PUBL0055: Introduction to Quantitative Methods

Lecture 2: Causality I

Michal Ovádek and Indraneel Sircar

Lecture Outline

Causality and counterfactuals

Why causality?

Cause-and-effect relationships are at the heart of some of the most interesting social science theories

Why causality?

Cause-and-effect relationships are at the heart of some of the most interesting social science theories

We can often express important research topics as a simple question: Does X cause Y?

Why causality?

Cause-and-effect relationships are at the heart of some of the most interesting social science theories

We can often express important research topics as a simple question: Does X cause Y?

- Does economic development lead to democracy?
- Does the race of a politician affect their chances of being elected?
- Does immigration increase support for right-wing parties?
- Does health insurance lead to better health?

Establishing causal effects using quantitative data is difficult!

Why causality?

Cause-and-effect relationships are at the heart of some of the most interesting social science theories

We can often express important research topics as a simple question: Does X cause Y?

- Does economic development lead to democracy?
- Does the race of a politician affect their chances of being elected?
- Does immigration increase support for right-wing parties?
- Does health insurance lead to better health?

Establishing causal effects using quantitative data is difficult!

Today we focus on defining what causal effects are, and how we might think about measuring them using quantitative data from experiments

Today's motivating example

Does health insurance improve health outcomes?

A critical issue of public policy is whether and how governments should provide health care to their citizens. But is there a *causal* effect of health insurance on actual levels of health? We will evaluate this question, and use it as an example to illustrate strengths and weaknesses of experimental and observational research.

- Y (Dependent variable, or "outcome"): *Health*
  - What is the self-assessed health of an individual?
- X (Independent variable, or "treatment"): *Health insurance*
  - Does the individual have health insurance, or not?

Counterfactuals

Whenever a person makes a causal statement, they are contrasting what they observe (something *factual*) with what they believe they would observe if a key condition was different (something *counterfactual*).

- "Austerity caused Brexit" (Fetzer, 2019)
  - Observation: Austerity policy in the UK and a vote to leave the EU
  - Belief: Without austerity, smaller Leave vote

Counterfactuals

Whenever a person makes a causal statement, they are contrasting what they observe (something *factual*) with what they believe they would observe if a key condition was different (something *counterfactual*).

- "Austerity caused Brexit" (Fetzer, 2019)
  - Observation: Austerity policy in the UK and a vote to leave the EU
  - Belief: Without austerity, smaller Leave vote
- "Oil wealth inhibits democracy" (Ross, 2013)
  - Observation: Oil-rich countries tend to be less democratic
  - Belief: If countries had less oil they would be more democratic

Counterfactuals

Whenever a person makes a causal statement, they are contrasting what they observe (something *factual*) with what they believe they would observe if a key condition was different (something *counterfactual*).

- "Austerity caused Brexit" (Fetzer, 2019)
  - Observation: Austerity policy in the UK and a vote to leave the EU
  - Belief: Without austerity, smaller Leave vote
- "Oil wealth inhibits democracy" (Ross, 2013)
  - Observation: Oil-rich countries tend to be less democratic
  - Belief: If countries had less oil they would be more democratic

Problem: We can't observe counterfactual outcomes! Causal inference requires *estimating* counterfactuals for comparison to realised outcomes.

Counterfactuals (example)

Health insurance $\rightarrow$ health

- Observation: People with health insurance are healthier
- Belief: If people had no health insurance, they would be less healthy

Treatment and outcome

We will think about causal relationships in terms of effects of treatments on outcomes.

Treatment and outcome

We will think about causal relationships in terms of effects of treatments on outcomes.

- Treatment: Where change originates
- Outcome: What is affected by change

Treatment and outcome

We will think about causal relationships in terms of effects of treatments on outcomes.

- Treatment: Where change originates
- Outcome: What is affected by change

We will focus on binary treatment variables:

- X_i is 1 if observation i is treated
- X_i is 0 if observation i is not treated

We will focus on continuous/interval outcome variables

Treatment and outcome (example)

In our example:

- X_i is 1 if individual i has health insurance
- X_i is 0 if individual does not have health insurance
- Y_i is individual's health

Problem: We can't observe counterfactual outcomes! Causal inference requires *estimating* counterfactuals for comparison to realised outcomes.

