

ECON4150 - Introductory Econometrics

Lecture 17: Experiments

Monique de Haan
(moniqued@econ.uio.no)

Stock and Watson Chapter 13

Lecture outline

- Why study experiments?
- The potential outcome framework.
- An ideal randomized experiment, potential outcomes & regression analysis
- Conditional mean independence vs conditional mean zero
- Randomized experiments in practice
 - Threats to internal validity in a randomized experiment
 - Threats to external validity in a randomized experiment

Why study experiments?

- Ideal randomized controlled experiments provide a conceptual benchmark for assessing observational studies.
- Experiments can overcome the threats to internal validity of observational studies, however they have their own threats to internal and external validity.
- Actual experiments are rare (\$\$\$) but influential.
- Thinking about experiments helps us to understand quasi-experiments, or “natural experiments,” if “natural” variation induces “as if” random assignment (topic of next week)

Terminology: experiments and quasi-experiments

An **experiment** is designed and implemented consciously by human researchers.

- An **experiment** randomly assigns subjects to treatment and control groups (think of clinical drug trials)

A **quasi-experiment** or natural experiment has a source of randomization that is “as if” randomly assigned.

- This variation was however not the result of an explicit randomized treatment and control design.

Program evaluation is the field of econometrics aimed at evaluating the effect of a program or policy, for example, an ad campaign to cut smoking, or a job training program.

Different Types of Experiments: Three Examples

- Clinical drug trial: does a proposed drug lower cholesterol?
 - Y = cholesterol level
 - X = treatment or control group (or dose of drug)
- Job training program
 - Y = has a job, or not (or Y = wage income)
 - X = went through experimental program, or not
- Class size effect (Tennessee class size experiment)
 - Y = test score (Stanford Achievement Test)
 - X = being in a small class

The Potential Outcome Framework

- Suppose we want to know the causal effect of a binary treatment X_i on the outcome Y_i
- For example let Y_i be health and the treatment is a new medicine with

$X_i = 1 \rightarrow$ *takes new medicine*

$X_i = 0 \rightarrow$ *does not take new medicine*

- For each individual there exist two *potential outcomes*

$Y_i(1)$ is the outcome of individual i if he takes the new medicine

$Y_i(0)$ is the outcome of individual i if he does not take the new medicine

- The causal effect of the treatment on the outcome of individual i is

$$\text{Causal effect}_i = Y_i(1) - Y_i(0)$$

The Potential Outcome Framework

- The observed outcome Y_i can be written in terms of the potential outcomes:

$$Y_i = Y_i(1) \cdot X_i + Y_i(0) \cdot (1 - X_i)$$

- If the individual received the treatment ($X_i = 1$):

$$Y_i = Y_i(1) \cdot 1 + Y_i(0) \cdot 0 = Y_i(1)$$

- If the individual did not receive the treatment ($X_i = 0$):

$$Y_i = Y_i(1) \cdot 0 + Y_i(0) \cdot 1 = Y_i(0)$$

The identification problem: We cannot identify the causal effect for individual i because we observe either $Y_i(1)$ or $Y_i(0)$ but never both!

The Potential Outcome Framework & A Randomized Experiment

- Although we can never observe the causal effect for individual i , we might be able to estimate the average causal effect in a population.
- The average causal effect/ average treatment effect:

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

- Suppose we set up a ideal randomized experiment
 - we take a random sample of the population
 - we randomly give half of the sample the treatment,
 - the other half does not get the treatment.

