The Experimental Ideal

EC 607, Set 02

Edward Rubin
Spring 2020

Prologue

Schedule

Last time

Research basics, our class, and `R`

Today

Admin: Zoom recordings on Canvas.
Material: The Rubin causal model (not mine), Chapter 2 MHE.
Assignment$_1$ Make sure `R` and `RStudio` are running on your computer.
Assignment$_2$ Take 15 minutes to quietly think about your interests.
Assignment$_3$ First step of project proposal due April 15th.

Future

Lab: Meet Colleen and start deepening R knowledge.
Long run: Deepen understandings/intuitions for causality and inference.

Review

Research fundamentals

Review

Research fundamentals

Angrist and Pischke provide four **fundamental questions for research:**

1. What is the **causal relationship of interest**?
2. How would an **ideal experiment** capture this causal effect of interest?
3. What is your **identification strategy**?
4. What is your **mode of inference**?

Seemingly straightforward questions can be fundamentally unanswerable.

General research recommendations

More unsolicited advice:

- Be curious.
- Ask questions.
- Attend seminars.
- Meet faculty (UO + visitors).
- Focus on learning—especially intuition.[†]
- **Be kind and constructive.**

[†] *Learning* is not always the same as getting good grades.

The experimental ideal

The experimental ideal

What's so great about experiments?

Science widely regards **experiments as the gold standard** for research.

But why? The costs can be substantial.

Costs

- slow and expensive
- heavily regulated by review boards
- can abstract away from the actual question/setting

Benefits

So the benefits need to be pretty large, right?

The experimental ideal

Example: Hospitals and health

Imagine we want to know the **causal effect of hospitals on health**.

Research question

Within the population of poor, elderly individuals, does visiting the emergency room for primary care improve health?

Empirical exercise

1. Collect data on *health status* and *hospital visits*.
2. Summarize health status by hospital-visit group.

The experimental ideal

Example: Hospitals and health

Our empirical exercise from the 2005 National Health Inteview Survey:

Group	Sample Size	Mean Health Status	Std. Error
Hospital	7,774	3.21	0.014
No hospital	90,049	3.93	0.003

We get a t statistic of 58.9 when testing a difference in groups' means (0.72).

Conclusion? Hospitals make folks worse. Hospitals make sick people sicker.

Alternative conclusion: Perhaps we're making a mistake in our analysis... maybe sick people go to hospitals?

The experimental ideal

Potential outcomes framework

Let's develop a framework to better discuss the problem here.

- Binary treatment variable (*e.g.*, hospitalized): $\mathbf{D}_i = 0, 1$
- Outcome for individual i (*e.g.*, health): $\mathbf{Y}_i$

This framework has a few names...

- Neyman potential outcomes framework
- Rubin causal model
- Neyman-Rubin `"potential outcome"|"causal" "framework"|"model"`

The experimental ideal

Potential outcomes framework

Research question: Does $\mathbf{D}_i$ affect $\mathbf{Y}_i$?

For each individual i, there are two **potential outcomes** (w/ binary $\mathbf{D}_i$)

1. $\mathbf{Y}_{1i}$ if $\mathbf{D}_i = 1$
   i's health outcome if she went to the hospital

2. $\mathbf{Y}_{0i}$ if $\mathbf{D}_i = 0$
   i's health outcome if she did not go to the hospital

The difference between these two outcomes gives us the **causal effect of hospital treatment**, *i.e.*,

$$\tau_i = \mathbf{Y}_{1i} - \mathbf{Y}_{0i}$$

The experimental ideal

#problems

This simple equation

$$\tau_i = \mathrm{Y}_{1i} - \mathrm{Y}_{0i}$$

leads us to ***the fundamental problem of causal inference.***

> We can never simultaneously observe Y_{1i} and Y_{0i}.

Most of applied econometrics focuses on addressing this simple problem.

Accordingly, our methods try to address the related question

> For each Y_{1i}, what is a (reasonably) good counterfactual?

The experimental ideal

Solutions?

Problem We cannot directly calculate $\tau_i = \mathrm{Y}_{1i} - \mathrm{Y}_{0i}$.