Potential outcomes and causal effects

We can define the causal effect of an independent/treatment variable X on an outcome variable Y by considering the **potential** outcomes of Y

Potential outcomes

The potential outcomes of Y are the values of Y that would be realised for different values of X. e.g.

- $Y_i(1) =$ the value Y_i *would have been* if X_i was equal to 1
- $Y_i(0) =$ the value Y_i *would have been* if X_i was equal to 0

Treatment effect

For any given individual, if we could observe both potential outcomes, the treatment effect of X on Y for that individual can be calculated as

$$Y_i(1) - Y_i(0)$$

Potential outcomes and causal effects (example)

What are the potential outcomes in our health insurance example?

- $Y_i(1) =$ The health individual i would have if the individual had health insurance
- $Y_i(0) =$ The health individual i would have if the individual did not have health insurance

Potential outcomes and causal effects (example)

What are the potential outcomes in our health insurance example?

- $Y_i(1) =$ The health individual i would have if the individual had health insurance
- $Y_i(0) =$ The health individual i would have if the individual did not have health insurance

What are the treatment effects?

- If $Y_i(1) > Y_i(0)$ then insurance improves health
- If $Y_i(1) < Y_i(0)$ then insurance worsens health
- If $Y_i(1) = Y_i(0)$ then insurance has no effect on health

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1			
2	1			
3	0			
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5		
2	1			
3	0			
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	
2	1			
3	0			
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1			
3	0			
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5		
3	0			
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	
3	0			
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0			
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3		
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0			

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4		

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4	3	

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4	3	1

Potential outcomes and causal effects (example)

- $X_i = 1$ if the individual is insured, and $X_i = 0$ if uninsured
- $Y_i(1)$ is the health of the individual if they were insured
- $Y_i(0)$ is the health of the individual if they were uninsured
- The treatment effect for an individual is $Y_i(1) - Y_i(0)$

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4	3	1

Average treatment effect (ATE) $= \frac{2+1+0+1}{4} = \frac{4}{4} = 1$

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1			
2	1			
3	0			
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5		
2	1			
3	0			
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	
2	1			
3	0			
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1			
3	0			
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5		
3	0			
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	
3	0			
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0			
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0	?		
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	?
4	0			

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	?
4	0	?		

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	?
4	0	?	3	

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	?
4	0	?	3	?

The Fundamental Problem of Causal Inference

But we cannot observe *both* potential outcomes for any individual!

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Causal effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	?
4	0	?	3	?

Average treatment effect (ATE) $= \frac{?+?+?+?}{4} = \frac{?}{4} = ?$

The Fundamental Problem of Causal Inference

The Fundamental Problem of Causal Inference

We only ever observe *one* potential outcome for a given individual, and our observed outcome depends on the status of our explanatory variable. We can never directly observe individual causal effects.

Consequences

- We cannot compute causal effects for individuals. *(We could perhaps in movies: Sliding Doors)*
- We have to *estimate* counterfactuals for comparison to realised outcomes.

Quantity of Interest

We want to estimate the **average treatment effect (ATE)**:

$$\frac{1}{N}\sum_{i=1}^{N}\{Y_i(1) - Y_i(0)\}$$

But we can't observe $Y_i(1)$ *and* $Y_i(0)$ for any given unit!

One alternative is to use the **difference-in-means** to *estimate* the ATE:

$$\bar{Y}_{X=1} - \bar{Y}_{X=0}$$

where $\bar{Y}_{X=1}$ and $\bar{Y}_{X=0}$ are the average *observed* outcomes for treated and control units, respectively.

Question: Is the difference-in-means equal to the ATE?

Using the difference-in-means to estimate the ATE

Question: Is the **difference-in-means** likely to equal the **ATE**?

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	?
4	0	?	3	?

Using the difference-in-means to estimate the ATE

Question: Is the **difference-in-means** likely to equal the **ATE**?

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	?	?
2	1	5	?	?
3	0	?	3	?
4	0	?	3	?

