ABDUL LATIF JAMEEL
Poverty Action Lab

TRANSLATING RESEARCH INTO ACTION

Threats and Analysis

Shawn Cole

Harvard Business School

povertyactionlab.org

Course Overview

- What is Evaluation?
- Measuring Impacts
- Why Randomise?
- How to Randomise
- Sampling and Sample Size
- Threats and Analysis
- Project from Start to Finish

Lecture Overview

- Attrition
- Spillovers
- Partial Compliance and Sample Selection Bias
- Intention to Treat & Treatment on Treated
- Choice of outcomes
- External validity
- Cost Effectiveness

Attrition

- Is it a problem if some of the people in the experiment vanish before you collect your data?
- It is a problem if the type of people who disappear is correlated with the treatment.
- Why is it a problem?
- Why should we expect this to happen?

Attrition bias: an example

- The problem you want to address:
  - Some children don't come to school because they are too weak (undernourished)
- You start a school feeding program and want to do an evaluation
  - You have a treatment and a control group
- Weak, stunted children start going to school more if they live next to a treatment school
- First impact of your program: increased enrollment.
- In addition, you want to measure the impact on child's growth
  - Second outcome of interest: Weight of children
- You go to all the schools (treatment and control) and measure everyone who is in school on a given day
- Will the treatment-control difference in weight be over-stated or understated?

	Before Treatment			After Treament	
	T	C		T	C
	20	20		**22**	20
	25	25		**27**	25
	30	30		**32**	30
Ave.	**25**	**25**		**27**	**25**
	Difference	0		Difference	2

What if only children > 20 Kg come to school?

What if only children > 20 Kg come to school?

	Before Treatment			After Treament	
	T	C		T	C
	[absent]	[absent]		**22**	[absent]
	25	25		**27**	25
	30	30		**32**	30
Ave.	**27.5**	**27.5**		**27**	**27.5**
	Difference	0		Difference	-0.5

Attrition Bias

- What source of attrition bias did they worry about in the de-worming case with regards to testing?
- If worms keep children out of school and have adverse cognitive consequences, then deworming medicine might induce the weaker students

Attrition Bias

- Devote resources to tracking participants after they leave the program
- If there is still attrition, check that it is not different in treatment and control. Is that enough?
- Also check that it is not correlated with observables.
- Try to bound the extent of the bias
  - suppose everyone who dropped out from the treatment got the lowest score that anyone got; suppose everyone who dropped out of control got the highest score that anyone got…
  - Why does this help?

Lecture Overview

- Attrition
- Spillovers
- Partial Compliance and Sample Selection Bias
- Intention to Treat & Treatment on Treated
- Choice of outcomes
- External validity
- Cost Effectiveness

Example: Deworming

- Previous studies randomize deworming treatment *within* schools
  - Suppose that deworming prevents the transmission of disease, what problems does this create for evaluation?
  - Suppose externalities are local? How can we measure total impact?

Externalities Within School

Without Externalities

School A	Treated?	Outcome
Pupil 1	Yes	no worms
Pupil 2	No	worms
Pupil 3	Yes	no worms
Pupil 4	No	worms
Pupil 5	Yes	no worms
Pupil 6	No	worms

Total in Treatment with Worms	0%
Total in Control with Worms	100%
Treament Effect	-100%

With Externalities

Suppose, because prevalence is lower, some children are not re-infected with worms

School A	Treated?	Outcome
Pupil 1	Yes	no worms
Pupil 2	No	no worms
Pupil 3	Yes	no worms
Pupil 4	No	worms
Pupil 5	Yes	no worms
Pupil 6	No	worms

Total in Treatment with Worms	0%
Total in Control with Worms	67%
Treatment Effect	-67%

How to measure program impact in the presence of spillovers?

- Design the unit of randomization so that it encompasses the spillovers

  If we expect externalities that are all within school:

  – Randomization at the level of the school allows for estimation of the overall effect

Measuring total impact in the presence of spillovers

GROUP	TREATMENT SCHOOLS	CONTROL SCHOOLS	PROGAM EFFECT
% of children with a moderate or heavy infection	27%	52%	-25%***
% of children who were sick the week before the survey	41%	45%	-4%**
% of children who are anemic	2%	4%	-2%*

Within-school health externalities

- What if we wanted to measure the spillovers?
- Deworming study
  - Because girls above 12 could not be treated in the treatment schools, we can compare girls above 12 in treatment schools to girls above 12 in comparison schools.
- More generally: need to randomize treatment within the unit so as to be able to learn about spillovers.