Proposed solution

Compare outcomes for people who visited the hospital $(\mathrm{Y}_{1i} \mid \mathrm{D}_i = 1)$ to outcomes for people who did not visit the hospital $\big(\mathrm{Y}_{0j} \mid \mathrm{D}_j = 0\big)$.

$$E[\mathrm{Y}_i \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_i \mid \mathrm{D}_i = 0]$$

which gives us the *observed difference in health outcomes*.

Q This comparison will return *an* answer, but is it *the* answer we want?

The experimental ideal

Selection

Q What does $E[\mathrm{Y}_i \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_i \mid \mathrm{D}_i = 0]$ actually tell us?

A First notice that we can write i's outcome Y_i as

$$\mathrm{Y}_i = \mathrm{Y}_{0i} + \mathrm{D}_i \underbrace{(\mathrm{Y}_{1i} - \mathrm{Y}_{0i})}_{\tau_i}$$

Now write out our expectation, apply this definition, do creative math.

$$\begin{aligned} &E[\mathrm{Y}_i \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_i \mid \mathrm{D}_i = 0] \\ &\quad = E[\mathrm{Y}_{1i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 0] \\ &\quad = \underbrace{E[\mathrm{Y}_{1i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 1]}_{\text{Average treatment effect on the treated 😀}} + \underbrace{E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 0]}_{\text{Selection bias 😞}} \end{aligned}$$

The experimental ideal

Selection

The **first term** is *good variation*—essentially the answer that we want.

$$
\begin{aligned}
&E[\mathrm{Y}_{1i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 1] \\
&\quad = E[\mathrm{Y}_{1i} - \mathrm{Y}_{0i} \mid \mathrm{D}_i = 1] \\
&\quad = E[\tau_i \mid \mathrm{D}_i = 1]
\end{aligned}
$$

The **average causal effect** of hospitalization *for hospitalized individuals*.

The **second term** is bad variation—preventing us from knowing the answer.

$$E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 0]$$

The difference in the average untreated outcome between the treatment and control groups.

Selection bias The extent to which the "control group" provides a bad counterfactual for the treated individuals.

The experimental ideal

Selection

Angrist and Pischke (MHE, p. 15),

> The goal of most empirical economic research is to overcome selection bias, and therefore to say something about the causal effect of a variable like $\mathbf{D}_i$.

Q So how do experiments—the gold standard of empirical economic (and scientific) research—accomplish this goal and overcome selection bias?

The experimental ideal

Back to experiments

Q How do experiments overcome selection bias?

A Experiments break the link between potential outcomes and treatment.

In other words: Randomly assigning $\mathbf{D}_i$ makes $\mathbf{D}_i$ independent of which outcome we observe (meaning Y_{1i} or Y_{0i}).

Difference in means with random assignment of $\mathbf{D}_i$

$$
\begin{aligned}
&E[\mathrm{Y}_i \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_i \mid \mathrm{D}_i = 0] \\
&\quad = E[\mathrm{Y}_{1i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 0] \\
&\quad = E[\mathrm{Y}_{1i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 1] && \text{from random assignment of } \mathbf{D}_i \\
&\quad = E[\mathrm{Y}_{1i} - \mathrm{Y}_{0i} \mid \mathrm{D}_i = 1] \\
&\quad = E[\tau_i \mid \mathrm{D}_i = 1] \\
&\quad = E[\tau_i] && \text{Random assignment of } \mathbf{D}_i \text{ breaks selection bias.}
\end{aligned}
$$

The experimental ideal

Randomly assigned treatment

The key to avoiding selection bias: **random assignment of treatment** (or *as-good-as random assignment*, *e.g.*, natural experiments).

Random assignment of treatment gives us

$$E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 0] = E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 1]$$

meaning the control group's mean now provides a good counterfactual for the treatment group's mean.

