J-PAL
ABDUL LATIF JAMEEL POVERTY ACTION LAB

Threats and Analysis

Course Overview

1. Why Evaluate
2. Theory of Change & Measurement
3. Why & When to Randomize
4. How to Randomize
5. Sample Size & Power
6. Randomized Evaluation from Start to Finish
7. Threats & Analysis
8. Ethical Considerations
9. Generalizing & Applying Evidence

Learning Objectives

- Identify the main **threats to validity** that can arise while implementing an intervention and evaluation
  - Main focus is internal validity (whether the estimated impact reflects a causal relationship between the treatment and the outcome)
- Discuss strategies to mitigate these threats during the **implementation phase**
- Learn some strategies to account for threats during the **analysis phase**

Introduction

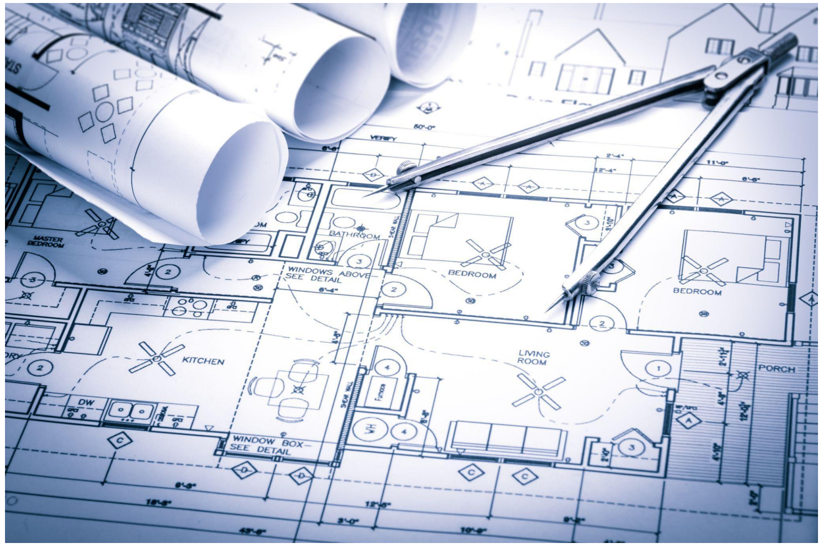

Photo credit: Shutterstock.com

During the **conception phase,** we design an evaluation that enables us to answer our research questions


Photo credit: Shutterstock.com

But the **implementation phase** of the evaluation is also extremely important: many things can go wrong

Lecture Overview

- Threats to validity
  - Spillovers
  - Attrition
  - Evaluation-driven effects
  - Partial compliance
- Generating impact estimates
  - Intention to Treat
  - Local Average Treatment Effect
  - Reporting results

Lecture Overview

- **Threats to validity**
  - **Spillovers**
  - Attrition
  - Evaluation-driven effects
  - Partial compliance
- Generating impact estimates
  - Intention to Treat
  - Local Average Treatment Effect
  - Reporting results

