# Making Causal Critiques

## Day 3 - Assessing Causal Evidence

Jonathan Phillips

January 30, 2019

## Solving the Problem of Causal Inference

- We cannot!
- But we can try and minimize the risks
- Selecting units that provide appropriate counterfactuals, avoiding:
  - Omitted variable bias
  - Selection Bias
  - Reverse Causation

# Solving the Problem of Causal Inference

- Experiments
  - Field Experiments
  - Lab Experiments
  - Survey Experiments
- Quasi-Experiments
  - Instrumental Variables
  - Regresssion Discontinuity
  - Difference-in-DIfferences

# Causal Inference

Types of Research Design:

| | Researcher controls the treatment assignment | Treatment assignment mechanism likely to create comparable potential outcomes ('Conditional Independence') |
|---|---|---|
| Controlled Experiments | Yes | Yes |
| Natural Experiments | No | Yes |
| Observable Studies | No | No |

# Field Experiments

- Field experiments provide confidence because treatment assignment is **controlled by the researcher**
- But still take place in real-world environments, so they identify (hopefully) meaningful treatment effects

# Field Experiments

- Why does randomization help us achieve causal inference?

# Field Experiments

- Why does randomization help us achieve causal inference?
  - A treatment assignment mechanism that **balances potential outcomes**
  - Every unit has **exactly the same** probability of treatment
  - If treatment is randomly distributed, **so are potential outcomes**
- Potential outcomes are - on average - the same for treated and control units
  - No omitted variable bias
  - No self-selection
  - No reverse causation

# Field Experiments

- Why does randomization help us achieve causal inference?
  - We want to estimate:

$$E(Y_1 - Y_0) \tag{1}$$

# Field Experiments

- Why does randomization help us achieve causal inference?
  - We want to estimate:

$$E(Y_1 - Y_0) \tag{1}$$

  - Our data provides:

$$E(Y_1|D=1)\ ,\ E(Y_0|D=0) \tag{2}$$

# Field Experiments

- Why does randomization help us achieve causal inference?
  - We want to estimate:

$$E(Y_1 - Y_0) \tag{1}$$

  - Our data provides:

$$E(Y_1|D=1)\ ,\ E(Y_0|D=0) \tag{2}$$

  - With randomization, $Y_1, Y_0 \perp D$:

## Field Experiments

- Why does randomization help us achieve causal inference?
  - We want to estimate:

$$E(Y_1 - Y_0) \tag{1}$$

  - Our data provides:

$$E(Y_1|D=1)\ ,\ E(Y_0|D=0) \tag{2}$$

  - With randomization, $Y_1, Y_0 \perp D$:

$$E(Y_1|D=1) \quad = \quad E(Y_1)$$

# Field Experiments

- Why does randomization help us achieve causal inference?
  - We want to estimate:

$$E(Y_1 - Y_0) \tag{1}$$

  - Our data provides:

$$E(Y_1|D=1) \text{ , } E(Y_0|D=0) \tag{2}$$

  - With randomization, $Y_1, Y_0 \perp D$:

$$\begin{aligned} E(Y_1|D=1) &= E(Y_1) \\ E(Y_0|D=0) &= E(Y_0) \end{aligned} \tag{3}$$

# Field Experiments

- Why does randomization help us achieve causal inference?
  - We want to estimate:

$$E(Y_1 - Y_0) \tag{1}$$

  - Our data provides:

$$E(Y_1|D = 1)\ ,\ E(Y_0|D = 0) \tag{2}$$

  - With randomization, $Y_1, Y_0 \perp D$:

$$\begin{aligned} E(Y_1|D = 1) &= E(Y_1) & (3)\\ E(Y_0|D = 0) &= E(Y_0) & (4)\\ E(Y_1|D = 1) - E(Y_0|D = 0) &= E(Y_1) - E(Y_0) \end{aligned}$$

## Field Experiments

- Why does randomization help us achieve causal inference?
  - We want to estimate:

$$E(Y_1 - Y_0) \tag{1}$$

  - Our data provides:

$$E(Y_1|D=1)\ ,\ E(Y_0|D=0) \tag{2}$$

  - With randomization, $Y_1, Y_0 \perp D$:

$$\begin{aligned}
E(Y_1|D=1) &= E(Y_1) && (3)\\
E(Y_0|D=0) &= E(Y_0) && (4)\\
E(Y_1|D=1) - E(Y_0|D=0) &= E(Y_1) - E(Y_0) && (5)\\
&= E(Y_1 - Y_0) && (6)
\end{aligned}$$

# Field Experiments

- But these are just **expectations** (averages)

# Field Experiments

- But these are just **expectations** (averages)
    - **On average**, potential outcomes will be balanced

# Field Experiments

- But these are just **expectations** (averages)
  - **On average**, potential outcomes will be balanced
  - More likely in larger samples

# Field Experiments

- But these are just **expectations** (averages)
  - **On average**, potential outcomes will be balanced
  - More likely in larger samples
  - We cannot measure potential outcomes

# Field Experiments

- But these are just **expectations** (averages)
  - **On average**, potential outcomes will be balanced
  - More likely in larger samples
  - We cannot measure potential outcomes
  - But we can assess balance in *observable* covariates

# Field Experiments

- But these are just **expectations** (averages)
  - **On average**, potential outcomes will be balanced
  - More likely in larger samples
  - We cannot measure potential outcomes
  - But we can assess balance in *observable* covariates
  - What if some covariates are imbalanced?

# Field Experiments

- Analysing field experiments
    - Comparison of means: t-test to test significance
    - Regression achieves the same thing
    - $Y_i \sim \alpha + \beta D_i + \epsilon_i$

# Field Experiments

- Assumptions
  - **Compliance with randomization** - Treatment was truly random and accepted
  - **No Spillovers (SUTVA)** - Treatment of one unit doesn't affect potential outcomes of other units

# Field Experiments

- Limitations of Field Experiments:

# Field Experiments

- Limitations of Field Experiments:
  - Small sample sizes still prevent inference
  - Ethics
  - Logistics/Finance
  - Some treatments can't be manipulated (history)
  - Lack of control over treatment content and context - is it informative?
  - Long-term/scale effects/adaptation?