Difference-in-means $= \frac{5+5}{2} - \frac{3+3}{2} = 5 - 3 = 2$

Using the difference-in-means to estimate the ATE

Question: Is the **difference-in-means** likely to equal the **ATE**?

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4	3	1

Difference-in-means $= \frac{5+5}{2} - \frac{3+3}{2} = 5 - 3 = 2$

Using the difference-in-means to estimate the ATE

Question: Is the **difference-in-means** likely to equal the **ATE**?

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4	3	1

Difference-in-means $= \frac{5+5}{2} - \frac{3+3}{2} = 5 - 3 = 2$

Average treatment effect (ATE) $= \frac{2+1+0+1}{4} = \frac{4}{4} = 1$

Using the difference-in-means to estimate the ATE

Question: Is the **difference-in-means** likely to equal the **ATE**?

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4	3	1

Difference-in-means $= \frac{5+5}{2} - \frac{3+3}{2} = 5 - 3 = 2$

Average treatment effect (ATE) $= \frac{2+1+0+1}{4} = \frac{4}{4} = 1$

No! The difference-in-means is larger than the ATE.

Using the difference-in-means to estimate the ATE

Question: Is the **difference-in-means** likely to equal the **ATE**?

Individual	X_i	$Y_i(1)$	$Y_i(0)$	Treatment effect
1	1	5	3	2
2	1	5	4	1
3	0	3	3	0
4	0	4	3	1

Difference-in-means $= \frac{5+5}{2} - \frac{3+3}{2} = 5 - 3 = 2$

Average treatment effect (ATE) $= \frac{2+1+0+1}{4} = \frac{4}{4} = 1$

No! The difference-in-means is larger than the ATE. Why?

Confounding

Confounding differences between treatment and control groups mean that the difference-in-means can give a **biased** estimate of the ATE.

Example 1: Causal effect of finance degree on income tax policy preferences

- Finding: People who study finance are more likely to prefer lower taxes
- Confounding: People who want to study finance might prefer lower income tax to begin with!
- Direction of bias: Positive, we overstate the effect of a finance degree

Confounding

Confounding differences between treatment and control groups mean that the difference-in-means can give a **biased** estimate of the ATE.

Example 2: Causal effect of aspirin on headache symptoms

- Finding: People who take aspirin report worse headache symptoms

Confounding

Confounding differences between treatment and control groups mean that the difference-in-means can give a **biased** estimate of the ATE.

Example 2: Causal effect of aspirin on headache symptoms

- Finding: People who take aspirin report worse headache symptoms
- Confounding: People who took aspirin had headaches to begin with!

Confounding

Confounding differences between treatment and control groups mean that the difference-in-means can give a **biased** estimate of the ATE.

Example 2: Causal effect of aspirin on headache symptoms

- Finding: People who take aspirin report worse headache symptoms
- Confounding: People who took aspirin had headaches to begin with!
- Direction of bias: Negative, we understate the effectiveness of aspirin

Confounding

Confounding differences between treatment and control groups mean that the difference-in-means can give a **biased** estimate of the ATE.

$\rightarrow$ Confounding bias can be positive or negative!

Confounding

Confounding differences between treatment and control groups mean that the difference-in-means can give a **biased** estimate of the ATE.

$\rightarrow$ Confounding bias can be positive or negative!

Question: Is confounding likely in the health care example?

Confounding bias (example)

Does health insurance improve health outcomes?

The National Health Interview Survey (NHIS) is an annual survey of the US population that asks questions about health and health insurance. We are interested in two questions in particular:

- Y (Dependent variable): *health*
  - "Would you say your health in general is excellent (5), very good (4), good (3), fair (2), or poor (1)?"
- X (Independent variable): *insured*
  - "Do you have health insurance?" `TRUE` = Insured, `FALSE` = Not insured

We will also use information from some of the other questions on the survey (gender, income, race, etc).

Confounding bias example

Here are the first six rows of the NHIS data:

Table 4: NHIS data

id	insured	health	age	female	years_educ	non_white	income
1	FALSE	4	29	TRUE	14	FALSE	19282.93
2	FALSE	4	35	FALSE	11	FALSE	19282.93
3	TRUE	3	32	FALSE	12	FALSE	167844.53
4	TRUE	3	34	TRUE	16	FALSE	167844.53
5	TRUE	4	45	FALSE	12	FALSE	85985.78
6	TRUE	4	44	TRUE	12	FALSE	85985.78

Subsetting data

To calculate the difference in means, we will have to **subset** our data.