Example: Price Information

- Providing farmers with spot and futures price information by mobile phone
- Should we expect spilloves?
- Randomize: individual or village level?
- Village level randomization
  - Less statistical power
  - “Purer control groups”
- Individual level randomization
  - More statistical power (if spillovers small)
  - Ability to measure spillovers

Example: Price Information

- Can we do both?
- Randomly assign villages into one of four groups, A, B, C, and D
- Group A Villages
  - SMS price information to all individuals with phones
- Group B Villages
  - SMS price information to randomly selected 75% of individuals with phones
- Group C Villages
  - SMS price information to randomly selected 25% of individuals with phones
- Group D Villages
  - No SMS price information

Lecture Overview

- Attrition
- Spillovers
- Partial Compliance and Sample Selection Bias
- Intention to Treat & Treatment on Treated
- Choice of outcomes
- External validity
- Cost Effectiveness

Basic setup of a randomized evaluation

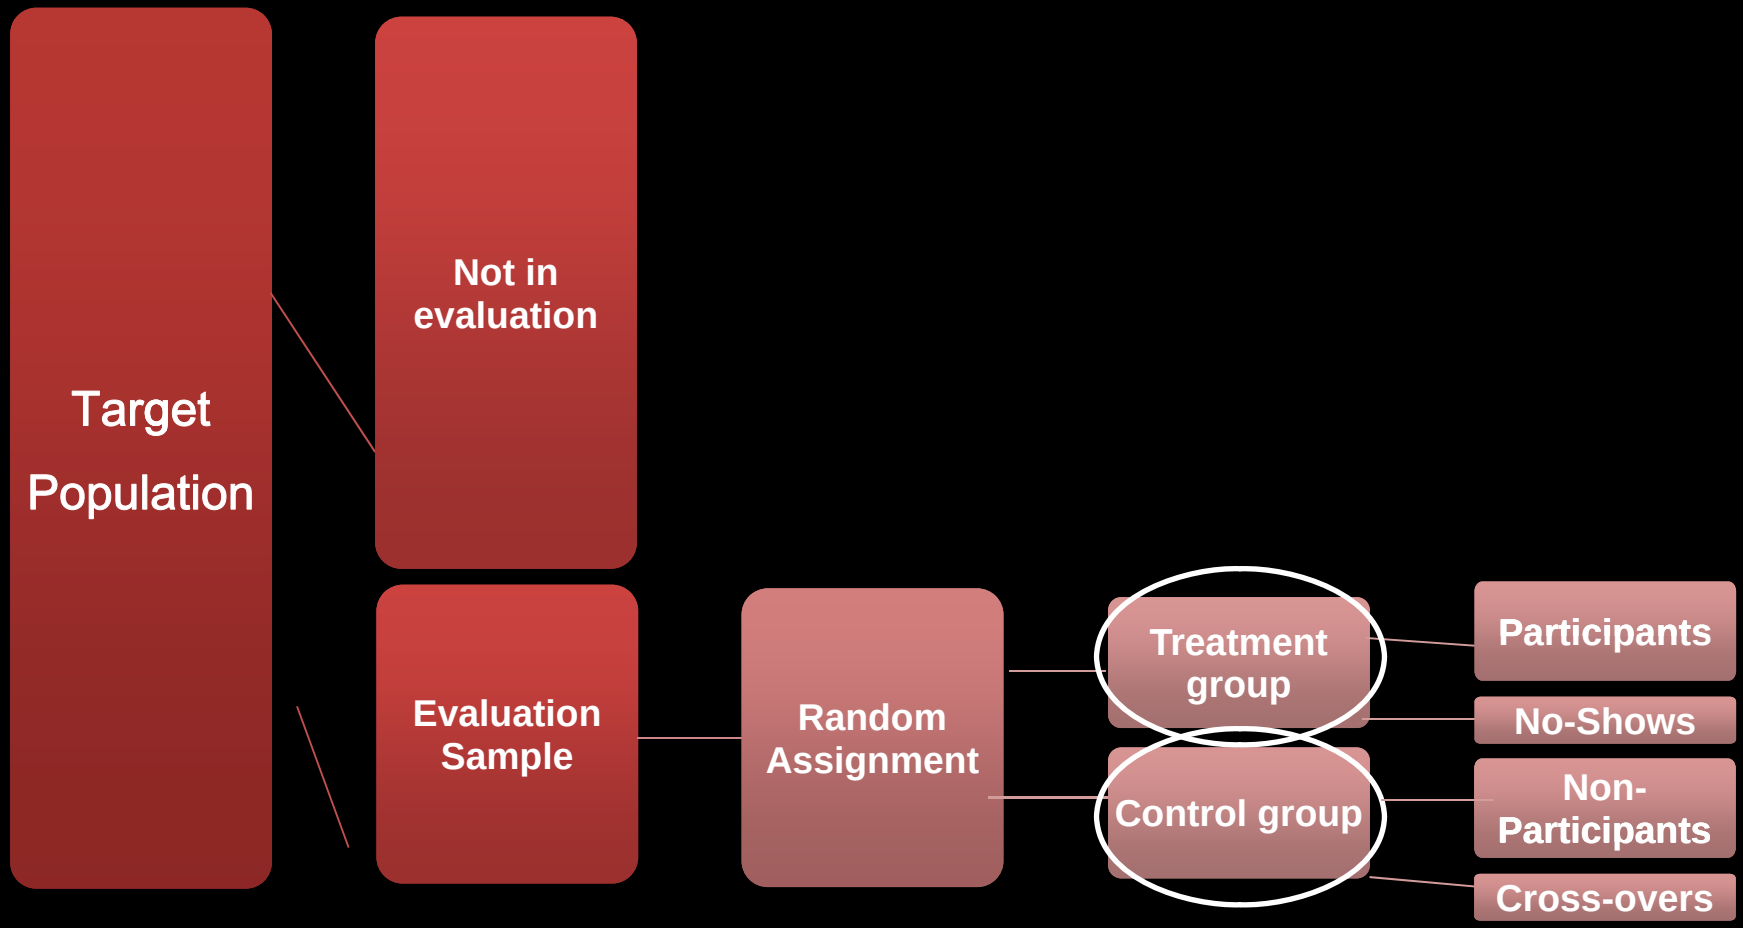

Sample selection bias

- Sample selection bias could arise if factors other than random assignment influence program allocation
  - Even if intended allocation of program was random, the actual allocation may not be

Sample selection bias

- Individuals assigned to comparison group could attempt to move into treatment group
  - De-worming program: parents could attempt to move their children from comparison school to treatment school
- Alternatively, individuals allocated to treatment group may not receive treatment
  - De-worming program: some students assigned to treatment in treatment schools did not receive medical treatment

Lecture Overview

- Spillovers
- Partial Compliance and Sample Selection Bias
- Intention to Treat & Treatment on Treated
- Choice of outcomes
- External validity
- Cost Effectiveness

ITT and ToT

- Vaccination campaign in villages
- Some people in treatment villages not treated
  - 78% of people assigned to receive treatment received some treatment
- What do you do?
  - Compare the beneficiaries and non-beneficiaries?
  - Why not?

Which groups can be compared ?

Treatment Group: vaccination	*Control Group*
TREATED	NON-TREATED
NON-TREATED	

What is the difference between the 2 random groups?

Treatment Group	*Control Group*
1: treated not infected 2: treated – not infected 3: treated – infected	5: non-treated infected 6: non-treated – not infected 7: non-treated – infected 8: non-treated – infected
4: non-treated – infected	

