23.1-23.7, 23.9 & 24.1-24.5, Ding Chapter 21.1-21.2

# Complications in real experiments

- Loss to follow-up: some assigned units have no observed outcome.
- Noncompliance: a control patient may seek a heart transplant outside the study.
- Measurement error: a PCR test may misclassify infection.
- Treatment versions: different doses may have different effects.
- Lack of blinding: patients or doctors may change behavior when they know the assignment.

Our focus is noncompliance. The other issues still require assumptions or analysis.

Does random assignment protect an analysis that drops everyone lost to follow-up?

# Random assignment and actual treatment

Lectures 3–8: random assignment justified comparisons of assigned groups.

Noncompliance:

- Some assigned to treatment refuse or cannot receive it.
- Some assigned to control obtain treatment elsewhere.
- Health, motivation, and access can affect both receipt and outcomes.

We now distinguish assignment $Z$ from treatment received $W$.

If we compare only people who followed their assignment, what protection from randomization might we lose?

# Vitamin A: assignment and the teaching model

- The experiment randomized villages to a vitamin A program or control.
- Each unit in our teaching analysis is a child: $Z_i$ is assignment, $W_i$ is receipt, and $Y_i$ is survival.
- Children in the same village share assignment, but they may differ in whether they receive the supplement.
- None of the individuals assigned to the control group received the supplements

Teaching simplification for the following calculations:

- We ignore the fact that treatment assignment is at the village level (clustered randomized experiment) and consider the experiment as from a completely randomized experiment.

# Vitamin A: observed data

Under the teaching model, $i$ indexes a child and $N$ = 23,682.
$Z$ = assignment, $W$ = vitamin A received, $Y$ = survival (1 = alive).

| Assignment Z | Receipt W | Survived Y = 1 | Died Y = 0 | Total |
|---|---|---|---|---|
| 0 | 0 | 11,514 | 74 | 11,588 |
| 1 | 0 | 2,385 | 34 | 2,419 |
| 1 | 1 | 9,663 | 12 | 9,675 |

- Three binary variables in principle give eight possible $(Z, W, Y)$ combinations.
- Some children assigned to the vitamin A program did not receive the supplement. This is **noncompliance with the assigned treatment**.
- Here no control child receives vitamin A. This is one-sided compliance.

Which two rows already suggest that receipt may be selective?

Sommer & Zeger (1991). The table combines the two survival outcomes within each (Z, W) row.

# Three comparisons, different causal questions

| Analysis | Groups compared | Survival difference |
|---|---|---|
| Intention-to-treat (ITT) | Z = 1 versus Z = 0 | 0.00258 |
| As-treated | W = 1 versus W = 0 | 0.00647 |
| Per-protocol | Z = 1, W = 1 versus Z = 0, W = 0 | 0.00515 |

- ITT estimates the programmatic effect of assignment to the vitamin A program.
- As-treated and per-protocol comparisons select on receipt after assignment and can suffer from selection bias.

Which comparison retains the randomized groups?

# Potential treatment and potential outcomes

- Treatment assignment (randomized encouragement): $Z_i \in \{0,1\}$

- Potential treatment variables: $(W_i(0), W_i(1))$
  - $W_i(z) = 1$: would receive the treatment if $Z_i = z$
  - $W_i(z) = 0$: would not receive the treatment if $Z_i = z$
- Observed treatment received: $W_i^{\text{obs}} = W_i(Z_i)$

- In the non-compliance setting, there are two "treatment": assignment to treatment and receipt of treatment
- Potential outcomes: $Y_i(z, w)$ potential outcome if unit is assigned to $z$ and receive $w$

- Observed outcome: $Y_i^{\text{obs}} = Y_i\big(Z_i, W_i(Z_i)\big)$
- We can also write the potential outcomes as $Y_i(z) = Y_i\big(z, W_i(z)\big)$

# Underlying assumptions

- No interference assumption for $W_i(z)$ and $Y_i(z, w)$
- Randomization of the treatment assignment

$$(Y_i(0,0), Y_i(0,1), Y_i(1,0), Y_i(1,1), W_i(0), W_i(1)) \perp Z_i$$

- We don't have

$$\big(Y_i(0,0), Y_i(0,1), Y_i(1,0), Y_i(1,1)\big) \perp W_i^{\text{obs}}$$

or

$$\big(Y_i(0,0), Y_i(0,1), Y_i(1,0), Y_i(1,1)\big) \perp W_i^{\text{obs}} | Z_i$$

Receipt need not be independent of potential outcomes, even conditional on Z.
We don't know why some units comply and some units don't

- Consistency connects these potential variables to observed receipt and outcomes.