# Field Experiments

- Limitations of Field Experiments: **Internal Validity**

# Field Experiments

- Limitations of Field Experiments: **Internal Validity**
  - No guarantee of actual balance
  - Unbiased but imprecise; variation still high if lots of other variables also affect Y
  - Hawthorne effect: participants adapt behaviour in experiments
  - Biased measurement if not double-blind
  - *Average* Treatment Effect can be skewed by Outliers
  - Complications of non-compliance, attrition

# Field Experiments

- All these complications mean we need lots of assumptions and background knowledge
- Just as with other methodologies

# Lab/Survey Experiments

- Why lab and survey experiments?

# Lab/Survey Experiments

- Why lab and survey experiments?
  - Treatments we cannot administer in reality
  - Outcome measurements that are hard to take in reality
  - Random treatment assignment not permitted in reality

# Lab/Survey Experiments

- **Treatment Assignment**: Same as a Field Experiment

# Lab/Survey Experiments

- **Treatment Assignment**: Same as a Field Experiment
- **Treatment**: Not a manipulation of real world political or economic processes, but establishing controlled 'lab' conditions

# Lab/Survey Experiments

- **Treatment Assignment**: Same as a Field Experiment
- **Treatment**: Not a manipulation of real world political or economic processes, but establishing controlled 'lab' conditions
  - The advantage: Control over context helps isolate mechanisms
  - The disadvantage: Can we generalize to the real world from this artificial context?

# Instrumental Variables

- What can we do when the treatment assignment mechanism is not random, and *cannot* be randomized?

# Instrumental Variables

- What can we do when the treatment assignment mechanism is not random, and *cannot* be randomized?
- An 'instrument' is a variable which assigns *part of* treatment in an 'as-if' random way

# Instrumental Variables

- What can we do when the treatment assignment mechanism is not random, and *cannot* be randomized?
- An 'instrument' is a variable which assigns *part of* treatment in an 'as-if' random way
  - Or at least in a way which is 'exogenous' - not related to omitted variables
  - Even if other variables **also** affect treatment

# Instrumental Variables

- We can use the instrument to isolate 'as-if' random variation in treatment, and use that to estimate the effect of treatment on the outcome

# Instrumental Variables

- We can use the instrument to isolate ’as-if’ random variation in treatment, and use that to estimate the effect of treatment on the outcome
- NOT the effect of the instrument on the outcome

# Instrumental Variables

- Example Instruments:
  - Rainfall for conflict
  - Sex-composition for effect of third child
  - Distance from the coast for exposure to slave trade

# Instrumental Variables

- Instrumental Variables Assumptions
  - **Strong First Stage:** The Instrument must **affect** the treatment

# Instrumental Variables

- Instrumental Variables Assumptions
  - **Strong First Stage:** The Instrument must **affect** the treatment
  - We can test this with a simple regression: $Treatment \sim Instrument$

# Instrumental Variables

- Instrumental Variables Assumptions
  - **Strong First Stage:** The Instrument must **affect** the treatment
  - We can test this with a simple regression: $Treatment \sim Instrument$
  - The instrument should be a significant predictor of treatment
  - Rule-of-thumb: $F - statistic > 10$

# Instrumental Variables

- Instrumental Variables Assumptions:
  1. **No Spillovers (SUTVA)**

# Instrumental Variables

- Instrumental Variables Assumptions:
  1. **No Spillovers (SUTVA)**
  2. **Strong First Stage (Instrument -> Treatment)**

# Instrumental Variables

- Instrumental Variables Assumptions:
  1. **No Spillovers (SUTVA)**
  2. **Strong First Stage (Instrument -> Treatment)**
  3. **Exclusion Restriction:** The Instrument **ONLY** affects the outcome through its effect on treatment, and not directly

# Instrumental Variables

- Instrumental Variables Assumptions:
  1. **No Spillovers (SUTVA)**
  2. **Strong First Stage (Instrument -> Treatment)**
  3. **Exclusion Restriction:** The Instrument **ONLY** affects the outcome through its effect on treatment, and not directly
     - **We cannot test or prove this assumption!**

# Instrumental Variables

- Instrumental Variables Assumptions:
  1. **No Spillovers (SUTVA)**
  2. **Strong First Stage (Instrument -> Treatment)**
  3. **Exclusion Restriction:** The Instrument **ONLY** affects the outcome through its effect on treatment, and not directly
     - **We cannot test or prove this assumption!**
- Theory and qualitative evidence needed to argue that the instrument is not correlated with any other factors affecting the outcome

# Instrumental Variables

- Instrumental Variables Methodology:

# Instrumental Variables

- Instrumental Variables Methodology:
  1. Use an all-in-one package, eg. *ivreg* in the *AER* package
     - Specify the formula: $Y\ D|Instrument$

# Instrumental Variables

- Instrumental Variables Methodology:
  1. Use an all-in-one package, eg. *ivreg* in the *AER* package
     - Specify the formula: $Y\ D|Instrument$
  2. Conduct 2-Stage Least Squares:

# Instrumental Variables

- Instrumental Variables Methodology:
  1. Use an all-in-one package, eg. *ivreg* in the *AER* package
     - Specify the formula: $Y\ D|Instrument$
  2. Conduct 2-Stage Least Squares:
     - Isolate the variation in treatment caused by the instrument: $D \sim Instrument$

