GOV 2000 Section 10: Inference with Randomized Experiments

Konstantin Kashin
Harvard University

November 7, 2012

Outline

Administrative Details

- Problem Set 7 returned; solutions will be posted this evening
- Problem Set 9 due next Tuesday

Outline

An example of a conditionally randomized experiment

- We want to study how information on unemployment insurance programs affects the likelihood of applying for benefits
- We can do so through a survey conditionally randomized experiment
- **Treatment:** receive information on unemployment insurance program (how it works) tailored to your state
  - $T = 1$ if receive information, $T = 0$ if do not receive information
- **Outcome:** likelihood of applying for benefits (measured from 1 to 5)
  - $Y \in \{1, 2, 3, 4, 5\}$, where 5 indicates very likely to apply
- Randomization is done *conditional* upon current employment status
  - $X = 1$ if employed, $X = 0$ if unemployed

Potential Outcomes Table

i	Name	X_i	$Y_i(1)$	$Y_i(0)$
1	Benjamin Franklin	1	3	2
2	Thomas Jefferson	1	3	3
3	George Washington	1	5	2
4	Alexander Hamilton	1	1	2
5	James Madison	1	2	1
6	John Jay	1	4	3
7	Roger Sherman	0	3	4
8	Charles Pinckney	0	4	5
9	John Hancock	0	4	3
10	Robert Morris	0	1	3

Assumptions of Causal Inference

What assumptions have we already made just by writing down the potential outcome table like this?

- SUTVA in Rubin framework: no interference between units
- Implicit in Hernan and Robins framework

Defining the Estimand

What are some potential causal quantities of interest we may be interested in?

- Average treatment effect: $\mathrm{E}[Y_i(1) - Y_i(0)]$
- Stratum-specific average treatment effect: $\mathrm{E}[Y_i(1) - Y_i(0)|X_i = x]$
- Generally, anything at all that can be written in terms of the potential outcomes in our table!

Potential Outcomes Table

i	Name	X_i	$Y_i(1)$	$Y_i(0)$	T_i
1	Benjamin Franklin	1	3	2	1
2	Thomas Jefferson	1	3	3	0
3	George Washington	1	5	2	1
4	Alexander Hamilton	1	1	2	0
5	James Madison	1	2	1	1
6	John Jay	1	4	3	0
7	Roger Sherman	0	3	4	1
8	Charles Pinckney	0	4	5	1
9	John Hancock	0	4	3	0
10	Robert Morris	0	1	3	1

Potential Outcomes Table

i	Name	X_i	$Y_i(1)$	$Y_i(0)$	T_i	Y_i^{obs}
1	Benjamin Franklin	1	3	2	1	3
2	Thomas Jefferson	1	3	3	0	3
3	George Washington	1	5	2	1	5
4	Alexander Hamilton	1	1	2	0	2
5	James Madison	1	2	1	1	2
6	John Jay	1	4	3	0	3
7	Roger Sherman	0	3	4	1	3
8	Charles Pinckney	0	4	5	1	4
9	John Hancock	0	4	3	0	3
10	Robert Morris	0	1	3	1	1

Assumptions of Causal Inference

What is the assumption that enabled us to write fill out Y_i^{obs} in terms of the corresponding potential outcomes?

- $Y_i^{obs} = T_i \cdot Y_i(1) + (1 - T_1) \cdot Y_i(0)$
  - Consistency in the Hernan and Robins framework
  - Part of SUTVA in the Rubin framework

Outline

Potential Outcomes Table

Here is what we actually observe:

i	Name	X	$Y(1)$	$Y(0)$	T	Y^{obs}
1	Benjamin Franklin	1	3	?	1	3
2	Thomas Jefferson	1	?	3	0	3
3	George Washington	1	5	?	1	5
4	Alexander Hamilton	1	?	2	0	2
5	James Madison	1	2	?	1	2
6	John Jay	1	?	3	0	3
7	Roger Sherman	0	3	?	1	3
8	Charles Pinckney	0	4	?	1	4
9	John Hancock	0	?	3	0	3
10	Robert Morris	0	1	?	1	1

Estimating the Estimand

Given what we observe, how do we estimate the estimand?

Estimating the ATE

Given what we observe, how do we estimate the ATE, defined as $E[Y_i(1) - Y_i(0)]$?

Estimating the ATE

What condition must hold to estimate the ATE using a simple difference-in-means between the treated and control units?

- **Exchangeability** in the Hernan and Robins framework: $Y(0), Y(1), X \perp\!\!\!\perp T$
- **Unconfoundedness** in the Rubin framework: $P(T|Y(0), Y(1), X) = P(T)$

This is satisfied when an experiment is marginally randomized.