We can denote subsets of a variable using subscripts. For instance:

$$\bar{Y}_{X=1}$$

means "the average value of Y when X is equal to 1."

We can then compare the average of Y in this subset to the average of Y in another subset (e.g. $\bar{Y}_{X=0}$).

Subsetting data

We can subset our data in R using the [,] parenthesis, which allow us to select certain rows and columns from the data.

To select **rows** use the space **before the comma**. Here we select rows 1, 2, and 3, and all columns.

```
nhis[1:3,]
```

```
##   insured health age female years_educ non_white    inc
## 1   FALSE      4  29   TRUE         14     FALSE  19282
## 2   FALSE      4  35  FALSE         11     FALSE  19282
## 3    TRUE      3  32  FALSE         12     FALSE 167844
```

Subsetting data

We can subset our data in R using the `[,]` parenthesis, which allow us to select certain rows and columns from the data.

To select **columns** use the space **after the comma**. Here we select columns 1, 2, and 3 and all rows.

```
head(nhis[,1:3])
```

```
##   insured health age
## 1   FALSE      4  29
## 2   FALSE      4  35
## 3    TRUE      3  32
## 4    TRUE      3  34
## 5    TRUE      4  45
## 6    TRUE      4  44
```

We can also use **logical values** and **logical and relational operators** to select rows and columns of interest.

For instance, we can ask R to tell us in which rows in our data the respondent's value for `female` is "TRUE":

```
nhis$female == TRUE
```

Where

- the `$` says that we would like to access the `female` variable from the `nhis` data
- the == says we would like the elements of that variable that **are equal to** the value TRUE

Which gives us this (these are just the first 6 elements):

```
## [1]  TRUE FALSE FALSE  TRUE FALSE  TRUE
```

We can also use **logical values** and **logical and relational operators** to select rows and columns of interest.

We can also ask R to return all rows in our data where the respondent's value for `age` is **not** 50:

```
nhis$age != 50
```

Where

- the `$` says that we would like to access the `age` variable from the `nhis` data
- the `!=` says we would like the elements of that variable that **are NOT equal to** the value 50

We learn more logical operators (such as <, >, >=) in the seminars.

Logical values and operators

We can combine == and [ ] to select **rows that match a criterion**:

```
nhis_insured <- nhis[nhis$insured == TRUE,]
head(nhis_insured)
```

```
##   insured health age female years_educ non_white    inc
## 3    TRUE      3  32  FALSE         12     FALSE 167844
## 4    TRUE      3  34   TRUE         16     FALSE 167844
## 5    TRUE      4  45  FALSE         12     FALSE  85985
## 6    TRUE      4  44   TRUE         12     FALSE  85985
## 7    TRUE      4  49  FALSE         16     FALSE 167844
## 8    TRUE      1  55   TRUE         11     FALSE 167844
```

Logical values and operators

We can combine == and [ ] to select **rows that match a criterion**:

```
nhis_uninsured <- nhis[nhis$insured == FALSE,]
head(nhis_uninsured)
```

```
##    insured health age female years_educ non_white   inco
## 1    FALSE      4  29   TRUE         14     FALSE 19282.
## 2    FALSE      4  35  FALSE         11     FALSE 19282.
## 21   FALSE      3  59  FALSE         12      TRUE 85985.
## 22   FALSE      5  53   TRUE          7      TRUE 85985.
## 45   FALSE      1  28  FALSE          7     FALSE 19282.
## 46   FALSE      1  27   TRUE          0     FALSE 19282.
```

Logical values and operators

We can even combine multiple conditions when using == and [ ] to select **rows that match a criterion** using &:

```
nhis_insured_female <- nhis[nhis$insured == TRUE &
                            nhis$female == TRUE,]
head(nhis_insured_female)
```

```
##    insured health age female years_educ non_white    inc
## 4     TRUE      3  34   TRUE         16     FALSE 167844
## 6     TRUE      4  44   TRUE         12     FALSE  85985
## 8     TRUE      1  55   TRUE         11     FALSE 167844
## 10    TRUE      4  43   TRUE         12     FALSE  61102
## 12    TRUE      4  34   TRUE         16     FALSE  85985
## 14    TRUE      3  54   TRUE         14     FALSE  70834
```

→ `nhis_insured_female` includes units who are insured and female

Difference-in-means

To calculate the difference in means between the two groups we will use:

- the `mean()` function
- the subsetting (`[]`) operator

Question: What do we want the difference of? Between which groups?