# Instrumental Variables

- Instrumental Variables Methodology:
  1. Use an all-in-one package, eg. *ivreg* in the *AER* package
     - Specify the formula: $Y\ D|Instrument$
  2. Conduct 2-Stage Least Squares:
     - Isolate the variation in treatment caused by the instrument: $D \sim Instrument$
     - Save the predicted values from this regression: $\hat{D} = D \sim Instrument$

# Instrumental Variables

- Instrumental Variables Methodology:
  1. Use an all-in-one package, eg. *ivreg* in the *AER* package
     - Specify the formula: $Y\ D|Instrument$
  2. Conduct 2-Stage Least Squares:
     - Isolate the variation in treatment caused by the instrument: $D \sim Instrument$
     - Save the predicted values from this regression: $\hat{D} = D \sim Instrument$
     - Estimate how the predicted values affect the outcome: $Y \sim \hat{D}$

# Instrumental Variables

- Instrumental Variables Methodology:
  1. Use an all-in-one package, eg. *ivreg* in the *AER* package
     - Specify the formula: $Y\ D|Instrument$
  2. Conduct 2-Stage Least Squares:
     - Isolate the variation in treatment caused by the instrument: $D \sim Instrument$
     - Save the predicted values from this regression: $\hat{D} = D \sim Instrument$
     - Estimate how the predicted values affect the outcome: $Y \sim \hat{D}$
     - Interpret the coefficient on $\hat{D}$

# Instrumental Variables

- IV Interpretation:

# Instrumental Variables

- IV Interpretation:
  - Your coefficient is a causal estimate ONLY for units that were actually treated **because of the instrument**

# Instrumental Variables

- IV Interpretation:
  - Your coefficient is a causal estimate ONLY for units that were actually treated **because of the instrument**
  - They don't tell us about the causal effect for other units that never responded to the instrument

# Instrumental Variables

- IV Interpretation:
  - Your coefficient is a causal estimate ONLY for units that were actually treated **because of the instrument**
  - They don’t tell us about the causal effect for other units that never responded to the instrument
  - We call our causal effect estimate a ’Local Average Treatment Effect’ (LATE)
  - ’Local’ to the units whose treatment status actually changed

# Regression Discontinuities

- As always, we need some ’as-if’ random variation in assignment to treatment to get plausible counterfactuals

## Regression Discontinuities

- As always, we need some 'as-if' random variation in assignment to treatment to get plausible counterfactuals
- Regression discontinuities take advantage of social rules that **treat similar people differently**

# Regression Discontinuities

- As always, we need some 'as-if' random variation in assignment to treatment to get plausible counterfactuals
- Regression discontinuities take advantage of social rules that **treat similar people differently**
- Specifically, similar people with slightly different 'scores' are assigned to treatment/control

# Regression Discontinuities

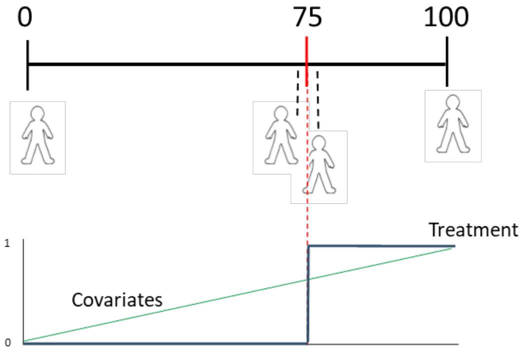

- Regression Discontinuity
  - Treatment assignment is ’as-if’ random only **really close to the threshold**

- Regression Discontinuity
  - Treatment assignment is 'as-if' random only **really close to the threshold**

$$D_i = \begin{cases} 1 & \text{if } x_i \geq \bar{x} \\ 0 & \text{if } x_i < \bar{x} \end{cases}$$

- Regression Discontinuity
  - Treatment assignment is 'as-if' random only **really close to the threshold**

$$D_i = \begin{cases} 1 & \text{if } x_i \geq \bar{x} \\ 0 & \text{if } x_i < \bar{x} \end{cases}$$

  - For units just above and below the threshold:
    - Their covariates are almost the same
    - Their potential outcomes are (on average) almost the same
    - They are plausible counterfactuals for each other

- Regression Discontinuity
  - Treatment assignment is 'as-if' random only **really close to the threshold**

$$D_i = \begin{cases} 1 & \text{if } x_i \geq \bar{x} \\ 0 & \text{if } x_i < \bar{x} \end{cases}$$

  - For units just above and below the threshold:
    - Their covariates are almost the same
    - Their potential outcomes are (on average) almost the same
    - They are plausible counterfactuals for each other
  - So we can compare them directly

- Example thresholds:
    - Exam cutoffs
    - Age cutoffs
    - Policy eligibility rules
    - Close elections
    - Adminsitrative boundaries

- Regresssion Discontinuity Variables:
  - **Running Variable, $x_i$:** The *continuous* variable to which the threshold/cutoff is applied, eg. exam score

- Regresssion Discontinuity Variables:
  - **Running Variable, $x_i$:** The *continuous* variable to which the threshold/cutoff is applied, eg. exam score
  - **Treatment, $D_i$:** Binary 0/1 depending on whether the running variable is above or below the threshold ($x_i >= \bar{x}$)

- Regresssion Discontinuity Variables:
  - **Running Variable,** $x_i$**:** The *continuous* variable to which the threshold/cutoff is applied, eg. exam score
  - **Treatment,** $D_i$**:** Binary 0/1 depending on whether the running variable is above or below the threshold ($x_i >= \bar{x}$)
  - **Outcome,** $Y_i$**:** Any subsequent outcome you have measured

- Regression Discontinuity Assumptions:
  1. No spillovers (SUTVA)

- Regression Discontinuity Assumptions:
    1. No spillovers (SUTVA)
    2. Potential outcomes vary continuously (are independent of treatment) at the threshold