Estimating the ATE

We find that $\bar{Y}|(T=1)=3$ and $\bar{Y}|(T=0)=2.75$.

Therefore, $\widehat{ATE}=\bar{Y}|T=1-\bar{Y}|T=0=3-2.75=0.25$.

Estimating the ATE

Estimating the ATE using a simple difference-in-means is equivalent to imputing the missing potential outcomes as:

i	Name	X	$Y(1)$	$Y(0)$	T	Y^{obs}
1	Benjamin Franklin	1	3	2.75	1	3
2	Thomas Jefferson	1	3	3	0	3
3	George Washington	1	5	2.75	1	5
4	Alexander Hamilton	1	3	2	0	2
5	James Madison	1	2	2.75	1	2
6	John Jay	1	3	3	0	3
7	Roger Sherman	0	3	2.75	1	3
8	Charles Pinckney	0	4	2.75	1	4
9	John Hancock	0	3	3	0	3
10	Robert Morris	0	1	2.75	1	1

Link to Bivariate Regression

The estimate of the ATE is equivalent to regression coefficient from regressing the outcome on the binary treatment indicator:

$$\begin{aligned} E[Y_i(1) - Y_i(0)] &= E[Y_i(1)] - E[Y_i(0)] \\ &= E[Y_i(1)|T_i = 1] - E[Y_i(0)|T_i = 0] \\ &= E[Y_i|T_i = 1] - E[Y_i|T_i = 0] \\ &= \beta \end{aligned}$$

Violation of Assumptions

Note that if exchangeability or unconfoundedness fail to hold as we defined them above, the simple difference-in-means estimate is no longer unbiased because:

$$E[Y_i(1)] - E[Y_i(0)] \neq E[Y_i(1)|T_i = 1] - E[Y_i(0)|T_i = 0]$$

Intuitively, it is wrong because we are comparing apples and oranges. We are imputing missing potential outcomes for apples using oranges, and vice versa.

Estimating the ATE

Estimating the ATE using a simple difference-in-means is equivalent to imputing the missing potential outcomes as:

i	Name	X	$Y(1)$	$Y(0)$	T	Y^{obs}
1	Benjamin Franklin	1	3	2.75	1	3
2	Thomas Jefferson	1	3	3	0	3
3	George Washington	1	5	2.75	1	5
4	Alexander Hamilton	1	3	2	0	2
5	James Madison	1	2	2.75	1	2
6	John Jay	1	3	3	0	3
7	Roger Sherman	0	3	2.75	1	3
8	Charles Pinckney	0	4	2.75	1	4
9	John Hancock	0	3	3	0	3
10	Robert Morris	0	1	2.75	1	1

Violation of this in Our Example

- We cannot analyze a conditionally randomized experiment as a marginally randomized experiment
- $P(T|X=1) \neq P(T|X=0)$ ($\frac{1}{2} \neq \frac{3}{4}$ in our example)
- Violation of exchangeability in the Hernan and Robins framework: $Y(0), Y(1), X$ are not independent from T
- Violation of unconfoundedness in Rubin framework: $P(T|Y(0), Y(1), X) \neq P(T)$ (instead it reduced to $P(T|X)$)

Outline

Assumptions

In a conditionally randomized experiment, where treatment was randomized within levels of X, we write conditional versions of the assumptions above:

- **Conditional exchangeability** in the Hernan and Robins framework: $Y(0), Y(1) \perp\!\!\!\perp T|X$
- In the Rubin framework, assignment mechanism (for who gets treatment) is no longer **unconfoundedness** but it is **known**: $P(T|Y(0), Y(1), X) = P(T|X)$

Assumptions

Conditional exchangeability implies that *within* strata defined by levels of X, we have exchangeability and can calculate ATEs using observed difference-in-means:

$$
\begin{aligned}
E[Y_i(1) - Y_i(0)|X_i = x] &= E[Y_i(1)|X_i] - E[Y_i(0)|X_i = x] \\
&= E[Y_i(1)|T_i = 1, X_i = x] - E[Y_i(0)|T_i = 0, X_i = x] \\
&= E[Y_i|T_i = 1, X_i = x] - E[Y_i|T_i = 0, X_i = x]
\end{aligned}
$$

$\bar{Y}|(T = 1, X = x) - \bar{Y}|(T = 0, X = x)$ is an unbiased estimate for $E[Y_i|T_i = 1, X_i = x] - E[Y_i|T_i = 0, X_i = x]$.