Reminder from How to Randomize Lecture: Spillovers

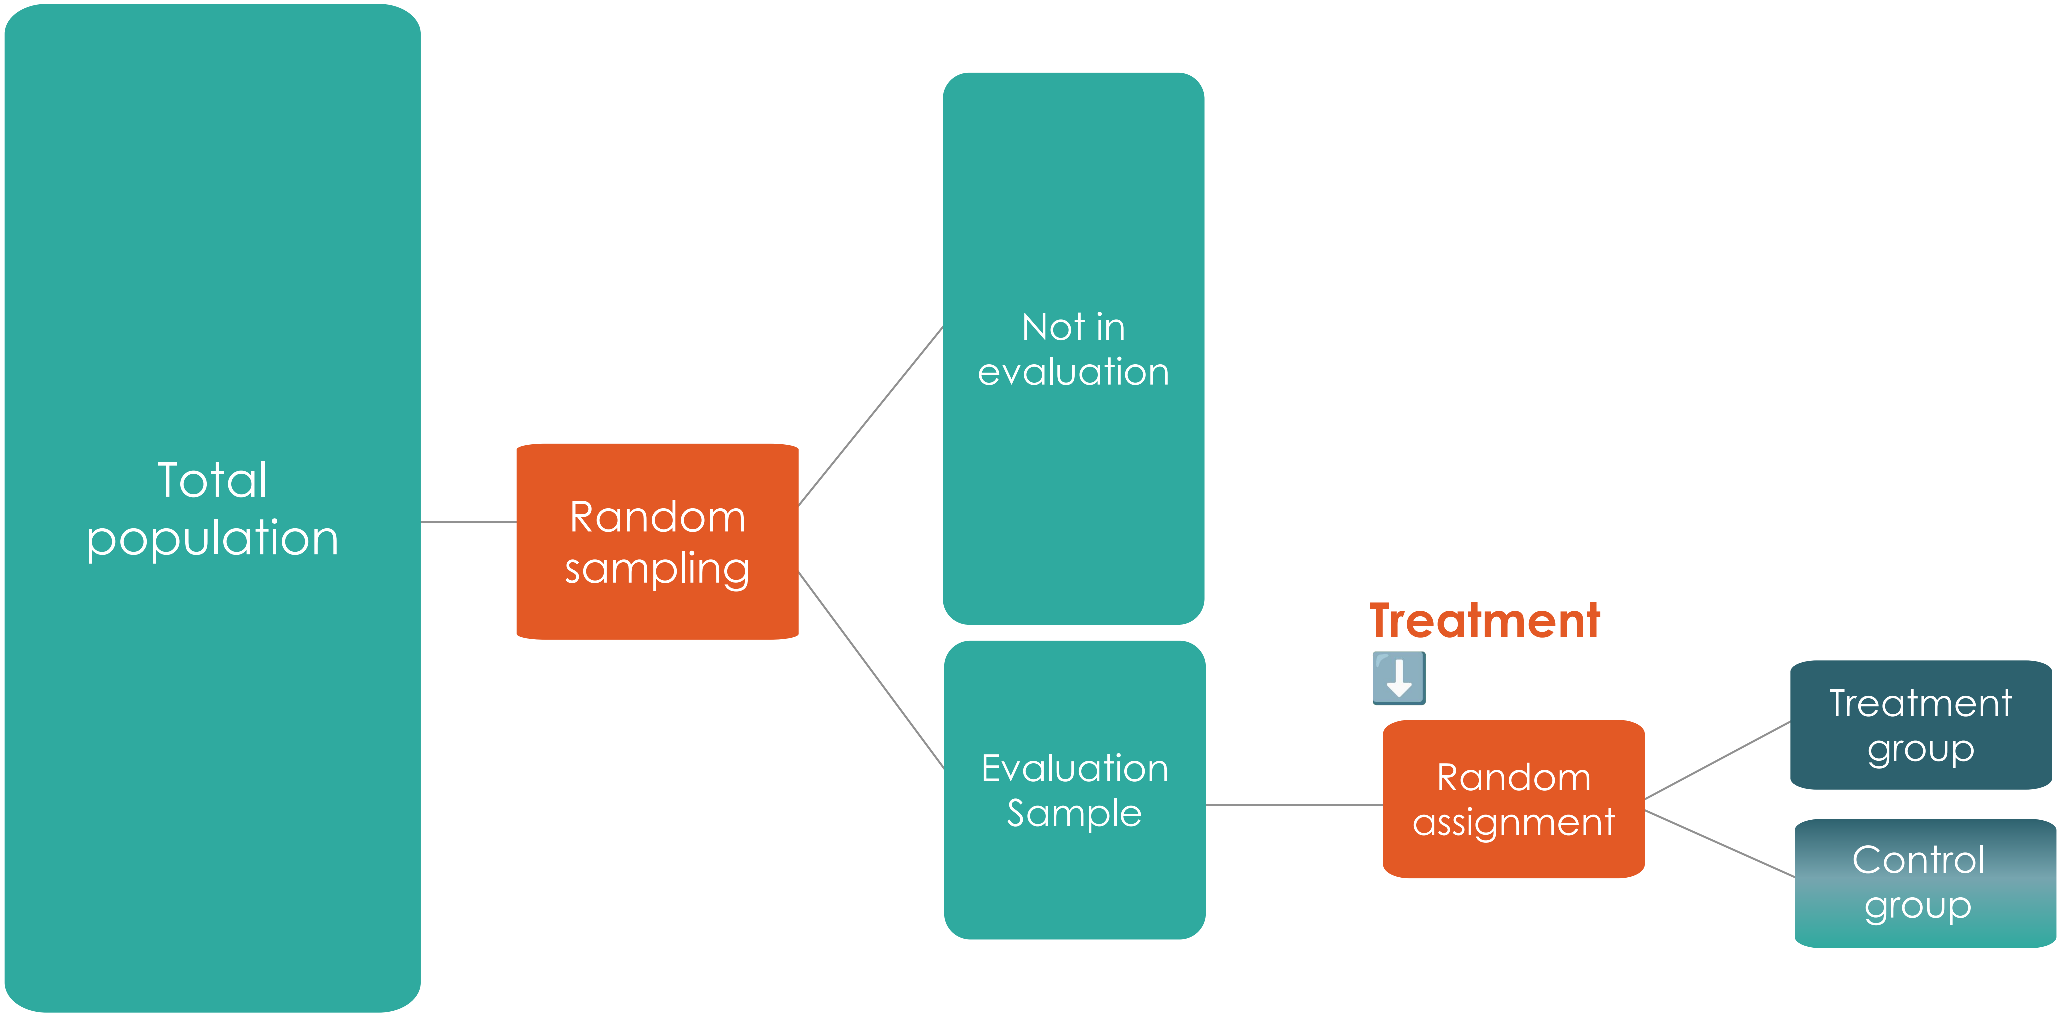

Spillovers

Spillovers occur when the outcomes of untreated units are indirectly affected by the treatment given to others.

- Spillovers **violate the key assumption** that one unit's treatment assignment has no effect on the outcomes of other units
- Spillovers are **not limited to subjects in the study sample**, but can affect anyone who is not treated
- Common causes: geographic proximity, social networks
- Make it difficult or impossible to measure the impact of the program
  - Comparison group no longer serves as a valid estimate of the counterfactual

Spillovers - Outcomes

- Spillovers may not put a study in jeopardy if they are contained or measured, but are problematic if they affect the comparison group
- Spillovers can be positive or negative
    - ➕ **Positive spillovers:** comparison group **benefits** from treatment group
    - ➖ **Negative spillovers:** comparison group is **harmed** by treatment group
- Spillovers can cause impact to be **underestimated <u>or</u> overestimated**
- Channels through which spillovers occur include physical, informational/behavioral, and marketwide/general equilibrium

Physical Spillover

Example: Cash transfer program

A member of the treatment group receives a cash transfer and gives some of the money to friends or relatives who are assigned to the control group

Control group benefits from treatment group

Treatment

Treatment group

Control group

Behavioral/Informational Spillover

Example: Handwashing promotion campaign

Control group imitates neighbors' hygiene practices or learns about the health benefits of handwashing