The Potential Outcome Framework & A Randomized Experiment

- The potential outcomes can differ between individuals

$$Y_i(1) \neq Y_j(1) \quad \text{and} \quad Y_i(0) \neq Y_j(0) \quad \text{for } i \neq j$$

- However if the treatment X_i is randomly assigned the *distribution of potential outcomes* will be the same in the treatment group ($X_i = 1$) and in the control group ($X_i = 0$)
- With random assignment the potential outcomes are independent of the treatment

$$Y_i(1), Y_i(0) \perp X_i$$

- We thus have that

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

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

The Potential Outcome Framework & A Randomized Experiment

- In a randomized experiment individuals are randomly assigned to a treatment and control group, we therefore have that

$$E[Y_i(1)] = E[Y_i(1) | X_i = 1] = E[Y_i | X_i = 1]$$

$$E[Y_i(0)] = E[Y_i(0) | X_i = 0] = E[Y_i | X_i = 0]$$

- This implies that

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

- We can thus estimate the average causal effect of the treatment by taking the difference in mean outcomes of the individuals in the treated group and control group

Example: project Star

- A large-scale and influential randomized experiment: Project STAR (Student-Teacher Achievement Ratio)
- Kindergarten students and their teachers were randomly assigned to one of three groups beginning in the 1985-1986 school year:
 - small classes (13-17 students per teacher),
 - regular-size classes (22-25 students),
 - regular/aide classes (22-25 students) which also included a full-time teacher's aide.
- Over all 4 years about 11,600 students from 80 schools participated in the experiment
- Project STAR was funded by the Tennessee legislature, at a total cost of approximately \$12 million over four years.

Example: project Star

- Kindergarten students were randomly assigned to 3 groups
- To simplify we combine the regular-size classes and the regular-size classes with an aide into 1 group
- This gives two groups:
 - A **treatment group** ($X_i = 1$): students assigned to a small class (13-17 students)
 - A **control group** ($X_i = 0$): students assigned to a regular class (22-25 students)
- We are interested in the causal effect of class size on student achievement.
- The outcome variable Y_i is the Stanford Achievement Test score at the end of kindergarten.

Example: project Star

- For each student i we have two potential outcomes:
 - $Y_i(1)$ is the test score in case student i would be in a small class
 - $Y_i(0)$ is the test score in case student i would be in a regular class
- The causal effect of class size on test score for pupil i is $Y_i(1) - Y_i(0)$
 - this is unobserved.
- Because students were randomly assigned to the treatment group (small class) and the control group (regular class)
 - we can estimate the mean causal effect $E[Y_i(1) - Y_i(0)]$
 - by comparing mean test scores of the students in a small class ($E[Y_i|X_i = 1]$)
 - with the mean test scores of students in a regular class ($E[Y_i|X_i = 0]$)

Example: project Star

- Mean test score students in regular class: $E[Y_i|X_i = 0] = 918.20$
- Mean test score students in small class: $E[Y_i|X_i = 1] = 931.94$
- Estimate of average causal effect: $E[Y_i|X_i = 1] - E[Y_i|X_i = 0] = 13.74$

```
. ttest testscore, by(small_class)
```

Two-sample t test with equal variances

Group	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf. Interval]	
0	4048	918.2013	1.135017	72.21422	915.9761	920.4266
1	1738	931.9419	1.831611	76.35863	928.3495	935.5343
combined	5786	922.3287	.9695111	73.7466	920.4281	924.2293
diff		-13.74055	2.107334		-17.87172	-9.609391

```
diff = mean( 0 ) - mean( 1)                                t = -6.5204
Ho: diff = 0                                               degrees of freedom = 5784
```

```
Ha: diff < 0
Pr(T < t) = 0.0000
```

```
Ha: diff != 0
Pr(|T| > |t|) = 0.0000
```

```
Ha: diff > 0
Pr(T > t) = 1.0000
```