- Regression Discontinuity Assumptions:
    1. No spillovers (SUTVA)
    2. Potential outcomes vary continuously (are independent of treatment) at the threshold
    3. Units cannot precisely 'manipulate' their score and sort either side of the threshold

- Regression Discontinuity Assumptions:
  1. No spillovers (SUTVA)
  2. Potential outcomes vary continuously (are independent of treatment) at the threshold
  3. Units cannot precisely ’manipulate’ their score and sort either side of the threshold
  4. The threshold is not chosen strategically

- Regression Discontinuity Assumptions:
  1. No spillovers (SUTVA)
  2. Potential outcomes vary continuously (are independent of treatment) at the threshold
  3. Units cannot precisely ’manipulate’ their score and sort either side of the threshold
  4. The threshold is not chosen strategically
  5. No compound treatments

- The threshold is more likely to be exogenous if:

- The threshold is more likely to be exogenous if:
  - Units are not aware of the threshold

- The threshold is more likely to be exogenous if:
  - Units are not aware of the threshold
  - The threshold is decided after units make choices

- The threshold is more likely to be exogenous if:
  - Units are not aware of the threshold
  - The threshold is decided after units make choices
  - The running variable is hard to manipulate precisely

- The threshold is more likely to be exogenous if:
  - Units are not aware of the threshold
  - The threshold is decided after units make choices
  - The running variable is hard to manipulate precisely
- We need qualitative evidence to support these assumptions

- We can check for sorting with a density test
- If units are bunched just above the threshold, this suggests manipulation

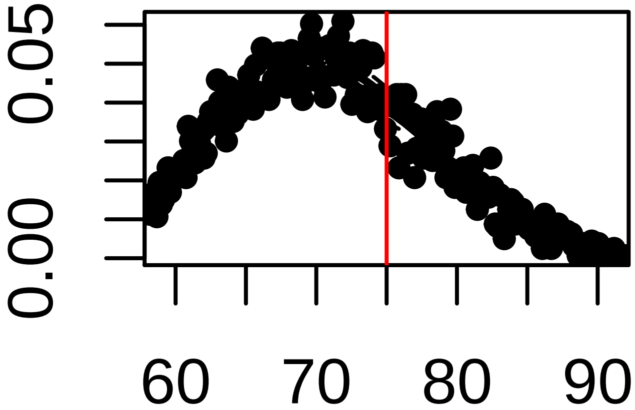

- **'Parametric' regression discontinuity:** Uses all the data and estimates:

$$Y_i = \alpha + \beta_1 Running\_Variable_i + \beta_2 Treatment_i + \epsilon_i$$

- **'Parametric' regression discontinuity:** Uses all the data and estimates:

$$Y_i = \alpha + \beta_1 Running\_Variable_i + \beta_2 Treatment_i + \epsilon_i$$

  - We just control for the 'smooth' variation in the running variable and estimate the 'jump' impact of treatment with a binary variable (dummy)

- **'Parametric' regression discontinuity:** Uses all the data and estimates:

$$Y_i = \alpha + \beta_1 Running\_Variable_i + \beta_2 Treatment_i + \epsilon_i$$

  - We just control for the 'smooth' variation in the running variable and estimate the 'jump' impact of treatment with a binary variable (dummy)
  - We may need to make the running variable non-linear