Intention to treat - ITT

Treatment Group: 50% infected

Control Group: 75% infected

- Y(T)= Average Outcome in Treatment Group
- Y(C)= Average Outcome in Control Group

ITT = Y(T) - Y(C)

- ITT = 50% - 75% = -25 percentage points

Intention To Treat (ITT)

- What does “intention to treat” measure?

  “What happened to the average child who is in a treated school in this population?”

- Is this difference the causal effect of the intervention?

When is ITT useful?

- May relate more to actual programs
- For example, we may not be interested in the medical effect of deworming treatment, but what would happen under an actual deworming program.
- If students often miss school and therefore don t get the deworming medicine, the intention to treat estimate may actually be most relevant.

School 1	Intention to Treat ?	Treated?	Observed Change in weight
Pupil 1	yes	**yes**	**4**
Pupil 2	yes	**yes**	**4**
Pupil 3	yes	**yes**	**4**
Pupil 4	yes	no	0
Pupil 5	yes	**yes**	**4**
Pupil 6	yes	no	2
Pupil 7	yes	no	0
Pupil 8	yes	**yes**	**6**
Pupil 9	yes	**yes**	**6**
Pupil 10	yes	no	0
Avg. Change among Treated A=			3

School 2			
Pupil 1	no	**no**	**2**
Pupil 2	no	**no**	**1**
Pupil 3	no	yes	3
Pupil 4	no	**no**	**0**
Pupil 5	no	**no**	**0**
Pupil 6	no	yes	3
Pupil 7	no	**no**	**0**
Pupil 8	no	**no**	**0**
Pupil 9	no	**no**	**0**
Pupil 10	no	**no**	**0**
Avg. Change among Not-Treated B=			0.9

School 1:
Avg. Change among Treated 3 (A)
School 2:
Avg. Change among not-treated 0.9 (B)

A-B 2.1

From ITT to effect of treatment on the treated (TOT)

- The point is that if there is leakage across the groups, the comparison between those originally assigned to treatment and those originally assigned to control is smaller

  But the difference in the probability of getting treated is also smaller
- Formally this is done by "instrumenting the probability of treatment by the original assignment

Treatment on the treated (TOT)

- The effect of the treatment on those who got the treatment:
  - Suppose children who got the treatment had a weight gain of A, irrespective of whether they were in a treatment or a control school
  - Suppose children who got no treatment had a weight gain of B, again in both kinds of schools
  - We want to know A-B, the difference between treated and non-treated students

Treatment on the treated (TOT)

- Then…
- Y(T)=A*Prob[treated|T]+B(1-Prob[treated|T])
- Y(C)=A*Prob[treated|C]+B(1-Prob[treated|C])
- A-B= (Y(T)-Y(C)) /

  (Prob[treated|T] – Prob[treated|C])

  = The “treatment on the treated” effect.