Why is adjusting for received treatment different from adjusting for baseline cholesterol in Lecture 6?

# Two intention-to-treat effects

Effect of assignment on receipt (the first stage):

$$\mathrm{ITT}_W = \frac{1}{N}\sum_{i=1}^{N}[W_i(1) - W_i(0)]$$

Effect of assignment on the outcome:

$$\mathrm{ITT}_Y = \frac{1}{N}\sum_{i=1}^{N}[Y_i(1) - Y_i(0)]$$

If nobody's receipt changes, can assignment still affect outcomes?
Which assumption would rule out that possibility?

# Estimating the ITT effects

$$\widehat{\mathrm{ITT}_W} = \overline{W_1^{\mathrm{obs}}} - \overline{W_0^{\mathrm{obs}}}, \qquad \widehat{\mathrm{ITT}_Y} = \overline{Y_1^{\mathrm{obs}}} - \overline{Y_0^{\mathrm{obs}}}$$

- Under complete randomization, these differences are unbiased for the ITT effects.
- Use the familiar variance formula within assigned groups:

$$\hat{V}(\widehat{\mathrm{ITT}_W}) = \frac{{s_{W,1}}^2}{N_1} + \frac{{s_{W,0}}^2}{N_0}, \qquad \hat{V}(\widehat{\mathrm{ITT}_Y}) = \frac{{s_{Y,1}}^2}{N_1} + \frac{{s_{Y,0}}^2}{N_0}$$

- We can also use regression analyses
- Drawback is that it estimates 'programmatic effectiveness' instead of 'biologic efficacy'

# Principal strata

Principal strata are latent types defined by both potential receipts ($W_i(0)$, $W_i(1)$).

| Type G | W(0) | W(1) | Effect of assignment on receipt |
|---|---|---|---|
| Never-taker (n) | 0 | 0 | 0 |
| Complier (c) | 0 | 1 | +1 |
| Always-taker (a) | 1 | 1 | 0 |
| Defier (d) | 1 | 0 | −1 |

Someone is assigned treatment and takes it. Must that person be a complier?

# Observed receipt only partly reveals the type

General two-sided noncompliance:

| **Observed receipt** | **Z = 0** | **Z = 1** |
|---|---|---|
| W = 0 | Never-taker or complier | Never-taker or defier |
| W = 1 | Always-taker or defier | Always-taker or complier |

One-sided noncompliance: $W_i(0) = 0$ for every unit.
Only compliers and never-takers remain.

Z = 1, W = 1 identifies compliers. Z = 1, W = 0 identifies never-takers.
The Z = 0 group still mixes compliers and never-takers.

1. Suppose we knew every child's true compliance type. Would comparing treatment-assigned and control-assigned children **within the complier group** preserve randomization? Why?
2. In the observed data, can we identify the compliers in each assignment group?

# Decomposing the assignment effects

- Denote the proportion of individuals that fall into each strata as $\pi_c, \pi_a, \pi_n, \pi_d$
  - For one-sided compliance data, $\pi_a = \pi_d = 0$
- Define the average ITT effect for each strata
  - For the treatment received $\mathrm{ITT}_{W,c}, \mathrm{ITT}_{W,a}, \mathrm{ITT}_{W,n}, \mathrm{ITT}_{W,d}$

$$\mathrm{ITT}_{W,c} = 1, \mathrm{ITT}_{W,a} = 0, \mathrm{ITT}_{W,n}=0, \mathrm{ITT}_{W,d} = -1$$

  - For the primary outcome $\mathrm{ITT}_c, \mathrm{ITT}_a, \mathrm{ITT}_n, \mathrm{ITT}_d$

- For the ITT effect on treatment received

$$\mathrm{ITT}_W = \frac{1}{N}\sum_{i=1}^{N} \mathrm{ITT}_{W,i} = \pi_c \mathrm{ITT}_{W,c} + \pi_a \mathrm{ITT}_{W,a} + \pi_n \mathrm{ITT}_{W,n} + \pi_d \mathrm{ITT}_{W,d} = \pi_c - \pi_d$$

- For the ITT effect on primary outcome

$$\mathrm{ITT}_Y = \frac{1}{N}\sum_{i=1}^{N} \mathrm{ITT}_{Y,i} = \pi_c \mathrm{ITT}_c + \pi_a \mathrm{ITT}_a + \pi_n \mathrm{ITT}_n + \pi_d \mathrm{ITT}_d$$

Which terms could disappear if assignment changes receipt only for compliers?

