GV300 Quantitative Political Analysis

Week 16
Causality and (laboratory) experiments

Dominik Duell (University of Essex)

January 15, 2020

Housekeeping

- Midterm
- Lab session: 1-on-1 R
- We conduct an experiment in the support session next Thursday, Jan 24, 12pm.
- Change in Syllabus, Run files, slides

Recap

- This is a class on research design and statistical methods
  - I started with the most widely used statistical tools: hypothesis testing, regression
  - Useful for description and inference but only causal inference under certain assumptions
- Now, shift focus on exclusively causality driven research design
- Later we come back to advanced regression models

Questions to ask to learn about the world

1. What is the causal relationship of interest?
2. What experiment would you want to run?
3. What is your identification strategy?
4. What is your mode of inference?

Causal inference

- We want to estimate the quantity

$$E[Y_{i,1} - Y_{i,0}]$$

  where $Y_{i,D}$ is the value of some outcome Y for individual i with some value of the treatment $D \in \{0, 1\}$.
- Note that we can't ever observe the value of $Y_{i,1}$ and $Y_{i,0}$ at the same time – **Fundamental problem of causal inference**
- Overcoming this will be the focus of the second semester, but we can't do anything with it until we understand some basic probability and statistics
- Then, we will look at tools that help us to identify causal effects

Statistical modeling and causality

- Data does not just speak for itself, we need a theory to express how we think the world looks like
- This theory is reflected in the statistical model we implement
  - Model: representation for a particular purpose
  - Statistical models: representation of data for a particular purpose
    1. Description
    2. Inference
    3. Measure of effect of a manipulation

Causation or correlation?

Definitions:

- Correlation: A relationship between two variables – gives us information, helps describe the world, does not answer most of the research questions we usually ask
- Causality:
  - A ceteris paribus change in one variable has an effect on another variable (Wooldridge)
  - We may investigate
    - **Causes of effects** (e.g., what causes turnout?) – derive predictions about behavior from a theory about the world deductively and see whether peoples behavior is consistent with those predictions – think Economics
    - **Effects of causes** (e.g., does increasing voter's information (cause) increase her probability of voting – isolate a single cause-effect relationship and built knowledge inductively – think Psychology

Causality

Rubin causal model

Some definitions:

- We obtain our information about the world from data generated by some process – call this the **data generating process (DGP)**
- Sometimes the DGP is given, sometimes researchers manipulate it (i.e, experiment)
- **Experiment:** when a researcher intervenes in the DGP by purposely manipulating elements of the DGP
- **Experimental data:** Data generated by nature and the intervention of an experimentalist
- **Observational data:** Data generated by nature without intervention from an experimentalist

see also Morton/Williams, chapter 2

Notation:

- Variables
  - Dependent variable / outcome: Y_i
  - Observables X_i
  - Unobservables U_i
  - Treatment T_i – think about this as your causal variable of interest
  - Experimental manipulation M_i
    Note, T_i need not be equal to M_i
- Causal effect $D_i = Y_i^T - Y_i^C$ or $D_i = Y_i^1 - Y_i^0$
- Fundamental problem of causal inference:
  $Y_i = T_i Y_i^1 + (1 - T_i) Y_i^0$
- RCM: compare outcomes in hypothetical states of the world on unit i!

Rubin causal model:

- Also called **potential outcomes model**
- One of the many ways to think about causality – discussion about how to think about causality gets philosophical quickly

see also Morgan/Winship, chapter 1 and Morton/Williams, chapter 3

Rubin causal model:

- Assume a population of N individuals, $i = 1, \ldots, N$
- We observe some outcome Y_i for each of them
- There is a treatment variable (independent variable), which either occurs for someone (X=Treatment or $X = T$) or does not (X=Control or $X = C$) – Important: X need not be binary, binary here to introduce the model
- Gives two **potential outcomes** for i:
  Y_i^T and Y_i^C
- Effect of variable X on Y is:
  $Y_i^T - Y_i^C$

Recall: Fundamental problem of causality:

- We do not observe both, Y_i^T and Y_i^C for each individual i
  - We observe Y_i^T if i is in the treatment group (i.e., $X_i = T$)
  - We observe Y_i^T if i is in the control group (i.e., $X_i = C$)
- Any causal claim involves making an assumption about the unobserved **counterfactual**!
- Note, you can think in counterfactual terms, in terms of treatment and control group for any kind of research question no matter the data/research design you use!
- But: always state clearly what is the counterfactual.

Way forward:

- **Estimate average causal effects under defendable assumptions**
- Formulate ideal experiment – posit counterfactuals
- State and defend identification assumptions

Estimating causal effects

Strategies to estimate causal effects:

- **Condition** causal variable of interest, T_i, on variables – **regression**, stratification, matching
- Use or induce random variation in T_i:
  - Experimental manipulation – generating a exhaustive, isolated mechanism of cause on effect
  - Instrumental variable estimation
  - Regression discontinuity design
  - ...

Conditioning on variables: **(Standard) Regression analysis**:

- Tell us about correlations
- True for any statistical method – for causal claims you need **additional assumptions**
- Regression tell us what is the best estimate of $E[Y|X]$ if we limit ourselves to a line
- We assume $E[e|X] = 0$ – defendable assumption? Why not? Mostly not! **OVB**

And the remaining strategies:

- We will look into
  - Experiments
  - Instrumental variables
  - Regression discontinuity design
  - Difference-in-difference estimator
  - ...
  - Fixed-effects regression
- Some of them are better, some worse or even incapable of solving the fundamental problem of causality ...