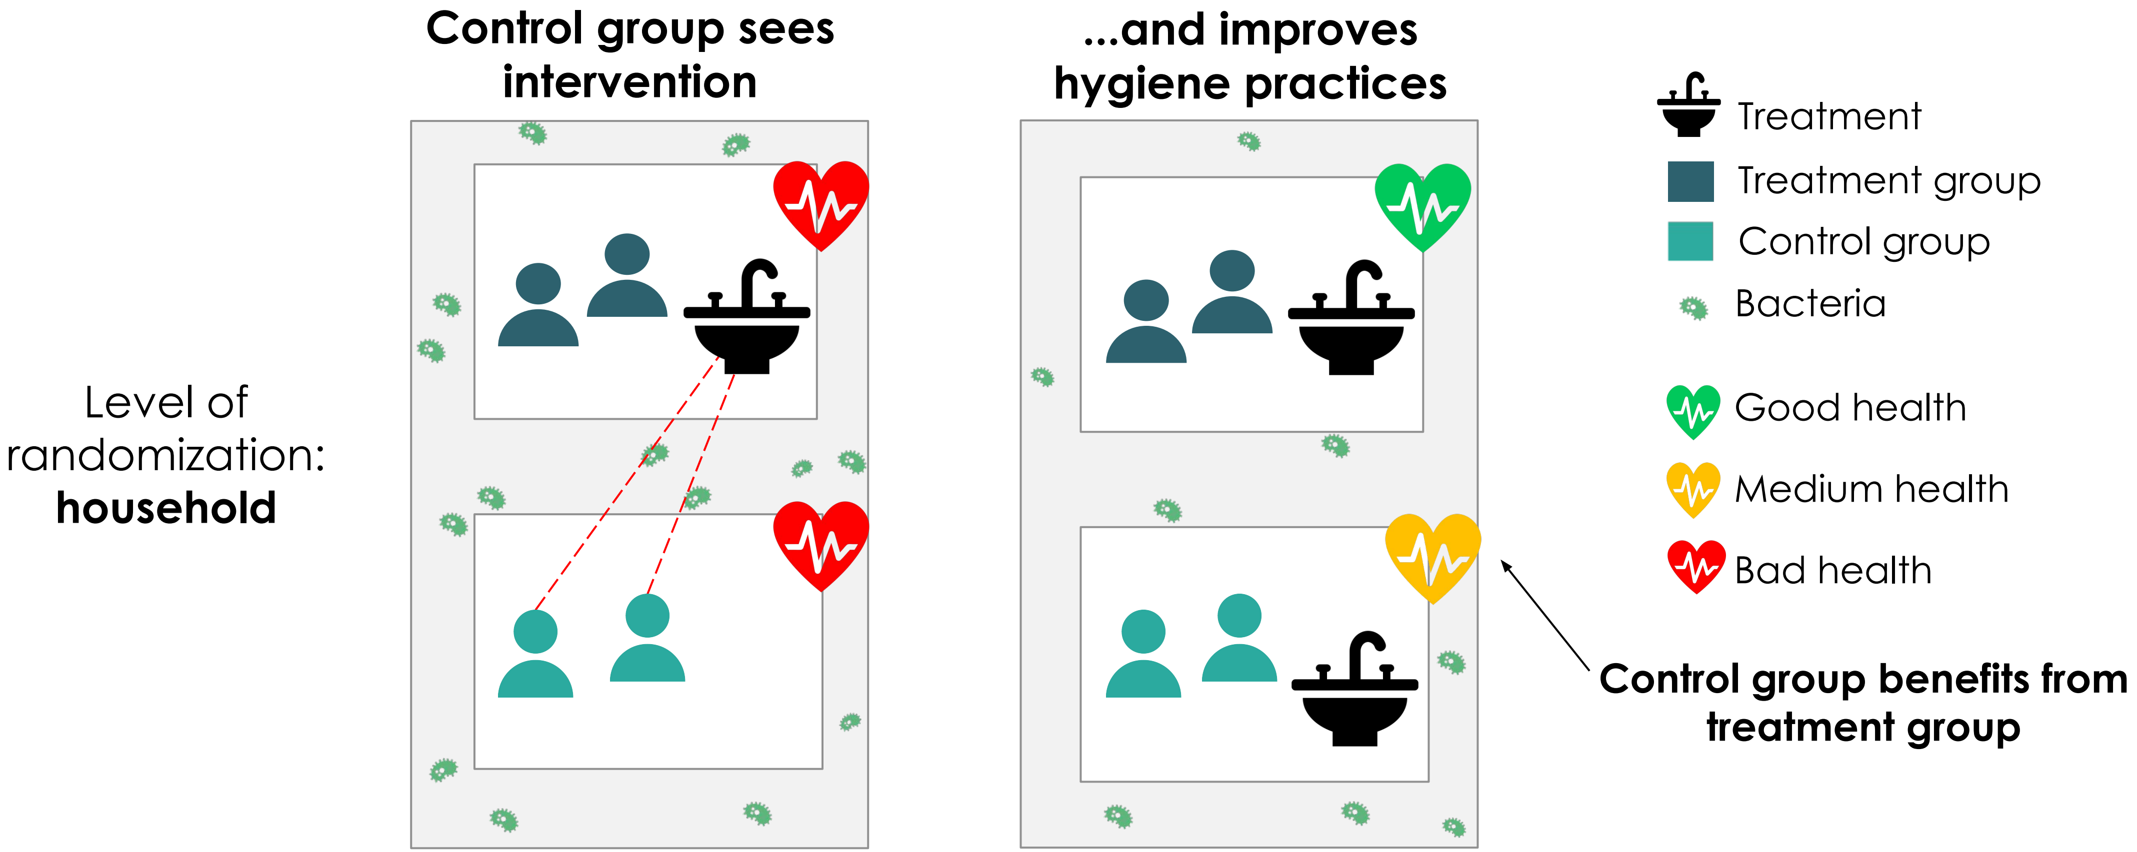

Marketwide/General Equilibrium Effects Spillover

Example: Microcredits

- **Program:** lending small amounts of money at low interest rates.

- **Threat:** delivering an intervention at scale can generate spillovers (positive or negatives) control group may affect the estimation of the impact of our program

  ➜ E.g., if we have positive effects, the spillovers can attenuate or even "wash-out" the effects

- Equilibrium effects can be relevant for policy decisions

  ➜ E.g., decisions regarding the regulatory treatment of microcredit

- **Outcome of interest:** impact of a microcredit over profits

Marketwide/General Equilibrium Effects Spillover

Example: Access to microcredit for entrepreneurs

Control group entrepreneurs are in competition with treatment group for recruiting from a limited number of workers in the local economy

Without intervention

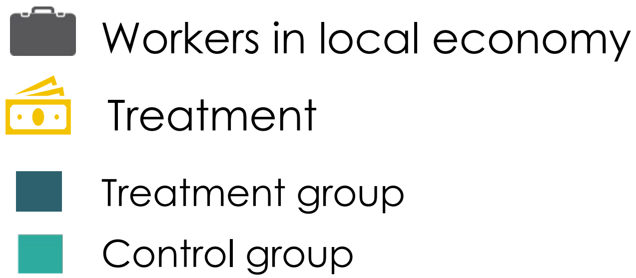

With intervention (if displacement occurs)

Control group is 'harmed' by treatment group

What can be done about spillovers?

Measure spillovers

- Build plans to **collect data** on spillovers into the experimental design
- **Measure** spillovers in the analysis phase

Avoid spillovers

- Choose **level of randomization** wisely, and randomize at a higher level if concerned about spillovers
- Incorporate **spatial buffers** between treatment and control units

Avoiding spillovers: Randomize at a different level

Randomizing at village level contains the positive spillovers within the treated households

Level of randomization: **village**

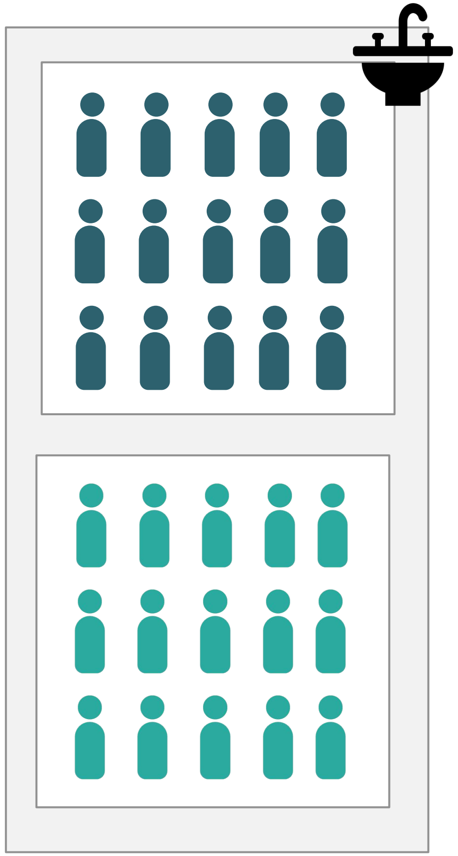

Intervention: **handwashing promotion campaign**

Treatment group

Control group

Avoiding spillovers: Build in buffers

Providing a buffer between the treatment and control subjects prevents treatment from spilling into control

Level of randomization: **individual**

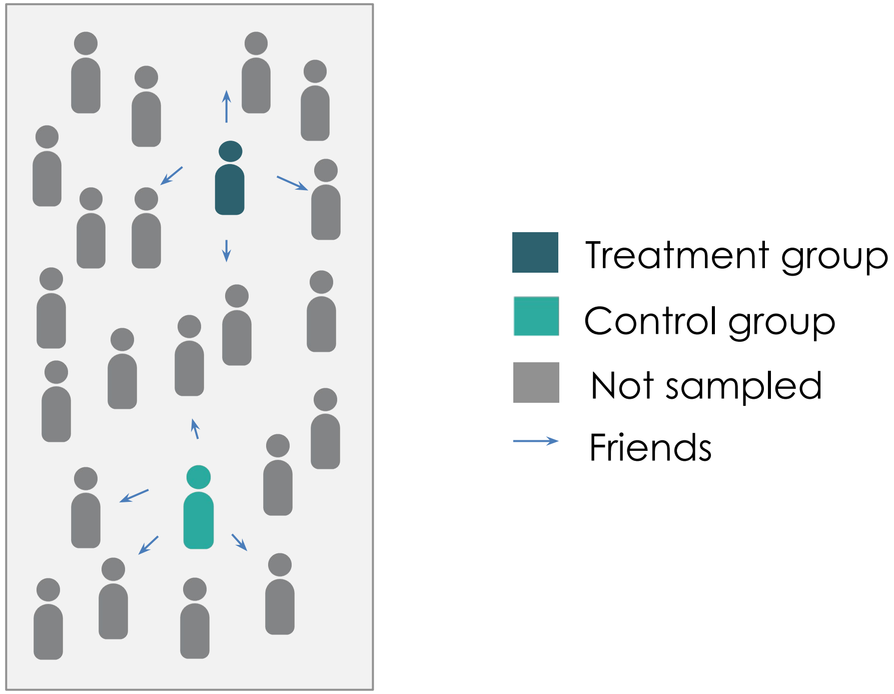

Thought exercise: Measuring informational spillovers

Imagine you are designing a randomized evaluation of a television program that features educational storylines about HIV/AIDs to understand the impact on viewers' knowledge, attitudes, and behaviors.

- **How could you design the evaluation to measure knowledge and behavior changes for viewers of the program—as well as the potentially positive informational spillovers to peers within their social networks?**

To learn more about the results of an HIV/AIDs edutainment intervention in Nigeria, **see** Banerjee, La Ferrara, and Orozco (2019), "The Entertaining Way to Behavioral Change: Fighting HIV with MTV."

Lecture Overview

- **Threats to validity**
  - Spillovers
  - **Attrition**
  - Evaluation-driven effects
  - Partial compliance
- Generating impact estimates
  - Intention to Treat
  - Local Average Treatment Effect
  - Reporting results

Attrition

Attrition occurs when study group members leave the study and data on their outcomes cannot be collected.