From Potential Outcomes to Regression

- Consider subject i drawn at random from a population and let:
 - $X_i = 1$ if subject is treated, $X_i = 0$ if not (binary treatment)
 - $Y_i(0)$ = potential outcome for subject i if untreated
 - $Y_i(1)$ = potential outcome for subject i if treated
- We saw on slide 7 that we can write the observed outcome as a function of the potential outcomes:

$$\begin{aligned}
 Y_i &= Y_i(1) \cdot X_i + Y_i(0) \cdot (1 - X_i) \\
 &= Y_i(0) + [Y_i(1) - Y_i(0)] \cdot X_i && \text{rewrite} \\
 &= E[Y_i(0)] + [Y_i(1) - Y_i(0)] \cdot X_i + [Y_i(0) - E[Y_i(0)]] && \text{add \& subtract } E[Y_i(0)]
 \end{aligned}$$

From Potential Outcomes to Regression

Let,

- $\beta_0 = E[Y_i(0)]$
- $\beta_{1i} = [Y_i(1) - Y_i(0)]$ is the causal effect for individual i
- $u_i = [Y_i(0) - E[Y_i(0)]]$

Then we have

$$Y_i = \underbrace{E[Y_i(0)]}_{\beta_0} + \underbrace{[Y_i(1) - Y_i(0)]}_{\beta_{1i}} \cdot X_i + \underbrace{[Y_i(0) - E[Y_i(0)]]}_{u_i}$$

If the causal effect is the same for all i , $\beta_{1i} = \beta_1$ for $i = 1, \dots, n$, we obtain the usual regression model

$$Y_i = \beta_0 + \beta_1 X_i + u_i$$

Potential Outcomes, Regression & a Randomized Experiment

In an ideal randomized experiment we have that

- the potential outcomes are independent of the treatment

$$Y_i(1), Y_i(0) \perp X_i$$

- We can thus estimate the average causal effect of the treatment by

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

In a regression framework this implies that:

- receiving the treatment is unrelated to the error term:

$$E[u_i|X_i] = 0$$

- We can thus estimate the average causal effect of the treatment by using OLS to estimate

$$Y_i = \beta_0 + \beta_1 \cdot X_i + u_i$$

Differences Estimator: $\hat{\beta}_1 = E[\widehat{Y_i X_i = 1}] - E[\widehat{Y_i X_i = 0}]$
--

Example: project Star

- We can therefore also estimate the average causal effect of class size by estimating a simple regression model

```
. regress testscore small_class
```

Source	SS	df	MS			
Model	229572.723	1	229572.723	Number of obs =	5786	
Residual	31232500	5784	5399.80983	F(1, 5784) =	42.51	
Total	31462072.8	5785	5438.56055	Prob > F =	0.0000	
				R-squared =	0.0073	
				Adj R-squared =	0.0071	
				Root MSE =	73.483	

testscore	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
small_class	13.74055	2.107334	6.52	0.000	9.609391	17.87172
_cons	918.2013	1.154965	795.00	0.000	915.9372	920.4655

- $\hat{\beta}_0 = E[\widehat{Y_i(0)}] = 918.20$
- $\hat{\beta}_1 = E[\widehat{Y_i(1)} - \widehat{Y_i(0)}] = 13.74$ is estimated average causal effect of being in a small class instead of a regular class

Randomization conditional on covariates

- In some experiments the treatment is randomly assigned conditional on individual characteristics
- For example, let Y_i be earnings and
 - $X_i = 1$ if individual is assigned to the treatment group that participates in a job training program
 - $X_i = 0$ if individual is assigned to the control group that does not participate in a job training program
- Suppose that the random assignment is conditional on the level of education where
 - 60% of low educated individuals are randomly assigned to the job training program,
 - 40% of high educated individuals are randomly assigned to the job training program

Randomization conditional on covariates

Education	i	X_i	$Y_i(0)$	$Y_i(1)$	causal effect	Y_i
high	1	1	10	20	10	20
high	2	1	10	20	10	20
high	3	0	10	20	10	10
high	4	0	10	20	10	10
high	5	0	10	20	10	10
low	6	1	0	10	10	10
low	7	1	0	10	10	10
low	8	1	0	10	10	10
low	9	0	0	10	10	0
low	10	0	0	10	10	0

$$E[Y_i|X_i = 1] - E[Y_i|X_i = 0] = \frac{20+20+10+10+10}{5} - \frac{10+10+10+0+0}{5} = 14 - 6 = 8$$