Difference-in-means

To calculate the difference in means between the two groups we will use:

- the `mean()` function
- the subsetting (`[]`) operator

```
## Mean health level for insured individuals
mean(nhis$health[nhis$insured == TRUE])
```

```
## [1] 3.943847
```

Difference-in-means

To calculate the difference in means between the two groups we will use:

- the `mean()` function
- the subsetting (`[]`) operator

```
## Mean health level for insured individuals
mean(nhis$health[nhis$insured == TRUE])
```

```
## [1] 3.943847
```

```
## Mean health level for non-insured individuals
mean(nhis$health[nhis$insured == FALSE])
```

```
## [1] 3.617647
```

Difference-in-means

To calculate the difference in means between the two groups we will use:

- the `mean()` function
- the subsetting (`[]`) operator

```r
## Mean health level for insured individuals
mean(nhis$health[nhis$insured == TRUE])
```

```
## [1] 3.943847
```

```r
## Mean health level for non-insured individuals
mean(nhis$health[nhis$insured == FALSE])
```

```
## [1] 3.617647
```

Insured individuals report health levels that are, on average, 0.33 (8.27%) better than non-insured individuals.

Difference-in-means

Questions:

1. Can we interpret the difference in means here as causal?
2. How might we assess the possibility of confounding bias here?

Checking "balance"

One implication of confounding: treatment and control units will be different with respect to characteristics other than the treatment. We can check this by evaluating the degree of **balance** for treated and control units on different pre-treatment variables.

1. Calculate $\bar{X}_{T=1}$ (average value of pre-treatment variable for treated)
2. Calculate $\bar{X}_{T=0}$ (average value of pre-treatment variable for control)
3. Large differences between $\bar{X}_{T=1}$ and $\bar{X}_{T=0}$ suggest confounding.

Checking "balance"

One implication of confounding: treatment and control units will be different with respect to characteristics other than the treatment. We can check this by evaluating the degree of **balance** for treated and control units on different pre-treatment variables.

1. Calculate $\bar{X}_{T=1}$ (average value of pre-treatment variable for treated)
2. Calculate $\bar{X}_{T=0}$ (average value of pre-treatment variable for control)
3. Large differences between $\bar{X}_{T=1}$ and $\bar{X}_{T=0}$ suggest confounding.

For example, for age:

```
mean(nhis$age[nhis$insured == TRUE]) -
  mean(nhis$age[nhis$insured == FALSE])
```

```
## [1] 2.432754
```

$\rightarrow$ The insured are, on average, 2.4 years older than the uninsured

Checking "balance"

insured	age	female	years_educ	income
Uninsured	40.9	49.0	11.3	42892.6
Insured	43.3	50.2	14.1	101315.4
Difference	2.4	1.2	2.9	58422.9

Our insured individuals are

- older (43) than non-insured (41)
- more educated (14.3 years) than the non-insured (12 years)
- significantly richer ($101315) than the non-insured ($42893)

Implications:

1. Several characteristics are *imbalanced* with respect to insurance status.
2. Confounding is very likely a problem in this data.

Break

Randomized experiments and observational studies

Randomized experiments

- Units are randomly assigned to different X values
- Researchers directly intervene in the world they study
- Gold standard for causal inference

Observational studies

- Units are assigned to X values ”by nature”
- Researchers observe, but do not intervene in, the world
- Very commonly used in social science research

Randomized experiments and causality

Randomization and selection bias

We saw previously that, in general, the difference-in-means will not be an unbiased estimate of the ATE because of confounding.

Question: How can we avoid this problem?