Discussion question: Why is it a problem if some of the people in the experiment leave the study before you finish collecting your data? Why might we expect this to happen?

- It may be a problem depending on how much of the study sample we lose
- It is a problem if the type of people who leave is correlated with the treatment
- Common drivers of attrition include mobility or migration, motivation, and mortality

Example: Microcredit

- Microfinance institutions expanded rapidly in recent decades and microcredit generated considerable enthusiasm and hope for fast poverty alleviation

- **Intervention:** Group-lending microcredit program for women in Hyderabad, India

Pair 104 identified neighborhoods based on per capita consumption and per-household debt

Pairwise randomization

Treatment
(52 neighborhoods)

Control
(52 neighborhoods)

- **Data collection:** Information collected at the household level from a random sample

Household business profits in Rupees before and after the microcredit program

	Before Treatment		After Treatment	
	T	C	T	C
	2,000	2,000	2,200	2,000
	2,500	2,500	2,700	2,500
	3,000	3,000	3,200	3,000
Avg.				
	Difference:		Difference:	

Household business profits in Rupees before and after the microcredit

	Before Treatment		After Treatment	
	T	C	T	C
	2,000	2,000	2,200	2,000
	2,500	2,500	2,700	2,500
	3,000	3,000	3,200	3,000
Avg.	2,500	2,500	2,700	2,500
	Difference:	0	Difference:	200

What if the households with the lowest income migrated to other regions?

Household business profits in Rs before and after the microcredit

Before Treatment		After Treatment	
T	C	T	C
2,000	2,000	2,200	[absent]
2,500	2,500	2,700	2,500
3,000	3,000	3,200	3,000

A. You will underestimate the impact
B. You will overestimate the impact
C. Neither
D. Ambiguous
E. Don’t know

What if the most disadvantaged households migrated to other regions?

Household business profits in Rs before and after the microcredit

	Before Treatment		After Treatment	
	T	C	T	C
	2,000	2,000	2,200	[absent]
	2,500	2,500	2,700	2,500
	3,000	3,000	3,200	3,000
Avg.	2,500	2,500	2,700	2,750
	Difference:	0	Difference:	-50

What can be done about attrition?

Implementation phase

- More **intensive follow-up** efforts with survey respondents
  - Account for follow-up costs in project planning and funding
  - For example: Tracking of respondents who moved to neighboring areas

Analysis phase

- Use **bounded estimates** to mitigate the effects of attrition on impact estimates
  - Bounded estimates: take the percentage difference between treatment and control and drop the top percentile and bottom percentile from the group with less attrition to bound the estimates, creating worst case and best case scenarios

When is attrition NOT a problem?

A. When the attrition rates are similar in both treatment and control groups
B. When the estimated treatment effect is zero (among those who remain in the study)
C. When the true treatment effect is zero
D. None of the above

When is attrition NOT a problem?

A. When the attrition rates are similar in both treatment and control groups

B. When the estimated treatment effect is zero (among those who remain in the study)

C. When the true treatment effect is zero

D. None of the above

Lecture Overview

- **Threats to validity**
  - Spillovers
  - Attrition
  - **Evaluation-driven effects**
  - Partial compliance
- Generating impact estimates
  - Intention to Treat
  - Local Average Treatment Effect
  - Reporting results

Evaluation-driven effects

Evaluation-driven effects occur when respondents change their behavior in response to the evaluation itself instead of the intervention.

Common causes: salience of being evaluated, social pressure

These include observer-driven effects and enumerator effects:

- **Hawthorne effects:** Behavior changes due to attention from the study or intervention
- **Anticipation effects:** Comparison group changes behavior because they expect to receive the treatment later (particular concern for phase-ins)
- **Resentment/demoralization effects:** Comparison group resents missing out on treatment and changes behavior
- **Demand effects:** Behavior changes due to perceptions of evaluator's objectives
- **Survey effects:** Being surveyed changes subsequent behavior

What can be done about evaluation-driven effects?

Evaluation design

- Use a **different level** of randomization
- Measure the evaluation-driven effects in a **subset** of the sample
  - Prime a subset of the sample by reminding them of the evaluation (e.g., Mummolo and Peterson 2019)
  - Supplement survey data with other measures of behavioral outcomes (e.g., Fearon, Humphreys, and Weinstein 2008)

Implementation phase

- **Minimize salience** of evaluation as much as possible
  - Do not announce phase-in
    - Downside is that this can be useful to reduce attrition!
  - Make sure staff are impartial and treat both groups similarly
    - E.g., do not share treatment assignment with data collection staff

Thought exercise: Feedback to teachers

Imagine you are designing a randomized evaluation of a program that provides feedback to teachers (based on students' testing performance) to help understand the impact on teacher effort and ultimately student learning outcomes. However, classroom observation and the presence of enumerators to measure teacher activity may drive teachers' behavior, rather than the treatment itself.

How could you disentangle program effects from potential Hawthorne effects?

Reminder: Hawthorne effects are behavior changes due to attention of the study or intervention.

To learn more about the results of a teacher feedback intervention in India, see Muralidharan and Sundararaman (2010), "The Impact of Diagnostic Feedback to Teachers on Student Learning."

Lecture Overview

- **Threats to validity**
  - Spillovers
  - Attrition
  - Evaluation-driven effects
  - **Partial compliance**
- Generating impact estimates
  - Intention to Treat
  - Local Average Treatment Effect
  - Reporting results

Partial Compliance

Noncompliance occurs when a unit's treatment assignment (to a treatment or comparison group) does not match their treatment status (took up or did not take up the program)

- ➔ Individuals assigned to the treatment group may not take up the program
- ➔ Individuals assigned to the comparison group may access the program
- ➔ Can be due to project **implementers** or the **participants** themselves

		Status (reality)	
		Treated	Not treated
Assignment	Treatment	Compliers	Non-compliers
	Control	Non-compliers	Compliers

When some participants are non-compliers, we say there is **partial compliance** in our study.

A study sample can be split into three distinct groups

Compliers	Always-takers	Never-takers
Follow their treatment assignment	Always take the treatment, even if assigned to the control group	Always refuse treatment, even if assigned to the treatment group

➜ We have the underlying assumption that there are no defiers,a fourth group who do the opposite of, or defy, their treatment assignment.

Potential Sources of Noncompliance

- Logistical or political challenges: For example, service providers may find it difficult to administer customized treatment alongside their other responsibilities
- Service providers might have trouble distinguishing between treatment and control, or may be unwilling to provide differential treatment

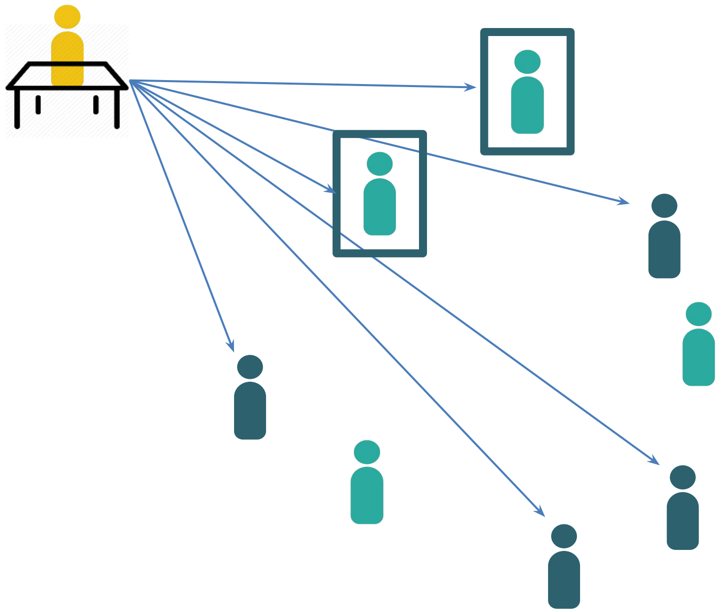

Service provider

Participant assigned to treatment group

Participant assigned to control group

Service provider treats this participant

Members of the control group who "cross over" into the treatment group and receive intervention