Estimating Stratum-Specific ATEs

Estimating the stratum-specific ATE using conditional difference-in-means is equivalent to imputing the missing potential outcomes as:

i	Name	X	$Y(1)$	$Y(0)$	T	Y^{obs}
1	Benjamin Franklin	1	3	$\frac{8}{3}$	1	3
2	Thomas Jefferson	1	$\frac{10}{3}$	3	0	3
3	George Washington	1	5	$\frac{8}{3}$	1	5
4	Alexander Hamilton	1	$\frac{10}{3}$	2	0	2
5	James Madison	1	2	$\frac{8}{3}$	1	2
6	John Jay	1	$\frac{10}{3}$	3	0	3
7	Roger Sherman	0	3	3	1	3
8	Charles Pinckney	0	4	3	1	4
9	John Hancock	0	$\frac{8}{3}$	3	0	3
10	Robert Morris	0	1	3	1	1

Estimating Stratum-Specific ATEs

For the stratum defined by $X = 1$:

$$\widehat{ATE}_{X=1} = \hat{\tau}_1 = \bar{Y}|(T = 1, X = 1) - \bar{Y}|(T = 0, X = 1) = \frac{10}{3} - \frac{8}{3} = \frac{2}{3}$$

For the stratum defined by $X = 0$:

$$\widehat{ATE}_{X=0} = \hat{\tau}_0 = \bar{Y}|(T = 1, X = 0) - \bar{Y}|(T = 0, X = 0) = \frac{8}{3} - 3 = -\frac{1}{3}$$

Estimating the ATE

But wait! We said we were interested in the ATE...

So can we combine stratum-specific ATEs to get one overall ATE?

Yes!

Estimating the ATE

There are several ways to calculate an overall ATE in a conditionally randomized experiment that capture the idea of combining strata-specific ATEs:

- Weight strata-specific ATEs by size of strata:
    - Standardization
    - IP weighting
    - Interactive regression
- Weight strata-specific ATEs by size of strata and within-strata variance of treatment indicator:
    - Additive regression

Standardization

The formula for standardization is:

$$\widehat{ATE} = \sum_x [\bar{Y}|(T = 1, X = x) \cdot Pr(X = x)] - \sum_w [\bar{Y}|(T = 0, X = x) \cdot Pr(X = x)]$$

With 2 strata, this simplifies to:

$$\widehat{ATE} = \tau_1 \cdot P(X = 1) + \tau_0 \cdot P(X = 0)$$

For our example:

$$\widehat{ATE} = \frac{2}{3} \cdot \frac{3}{5} + \frac{-1}{3} \cdot \frac{2}{5} = \frac{4}{15}$$

Standardization

Note that the answer given by standardization is identical to that provided if we just take the difference-in-means in *imputed* potential outcomes:

i	Name	X	$Y(1)$	$Y(0)$	T	Y^{obs}
1	Benjamin Franklin	1	3	$\frac{8}{3}$	1	3
2	Thomas Jefferson	1	$\frac{10}{3}$	3	0	3
3	George Washington	1	5	$\frac{8}{3}$	1	5
4	Alexander Hamilton	1	$\frac{10}{3}$	2	0	2
5	James Madison	1	2	$\frac{8}{3}$	1	2
6	John Jay	1	$\frac{10}{3}$	3	0	3
7	Roger Sherman	0	3	3	1	3
8	Charles Pinckney	0	4	3	1	4
9	John Hancock	0	$\frac{8}{3}$	3	0	3
10	Robert Morris	0	1	3	1	1
			3.67	2.80		

Standardization and Interactive Regression

It turns out that an interactive regression model will give you the same results as standardization:

$$E[Y|T = t, X = x] = \beta_0 + \beta_1 \cdot t + \beta_2 \cdot x + \beta_3 \cdot t \cdot x$$

We can interpret $\beta_1 + \beta_3 \cdot x$ as a conditional difference-in-means between treated and control groups:

$$E[Y|T = 1, X = x] - E[Y|T = 0, X = x]$$

Note that for $X = x$, we get out stratum-specific ATEs:

- $X = 0 \Rightarrow \tau_0 = \beta_1$
- $X = 1 \Rightarrow \tau_1 = \beta_1 + \beta_3$

Standardization and Interactive Regression

To get the overall ATE, we can use the law of total expectation to obtain that:

$$E[Y_i(1) - Y_i(0)] = \beta_1 + \beta_3 \cdot E_X[X]$$

Therefore:

$$\widehat{ATE} = \hat{\beta}_1 + \hat{\beta}_3 \cdot \bar{X}$$

In practice:

- Recenter X by subtracting off $\bar{X}$ (de-meaning)
- Re-run interactive regression above using demeaned X
- $\widehat{ATE}$ will now be $\hat{\beta}_1$

Conclusion

ANY QUESTIONS?