However, if we estimate effect conditional on education:

$$E[Y_i|X_i = 1, high] - E[Y_i|X_i = 0, high] = \frac{20+20}{2} - \frac{10+10+10}{3} = 20 - 10 = 10$$

$$E[Y_i|X_i = 1, low] - E[Y_i|X_i = 0, low] = \frac{10+10+10}{3} - \frac{0+0}{2} = 10 - 0 = 10$$

Randomization conditional on covariates

In a regression framework:

- If we estimate $Y_i = \beta_0 + \beta_1 X_i + u_i$, the conditional mean zero assumption ($E[u_i|X_i] = 0$) will be violated.
- The individuals in the control group are on average higher educated than the individuals in the treatment group
- High educated individuals generally have higher earnings.
- $\hat{\beta}_1$ will be a biased estimate of the average causal effect of the job training program due to omitted variable bias.

Randomization conditional on covariates

- However, conditional on education assignment to the treatment group is random
- If we include education as control variable we can obtain an unbiased estimate of the average causal effect of the job training program
- We will however not obtain an unbiased estimate of the effect of education,
 - because education is likely correlated with unobserved characteristics (ability, motivation)

Conditional Mean Independence (S&W appendix 7.1)

Suppose we have the following regression model

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_i + u_i$$

- Y_i is earnings
- X_i equals 1 if individual participated in job training program
- W_i equals 1 for high educated individuals 0 for low educated individuals
- Until now we have always defined the first OLS assumption to be

$$E[u_i | X_i, W_i] = 0$$

- This means that both X_i and W_i are uncorrelated with the error term
- In the example W_i is likely correlated with u_i
- But conditional on education treatment assignment is random, so conditional on W_i , X_i is uncorrelated with u_i

Conditional Mean Independence (S&W appendix 7.1)

Conditional Mean Independence : $E[u_i|X_i, W_i] = E[u_i|W_i] \neq 0$

- Under Conditional Mean Independence, OLS will give an **unbiased** estimate of the causal effect of X_i :

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_i + u_i$$

- $E[Y_i|X_i = 1, W_i] = \beta_0 + \beta_1 + \beta_2 W_i + E[u_i|X_i = 1, W_i]$
- $E[Y_i|X_i = 0, W_i] = \beta_0 + \quad + \beta_2 W_i + E[u_i|X_i = 0, W_i]$

$$\begin{aligned} E[Y_i|X_i = 1, W_i] - E[Y_i|X_i = 0, W_i] &= \beta_1 + E[u_i|X_i = 1, W_i] - E[u_i|X_i = 0, W_i] \\ &= \beta_1 + E[u_i|W_i] - E[u_i|W_i] \\ &= \beta_1 \end{aligned}$$

Conditional Mean Independence (S&W appendix 7.1)

Conditional Mean Independence : $E[u_i|X_i, W_i] = E[u_i|W_i] \neq 0$

- Under Conditional Mean Independence, OLS will give a **biased** estimate of the causal effect of W_i

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_i + u_i$$

- $E[Y_i|X_i, W_i = 1] = \beta_0 + \beta_1 X_i + \beta_2 + E[u_i|X_i, W_i = 1]$
- $E[Y_i|X_i, W_i = 0] = \beta_0 + \beta_1 X_i + \quad + E[u_i|X_i, W_i = 0]$

$$\begin{aligned} E[Y_i|X_i, W_i = 1] - E[Y_i|X_i, W_i = 0] &= \beta_2 + E[u_i|X_i, W_i = 1] - E[u_i|X_i, W_i = 0] \\ &= \beta_2 + E[u_i|W_i = 1] - E[u_i|W_i = 0] \\ &\neq \beta_2 \end{aligned}$$

