Lecture 03:
Randomized controlled trials I

PPHA 34600

Prof. Fiona Burlig

Harris School of Public Policy
University of Chicago

From last time: selection is an issue

Recall that there are lots of things we want to estimate.

We need to get around selection bias to do this.

In other words, we need:

$$E[Y_i(1)] = E[Y_i(1)|D_i = 1] = E[Y_i(1)|D_i = 0]$$

and

$$E[Y_i(0)] = E[Y_i(0)|D_i = 0] = E[Y_i(0)|D_i = 1]$$

Regression equivalent:

$$E[\varepsilon_i|D_i] = 0$$

Random assignment as a solution

When treatment status is randomly assigned,

$$F(X, \varepsilon | D = 1) = F(X, \varepsilon | D = 0) = F(X, \varepsilon)$$

In words:

The distribution of **both** observables (Xs) **and** unobservables (εs) is the same for treated and untreated units!

There is **no selection problem** by construction!

Again, but mathier

When D, treatment, is **randomly assigned**:

- D is independent of $Y(0)$ and $Y(1)$
- The distribution of $Y_i(0)|D_i$ is equal to the unconditional distribution
- The distribution of $Y_i(1)|D_i$ is equal to the unconditional distribution
- $E[Y_i(1)|D_i = 1] = E[Y_i(1)]$
- $E[Y_i(0)|D_i = 0] = E[Y_i(0)]$

Again, but mathier

When D, treatment, is **randomly assigned**:

- D is independent of $Y(0)$ and $Y(1)$
- The distribution of $Y_i(0)|D_i$ is equal to the unconditional distribution
- The distribution of $Y_i(1)|D_i$ is equal to the unconditional distribution
- $E[Y_i(1)|D_i = 1] = E[Y_i(1)]$
- $E[Y_i(0)|D_i = 0] = E[Y_i(0)]$

As a result:

$$\tau^{ATE} = E[Y_i(1)] - E[Y_i(0)]$$

$$= E[Y_i(1)|D_i = 1] - E[Y_i(0)|D_i = 0]$$

$$= E[Y_i|D_i = 1] - E[Y_i|D_i = 0]$$

This bears repeating

Under randomization:

$$\tau^{ATE} = E[Y_i|D_i = 1] - E[Y_i|D_i = 0]$$

This bears repeating

Under randomization:

$$\tau^{ATE} = E[Y_i|D_i = 1] - E[Y_i|D_i = 0]$$

We can easily estimate this from data:

$$\hat{\tau}^{ATE} = \overline{Y(1)} - \overline{Y(0)}$$

We can estimate the ATE simply from the difference in means between treated and "control" group.

This bears repeating

Under randomization:

$$\tau^{ATE} = E[Y_i|D_i = 1] - E[Y_i|D_i = 0]$$

We can easily estimate this from data:

$$\hat{\tau}^{ATE} = \overline{Y(1)} - \overline{Y(0)}$$

We can estimate the ATE simply from the difference in means between treated and "control" group.

Obvious (?) caveat: We still can't get τ_i, because we only observe i once.

Evaluating an RCT

This is not a class on how to do RCTs

- As always, the devil is in the details
- Field experiments are *hard!*
- But supposing you've got one...

Evaluating an RCT

This is not a class on how to do RCTs

- As always, the devil is in the details
- Field experiments are *hard!*
- But supposing you've got one...

Basic RCT checklist

- ☐ Verify random assignment
- ☐ Check compliance with treatment
- ☐ Estimate the ATE (or other things...)

What is this experiment trying to learn?

When running an RCT, you want to have a "research question" in mind:

What is the causal effect of [program x] on [outcome y]?

What is this experiment trying to learn?

When running an RCT, you want to have a "research question" in mind:

What is the causal effect of [program x] on [outcome y]?

Why do we need an RCT to study this?

What is this experiment trying to learn?

When running an RCT, you want to have a "research question" in mind:

What is the causal effect of [program x] on [outcome y]?

Why do we need an RCT to study this?

- Program X targets certain individuals
- Individuals who choose to participate look different than non-participants
- Others?

Understanding RCTs

Basic ingredients for an RCT:

- What is the research design?
  - What is the unit of randomization?
  - How was randomization performed?
- What are the outcomes of interest?

Verifying random assignment

Did randomization "work"?