Randomization and selection bias

We saw previously that, in general, the difference-in-means will not be an unbiased estimate of the ATE because of confounding.

Question: How can we avoid this problem? Answer: Conduct a randomised experiment!

Randomization and selection bias

We saw previously that, in general, the difference-in-means will not be an unbiased estimate of the ATE because of confounding.

Question: How can we avoid this problem? Answer: Conduct a randomised experiment!

Key intuition

Randomisation of treatment doesn't eliminate differences between individuals, but ensures that the mix of individuals being compared is the same on average.

Consequence

We cannot use randomisation to calculate individual causal effects, but we can use it to estimate average causal effects.

Randomization and selection bias

When the treatment, X, is randomly assigned to units...

- ... treated and untreated groups will be similar, on average, in terms of all characteristics (both observed and unobserved)
- ... the only *systematic* difference between the two will be the receipt of treatment
- ... the average outcome for controls will be similar to what *would have happened* for the treatment group if they had not been treated

$\rightarrow$ we can use the difference in means to estimate the ATE!

Randomization and selection bias (visualization)

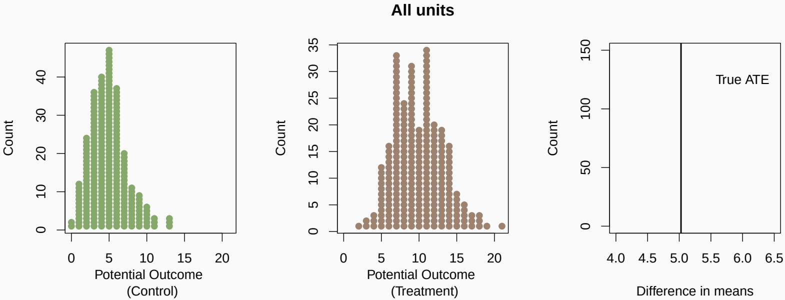

Randomization and selection bias (visualization)

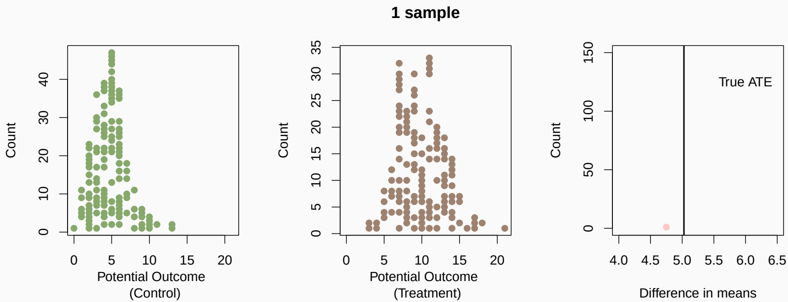

Randomization and selection bias (visualization)

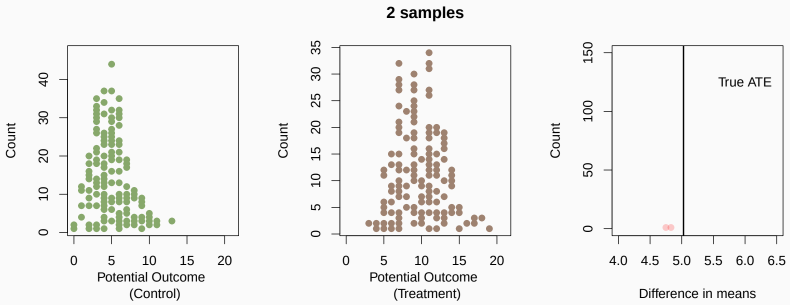

Randomization and selection bias (visualization)

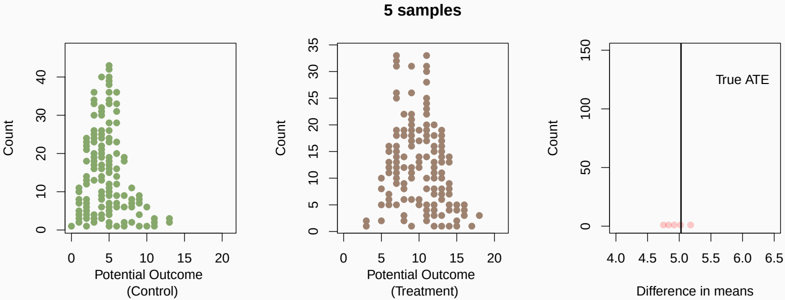

Randomization and selection bias (visualization)