## Raw Data

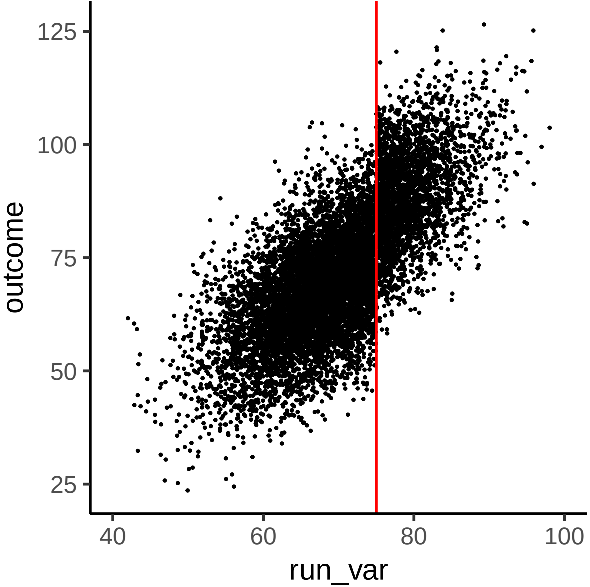

## 'Binned' Data

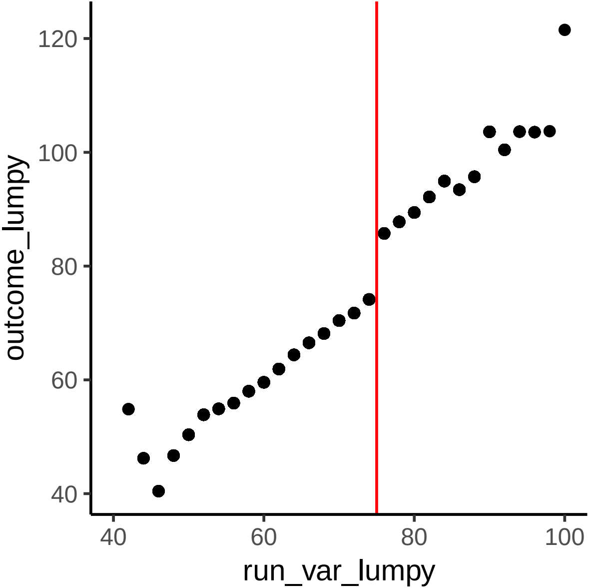

# Difference-in-Means

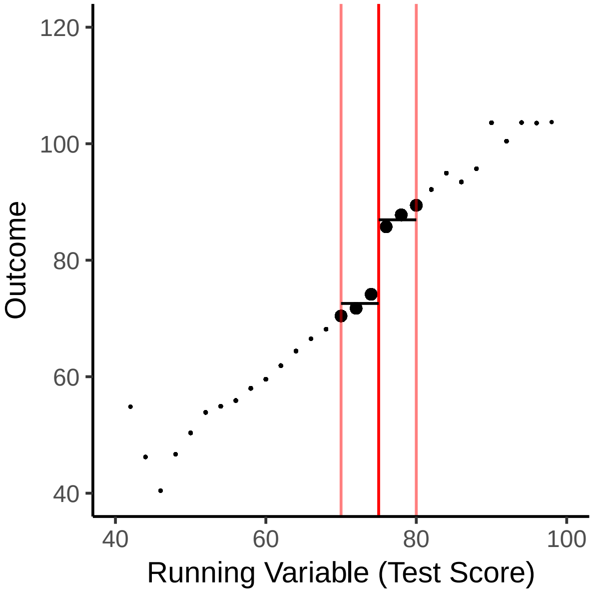

# Parametric Regression - Linear

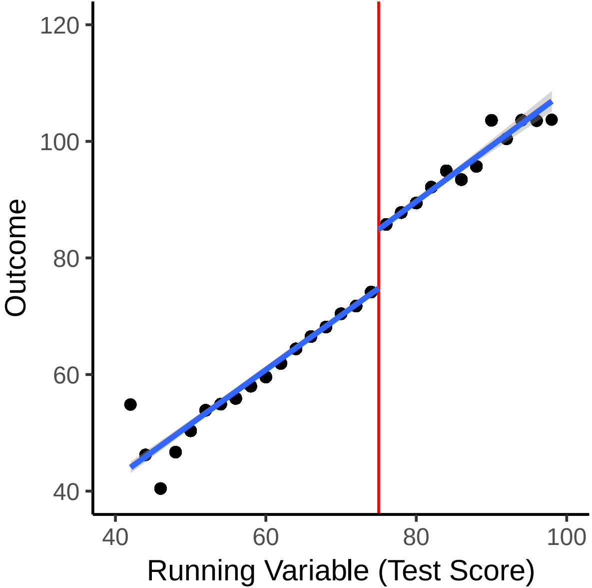

# Parametric Regression - Non-linear

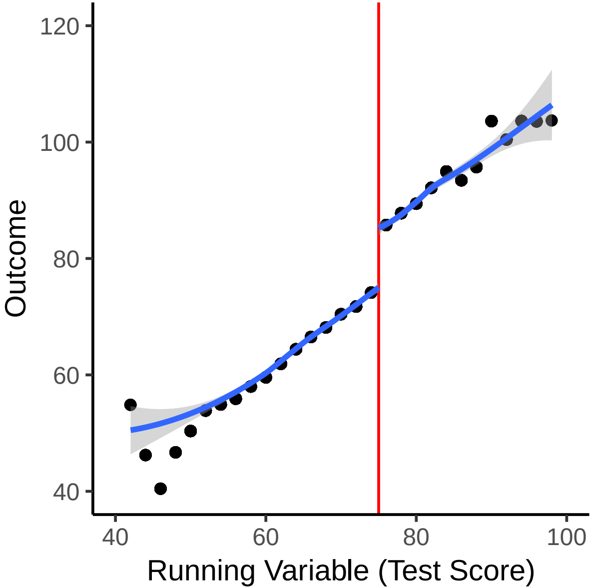

- Why does RD estimate a **Local** Average Treatment Effect?

- Why does RD estimate a **Local** Average Treatment Effect?
  - Treatment assignment is only random at the threshold

- Why does RD estimate a **Local** Average Treatment Effect?
  - Treatment assignment is only random at the threshold
  - Our estimates only apply to units close to the threshold

- Why does RD estimate a **Local** Average Treatment Effect?
    - Treatment assignment is only random at the threshold
    - Our estimates only apply to units close to the threshold
    - Units far from the threshold are very different for a reason, and causal effects are likely to be different

- Limitations:
  - Risk of sorting/manipulation

- Limitations:
  - Risk of sorting/manipulation
  - Lots of alternative specifications so no single simple test

- Limitations:
  - Risk of sorting/manipulation
  - Lots of alternative specifications so no single simple test
  - Opportunistic regression discontinuities may not identify a useful causal effect or for a relevant group

- Limitations:
  - Risk of sorting/manipulation
  - Lots of alternative specifications so no single simple test
  - Opportunistic regression discontinuities may not identify a useful causal effect or for a relevant group

- Close elections are one type of regression discontinuity in which political office is 'as-if' randomized

- Close elections are one type of regression discontinuity in which political office is 'as-if' randomized
- Particularly useful for understanding the effects of political power

- Close elections are one type of regression discontinuity in which political office is 'as-if' randomized
- Particularly useful for understanding the effects of political power
  - **Running Variable:** Margin of victory
  - **Treatment:** Winning a close election
  - **Control:** Losing a close election
  - **Outcome:** Anything that happens later...

- How much faith should we have in 'close election' regression discontinuities?

- How much faith should we have in ’close election’ regression discontinuities?
- Eggers et al (2013):

- How much faith should we have in 'close election' regression discontinuities?
- Eggers et al (2013):
  - US House of Representatives elections show sorting in very close elections (<1%)

- How much faith should we have in ’close election’ regression discontinuities?
- Eggers et al (2013):
    - US House of Representatives elections show sorting in very close elections (<1%)
    - Politicians (incumbents, the wealthy) can control whether they win, even when it’s a tight race

- How much faith should we have in ’close election’ regression discontinuities?
- Eggers et al (2013):
  - US House of Representatives elections show sorting in very close elections (<1%)
  - Politicians (incumbents, the wealthy) can control whether they win, even when it’s a tight race
  - They have extremely detailed information to predict vote results

- How much faith should we have in ’close election’ regression discontinuities?
- Eggers et al (2013):
  - US House of Representatives elections show sorting in very close elections (<1%)
  - Politicians (incumbents, the wealthy) can control whether they win, even when it’s a tight race
  - They have extremely detailed information to predict vote results
  - So potential outcomes are not balanced

- How much faith should we have in 'close election' regression discontinuities?
- Eggers et al (2013):
  - US House of Representatives elections show sorting in very close elections (<1%)
  - Politicians (incumbents, the wealthy) can control whether they win, even when it's a tight race
  - They have extremely detailed information to predict vote results
  - So potential outcomes are not balanced
  - But no other case (9 countries) has this problem

# Difference-in-Differences

- Some treatments happen at a specific point in time
  - Can't we compare the same unit before and after treatment?