In other words, there is no selection bias, *i.e.*,

$$\text{Selection bias} = E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_{0i} \mid \mathrm{D}_i = 0] = 0$$

The experimental ideal

Randomly assigned treatment

Additional benefit of randomization:

The *average treatment effect* is now representative of the *population average*, rather than the *treatment-group average*.

$$E[\tau_i \mid \mathrm{D}_i = 1] = E[\tau_i \mid \mathrm{D}_i = 0] = E[\tau_i]$$

The experimental ideal

Example: Training programs

Governments subsidize training programs to assist disadvantaged workers.

Q Do these programs have the desired effects (*i.e.*, increase wages)?

A Observational studies—comparing wage data from participants and non-participants—often find that people who complete these programs actually make lower wages.

Challenges Participants self select. **+** Programs target lower-wage workers.

The experimental ideal

Example: Training programs

How do we formalize these concerns in our framework?

Observational program evaluations

$$E[\text{Wage}_i \mid \text{Program}_i = 1] - E[\text{Wage}_i \mid \text{Program}_i = 0] =$$

$$\underbrace{E[\text{Wage}_{1i} \mid \text{Program}_i = 1] - E[\text{Wage}_{0i} \mid \text{Program}_i = 1]}_{\text{Average causal effect of training program on wages for participants, } i.e., \bar{\tau}_1} +$$

$$\underbrace{E[\text{Wage}_{0i} \mid \text{Program}_i = 1] - E[\text{Wage}_{0i} \mid \text{Program}_i = 0]}_{\text{Selection bias}}$$

If the program attracts/selects individuals who, on average, have lower wages without the program (sort of the point of the program), then we have negative selection bias.

The experimental ideal

Example: Training programs

$$E[\text{Wage}_i \mid \text{Program}_i = 1] - E[\text{Wage}_i \mid \text{Program}_i = 0] =$$

$$E[\text{Wage}_{1i} \mid \text{Program}_i = 1] - E[\text{Wage}_{0i} \mid \text{Program}_i = 1] +$$

$$E[\text{Wage}_{0i} \mid \text{Program}_i = 1] - E[\text{Wage}_{0i} \mid \text{Program}_i = 0]$$

So even if the program, on average, has an positive wage effect (in the participant group), *i.e.*, $\bar{\tau}_1 > 0$, we will detect a lower effect due to the negative selection bias.

If the bias is sufficiently large (relative to the treatment effect), our estimate will even get the sign of the effect wrong.

Related While observational studies typically found negative program effects, several experiments found positive program effects.

The experimental ideal

Example: The STAR experiment

The Tennessee STAR experiment is a famous/popular example of an experiment that allows us to answer an important social/policy question.

Research question Do classroom resources affect student performance?

- Statewide(-ish) in Tennessee for the 1985–1986 kindergarten cohort
- Ran for 4 years with ~11,600 children. Cost ~$12 million.

Treatments

1. *Small* classes (13–17 students)
2. *Regular* classes (22–35 students) plus part-time teacher's aide
3. *Regular* classes (22–35 students) plus full-time teacher's aide

The experimental ideal

Example: The STAR experiment

First question Did the randomization balance participants' characteristics across the treatment groups?

Ideally, we would have pre-experiment data on outcome variable.

Unfortunately, we only have a few demographic attributes.

Table 2.2.1, MHE

	Treatment: Class Size			
Variable	**Small**	**Regular**	**Regular + Aide**	**P-value**
Free lunch	0.47	0.48	0.50	0.09
White/Asian	0.68	0.67	0.66	0.26
Age in 1985	5.44	5.43	5.42	0.32
Attrition rate	0.49	0.52	0.53	0.02
K. class size	15.10	22.40	22.80	0.00
K. test percentile	54.70	48.90	50.00	0.00

Table 2.2.1, MHE

	Treatment: Class Size			
Variable	**Small**	**Regular**	**Regular + Aide**	**P-value**
Free lunch	**0.47**	**0.48**	**0.50**	**0.09**
White/Asian	**0.68**	**0.67**	**0.66**	**0.26**
Age in 1985	**5.44**	**5.43**	**5.42**	**0.32**
Attrition rate	0.49	0.52	0.53	0.02
K. class size	15.10	22.40	22.80	0.00
K. test percentile	54.70	48.90	50.00	0.00

Demographics appear balanced across the three treatment groups.