# Assumptions for valid instrumental variables (IV)

| Assumption | What it requires |
|---|---|
| Independence | $Z_i$ is independent of all potential receipts and outcomes. |
| Relevance | Assignment changes the average probability of receipt. |
| Exclusion restriction | At fixed receipt, changing Z does not change the outcome. $Y_i(1, w) = Y_i(0, w)$ |
| Monotonicity | $W_i(1) \geq W_i(0)$ for every unit: no defiers. $\pi_d = 0$ |

With monotonicity, relevance means a positive fraction of compliers: $\pi_c > 0$

Which assumption follows from randomizing Z?
Which still need evidence or substantive reasoning?

# Monotonicity

$$W_i(1) \geq W_i(0) \quad \text{for every unit}$$

- Encouragement may leave receipt unchanged or increase it for each person.
- Requires no defier.
- Structural one-sided access (one-sided compliance), guarantees monotonicity for binary receipt.
- Check the monotonicity assumption in general:
  - $\mathrm{ITT}_W = \pi_c - \pi_d > 0$ if $\pi_d = 0$, so if we can reject the null that $\mathrm{ITT}_W \geq 0$, then monotonicity assumption must fail
  - Otherwise, the monotonicity assumption is not testable
  - Need to decide whether the monotonicity assumption is reasonable or not based on domain knowledge

Suppose 25% are compliers and 5% are defiers. What is $\mathrm{ITT}_W$? Does a positive $\mathrm{ITT}_W$ establish monotonicity?

# Exclusion restriction

$$Y_i(1, w) = Y_i(0, w) \equiv Y_i(w), \qquad w \in \{0,1\}$$

At fixed receipt, assignment has no additional effect on the outcome.

- Always-takers and never-takers keep the same receipt, so their assignment effects are zero.
- Blinding supports exclusion but does not guarantee it.
- Observed data alone cannot establish exclusion. We need subject-matter reasoning.
- An offer may change information, access to other services, or behavior.

In Seguro Popular, could improving health services affect outcomes even for households whose insurance enrollment does not change?

# Identification of the complier average effect

After exclusion, define the complier average causal effect (CACE):

$$\tau_c = \mathrm{E}[Y_i(1) - Y_i(0) \mid W_i(1) > W_i(0)] = \mathrm{ITT}_c$$

Monotonicity and exclusion simplify the two mixtures:

- $\mathrm{ITT}_W = \pi_c, \mathrm{ITT}_a = \mathrm{ITT}_n = 0$
- $\mathrm{ITT}_Y = \pi_c \mathrm{ITT}_c + \pi_a \mathrm{ITT}_a + \pi_n \mathrm{ITT}_n = \pi_c \tau_c$

$$\tau_c = \frac{\mathrm{ITT}_Y}{\mathrm{ITT}_W} \quad \text{(also called the local average treatment effect, LATE)}$$

# Why the effect is local

$$\text{ATE} = \pi_c \tau_c + \pi_a \tau_a + \pi_n \tau_n$$

- IV identifies the effect for people whose receipt changes with this instrument.
- It does not generally identify effects for always-takers or never-takers.
- CACE equals ATE if the average effect among noncompliers also equals CACE.
- With one-sided access and randomized Z, CACE also equals the population ATT.

If 30% are compliers with effect 0.10 and everyone else has effect 0, what are CACE and ATE?

# Method-of-moments (Wald) estimation

- Causal estimand assuming a super population

$$\text{CATE} = \frac{\text{ITT}_Y}{\text{ITT}_W} = \frac{\mathbb{E}(Y_i(1) - Y_i(0))}{\mathbb{E}(W_i(1) - W_i(0))}$$

- Method-of-moment estimator:

$$\hat{\tau}_{\text{IV}} = \frac{\widehat{\text{ITT}}_Y}{\widehat{\text{ITT}}_W} = \frac{\overline{Y_1{}^{\text{obs}}} - \overline{Y_0{}^{\text{obs}}}}{\overline{W_1{}^{\text{obs}}} - \overline{W_0{}^{\text{obs}}}}$$

Simplification under one-sided compliance:

- As $W_i(0) \equiv 0$, we have

$$\widehat{\text{ITT}}_W = \bar{W}_1^{\text{obs}} - \bar{W}_0^{\text{obs}} = \bar{W}_1^{\text{obs}}$$

proportions of units who follow the assignment in the treated group

# Result in Sommer-Zeger Vitamin Supplement data

ITT Estimates:

- $N_1 = 12 + 9663 + 34 + 2385 = 12094$, $N_0 = 74 + 11514 = 11588$
- $\widehat{\mathrm{ITT}}_W = \bar{W}_1^{\mathrm{obs}} = \frac{12+9663}{N_1} = 0.8$
- $\widehat{\mathrm{ITT}}_Y = \frac{2385+9663}{N_1} - \frac{11514}{N_0} = 0.0026$