School 1	Intention to Treat ?	Treated?	Observed Change in weight
Pupil 1	yes	yes	4
Pupil 2	yes	yes	4
Pupil 3	yes	yes	4
Pupil 4	yes	no	0
Pupil 5	yes	yes	4
Pupil 6	yes	no	2
Pupil 7	yes	no	0
Pupil 8	yes	yes	6
Pupil 9	yes	yes	6
Pupil 10	yes	no	0
		Avg. Change Y(T)=	3

School 2			
Pupil 1	no	no	2
Pupil 2	no	no	1
Pupil 3	no	yes	3
Pupil 4	no	no	0
Pupil 5	no	no	0
Pupil 6	no	yes	3
Pupil 7	no	no	0
Pupil 8	no	no	0
Pupil 9	no	no	0
Pupil 10	no	no	0
		Avg. Change Y(C) =	0.9

A = Gain if Treated
B = Gain if not Treated

ToT Estimator: A-B

$$\textbf{A-B} = \frac{Y(T)-Y(C)}{Prob(Treated|T)-Prob(Treated|C)}$$

Y(T)	3
Y(C)	0.9
Prob(Treated|T)	60%
Prob(Treated|C)	20%
Y(T)-Y(C)	2.1
Prob(Treated|T)-Prob(Treated|C)	40%
A-B	5.25

Estimating TOT

- What values do we need?
- Y(T)
- Y(C)

- Prob[treated|T]
- Prob[treated|C]

TOT not always appropriate…

- Example: send 50% of MIT staff a letter warning of flu season, encourage them to get vaccines
- Suppose 50% in treatment, 0% in control get vaccines
- Suppose incidence of flu in treated group drops 35% relative to control group
- Is (.35) / (.5 – 0) = 70% the correct estimate?
- What effect might letter alone have?
- Excel Example

Lecture Overview

- Spillovers
- Partial Compliance and Sample Selection Bias
- Intention to Treat & Treatment on Treated
- Choice of outcomes
- External validity
- Cost Effectiveness

Multiple outcomes

- Can we look at various outcomes?

- The more outcomes you look at, the higher the chance you find at least one significantly affected by the program
  - Pre-specify outcomes of interest
  - Report results on all measured outcomes, even null results
  - Correct statistical tests (Bonferroni)

Covariates

- Why include covariates?
  - May explain variation, improve statistical power
- Why not include covariates?
  - Appearances of "specification searching"
- What to control for?
  - If stratified randomization: add strata fixed effects
  - Other covariates

Rule: Report both "raw" differences and regression-adjusted results

Lecture Overview

- Spillovers
- Partial Compliance and Sample Selection Bias
- Intention to Treat & Treatment on Treated
- Choice of outcomes
- External validity
- Cost Effectiveness

Threat to external validity:

- Behavioral responses to evaluations

- Generalizability of results

Threat to external validity: Behavioral responses to evaluations

- One limitation of evaluations is that the evaluation itself may cause the treatment or comparison group to change its behavior
  - Treatment group behavior changes: Hawthorne effect
  - Comparison group behavior changes: John Henry effect

- Minimize salience of evaluation as much as possible
- Consider including controls who measured at endline only

Generalizability of results

- Depend on three factors:
  - Program Implementation: can it be replicated at a large (national) scale?
  - Study Sample: is it representative?
  - Sensitivity of results: would a similar, but slightly different program, have same impact?

Lecture Overview

- Spillovers
- Partial Compliance and Sample Selection Bias
- Intention to Treat & Treatment on Treated
- Choice of outcomes
- External validity
- Conclusion

Conclusion

- There are many threats to the internal and external validity of randomized evaluations…
- …as are there for every other type of study
- Randomized trials:
  - Facilitate simple and transparent analysis
    - Provide few "degrees of freedom" in data analysis (this is a good thing)
  - Allow clear tests of validity of experiment

Further resources

- Using Randomization in Development Economics Research: A Toolkit (Duflo, Glennerster, Kremer)
- Mostly Harmless Econometrics (Angrist and Pischke)
- Identification and Estimation of Local Average Treatment Effects (Imbens and Angrist, Econometrica, 1994)

Resource: Abdul Latif Jameel Poverty Action Lab Executive Training: Evaluating Social Programs

Dr. Rachel Glennerster, Prof. Abhijit Banerjee, Prof. Esther Duflo