- Fundamental problem of causal inference tackled by comparisons of averages
- Which causal effects are we estimating?
  - **Average treatment effect**: $ATE = E[D_i]$
  - **Conditional average treatment effect**: $CATE = E[D_i|X_i]$
  - **Local average treatment effect, LATE**: ATE estimated on subjects who comply with (take) the treatment
  - ...

Experiments

(Laboratory) Experiments are the gold standard

- Experiments are superior in testing causal claims
- Why? Experiments . . .
  - allow tighter **control** of confounding observable and unobservable variables – think about this as analytic decomposition of the DGP
  - allow for **randomization** – sidestepping the confounding variables problem altogether

Types of experiments

- Amazingly divers set of experiments out there
- We will talk about
  - laboratory experiments
  - survey experiments – week 17
  - field experiments – week 18
- Types of experiments differ in tightness of control over DGP and sample

What is not an experiments

- Natural experiments– we got some debate here
- Interviews, surveys
- Policy experiments
- Computer simulations
- Observational data studies

Controlling for observables and unobservables

Experiments control the DGP:

- We want to model how an outcome is generated:

  $Y_i \leftarrow T_i, X_i, U_i$
- Issues: which treatment, which observables, what's still unobserved?
- Then, how good is the treatment:

  $T_i \leftarrow M_i, Z_i, V_i$

  where Z_i are observables and V_i are unobservables

Experiments control the DGP:

- picking up elements of X_i and U_i in a manipulation (M_i, potentially in T_i)
- fixing observables Z_i and V_i in a script
- repeating observations within subject $\rightarrow$ holding elements of U_i and V_i constant
- (monetary) incentives $\rightarrow$ moving elements of U_i into X_i

$\rightarrow$ most of this is easier in the laboratory

Controlling unobservables through M_i:

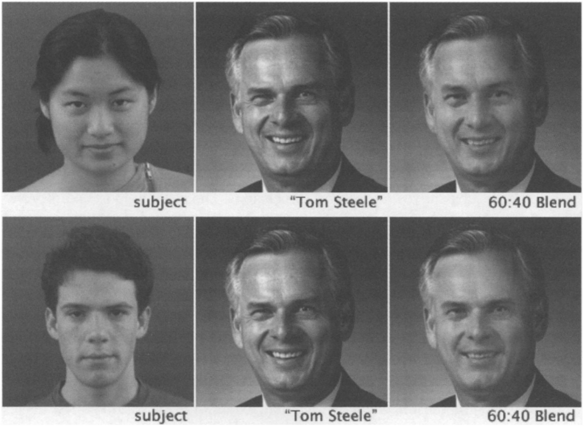

Figure 1. The original and morphed faces.

Bailenson et al (2006), "Transformed Facial Similarity as a Political Cue: A Preliminary Investigation", Political Psychology 27 (3), 373-85

Controlling unobservables through M_i:

Welcome to this experiment in decision-making during which you may earn some money that will be paid to you, privately and in cash, at the end.

The interaction in the experiment will be in pairs. You are

called Student A and the student you are matched with is called Student B. The participants in the role of Student B are from the University of Haifa.

At the beginning of the experiment you will receive NIS 20, and Student B will not receive any money. You are asked to decide whether you wish to transfer any amount of the NIS 20 to the student you are matched with and if so, how much. We will triple the amount you transfer and give it to Student B; that is, for every NIS 1 that you transfer, Student B will receive NIS 3.

In a few days time, we will ask Student B to decide if (s)he wants to return any of the money (s)he received (three times what you sent); and if so, how much. This amount will not be tripled. This will conclude the experiment, and the money will be paid.
Name of the student you are matched with (Student B): ____
Your name: ____
Amount of money you wish to transfer to Student B: ____ (Please remember that this amount should be between NIS 0 and NIS 20.)

Fersthman/Gneezy (2001), "Discrimination in a Segmented Society: An Experimental Approach", Quarterly Journal of Economics 116 (1), 351-77

Randomization

- Randomization implies:
  $$E[Y_i^T - Y_i^C] = \frac{\sum_{i=1}^{N_T} Y_i^T}{N_T} - \frac{\sum_{j=1}^{N_C} Y_j^C}{N_C}$$
  Where N_T and N_C are the number of subjects in treatment and control group, respectively
- To establish a causal effect of M (why not T?) on Y by random assignment of M we need:
  - Independence: M is statistically independent of Y
  - Perfect substitute: M is a perfect determinant of who receives treatment (T)
  - No missing data: We can perfectly observe Y associated with those affected by M.

- (Laboratory) experiments satisfy those assumptions in many cases
- Important: computation of ATE estimate as given holds in expectation! In a particular realization of random assignment ...
  - all subjects may be assigned to one group
  - may be correlated with observables or unobservables

$\leftrightarrow$ how does randomization help us with confounding variables?
Let's draw again some graphs ...

Example

Consider this experiment linked here $\Rightarrow$ and answer the following questions:

- What is the puzzle?
- What is the research question?
- What type of experiment?
- Sketch the DGP?
- What is the randomization?
- How are observable and unobservable confounds controlled?

Where to go from here?

Start your research with two planets

Whenever you pose a research question ask yourself: **What would be the ideal experiment?**

Which experiment do you want?

Experiments may be assessed by whether,

- no other factor we could have controlled for drives results $\rightarrow$ **internal validity**
- More generally,
- Internal validity is the he approximate truth of the inference or knowledge claim within a target population studied
- Our claim is **one** potential explanation, there will be others!
- Do not mix up with, **external validity**, which is the approximate truth of the inference or knowledge claim for observations beyond the target population studied.