# Difference-in-Differences

- Some treatments happen at a specific point in time
  - Can’t we compare the same unit before and after treatment?
  - Surely this limits the number of omitted variables - Chile today is very similar to Chile tomorrow
- But No!
  - Other factors influencing the outcome might also have changed between our measurements (eg. any news event!)
  - Eg. a worldwide recession

## Difference-in-Differences

- But what if we combine the time-series and cross-section variation?

# Difference-in-Differences

- But what if we combine the time-series and cross-section variation?
- We can keep lots of variables fixed if we compare the same unit before and after treatment

# Difference-in-Differences

- But what if we combine the time-series and cross-section variation?
- We can keep lots of variables fixed if we compare the same unit before and after treatment
- We can measure how much other factors changed over time if we have units that were not exposed to treatment

# Difference-in-Differences

- But what if we combine the time-series and cross-section variation?
- We can keep lots of variables fixed if we compare the same unit before and after treatment
- We can measure how much other factors changed over time if we have units that were not exposed to treatment
- There is nothing 'random' here, but we are more easily able to limit the risk of omitted variables

# Difference-in-Differences

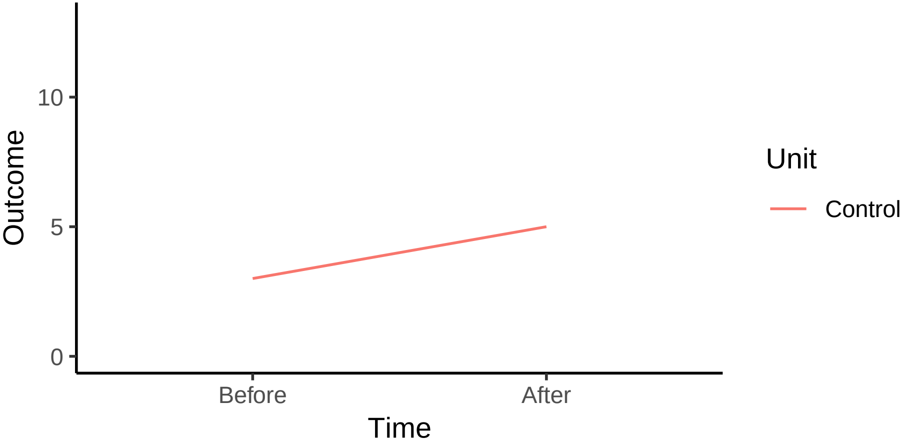

# Difference-in-Differences

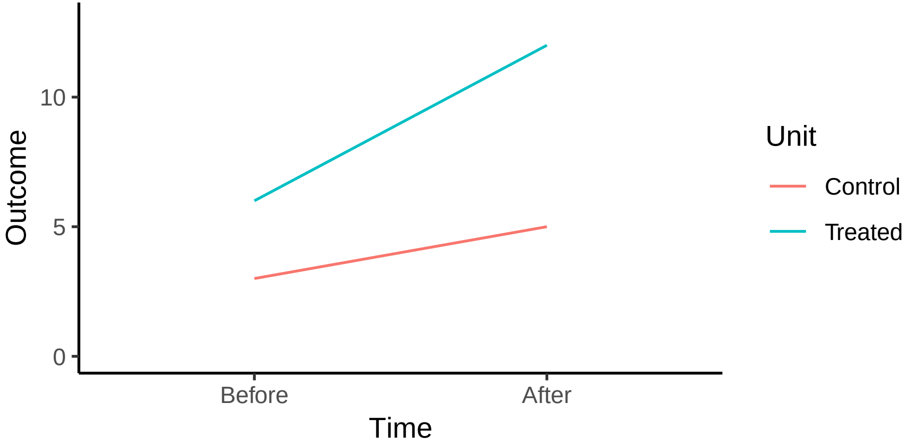

# Difference-in-Differences

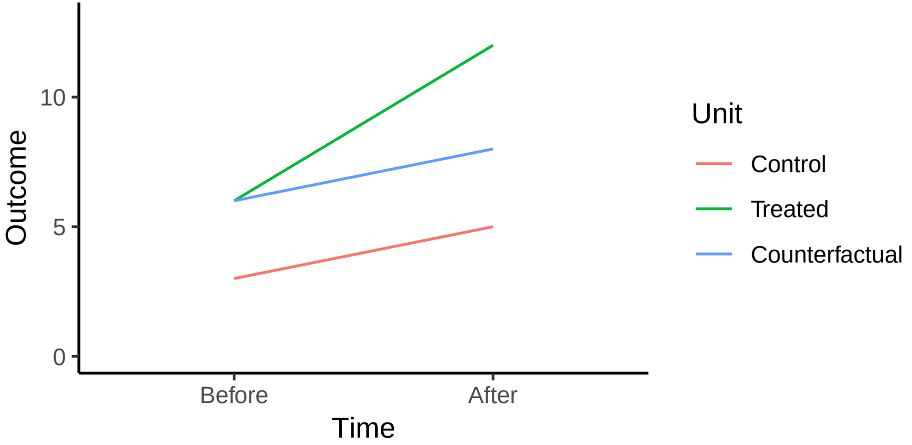

# Difference-in-Differences

- Example: How has Brexit affected the UK's growth rate?