CATE estimate:

- $\hat{\tau}_{\mathrm{IV}} = \frac{0.0026}{0.8} = 0.0032$
- CACE: about 3.23 additional survivors per 1,000 compliers

| Assignment $Z_i$ | Vitamin Supplements $W_i^{\mathrm{obs}}$ | Survival $Y_i^{\mathrm{obs}}$ | Number of Units ($N = 23{,}682$) |
|---|---|---|---|
| 0 | 0 | 0 | 74 |
| 0 | 0 | 1 | 11,514 |
| 1 | 0 | 0 | 34 |
| 1 | 0 | 1 | 2385 |
| 1 | 1 | 0 | 12 |
| 1 | 1 | 1 | 9663 |

- ITT = 0.8 × CACE: nonreceipt dilutes the assignment effect.
- The as-protocol or as-treated estimates are possibly biased up

| ITT | 0.0026 |
|---|---|
| As-treated | 0.0065 |
| Per-protocol | 0.0052 |

# Uncertainty for a ratio estimator

- Method-of-moment estimator: $\hat{\tau}_{\mathrm{IV}} = \frac{\widehat{\mathrm{ITT}}_Y}{\widehat{\mathrm{ITT}}_W}$

- How to estimate the variance of $\hat{\tau}_{\mathrm{IV}}$?
  - Estimation of $\widehat{\mathrm{ITT}}_Y$ and $\widehat{\mathrm{ITT}}_W$ are correlated because they use the same dataset

- When the number of units $N$ is large
  - $\widehat{\mathrm{ITT}}_Y$ and $\widehat{\mathrm{ITT}}_W$ are close to the true values $\mathrm{ITT}_Y$ and $\mathrm{ITT}_W$

$$\widehat{\mathrm{ITT}}_Y = \mathrm{ITT}_Y + O\left(\frac{1}{\sqrt{N}}\right), \qquad \widehat{\mathrm{ITT}}_W = \mathrm{ITT}_W + O\left(\frac{1}{\sqrt{N}}\right)$$

  - Perform a first-order Taylor expansion of $\hat{\tau}_{\mathrm{IV}}$ at $\mathrm{ITT}_Y$ and $\mathrm{ITT}_W$:

$$\frac{\widehat{\mathrm{ITT}}_Y}{\widehat{\mathrm{ITT}}_W} = \frac{\mathrm{ITT}_Y}{\mathrm{ITT}_W} + \frac{1}{\mathrm{ITT}_W}\left(\widehat{\mathrm{ITT}}_Y - \mathrm{ITT}_Y\right) - \frac{\tau_c}{\mathrm{ITT}_W}\left(\widehat{\mathrm{ITT}}_W - \mathrm{ITT}_W\right) + O\left(\frac{1}{N}\right)$$

Why is it usually wrong to treat the estimated first stage as a fixed number when computing the standard error?

# The delta-method variance

$$\mathrm{Var}(\hat{\tau}_{\mathrm{IV}}) \approx \frac{\mathrm{Var}(\widehat{\mathrm{ITT}_Y}) + {\tau_c}^2 \mathrm{Var}(\widehat{\mathrm{ITT}_W}) - 2\tau_c \mathrm{Cov}(\widehat{\mathrm{ITT}_Y}, \widehat{\mathrm{ITT}_W})}{{\mathrm{ITT}_W}^2}$$

The numerator is the variance of a residualized assignment contrast:

$$\mathrm{Var}(\hat{\tau}_{\mathrm{IV}}) \approx \frac{\mathrm{Var}(\widehat{\mathrm{ITT}_Y - \tau_c \mathrm{ITT}_W})}{{\mathrm{ITT}_W}^2}$$

If $\mathrm{ITT}_W$ is halved while the residual variance stays the same, what happens approximately to the IV standard error?

# Estimating the covariance

Compute the sample covariance within each assigned group Z = z:

$$s_{YW,z} = \frac{1}{N_z - 1} \sum_{i:Z_i=z} (Y_i^{\text{obs}} - \overline{Y_z^{\text{obs}}})(W_i^{\text{obs}} - \overline{W_z^{\text{obs}}})$$

The covariance between the two ITT estimators is estimated by:

$$\widehat{\text{Cov}}(\widehat{\text{ITT}_Y}, \widehat{\text{ITT}_W}) = \frac{s_{YW,1}}{N_1} + \frac{s_{YW,0}}{N_0}$$

This equality uses independent sampled units. The finite-population distinction appears in the reference slide.