- This is unproblematic as long as we are only interested in the causal effect of X_i and not in the causal effect of W_i

Conditional Mean Independence (S&W appendix 7.1)

- Concept of Conditional Mean Independence also relevant in studies with observational data.
- Often we are interested in obtaining an unbiased & consistent estimate of 1 particular variable X_i on an outcome Y_i
- We generally include control variables W_{1i}, \dots, W_{ri} to eliminate omitted variable bias
- This will give an unbiased & consistent estimate of the effect of X_i if

$$E[u_i | X_i, W_{1i}, \dots, W_{ri}] = E[u_i | W_{1i}, \dots, W_{ri}]$$

- But often we don't obtain unbiased & consistent estimates of W_{1i}, \dots, W_{ri} because

$$E[u_i | W_{1i}, \dots, W_{ri}] \neq 0$$

The differences estimator with additional regressors

- One reason to include control variables is when assignment is random conditional on observed characteristics
- Another reason is to increase the precision of the estimate of the average treatment effect.
- If you observe pre-treatment characteristics that affect the outcome variable, you can include these to reduce the variance of the error term

$$Y_i = \beta_0 + \beta_1 X_i + \varepsilon_i$$

$$Y_i = \beta_0 + \beta_1 X_i + \delta_1 W_{1i} + \dots + \delta_r W_{ri} + u_i$$

$$\text{Var}(\varepsilon_i) > \text{Var}(u_i)$$

- This will reduce the standard error of the estimated effect of the treatment.
- But never include post-treatment characteristics, these are “bad controls”!

Example: project Star

```
1 . regress testscore small_class
```

Source	SS	df	MS	Number of obs = 5786		
Model	229572.723	1	229572.723	F(1, 5784) =	42.51	
Residual	31232500	5784	5399.80983	Prob > F =	0.0000	
				R-squared =	0.0073	
				Adj R-squared =	0.0071	
Total	31462072.8	5785	5438.56055	Root MSE =	73.483	

testscore	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
small_class	13.74055	2.107334	6.52	0.000	9.609391	17.87172
_cons	918.2013	1.154965	795.00	0.000	915.9372	920.4655

```
2 . regress testscore small_class boy
```

Source	SS	df	MS	Number of obs = 5786		
Model	502140.464	2	251070.232	F(2, 5783) =	46.90	
Residual	30959932.3	5783	5353.61098	Prob > F =	0.0000	
				R-squared =	0.0160	
				Adj R-squared =	0.0156	
Total	31462072.8	5785	5438.56055	Root MSE =	73.168	

testscore	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
small_class	13.76958	2.098303	6.56	0.000	9.656121	17.88304
boy	-13.73209	1.924522	-7.14	0.000	-17.50488	-9.959309
_cons	925.2438	1.515476	610.53	0.000	922.2729	928.2147

Threats to internal validity in a randomized experiment

- Analyzing data from an ideal randomized experiment will give an unbiased & consistent estimate of the causal effect of the treatment.
- In practice, setting up an ideal randomized experiment is not easy and often things do not go as planned
- This can lead to the following threats to internal validity
 - Failure to randomize
 - Failure to follow the treatment protocol
 - Attrition
 - Experimental effects/ Hawthorne effect
 - Small samples

Failure to randomize

- The treatment might not be assigned randomly but instead is based on characteristics or preferences of the subjects
- If this is due to the fact that the experimenter assigned the treatment randomly conditional on observed characteristics...
- ...we can estimate the causal effect by including these observed characteristics in the regression (conditional mean independence)
- If the treatment is randomly assigned conditional on unobserved characteristics or preferences...
- ...the estimated treatment “effect” will reflect both the effect of the treatment and the effect of these unobserved characteristics.