- Randomization should mean treated and control units are similar
- This is true *in expectation*, not necessarily for one draw

Verifying random assignment

Did randomization "work"?

- Randomization should mean treated and control units are similar
- This is true *in expectation*, not necessarily for one draw

Testing whether randomization was effective:

- We want T and C to be similar on observables **and** unobservables
- We can only test this for observables
- To check this, we "test for balance":
- Compare mean outcomes for T vs. C *at baseline* (before treatment) or in fixed characteristics

  → Implementation: Regress $Y_i^{baseline} = \alpha + \tau D_i + \nu_i$

Checking for balance

Three things to check for:

1. Did they test for all outcome variables?
2. Are differences statistically significant?
3. Are magnitudes economically meaningful?

Checking compliance with treatment

Did assignment to treatment affect treatment status?

Trying to verify whether...

- Units assigned to treatment were actually treated
- Units assigned to control were *not* treated

There is often substantial non-compliance. We'll talk more about exactly how to deal with this issue next time.

Thinking about non-compliance

We will treat this more formally next time

For now, non-compliance changes the interpretation of our estimates:

> Rather than asking "What does treatment do to our outcome activities?"...
>
> ... we're asking "What does offering treatment do to our outcome?"

This may be the policy-relevant quantity

We want to estimate the ATE

Recall that the ATE is just:

$$\tau^{ATE} = E[Y_i(1)] - E[Y_i(0)]$$

Since we have random assignment, we can estimate this as:

$$\hat{\tau}^{ATE} = \overline{Y(1)} - \overline{Y(0)}$$

We want to estimate the ATE

Recall that the ATE is just:

$$\tau^{ATE} = E[Y_i(1)] - E[Y_i(0)]$$

Since we have random assignment, we can estimate this as:

$$\hat{\tau}^{ATE} = \overline{Y(1)} - \overline{Y(0)}$$

Regression is a convenient way to do this:

$$Y_i = \alpha + \tau D_i + \varepsilon_i$$

Since our $E[\varepsilon|D_i] = 0$ assumption is satisfied (why?), $\hat{\tau} = \hat{\tau}^{ATE}$

Estimating treatment effects

We'll often see things that look like this:

$$y_{ia} = \alpha + \tau Treat_{ia} + \gamma \mathbf{X}_a^{\text{baseline}} + \varepsilon_{ia}$$

where:

- y_{ia} are outcomes for household i in area a
- α is a constant
- $Treat_{ia}$ is a treatment dummy (think D_i)
- $\mathbf{X}_a^{\text{baseline}}$ is a set of baseline area controls
- ε_{ia} is an error term

What is this equation estimating?

$$y_{ia} = \alpha + \tau Treat_{ia} + \gamma \mathbf{X}_a^{\text{baseline}} + \varepsilon_{ia}$$

This differs from our basic regression a bit:

- There's an *i and* an *a*
- We have $\gamma \mathbf{X}_a^{\text{baseline}}$

Let's unpack each of these in turn...

Randomization by area, data on individuals

We have *i*-ndividual level data, but *a*-rea level randomization

Randomizing at a higher level of aggregation is common:

- Some questions can't be answered at *i* level (no personal bank branches)
- Ethics concerns: can sometimes delay implementation for a whole group; hard for individuals
- Reduce spillovers (more on this later)

Randomization by area, data on individuals

We have *i*-ndividual level data, but *a*-rea level randomization

Randomizing at a higher level of aggregation is common:

- Some questions can't be answered at *i* level (no personal bank branches)
- Ethics concerns: can sometimes delay implementation for a whole group; hard for individuals
- Reduce spillovers (more on this later)

Randomizing at a higher level affects the analysis:

- Interpretation is different (what exactly is treatment?)
- Getting standard errors right requires either:
  1. Estimate *i*-level effects, but cluster at *a*-level
     or
  2. Averaging outcomes at the group level (weight by individuals per group)

Adding controls

If D_i is randomly assigned, we don't need X_i!

We often add controls anyway:

Adding controls

If D_i is randomly assigned, we don't need X_i!

We often add controls anyway:

- Controlling for X_i should not affect $\hat{\tau}$

  $\rightarrow$ Why?

Adding controls

If D_i is randomly assigned, we don't need X_i!

We often add controls anyway:

- Controlling for X_i should not affect $\hat{\tau}$

  $\rightarrow$ Why?