Sample Selection Bias

Noncompliance can lead to **sample selection bias**

This threatens internal validity if not properly accounted for in the analysis

➔ **Selection bias** occurs when individuals who receive or opt into the program are systematically different from those who do not

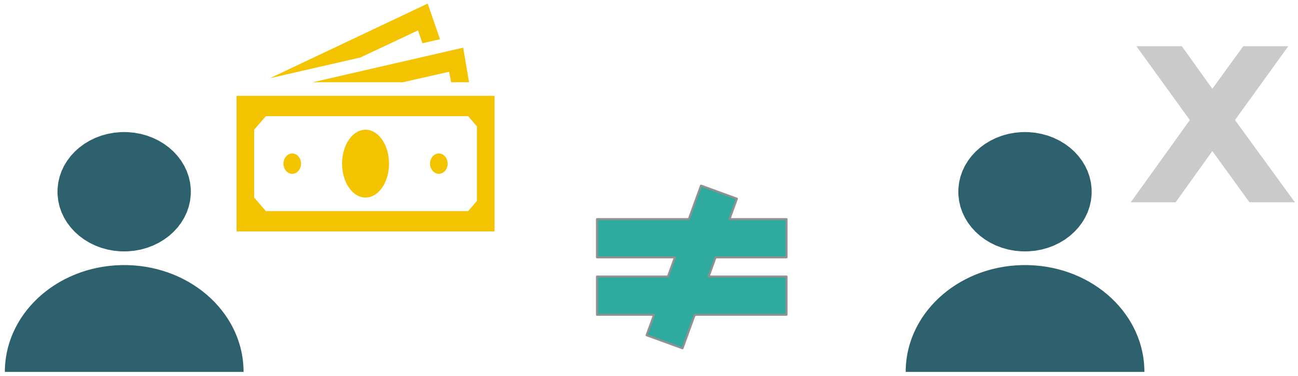

Entrepreneurs who take up microcredit

Entrepreneurs who do not take up microcredit

Non-compliers

What can you do?

Can you switch them?

No!

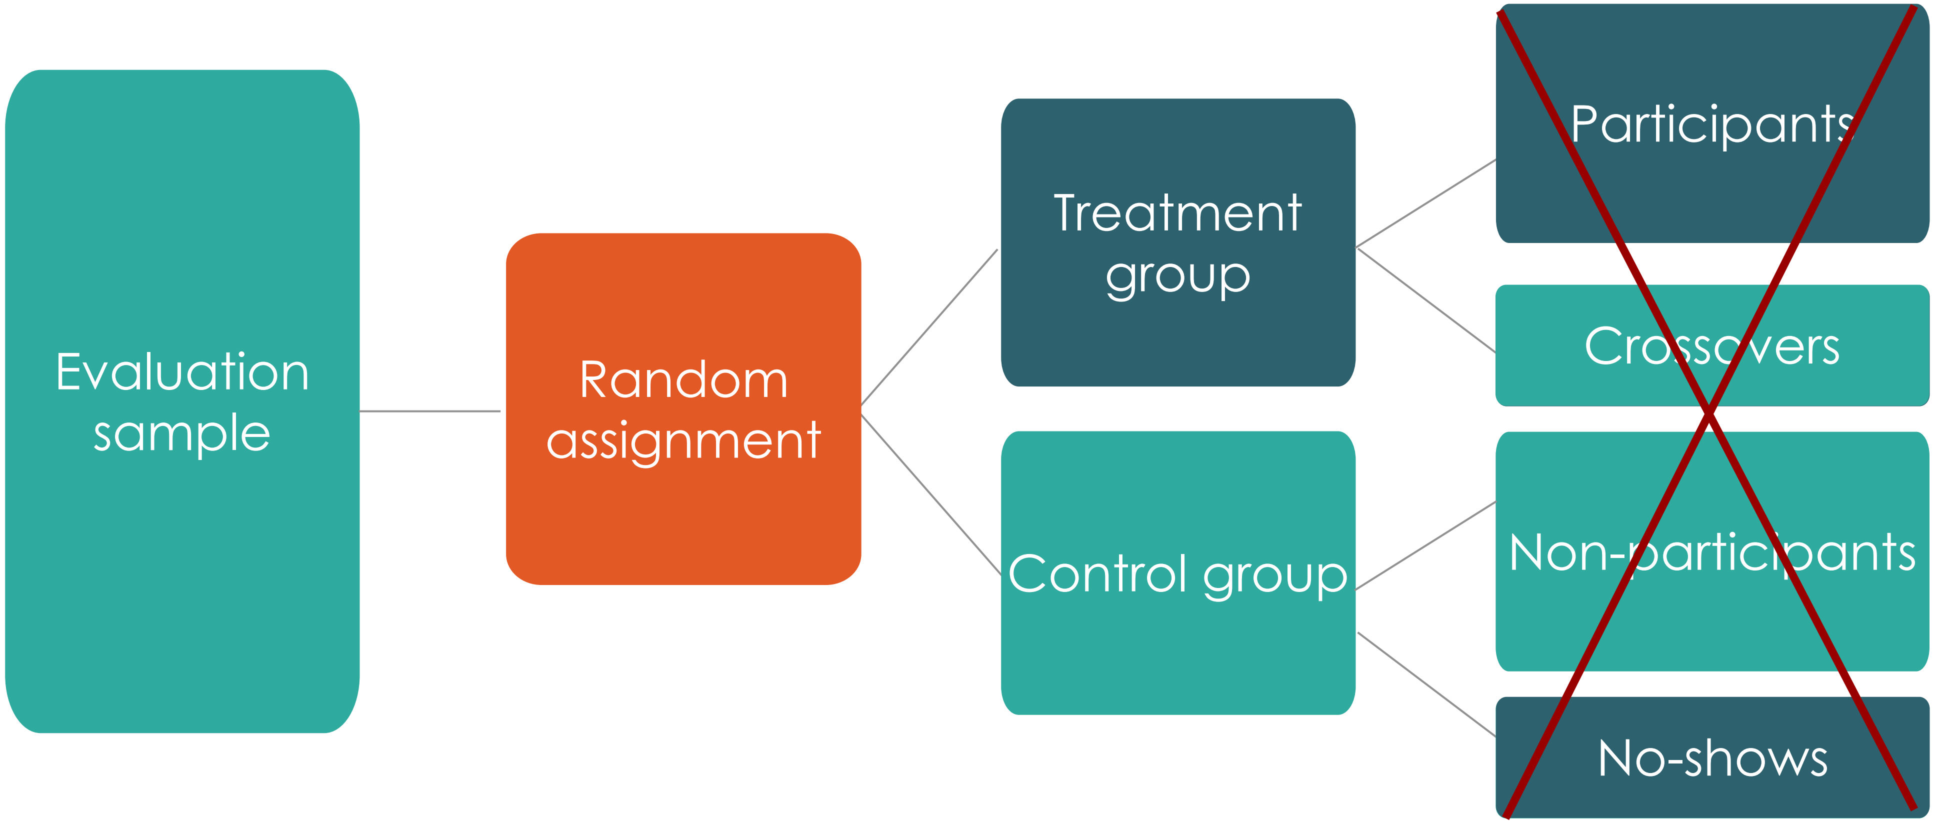

Non-compliers

What can you do?

Can you drop them?

No!