Failure to randomize

- We can “check” whether the treatment was randomly assigned by comparing observed characteristics between the treatment and control group.
- Table shows mean characteristics of students assigned to small vs regular classes in project STAR

	Mean small class	Mean regular class	Difference	p-value
Gender (boy=1)	0.514	0.513	0.001	0.969
Race (black=1)	0.312	0.331	-0.019	0.140
Eligible for free lunch	0.471	0.490	-0.019	0.162

- No significant difference in the observed characteristics between those assigned to the treatment group (small class) and the control group (regular class)

Failure to follow the treatment protocol

- Even if treatment assignment is random, treatment receipt might not be.
- Individuals assigned to the treatment group might not receive the treatment
 - for example if individuals assigned to a job training program do not show up for the training sessions
- Individuals assigned to the control group might receive the treatment
 - for example if individuals assigned to the control group manage to convince the instructor and attend the training sessions.
- If X_i equals 1 if an individual *received* the treatment and 0 otherwise....
- ...regressing Y_i on X_i will give a biased estimate of the treatment effect.
- Treatment received is related to (unobserved) characteristics and preferences!

Failure to follow the treatment protocol

- If we have data on the treatment actually received X_i and on the initial random assignment Z_i ...
- ..we can use the instrumental variable approach to estimate the treatment effect.

$$Y_i = \beta_0 + \beta_1 X_i + u_i \quad X_i = \pi_0 + \pi_1 Z_i + v_i$$

- We can use the initial random assignment as instrument for the treatment actually received!

Instrument relevance: $Cov(Z_i, X_i) \neq 0$ as long as treatment assignment partially determines the treatment received, this condition holds.

Instrument exogeneity: $Cov(Z_i, u_i) = 0$ as long as treatment assignment is random, this condition holds.

Attrition

- Attrition refers to subjects dropping out of the study after being randomly assigned to the treatment or control group
- Not problematic if attrition is unrelated to the treatment.
- If attrition is related to the treatment, the OLS estimator of the treatment effect will be biased
- For example if individuals that participated in the job training program moved out of town because they found a better job (due to the training)
- This is a reincarnation of sample selection bias from Ch. 9 (the sample is selected in a way related to the outcome variable).

Experimental effect/ Hawthorne effect

Hawthorne effect: Human subjects might change their behavior, merely because they are part of an experiment.

- For example, teachers assigned to small classes might put in extra effort
- They would like the researchers to find a positive effect of small class size.
- Teachers like to teach small classes.

Experimental effect/ Hawthorne effect

- In some experiments, a “double-blind” protocol can mitigate the Hawthorne effect
 - subjects and experimenters know that they are in an experiment....
 - ...but neither knows which subjects are in the treatment group and which in the control group.
 - In this case the treatment & control group experience the same experimental effects...
 - ... and differences in outcomes can be attributed to the treatment.
- Unfortunately double-blind experiments are often not feasible within the field of economics.

Small samples

- Experiments with human subjects can be expensive.
- The sample size in experiments is therefore sometimes (too) small.
- Small samples do not produce biased estimates, but often produce imprecise estimates (large se's).
- In addition large-sample approximations might not be justified and confidence intervals and hypothesis test might not be valid.

Threats to external validity in a randomized experiment

- Can we generalize the results based on the randomized experiment to other settings and populations?

Nonrepresentative sample: The population studied and the population of interest might differ.

- Often experiments use subjects that signed up for participation in the experiment (volunteers)
- These volunteers are often more motivated.
- Even if these volunteers are randomly assigned to treatment and control group...
- ...the estimated average treatment effect might not be informative for a general population.

Threats to external validity in a randomized experiment

Nonrepresentative program or policy: Program or policy of interest might differ from the program studied.

- Experimental program is often small scale and tightly monitored.
- The quality of the actual program, when widely implemented, might therefore be lower than the experimental program.

General equilibrium effects: Turning a small experimental program into a widespread, permanent program might change the economic environment.

- An experiment testing a small scale job training program might find positive effects on earnings.
- A large scale government funded job training program might not be beneficial if it crowds out employer funded training.