- Controlling for X_i will affect the standard error on $\hat{\tau}$

  $\rightarrow$ Why?

  do **not** control for post-treatment outcomes

Adding bad controls

First rule of RCT club:

- Do **not** control for post-treatment outcomes
- Do **not** control for post-treatment outcomes

$\rightarrow$ If treatment affects these outcomes, you can get bias!

Simple example:

- Suppose microfinance impacts business ownership
- By random assignment, households with and without loans have the same potential income
- Once we condition on business ownership, this is no longer true!

We can use simulated data to think about this

Type of household	Potential business ownership		Potential income		Average earnings by ownership	
	Without MF	With MF	Without MF	With MF	Without MF	With MF
Never owner	No	No	1,000	1,500		
Moved by MF	No	Yes	2,000	2,500		
Always owner	Yes	Yes	3,000	3,500		

We can use simulated data to think about this

Type of household	Potential business ownership		Potential income		Average earnings by ownership	
	Without MF	With MF	Without MF	With MF	Without MF	With MF
Never owner	No	No	1,000	1,500	Don't own: 1,500	Don't own: 1,500
Moved by MF	No	Yes	2,000	2,500		Own: 3,000
Always owner	Yes	Yes	3,000	3,500	Own: 3,000	

We can use simulated data to think about this

Type of household	Potential business ownership		Potential income		Average earnings by ownership	
	Without MF	With MF	Without MF	With MF	Without MF	With MF
Never owner	No	No	1,000	1,500	Don't own: 1,500	Don't own: 1,500
Moved by MF	No	Yes	2,000	2,500		Own: 3,000
Always owner	Yes	Yes	3,000	3,500	Own: 3,000	

- The return to MFI is 500 for everyone...
- But once we condition on ownership, it looks like the return is 0!

→ This is because we don't have random assignment **within** ownership!

We can use simulated data to think about this

Type of household	Potential business ownership		Potential income		Average earnings by ownership	
	Without MF	With MF	Without MF	With MF	Without MF	With MF
Never owner	No	No	1,000	1,500	Don't own: 1,500	Don't own: 1,500
Moved by MF	No	Yes	2,000	2,500		Own: 3,000
Always owner	Yes	Yes	3,000	3,500	Own: 3,000	

- The return to MFI is 500 for everyone...
- But once we condition on ownership, it looks like the return is 0!

→ This is because we don't have random assignment **within** ownership!

Do not control for post-treatment outcomes!

We can also estimate heterogeneous effects

Heterogeneous effects are straightforward:

$$\tau(X_1 = x_1) = E[Y_i(1)|X_1 = x_1] - E[Y_i(0)|X_1 = x_1]$$

We can also estimate heterogeneous effects

Heterogeneous effects are straightforward:

$$\tau(X_1 = x_1) = E[Y_i(1)|X_1 = x_1] - E[Y_i(0)|X_1 = x_1]$$

We typically estimate these in two ways:

1. Add an **interaction term** to the regression:

$$y_i = \alpha + \tau Treat_i + \gamma Treat_i \cdot X_i + \delta X_i + \varepsilon_i$$

$\rightarrow$ Make sure to add both the interaction and the base term

We can also estimate heterogeneous effects

Heterogeneous effects are straightforward:

$$\tau(X_1 = x_1) = E[Y_i(1)|X_1 = x_1] - E[Y_i(0)|X_1 = x_1]$$

We typically estimate these in two ways:

1. Add an **interaction term** to the regression:

$$y_i = \alpha + \tau Treat_i + \gamma Treat_i \cdot X_i + \delta X_i + \varepsilon_i$$

   → Make sure to add both the interaction and the base term

2. Estimate the regression **separately** by heterogeneity

   → Equivalent to a *fully* interacted model

Estimate heterogeneity by pre-determined characteristics only!