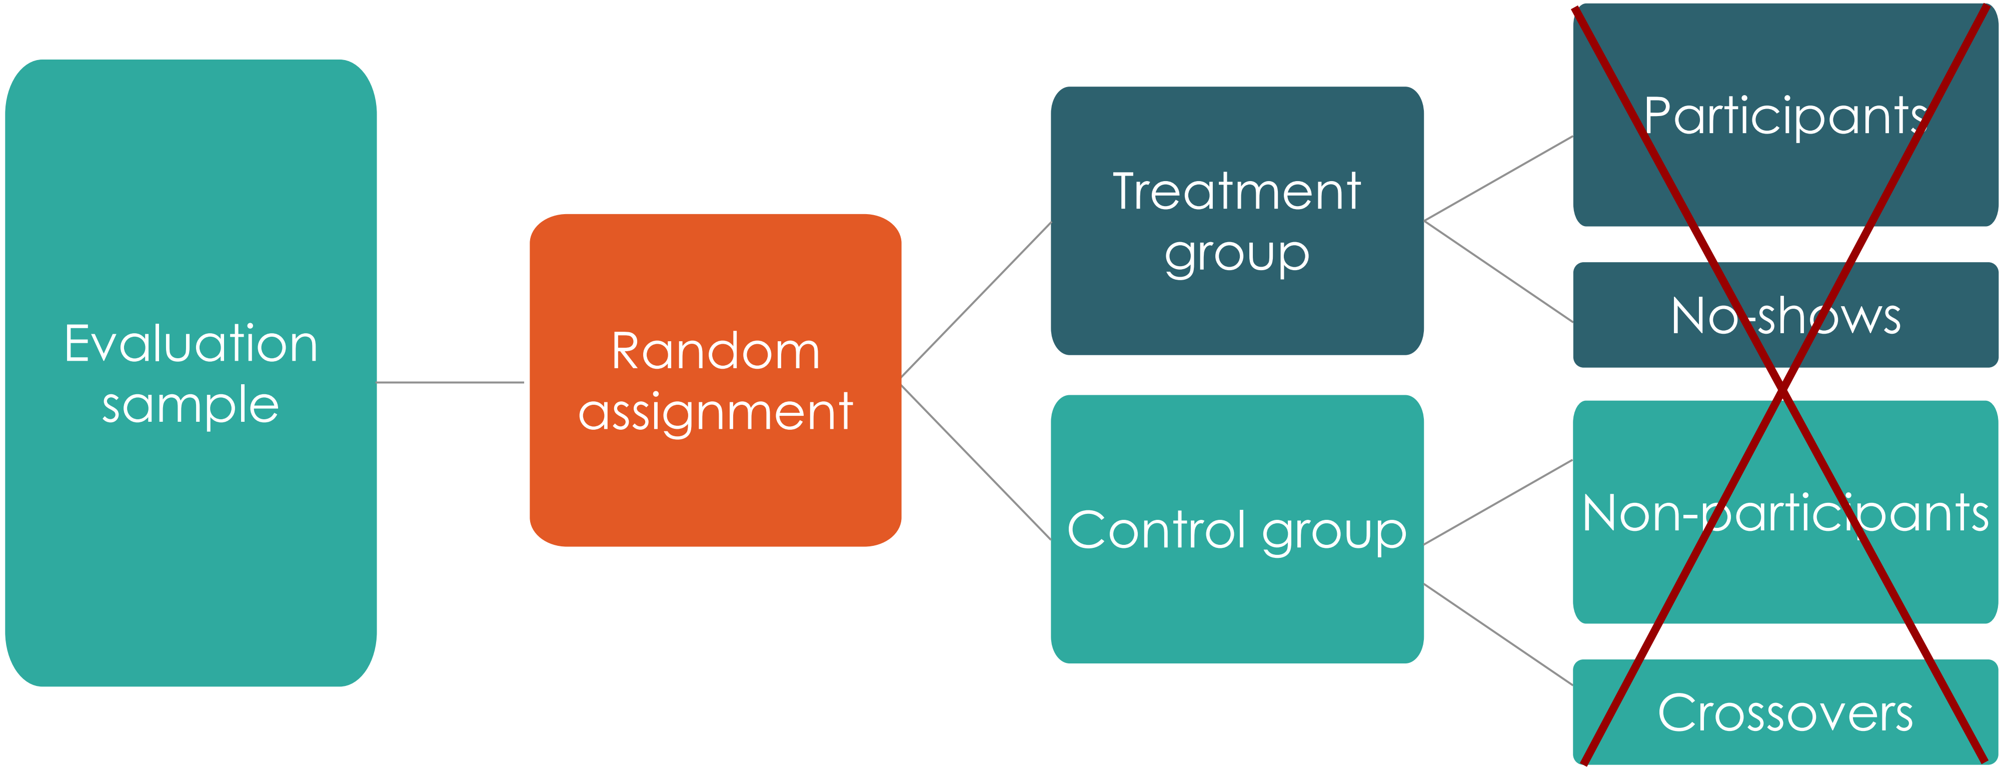

Non-compliers

You can compare the original groups

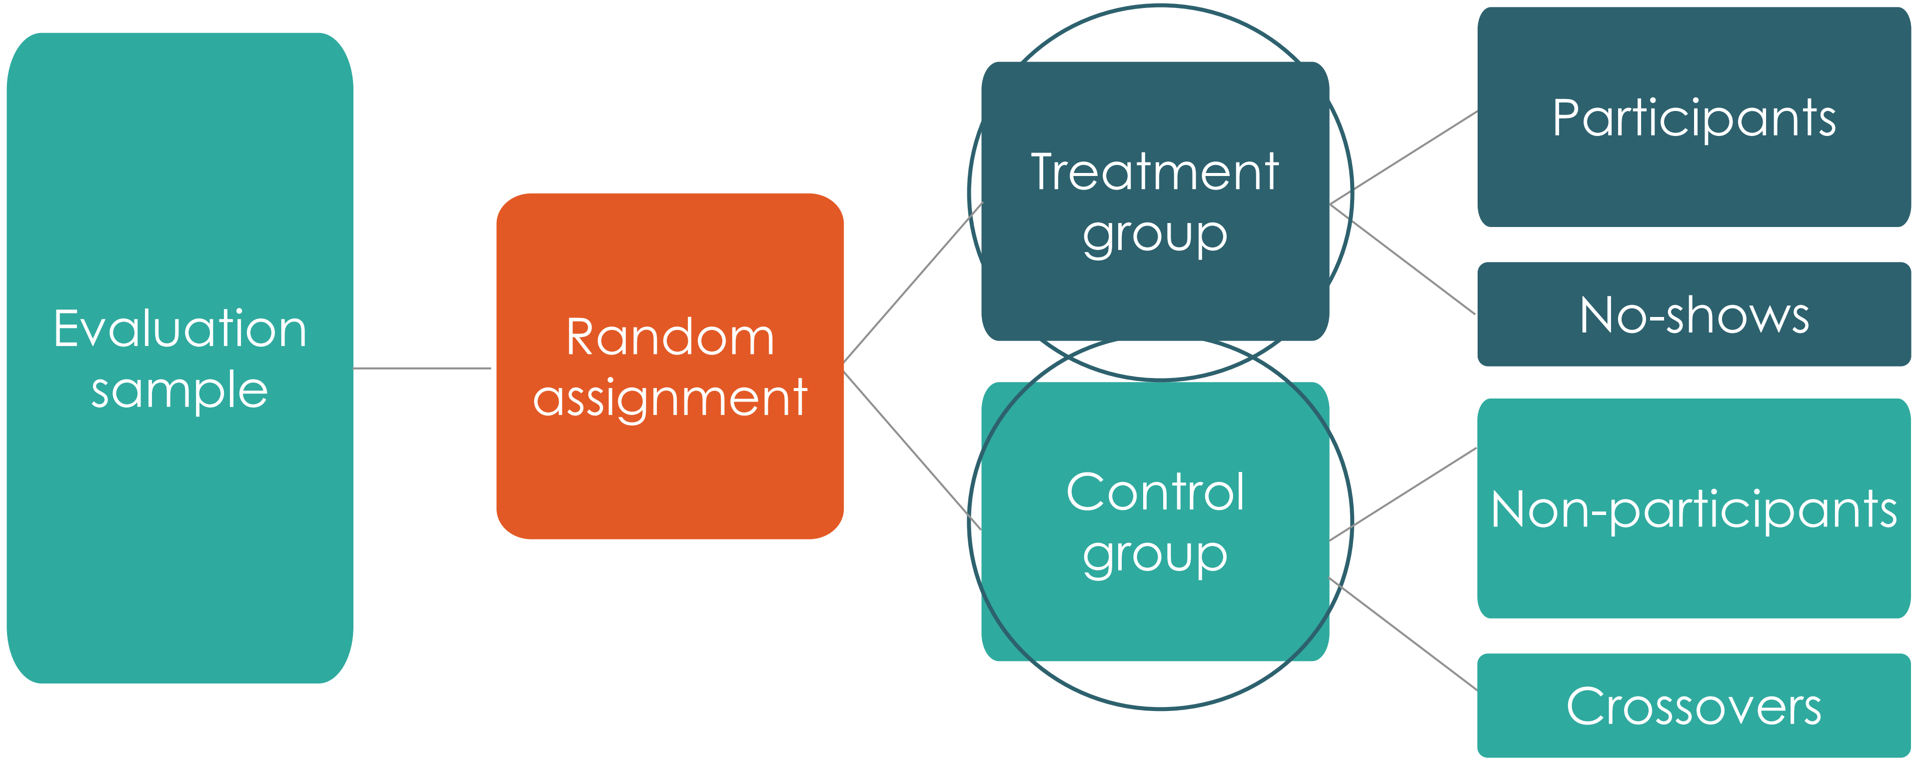

Based on what we just discussed, our treatment group for analysis is…

A. Individuals assigned to treatment who were *actually* treated

B. All individuals who were *actually* treated

C. Individuals assigned to treatment, regardless of whether or not they were treated

D. Don't know

Based on what we just discussed, our treatment group for analysis is…

A. Individuals assigned to treatment who were *actually* treated

B. All individuals who were *actually* treated

C. Individuals assigned to treatment, regardless of whether or not they were treated

D. Don't know

Example: Measuring Take-Up

Fazzio et al. (2021) studied the impact of an alternative to government-run primary schools in isolated rural areas in Guinea-Bissau:

- The intervention provides 4 years of primary education classes
- Randomized at the village level: comparison villages continue with existing school options, and treatment villages receive the intervention

Do enrolled children in intervention villages attend classes?

Attendance level	Percent of students in treatment villages
Mean attendance	**85.72%**
Attend 0% of classes	9.27%
Attend >0 to 25% of classes	1.24%
Attend >25 to 50% of classes	2.32%
Attend >50 to 75% of classes	2.01%
Attend >75% to 100% of classes	85.16%

Example: Measuring Take-Up

Discussion question: What steps would you take in the design or implementation phases to maximize take-up of the intervention?

Attendance level	Percent of students in treatment villages
Mean attendance	**85.72%**
Attend 0% of classes	9.27%
Attend >0 to 25% of classes	1.24%
Attend >25 to 50% of classes	2.32%
Attend >50 to 75% of classes	2.01%
Attend >75% to 100% of classes	85.16%

What can be done about noncompliance?

Design phase

- **Randomize** at a higher level to enable providers to treat clusters the same

Implementation phase

- **Prevent** noncompliance, e.g., by making take up easy or by incentivizing take up
  ➜ cannot always be done

- **Monitor** noncompliance to be aware if/when it happens

Analysis phase

- **Interpret** it during analysis phase
  ➜ see next section

Lecture Overview

- Threats to validity
  - Spillovers
  - Attrition
  - Evaluation-driven effects
  - Partial compliance
- **Generating impact estimates**