# Difference-in-Differences

- Example: How has Brexit affected the UK's growth rate?
  - Comparing with European growth rates is biased - UK growth is influenced by oil, different labour laws etc.

# Difference-in-Differences

- Example: How has Brexit affected the UK's growth rate?
  - Comparing with European growth rates is biased - UK growth is influenced by oil, different labour laws etc.
  - Comparing before and after Brexit is biased - the world economy improved around the same time as Brexit (coincidentally)

# Difference-in-Differences

- Example: How has Brexit affected the UK's growth rate?
  - Comparing with European growth rates is biased - UK growth is influenced by oil, different labour laws etc.
  - Comparing before and after Brexit is biased - the world economy improved around the same time as Brexit (coincidentally)
  - But compare how European growth changes (-0.05%) and UK growth changed (-0.4%)

# Difference-in-Differences

- Example: How has Brexit affected the UK's growth rate?
  - Comparing with European growth rates is biased - UK growth is influenced by oil, different labour laws etc.
  - Comparing before and after Brexit is biased - the world economy improved around the same time as Brexit (coincidentally)
  - But compare how European growth changes (-0.05%) and UK growth changed (-0.4%)
  - The net effect of Brexit is -0.35%

# Difference-in-Differences

- Example: How has Brexit affected the UK's growth rate?
  - Comparing with European growth rates is biased - UK growth is influenced by oil, different labour laws etc.
  - Comparing before and after Brexit is biased - the world economy improved around the same time as Brexit (coincidentally)
  - But compare how European growth changes (-0.05%) and UK growth changed (-0.4%)
  - The net effect of Brexit is -0.35%
  - That's two differences
    - **Difference 1:** Between before and after (over time)
    - **Difference 2:** Between treated and control units

## Difference-in-Differences

- Example: How has Brexit affected the UK's growth rate?
  - Comparing with European growth rates is biased - UK growth is influenced by oil, different labour laws etc.
  - Comparing before and after Brexit is biased - the world economy improved around the same time as Brexit (coincidentally)
  - But compare how European growth changes (-0.05%) and UK growth changed (-0.4%)
  - The net effect of Brexit is -0.35%
  - That's two differences
    - **Difference 1:** Between before and after (over time)
    - **Difference 2:** Between treated and control units

# Difference-in-Differences

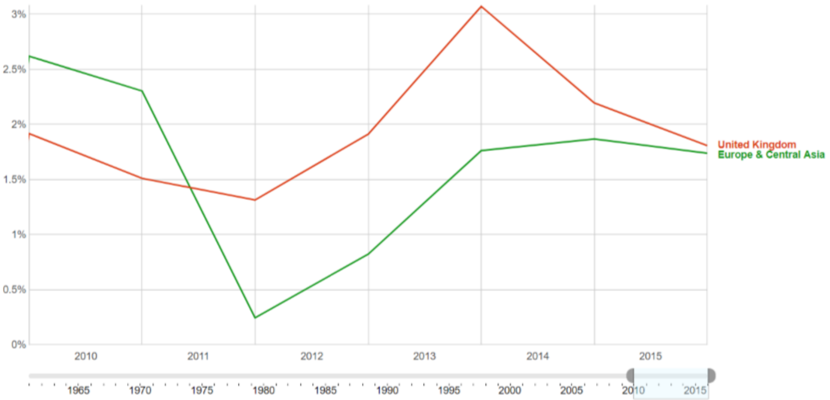

# Difference-in-Differences

- We’re now comparing *changes* (differences), not *levels* of the outcome
  - Most omitted variables affect ’levels’, so this makes our counterfactuals more plausible
    - Eg. different laws affect growth rates, not the change in growth over time
  - And crucially, we can remove omitted variables even for *unobserved* confounders

# Difference-in-Differences

- Difference-in-differences only removes **time-invariant ('levels') confounders**

## Difference-in-Differences

- Difference-in-differences only removes **time-invariant (’levels’) confounders**
  - Most omitted variables affect ’levels’, so this makes our counterfactuals more plausible

# Difference-in-Differences

- Difference-in-differences only removes **time-invariant ('levels') confounders**
  - Most omitted variables affect 'levels', so this makes our counterfactuals more plausible
  - Eg. different laws affect growth rates, not the change in growth over time

# Difference-in-Differences

- Difference-in-differences only removes **time-invariant ('levels') confounders**
  - Most omitted variables affect 'levels', so this makes our counterfactuals more plausible
  - Eg. different laws affect growth rates, not the change in growth over time
- We still need to *make the assumption or argument* that there are **no time-varying confounders**

# Difference-in-Differences

- Difference-in-differences only removes **time-invariant ('levels') confounders**
  - Most omitted variables affect 'levels', so this makes our counterfactuals more plausible
  - Eg. different laws affect growth rates, not the change in growth over time
- We still need to *make the assumption or argument* that there are **no time-varying confounders**
- Factors that affect the **trend** in the outcome *differentially* in treated and control units

# Difference-in-Differences

- Difference-in-differences only removes **time-invariant ('levels') confounders**
  - Most omitted variables affect 'levels', so this makes our counterfactuals more plausible
  - Eg. different laws affect growth rates, not the change in growth over time
- We still need to *make the assumption or argument* that there are **no time-varying confounders**
- Factors that affect the **trend** in the outcome *differentially* in treated and control units
- Eg. Even before Brexit, the UK had falling growth while growth in the eurozone was improving

# Difference-in-Differences

- Estimating Difference-in-Differences

## Difference-in-Differences

- Estimating Difference-in-Differences
- Time (Before and after) and treatment status (treated and control) are just variables in our data