A note on assumptions for the RCT

We still need several assumptions for the RCT to work:

- $E[Y_i(1)|D_i = 1] = E[Y_i(1)|D_i = 0]$
  and
  $E[Y_i(0)|D_i = 1] = E[Y_i(0)|D_i = 0]$

  $\rightarrow$ We "get this" via randomization, but only in expectation

A note on assumptions for the RCT

We still need several assumptions for the RCT to work:

- $E[Y_i(1)|D_i = 1] = E[Y_i(1)|D_i = 0]$
  and
  $E[Y_i(0)|D_i = 1] = E[Y_i(0)|D_i = 0]$

  $\rightarrow$ We "get this" via randomization, but only in expectation

- Perfect compliance

  $\rightarrow$ Kinda. More on this next class

A note on assumptions for the RCT

We still need several assumptions for the RCT to work:

- $E[Y_i(1)|D_i = 1] = E[Y_i(1)|D_i = 0]$
  and
  $E[Y_i(0)|D_i = 1] = E[Y_i(0)|D_i = 0]$

  $\rightarrow$ We "get this" via randomization, but only in expectation

- Perfect compliance

  $\rightarrow$ Kinda. More on this next class

- No spillovers: "SUTVA"

  - Stable Unit Treatment Value Assumption: D_i doesn't affect j's potential outcomes

  $\rightarrow$ Kinda. More on this in two classes

Application: Health worker performance in Sierra Leone

Christensen, Dube, Haushofer, Siddiqi, and Voors (2020 NBER WP)

Policy challenge:

- Health in developing countries is low
- Healthcare systems are thought to be of low quality...
- ...and are therefore potentially not used as much as they could be

Intervention(s):

- **#1:** Community monitoring: rankings and public meetings
- **#2:** Staff competition for the best (and most improved) clinic
- No additional inputs to clinics

Improving health services: The experiment

$\rightarrow$ **Lesson for you as MPPs:** RCTs are doable in high-stakes contexts!

This is a **blocked randomization design**:

- 254 government-run health clinics in Sierra Leone serving $\sim$1M
- Clinics grouped into threes
- Randomization at the clinic level
- Clinics selected to minimize spillovers
- Households in clinic catchment areas surveyed

Outcomes of interest

Outcome data: 2011 (baseline) and 2013 (endline); Ebola cases (2014–16)

Outcomes of interest:

- Health utilization (plus maternal utilization)
- Household-level health outcomes
- Household satisfaction
- Clinic quality
- Economic outcomes
- (Plus more)
- Ebola cases

Balance?

	(1) Control Mean	(2) CM	(3) NFA	(4) Difference	(5) N
Village characteristics					
Motorable road	0.891	−0.009	0.005	−0.014	503
	(0.313)	(0.036)	(0.035)	(0.035)	
Mobile phone coverage	0.812	0.058	0.096	−0.038	504
	(0.392)	(0.044)	(0.041)**	(0.037)	
Distance to the closest clinic	1.362	−0.204	0.338	−0.542	504
	(2.217)	(0.329)	(0.481)	(0.463)	
Travel cost to closest clinic	94.225	−24.273	−24.389	0.116	503
	(869.677)	(72.811)	(74.502)	(65.453)	
Household characteristics and questions to household head					
Household size	3.369	−0.061	0.007	−0.068	4774
	(2.979)	(0.056)	(0.058)	(0.059)	
Number of illness or injury cases per household	0.054	−0.039	−0.026	−0.013	4774
	(0.237)	(0.011)***	(0.012)**	(0.013)	
Birth in household last year	0.157	−0.028	0.009	−0.037	2127
	(0.363)	(0.014)**	(0.015)	(0.014)**	
Child under 2 in household	0.230	−0.013	0.027	−0.040	2126
	(0.421)	(0.018)	(0.018)	(0.018)**	
Prominent village member in household	0.042	−0.007	−0.002	−0.005	2090
	(0.200)	(0.010)	(0.010)	(0.009)	
Believes doctor's advice	0.995	0.000	−0.007	0.007	1977
	(0.072)	(0.004)	(0.005)	(0.004)*	
Health care fees unaffordable	2.307	0.023	0.030	−0.007	2057
	(0.784)	(0.045)	(0.050)	(0.046)	
Trust in the community	1.856	−0.032	−0.010	−0.021	2127
	(0.663)	(0.051)	(0.049)	(0.052)	
Community cohesion	2.420	−0.018	0.009	−0.026	2122
	(0.610)	(0.034)	(0.036)	(0.035)	
Believe VHC members represent your interest	2.743	0.094	0.159	−0.066	984
	(1.061)	(0.107)	(0.122)	(0.104)	
The VHC can be trusted	2.453	−0.171	−0.089	−0.083	1148
	(0.967)	(0.103)*	(0.106)	(0.101)	

Compliance?