  - **Intention to Treat**
  - Local Average Treatment Effect
  - Reporting results

Intention to Treat (ITT)

- Easiest way to deal with partial compliance - calculate the Intent to Treat (ITT):
  - The difference between the average outcome of the group that was **randomly assigned to treatment** and the group that was **randomly assigned to comparison**, ***regardless of whether units within those groups actually received the treatment***.

ITT = (avg. outcome in group assigned to treatment) – (avg. outcome in group assigned to control)

- What does "intention to treat" measure?

  "What happened to the average treated unit in this population?"

- Is this difference the causal effect of the intervention?

Let’s go back to our microcredit example

Photo: Small businesses in Hyderabad, India | Sean Hallisey

- **Intervention:** Group-lending microcredit program that targets women in Hyderabad, India
- 104 neighborhoods
  - 52 assigned to the control group
  - 52 assigned to the treatment group

Neighborhood 1: Treatment

	Treatment assignment	Treated status	Change in profits (in Rupees)
Household 1	Yes	Yes	400
Household 2	Yes	Yes	400
Household 3	Yes	Yes	400
Household 4	Yes	No	0
Household 5	Yes	Yes	400
Household 6	Yes	No	200
Household 7	Yes	No	0
Household 8	Yes	Yes	600
Household 9	Yes	Yes	600
Household 10	Yes	No	0

Neighborhood 2: Control

	Treatment assignment	Treated status	Change in profits (in Rupees)
Household 1	No	No	200
Household 2	No	No	100
Household 3	No	Yes	300
Household 4	No	No	0
Household 5	No	No	0
Household 6	No	Yes	300
Household 7	No	No	0
Household 8	No	No	0
Household 9	No	No	100
Household 10	No	No	0

Neighborhood 1: Treatment

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	Yes	Yes	400
Household 2	Yes	Yes	400
Household 3	Yes	Yes	400
Household 4	Yes	No	0
Household 5	Yes	Yes	400
Household 6	Yes	No	200
Household 7	Yes	No	0
Household 8	Yes	Yes	600
Household 9	Yes	Yes	600
Household 10	Yes	No	0

Neighborhood 2: Control

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	No	No	200
Household 2	No	No	100
Household 3	No	Yes	300
Household 4	No	No	0
Household 5	No	No	0
Household 6	No	Yes	300
Household 7	No	No	0
Household 8	No	No	0
Household 9	No	No	100
Household 10	No	No	0

Mean change in profits:

Treated households neighborhood 1	466.7
Non-treated households neighborhood 2	50
Difference	416.7

Effect of treatment on profits?

NOT CORRECT!