Table 2.2.1, MHE

	Treatment: Class Size			
Variable	**Small**	**Regular**	**Regular + Aide**	**P-value**
Free lunch	0.47	0.48	0.50	0.09
White/Asian	0.68	0.67	0.66	0.26
Age in 1985	5.44	5.43	5.42	0.32
Attrition rate	**0.49**	**0.52**	**0.53**	**0.02**
K. class size	15.10	22.40	22.80	0.00
K. test percentile	54.70	48.90	50.00	0.00

The three groups differ significantly on attrition rate.

Table 2.2.1, MHE

	Treatment: Class Size			
Variable	**Small**	**Regular**	**Regular + Aide**	**P-value**
Free lunch	0.47	0.48	0.50	0.09
White/Asian	0.68	0.67	0.66	0.26
Age in 1985	5.44	5.43	5.42	0.32
Attrition rate	0.49	0.52	0.53	0.02
K. class size	**15.10**	**22.40**	**22.80**	**0.00**
K. test percentile	54.70	48.90	50.00	0.00

The randomization generated variation in the treatment.

Table 2.2.1, MHE

	Treatment: Class Size			
Variable	**Small**	**Regular**	**Regular + Aide**	**P-value**
Free lunch	0.47	0.48	0.50	0.09
White/Asian	0.68	0.67	0.66	0.26
Age in 1985	5.44	5.43	5.42	0.32
Attrition rate	0.49	0.52	0.53	0.02
K. class size	15.10	22.40	22.80	0.00
K. test percentile	**54.70**	**48.90**	**50.00**	**0.00**

The small-class treatment significantly increased test scores.

The experimental ideal

The STAR experiment

The previous table estimated/compared the treatment effects using simple differences in means.

We can make the same comparisons using regressions.

Specifically, we regress our outcome (test percentile) on dummy variables (binary indicator variables) for each treatment group.

The experimental ideal

Example of our three treatment dummies.

i	y_i	Trt_{1i}	Trt_{2i}	Trt_{3i}
1	y_1	1	0	0
2	y_2	1	0	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
ℓ	y_ℓ	1	0	0
$\ell+1$	$y_{\ell-1}$	0	1	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
p	y_p	0	1	0
$p+1$	y_{p+1}	0	0	1
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
N	y_N	0	0	1

The experimental ideal

Regression analysis

Assume for the moment that the treatment effect is constant[†], *i.e.*,

$$\mathrm{Y}_{1i} - \mathrm{Y}_{0i} = \rho \quad \forall i$$

then we can rewrite

$$\mathrm{Y}_i = \mathrm{Y}_{0i} + \mathrm{D}_i \left(\mathrm{Y}_{1i} - \mathrm{Y}_{0i} \right)$$

as

$$\mathrm{Y}_i = \underbrace{\alpha}_{=E[\mathrm{Y}_{0i}]} + \mathrm{D}_i \underbrace{\rho}_{\mathrm{Y}_{1i}-\mathrm{Y}_{0i}} + \underbrace{\eta_i}_{\mathrm{Y}_{0i}-E[\mathrm{Y}_{0i}]}$$

[†]You'll often hear econometricians say "homogeneous" (*vs.* "hetergeneous").

The experimental ideal

Regression analysis

$$\mathbf{Y}_i = \alpha + \mathbf{D}_i \rho + \eta_i$$

Now write out the conditional expectation of $\mathbf{Y}_i$ for both levels of $\mathbf{D}_i$

$$E[\mathbf{Y}_i \mid \mathrm{D}_i = 1] = E[\alpha + \rho + \eta_i \mid \mathrm{D}_i = 1] = \alpha + \rho + E[\eta_i | \mathrm{D}_i = 1]$$

$$E[\mathbf{Y}_i \mid \mathrm{D}_i = 0] = E[\alpha + \eta_i \mid \mathrm{D}_i = 0] = \alpha + E[\eta_i \mid \mathrm{D}_i = 0]$$