Regression specification and parameters of interest

These authors estimate a (slightly more complicated) version of:

$$Y_{ivc} = \alpha + \tau^{CM} D_c^{CM} + \tau^{NFA} D_c^{NFA} + \gamma_b + \delta \mathbf{X}_{vc} + \varepsilon_{ivc}$$

where:

- Y_{ivc} is an outcome for (household i in village v) in catchment area c
- D_c^{CM}, D_c^{NFA} are treatment indicators for the two treatments
- γ_b is a triad fixed effect
- $\mathbf{X}_{vc}$ are controls: baseline outcomes
- ε_{ivc} is an error term

Findings

	(1) Control Mean	(2) Pooled	(3) CM	(4) NFA	(5) Difference	(6) Joint F-test (p)	(7) N
General utilization	0.000	0.112	0.126	0.099	0.026	7.054	4496
	(1.000)	(0.031)***	(0.034)***	(0.037)***	(0.033)	(0.001)***	
		[0.005]***	[0.003]***	[0.032]**			
Maternal utilization	0.000	0.061	0.175	−0.043	0.218	4.128	888
	(1.000)	(0.064)	(0.077)**	(0.076)	(0.081)***	(0.017)**	
		[0.327]	[0.068]*	[0.548]			
Health outcomes	0.000	0.053	0.146	−0.037	0.184	6.318	5053
	(1.000)	(0.051)	(0.056)***	(0.059)	(0.056)***	(0.002)***	
		[0.327]	[0.045]**	[0.548]			
Satisfaction	0.000	0.101	0.091	0.109	−0.018	2.876	5052
	(1.000)	(0.042)**	(0.049)*	(0.049)**	(0.048)	(0.058)*	
		[0.038]**	[0.095]*	[0.048]**			

Findings

	(1) Control Mean	(2) Pooled	(3) CM	(4) NFA	(5) Difference	(6) Joint F-test (p)	(7) N
Clinic quality	0.000	0.104	−0.004	0.213	−0.216	0.929	254
	(1.000)	(0.149)	(0.175)	(0.176)	(0.184)	(0.397)	
		[0.395]	[0.649]	[0.237]			
Health service delivery	0.000	0.039	0.070	0.027	0.043	0.507	2877
	(1.000)	(0.059)	(0.071)	(0.062)	(0.059)	(0.603)	
		[0.395]	[0.266]	[0.583]			
Community support	0.000	0.007	0.027	−0.013	0.040	0.062	508
	(1.000)	(0.095)	(0.112)	(0.109)	(0.116)	(0.940)	
		[0.891]	[0.564]	[0.619]			
CDPE	0.000	0.231	0.202	0.261	−0.059	3.849	508
	(1.000)	(0.085)***	(0.102)**	(0.101)**	(0.110)	(0.023)**	
		[0.034]**	[0.095]*	[0.032]**			