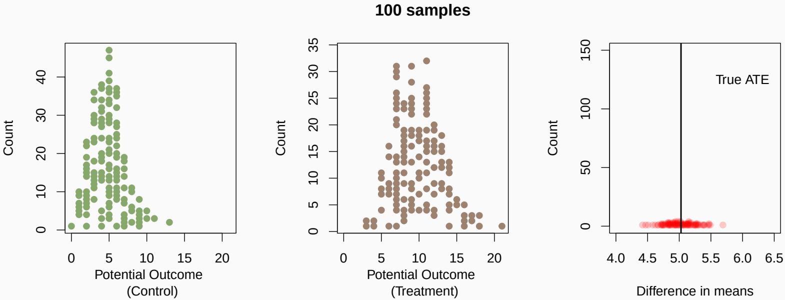

Randomization and selection bias (visualization)

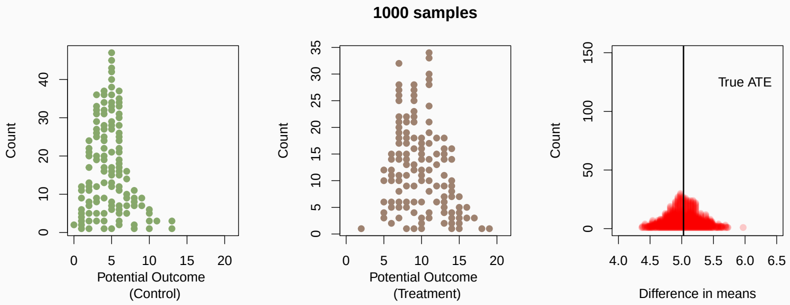

Randomization and selection bias (visualization)

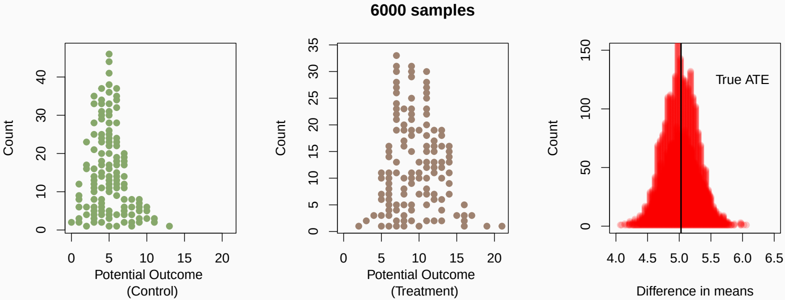

Randomization and selection bias

- Randomization of the treatment makes the difference in group means an **unbiased** estimator of the true ATE
- On average, you can expect a randomized experiment to get the right answer
- This does not guarantee that the answer you get from any particular randomization will be exactly correct!
- Considered the gold standard of causal inference

Example – Experimental data

Does health insurance improve health outcomes?

The RAND Health Insurance Experiment (RAND) was an experiment conducted between 1974 and 1982 in the US. In this experiment, researchers *randomly allocated* individuals to receive health insurance.

- Y (Dependent variable): *health*
  - "Would you say your health in general is excellent (5), very good (4), good (3), fair (2), or poor (1)?"
- X (Independent variable): *insured*
  - Was the participant *randomly allocated* to receive health insurance? `TRUE` = Insured, `FALSE` = Not insured

We will again use information from some of the other questions on the survey (gender, income, race, etc).