Take the difference...

$$\begin{aligned} &E[\mathbf{Y}_i \mid \mathrm{D}_i = 1] - E[\mathbf{Y}_i \mid \mathrm{D}_i = 0] \\ &\quad = \rho + \underbrace{E[\eta_i | \mathrm{D}_i = 1] - E[\eta_i \mid \mathrm{D}_i = 0]}_{\text{Selection bias}} \end{aligned}$$

The experimental ideal

Regression analysis

$$E[\mathrm{Y}_i \mid \mathrm{D}_i = 1] - E[\mathrm{Y}_i \mid \mathrm{D}_i = 0] = \rho + E[\eta_i | \mathrm{D}_i = 1] - E[\eta_i \mid \mathrm{D}_i = 0]$$

Again, our estimate of the **treatment effect** (ρ) is only going to be as good as our ability to shut down the **selection bias**.

Selection bias in regression model: $E[\eta_i | \mathrm{D}_i = 1] - E[\eta_i \mid \mathrm{D}_i = 0]$

Selection bias here should remind you a lot of **omitted-variable bias.**

There is something in our disturbance η_i that is affecting Y_i and is also correlated with D_i.

In other metrics-*y* words: Our treatment D_i is endogenous.

The experimental ideal

Solutions and covariates

Selection bias in regression model: $E[\eta_i|\mathrm{D}_i = 1] - E[\eta_i \mid \mathrm{D}_i = 0]$

As before, if we randomly assign $\mathbf{D}_i$, then selection bias disappears.

Another potential route to identification is to condition on covariates in the hopes that they "take care of" the relationship between $\mathbf{D}_i$ and whatever is in our disturbance η_i.

Without very clear reasons explaining how you know you've controlled for the "bad variation", clean and convincing identification on this path is going to be challenging.

The experimental ideal

Covariates

That said, covariates can help with two things:

1. Even experiments may need **conditioning/controls**: The STAR experiment was random *within school*—not across schools.

2. Covariates can soak up unexplained variation—**increasing precision.**

Now that we've seen regression can analyze experiments, let's estimate the STAR example...

Table 2.2.2, MHE

Explanatory variable	**1**	**2**	**3**
Small class	4.82	5.37	5.36
	(2.19)	(1.26)	(1.21)
Regular + aide	0.12	0.29	0.53
	(2.23)	(1.13)	(1.09)
White/Asian			8.35
			(1.35)
Female			4.48
			(0.63)
Free lunch			-13.15
			(0.77)
School F.E.	F	T	T

The omitted level is *Regular* (with part-time aide).

Table 2.2.2, MHE

Explanatory variable	**1**	**2**	**3**
Small class	**4.82**	5.37	5.36
	(2.19)	(1.26)	(1.21)
Regular + aide	**0.12**	0.29	0.53
	(2.23)	(1.13)	(1.09)
White/Asian			8.35
			(1.35)
Female			4.48
			(0.63)
Free lunch			-13.15
			(0.77)
School F.E.	**F**	T	T

Results without other controls are very similar to the difference in means.

Table 2.2.2, MHE

Explanatory variable	**1**	**2**	**3**
Small class	4.82	**5.37**	5.36
	(2.19)	**(1.26)**	(1.21)
Regular + aide	0.12	**0.29**	0.53
	(2.23)	**(1.13)**	(1.09)
White/Asian			8.35
			(1.35)
Female			4.48
			(0.63)
Free lunch			-13.15
			(0.77)
School F.E.	F	**T**	T

School FEs enforce the experiment's design and increase precision.

Table 2.2.2, MHE

Explanatory variable	**1**	**2**	**3**
Small class	4.82	5.37	**5.36**
	(2.19)	(1.26)	**(1.21)**
Regular + aide	0.12	0.29	**0.53**
	(2.23)	(1.13)	**(1.09)**
White/Asian			**8.35**
			(1.35)
Female			**4.48**
			(0.63)
Free lunch			**-13.15**
			(0.77)
School F.E.	F	T	**T**

Additional controls slightly increase precision.

Table of contents