Findings

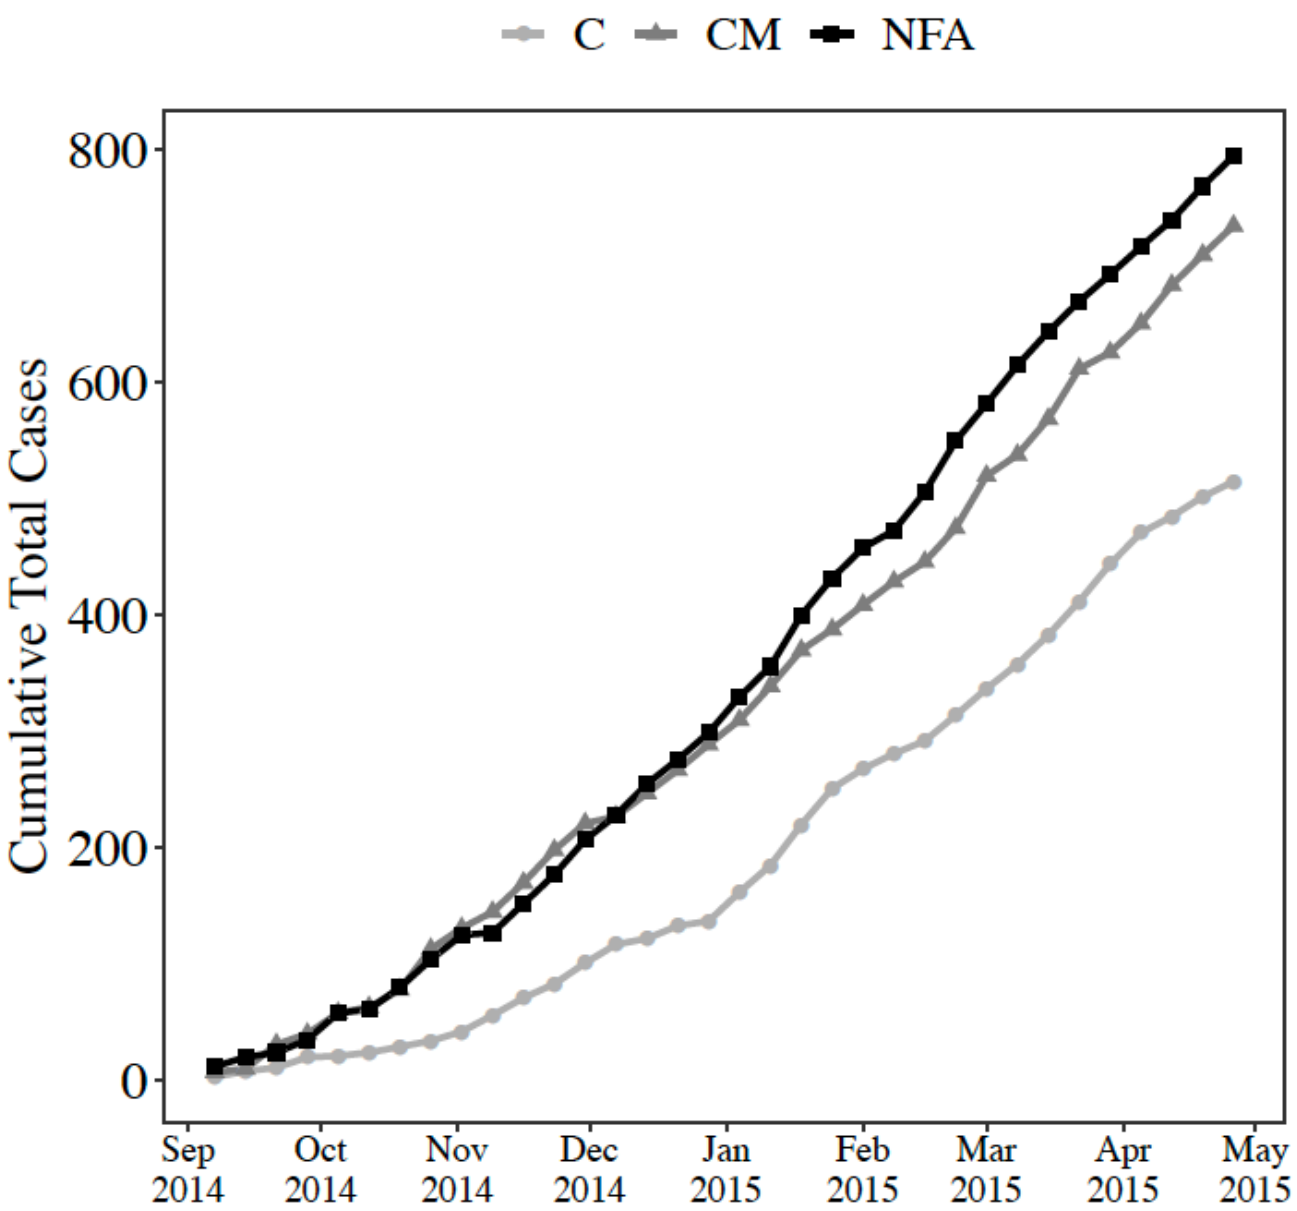
C
CM
NFA
Cumulative Total Cases
800
600
400
200
0
Sep 2014
Oct 2014
Nov 2014
Dec 2014
Jan 2015
Feb 2015
Mar 2015
Apr 2015
May 2015

Findings

Total Cases in Last 2 Weeks	Predicted Deaths in Control	Predicted Deaths in Pooled	Difference
2 reported cases	0.49	0.36	0.13
	(0.04)	(0.05)	(0.06)**
5 reported cases	1.23	0.80	0.43
	(0.11)	(0.17)	(0.19)**
10 reported cases	2.45	1.53	0.92
	(0.21)	(0.36)	(0.40)**

Recap

TL;DR:

1. RCTs are great!
2. Experiments solve our selection problem
3. Be very careful with adding controls