Neighborhood 1: Treatment

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	Yes	Yes	400
Household 2	Yes	Yes	400
Household 3	Yes	Yes	400
Household 4	Yes	No	0
Household 5	Yes	Yes	400
Household 6	Yes	No	200
Household 7	Yes	No	0
Household 8	Yes	Yes	600
Household 9	Yes	Yes	600
Household 10	Yes	No	0

Neighborhood 2: Control

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	No	No	200
Household 2	No	No	100
Household 3	No	Yes	300
Household 4	No	No	0
Household 5	No	No	0
Household 6	No	Yes	300
Household 7	No	No	0
Household 8	No	No	0
Household 9	No	No	100
Household 10	No	No	0

The ***Intent to Treat (ITT)*** estimate:

Mean in neighborhood 1	300
Mean in neighborhood 2	100
Difference:	200

Lecture Overview

- Threats to validity
  - Spillovers
  - Attrition
  - Evaluation-driven effects
  - Partial compliance
- **Generating impact estimates**
  - Intention to Treat
  - **Local Average Treatment Effect**
  - Reporting results

Local Average Treatment Effect (LATE)

- The Local Average Treatment Effect (LATE) is:

$$\textbf{LATE} = \frac{\textbf{ITT}}{\textbf{(take-up in treatment group) - (take-up in control group)}}$$

- What does the LATE estimate? ***The effect of the program on those who complied with their treatment assignment***

- Note: Effects on people who didn't take it up might have been quite different

- Very similar: "Treatment on the Treated" (TOT)

Local Average Treatment Effect (LATE)

- <u>The intuitive idea:</u>
  - Let's say the ITT effect of a microcredit program on profits is a 300 Rupee (Rs) difference between treatment and control villages
  - But imagine only 50% of the households in the treatment neighborhood actually participate (let's assume no household in control neighborhood participate)

- If the effect of 50% take-up is an increase in profits of 300 Rs, then we can say that if everyone were to take up the microcredit program, the effect would be:

$$\text{LATE} = \frac{\text{ITT}}{\text{(take-up in treatment group)} - \text{(take-up in control group)}}$$

$$= \frac{300}{(0.5) - (0)} = 300 * 2 = 600 \text{ Rs}$$

Neighborhood 1: Treatment

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	Yes	Yes	400
Household 2	Yes	Yes	400
Household 3	Yes	Yes	400
Household 4	Yes	**No**	0
Household 5	Yes	Yes	400
Household 6	Yes	**No**	200
Household 7	Yes	**No**	0
Household 8	Yes	Yes	600
Household 9	Yes	Yes	600
Household 10	Yes	**No**	0

Neighborhood 2: Control

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	No	No	200
Household 2	No	No	100
Household 3	No	**Yes**	300
Household 4	No	No	0
Household 5	No	No	0
Household 6	No	**Yes**	300
Household 7	No	No	0
Household 8	No	No	0
Household 9	No	No	100
Household 10	No	No	0

The ***Intent to Treat (ITT)*** estimate:

Mean in neighborhood 1	300
Mean in neighborhood 2	100
Difference:	200

Treatment probability:

Fraction treated in neighborhood 1	0.6
Fraction treated in neighborhood 2	0.2
Difference:	0.4

Neighborhood 1: Treatment

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	Yes	Yes	400
Household 2	Yes	Yes	400
Household 3	Yes	Yes	400
Household 4	Yes	No	0
Household 5	Yes	Yes	400
Household 6	Yes	No	200
Household 7	Yes	No	0
Household 8	Yes	Yes	600
Household 9	Yes	Yes	600
Household 10	Yes	No	0

Neighborhood 2: Control

	Treatment assignment	Treated status	Change in profits (in Rs)
Household 1	No	No	200
Household 2	No	No	100
Household 3	No	Yes	300
Household 4	No	No	0
Household 5	No	No	0
Household 6	No	Yes	300
Household 7	No	No	0
Household 8	No	No	0
Household 9	No	No	100
Household 10	No	No	0

The ***Intent to Treat (ITT)*** estimate:

Mean in neighborhood 1	300
Mean in neighborhood 2	100
Difference:	200

Treatment probability:

Fraction treated in neighborhood 1	0.6
Fraction treated in neighborhood 2	0.2
Difference:	0.4

Local Average Treatment Effect (LATE): 200/0.4 = 500 Rs

ITT vs LATE

If obtaining the estimate is easy, why not always use LATE?

- In order to estimate LATE we need data on compliance
- ITT may be the policy-relevant parameter of interest
  - For example, we may not be interested in the medical effect of deworming treatment, but what would **happen under actual implementation** of the program
  - If students often miss school and therefore don't get the deworming medicine, the Intention to Treat estimate may be most relevant

ITT & LATE: Conclusions

- Both ITT and LATE can provide valuable information to decision-makers

ITT

Estimates the **overall effect of the intervention**, admitting that noncompliance can happen (which is inherent to any policy).

LATE

Estimates the **effect of the intervention for those who comply** with their assignment to treatment or control.

Lecture Overview

- Threats to validity
  - Spillovers
  - Attrition
  - Evaluation-driven effects
  - Partial compliance
- **Generating impact estimates**
  - Intention to Treat
  - Local Average Treatment Effect
  - **Reporting results**

Reporting Results

Reporting bias occurs when the decision on whether and how to report impact estimates depends on the direction and significance of the estimate

How researchers report their results can also threaten validity

Potential sources of reporting bias:

- **Specification searching:** trying different analyses to find one that is statistically significant
  - The more outcomes and adjustments to covariates you look at, the higher the chance you find at least one significant effect
- **File drawer problem:** significant results are more likely to be published

What can be done about reporting bias?

- Pre-specify outcomes of interest
  - Pre-analysis plans are becoming more common, and pre-specified analyses may be given more weight
  - Differentiate between pre-specified and exploratory analysis
- Report raw differences between treatment and control as well as regression estimates (adjusted based on covariates)
- Share data and code along with research papers
- Report all results, not just the most impressive or significant ones
  - There is a lot we can learn from papers reporting no significant impact of the intervention on the outcomes of interest

Conclusions

- **Internal validity is a strength of well-designed randomized evaluations…**

  …so everything undermining it must be carefully considered

- **Consider which threats are likely factors for a given evaluation…**

  …and plan to mitigate and monitor these

- **The design phase and project planning are important…**

  …but so is the ability to face challenges during implementation phase

- **Analyzing impact using different estimators can teach us different things…**

  … so think critically about which results are reported and how to interpret

Further Resources

- "Using Randomization in Development Economics Research: A Toolkit" (Duflo, Glennerster, and Kremer 2006)
- *Mostly Harmless Econometrics* (Angrist and Pischke 2008)
- "Identification and Estimation of Local Average Treatment Effects" (Imbens and Angrist 1994)
- *Impact Evaluation in Practice*, Chapter 9 (Gertler et al. 2016)
- "On Minimizing the Risk of Bias in Randomized Controlled Trials in Economics" (Eble, Boone, and Elbourne 2016)
- J-PAL Reseach Resource. "Data analysis."

References

Bagga, Aanchal, and Cynthia Kinnan. 2024. "Microcredit: Equilibrium effects." *Oxford Review of Economic Policy* 40(1): 54–70.

Banerjee, Abhijit, Eliana La Ferrara, and Victor H. Orozco-Olvera. 2019. "The entertaining way to behavioral change: Fighting HIV with MTV." National Bureau of Economic Research Working Paper No. w26096.

Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. 2015. "The Miracle of Microfinance? Evidence from a Randomized Evaluation." *American Economic Journal: Applied Economics* 7 1): 22-53.

Fazzio, Ila, Alex Eble, Robin L. Lumsdaine, Peter Boone, Baboucarr Bouy, Pei-Tseng Jenny Hsieh, Chitra Jayanty, Simon Johnson, and Ana Filipa Silva. 2021. "Large learning gains in pockets of extreme poverty: Experimental evidence from Guinea Bissau." *Journal of Public Economics* 199, 104385.

Fearon, James D., Macartan Humphreys, and Jeremy M. Weinstein. 2009. "Can development aid contribute to social cohesion after civil war? Evidence from a field experiment in post-conflict Liberia." *American Economic Review* 99(2): 287-291.

J-PAL. "Measuring the Impact of Microfinance in Hyderabad, India." J-PAL Evaluation Summary.

J-PAL. "The Impact of Diagnostic Feedback for Teachers on Student Learning in India." J-PAL Evaluation Summary. 2010.

Mummolo, Jonathan, and Erik Peterson. 2019. "Demand effects in survey experiments: An empirical assessment." *American Political Science Review* 113(2): 517-529.

Reuse and citation

To reference this lecture, please cite as:

J-PAL. 2024. “Lecture: Threats & Analysis.” Abdul Latif Jameel Poverty Action Lab. Cambridge, MA