Experimental data

Here are the first six rows of the RAND data:

Table 6: RAND data

id	insured	health	age	female	years_educ	non_white	income
1	FALSE	4	42	FALSE	12	TRUE	67486.484
2	FALSE	4	43	TRUE	12	TRUE	67486.484
3	TRUE	2	38	TRUE	12	TRUE	27608.107
4	TRUE	3	24	FALSE	12	TRUE	4322.203
5	TRUE	4	60	TRUE	9	TRUE	24540.541
6	TRUE	4	59	FALSE	4	TRUE	24540.541

Difference-in-means - Experimental data

Let's calculate the difference in means again using the experimental data:

```
## Mean health level for insured individuals
mean(rand$health[rand$insured == TRUE])
```

```
## [1] 3.405373
```

```
## Mean health level for non-insured individuals
mean(rand$health[rand$insured == FALSE])
```

```
## [1] 3.423304
```

The difference between insured and non-insured individuals almost completely disappear in the experimental data!

Results - Experimental data

insured	age	female	years_educ	income
Uninsured	33.1	56.0	12.1	32597.2
Insured	33.2	53.4	12.0	31220.9
Difference	0.1	-2.6	-0.1	-1376.3

Insured and non-insured individuals

- are roughly equally likely to be female
- have similar ages
- have similar years of education
- have similar income levels.

Implications:

1. There is less evidence of imbalance in the experimental data
2. The experimental effect of insurance is much smaller

Observational and Experimental Results Compared

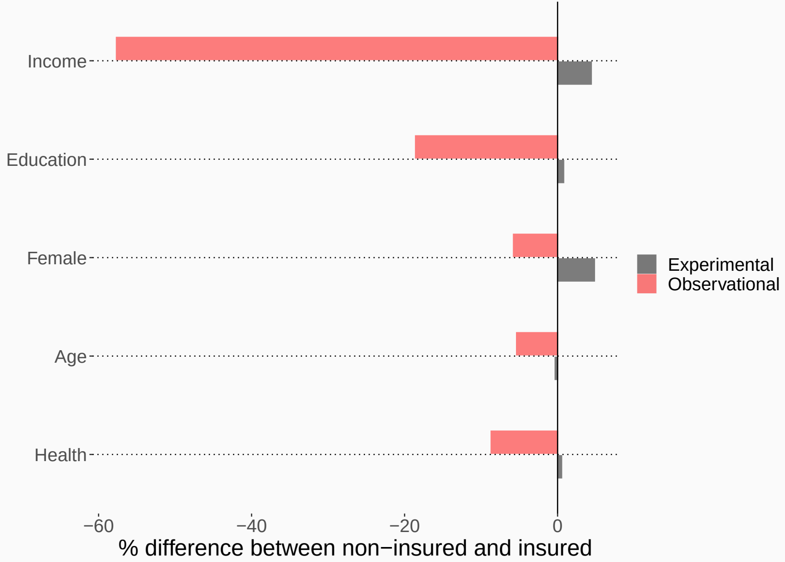

Why not always use experiments?

If randomized experiments are the "gold standard" for causal inference, why not always use them?

1. **Practical concerns**
   - It is often unfeasible to randomly vary the "treatment" you care about
   - e.g. Could you randomize the type of electoral system in a country?
2. **Ethical concerns**
   - Experiments involving real people can be ethically dubious
   - e.g. Manipulating the type of messages people see on Facebook
3. **Cost concerns**
   - Experiments can be expensive in both money and time
   - Particularly relevant to planning dissertation projects

External and internal validity

Internal validity: Can we interpret estimates in a study as causal?

- Randomized experiments tend to have high internal validity $\rightarrow$ randomization makes treatment and control groups similar on average
- Observational studies tend to have lower internal validity $\rightarrow$ pre-treatment variables may differ between treatment and control

External validity: How generalizable are the conclusions of a study?

- Randomized experiments tend to have lower external validity $\rightarrow$ difficult to conduct on representative samples
- Observational studies tend to have higher external validity $\rightarrow$ easier to use representative samples or the population itself

Picking a research design requires making trade-offs between these 2 goals.

Conclusions

What have we covered?

- Causality can be understood as a process of counterfactual reasoning
- We can understand causal effects as being the differences in *potential* outcomes for units with different treatment statuses
- Randomized experiments are the "gold standard" for causal inference
- Causal inference is hard. *But at least now you know that it is hard.*

Seminar

In seminars this week, you will learn about ...

1. ... working directories
2. ... loading data into R
3. ... calculating causal effects using experimental data
4. ... thinking critically about causality in empirical research