# Plug-in variance and interval

Substitute the estimated ITT variances, covariance, and complier effect:

$$\hat{V}(\hat{\tau}_{\mathrm{IV}}) = \frac{\hat{V}(\widehat{\mathrm{ITT}_Y}) + \hat{\tau}_{\mathrm{IV}}{}^2 \hat{V}(\widehat{\mathrm{ITT}_W}) - 2\hat{\tau}_{\mathrm{IV}}\widehat{\mathrm{Cov}}(\widehat{\mathrm{ITT}_Y}, \widehat{\mathrm{ITT}_W})}{\widehat{\mathrm{ITT}_W}^2}$$

$$\text{Approximate 95\% CI: } \hat{\tau}_{\mathrm{IV}} \pm 1.96\sqrt{\hat{V}(\hat{\tau}_{\mathrm{IV}})}$$

The normal approximation requires large samples and a sufficiently strong $\mathrm{ITT}_W$.

- One-sided noncompliance: W = 0 in controls, hence $s_{W,0}{}^2 = s_{YW,0} = 0$

$$\hat{V}(\widehat{\mathrm{ITT}_W}) = \frac{\overline{W_1}^{\mathrm{obs}}(1 - \overline{W_1}^{\mathrm{obs}})}{N_1 - 1}, \qquad \widehat{\mathrm{Cov}}(\widehat{\mathrm{ITT}_Y}, \widehat{\mathrm{ITT}_W}) = \frac{s_{YW,1}}{N_1}$$

Does that mean the control group contributes no uncertainty to the IV estimate?

# Vitamin A: estimated effects and uncertainty

ITT Estimates:

- $N_1 = 12094, N_0 = 11588$
- $\widehat{\mathrm{ITT}}_W = \bar{W}_1^{\mathrm{obs}} = \frac{12+9663}{N_1} = 0.8, \widehat{\mathbb{V}}\left(\widehat{\mathrm{ITT}}_W\right) = \frac{\bar{W}_1^{\mathrm{obs}}(1-\bar{W}_1^{\mathrm{obs}})}{N_1-1} = \frac{0.2*0.8}{12093} = 0.0036^2$
- $\widehat{\mathrm{ITT}}_Y = \frac{2385+9663}{N_1} - \frac{11514}{N_0} = 0.0026, \widehat{\mathbb{V}}\left(\widehat{\mathrm{ITT}}_Y\right) = \sum_{z=0}^{1} \frac{\bar{Y}_z^{\mathrm{obs}}(1-\bar{Y}_z^{\mathrm{obs}})}{N_z-1} = 0.0009^2$
- 95% CI of $\mathrm{ITT}_Y$: $(0.0008, 0.0044)$

CATE estimate:

- $\hat{\tau}^{iv} = \frac{0.0026}{0.8} = 0.0032$
- $\widehat{\mathrm{Cov}}\left(\widehat{\mathrm{ITT}}_W, \widehat{\mathrm{ITT}}_Y\right) = 0.0000017$ (correlation 0.05)
- $\widehat{\mathbb{V}}\left(\hat{\tau}^{iv}\right) = 0.0012^2$
- 95% CI of CATE: $(0.0010, 0.0055)$

| Assignment $Z_i$ | Vitamin Supplements $W_i^{\mathrm{obs}}$ | Survival $Y_i^{\mathrm{obs}}$ | Number of Units ($N = 23{,}682$) |
|---|---|---|---|
| 0 | 0 | 0 | 74 |
| 0 | 0 | 1 | 11,514 |
| 1 | 0 | 0 | 34 |
| 1 | 0 | 1 | 2385 |
| 1 | 1 | 0 | 12 |
| 1 | 1 | 1 | 9663 |

# What IV adds to the causal analysis

- Randomization identifies assignment effects even when people do not comply.
- The IV assumptions turn an assignment contrast into an effect of receipt for compliers.
- The same logic can address observational treatment selection when a credible instrument exists.
- The ratio requires uncertainty for both estimated contrasts and their covariance.

1. What remains identified if exclusion restriction fails in a randomized trial?
2. Does a strong $\mathrm{ITT}_W$ establish that an instrument is valid?
3. What new justification is needed in an observational study?

# IV beyond randomized experiments

- With noncompliance, receipt W is not randomized even when assignment Z is.
- In observational studies, seek an instrument Z that changes receipt and is independent of potential receipts and outcomes, possibly conditional on baseline covariates.
- Exclusion and relevance remain necessary. Monotonicity gives a complier-average interpretation.

Lecture 10: the draft lottery, covariate adjustment, and weak instruments.

What would justify treating (Z) as if randomly assigned, possibly within groups defined by baseline covariates?