## Difference-in-Differences

- Estimating Difference-in-Differences
- Time (Before and after) and treatment status (treated and control) are just variables in our data
- We know how to do a regression for the effect of treatment status on the outcome

$$Y_{it} = \alpha + \gamma D_i$$

## Difference-in-Differences

- Estimating Difference-in-Differences
- Time (Before and after) and treatment status (treated and control) are just variables in our data
- We know how to do a regression for the effect of treatment status on the outcome

$$Y_{it} = \alpha + \gamma D_i$$

- The difference-in-differences estimate is just the *interaction* of time and treatment status

$$Y_{it} = \alpha + \gamma D_i + \delta T_t + \beta D_i * T_t$$

- $\beta$ is our causal effect estimate

# Difference-in-Differences

- Assumptions Required:
  1. **No Spillovers** (SUTVA)
  2. **No time-varying confounders** (Parallel trends)
  3. **Well-defined treatment** (many things changed at the same time!)
     - Eg. The UK also announced new rules to regulate the banking sector on the same day as Brexit
  4. **Groups are stable** (eg. no migration due to treatment)

# Difference-in-Differences

- No time-varying confounders is a difficult assumption

# Difference-in-Differences

- No time-varying confounders is a difficult assumption
- Selection into treatment is usually not just due to ’fixed’ variables (eg. gender) but due to ’time-varying’ variables (eg. income, employment etc.)

# Difference-in-Differences

- No time-varying confounders is a difficult assumption
- Selection into treatment is usually not just due to ’fixed’ variables (eg. gender) but due to ’time-varying’ variables (eg. income, employment etc.)
- Eg. Participants who join a training program usually experience income falls in the previous few months

# Difference-in-Differences

- How do we know if there are time-varying confounders?

## Difference-in-Differences

- How do we know if there are time-varying confounders?
- We really want the outcome for the treated group to have the same trend as the control group

## Difference-in-Differences

- How do we know if there are time-varying confounders?
- We really want the outcome for the treated group to have the same trend as the control group
  - So any difference in trend is only due to treatment

# Difference-in-Differences

- How do we know if there are time-varying confounders?
- We really want the outcome for the treated group to have the same trend as the control group
  - So any difference in trend is only due to treatment
- One test of this is to check if **pre-treatment trends are parallel**

## Difference-in-Differences

- How do we know if there are time-varying confounders?
- We really want the outcome for the treated group to have the same trend as the control group
  - So any difference in trend is only due to treatment
- One test of this is to check if **pre-treatment trends are parallel**
- Then our counterfactual makes sense

# Difference-in-Differences

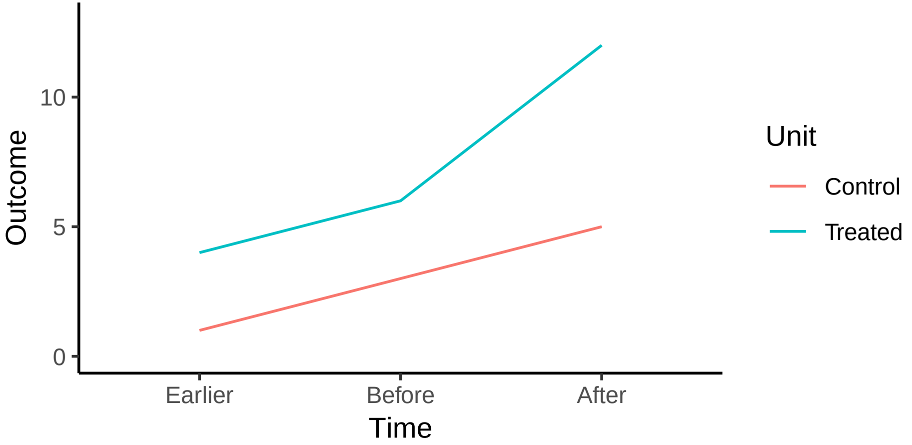

# Difference-in-Differences

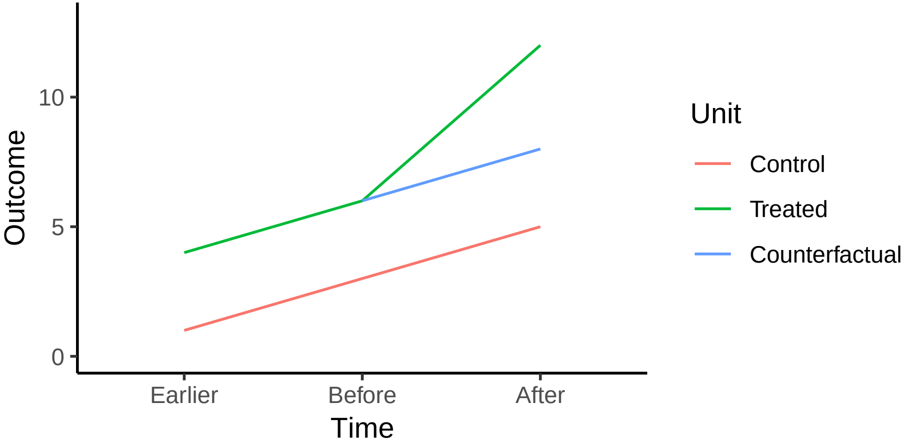

# Assumptions

Causal Methodology Assumptions

| Research Design | Assumptions required for valid causal inference |
|---|---|
| Field/Lab/Survey Experiments | No spillovers, Randomization implemented correctly, Randomization complied with, No Hawthorne Effects |
| Instrumental Variables | No Spillovers, First stage predicts treatment, Exclusion restriction |
| Regression Discontinuities | No Spillovers, Continuity (balance) of covariates, No precise manipulation, No strategic threshold, No compounding discontinuities |
| Difference-in-Differences | No Spillovers, No time-varying confounders (parallel trends), Well-defined treatment, Stable groups |