Empirical Strategies

October 30, 2012

Empirical Labor Economics

- Labor economists are typically interested in questions such as the *impact* of
  - welfare programs on labor supply
  - minimum wages on employment
  - immigration on native wages
  - education on lifetime earnings
  - training program on the re-employment probability of an unemployed worker
- These are causal question, i.e. they concern *differences between counterfactual states of the world*

Empirical Labor Economics

- Labor economists are typically interested in questions such as the *impact* of
  - welfare programs on labor supply
  - minimum wages on employment
  - immigration on native wages
  - education on lifetime earnings
  - training program on the re-employment probability of an unemployed worker
- These are causal question, i.e. they concern *differences between counterfactual states of the world*
- For example: How does a MSc degree affect *your* lifetime income (in comparison to taking a vocational degree)?
  - difference between your true lifetime income and what you would have earned, had you taken a vocational degree

Potential Outcomes

A binary treatment

$$\text{potential outcome} = \begin{cases} y_{1i} & \text{if } D_i = 1 \\ y_{0i} & \text{if } D_i = 0 \end{cases}$$

y_{1i} =outcome for individual i if she receives the "treatment" D_i (e.g. a MSc degree);
y_{0i} =outcome for the *same* individual if she does not receive the treatment

Potential Outcomes

A binary treatment

$$\text{potential outcome} = \begin{cases} y_{1i} & \text{if } D_i = 1 \\ y_{0i} & \text{if } D_i = 0 \end{cases}$$

y_{1i} =outcome for individual i if she receives the "treatment" D_i (e.g. a MSc degree);
y_{0i} =outcome for the *same* individual if she does not receive the treatment

The causal effect (for indvidual i)

$$[y_{1i} - y_{0i}]$$

the difference in the potential outcomes with and without the treatment for individual i

Potential Outcomes

A binary treatment

$$\text{potential outcome} = \begin{cases} y_{1i} & \text{if } D_i = 1 \\ y_{0i} & \text{if } D_i = 0 \end{cases}$$

y_{1i} =outcome for individual i if she receives the "treatment" D_i (e.g. a MSc degree);
y_{0i} =outcome for the *same* individual if she does not receive the treatment

The causal effect (for indvidual i)

$$[y_{1i} - y_{0i}]$$

the difference in the potential outcomes with and without the treatment for individual i

But we only observe

$$y_i = y_{oi} + [y_{1i} - y_{0i}]\, D_i$$

Treatment Effects

ATE and ATT

$$\begin{aligned} \text{Average treatment effect (ATE)} &= \mathbb{E}\left[y_{i1} - y_{0i}\right] \\ \text{ATE for the treated (ATT)} &= \mathbb{E}\left[y_{i1} - y_{0i} | D_i = 1\right] \end{aligned}$$

$\mathbb{E}\left[a|b\right]$ =expectation of a conditional on b

y_{1i} =outcome for individual i if she receives the treatmentm (e.g. MSc degree)
y_{0i} =outcome for the *same* individual if she does not receive the treatment

Treatment Effects

ATE and ATT

$$\begin{aligned} \text{Average treatment effect (ATE)} &= \mathbb{E}\left[y_{i1} - y_{0i}\right] \\ \text{ATE for the treated (ATT)} &= \mathbb{E}\left[y_{i1} - y_{0i} | D_i = 1\right] \end{aligned}$$

$\mathbb{E}[a|b]$ =expectation of a conditional on b

y_{1i} =outcome for individual i if she receives the treatmentm (e.g. MSc degree)
y_{0i} =outcome for the *same* individual if she does not receive the treatment

- Why ATE *and* ATT?
    - the treatment group may be a non-random sample of the population, e.g. if people choose education based on their individual returns, ATT $>$ ATE

Treatment Effects

ATE and ATT

$$\begin{aligned} \text{Average treatment effect (ATE)} &= \mathbb{E}\left[y_{i1} - y_{0i}\right] \\ \text{ATE for the treated (ATT)} &= \mathbb{E}\left[y_{i1} - y_{0i} | D_i = 1\right] \end{aligned}$$

$\mathbb{E}[a|b]$ =expectation of a conditional on b

y_{1i} =outcome for individual i if she receives the treatmentm (e.g. MSc degree)
y_{0i} =outcome for the *same* individual if she does not receive the treatment

- Why ATE *and* ATT?
  - the treatment group may be a non-random sample of the population, e.g. if people choose education based on their individual returns, ATT > ATE
- **Internal and external validity**
  - Do we identify the true effect for some population (internal)?
  - Can we extrapolate to other populations (external)?

Selection Bias

What we see in the data (and what we don't)

$$\underbrace{\mathbb{E}\left[y_i|D_i = 1\right] - \mathbb{E}\left[y_i|D = 0\right]}_{\text{Observed difference}} =$$

Selection Bias

What we see in the data (and what we don't)

$$\underbrace{\mathbb{E}\left[y_i|D_i=1\right]-\mathbb{E}\left[y_i|D=0\right]}_{\text{Observed difference}} =$$

-
-

Selection Bias

What we see in the data (and what we don't)

$$\underbrace{\mathbb{E}\left[y_i|D_i=1\right]-\mathbb{E}\left[y_i|D=0\right]}_{\text{Observed difference}} = \underbrace{\mathbb{E}\left[y_{1i}|D_i=1\right]}_{} \qquad \underbrace{-\,\mathbb{E}\left[y_{0i}|D_i=0\right]}_{}$$

$\mathbb{E}\left[y_{1i}|D_i=1\right]$: exp. outcome with the treatment for those who got it (**observed**)

Note that $\mathbb{E}\left[y_i|D_i=1\right]=\mathbb{E}\left[y_{1i}|D_i=1\right]$ and

Selection Bias

What we see in the data (and what we don't)

$$\underbrace{\mathbb{E}\left[y_i|D_i = 1\right] - \mathbb{E}\left[y_i|D = 0\right]}_{\text{Observed difference}} = \underbrace{\mathbb{E}\left[y_{1i}|D_i = 1\right] \qquad\qquad}_{} \underbrace{\qquad\qquad - \mathbb{E}\left[y_{0i}|D_i = 0\right]}_{}$$

$\mathbb{E}\left[y_{1i}|D_i = 1\right]$: exp. outcome with the treatment for those who got it (**observed**)

Note that $\mathbb{E}\left[y_i|D_i = 1\right] = \mathbb{E}\left[y_{1i}|D_i = 1\right]$ and

-
-

Selection Bias

What we see in the data (and what we don't)

$$\underbrace{\mathbb{E}\left[y_i|D_i=1\right]-\mathbb{E}\left[y_i|D=0\right]}_{\text{Observed difference}} = \underbrace{\mathbb{E}\left[y_{1i}|D_i=1\right]-\mathbb{E}\left[y_{0i}|D_i=1\right]}_{\text{ATT}}$$
$$+\underbrace{\mathbb{E}\left[y_{0i}|D_i=1\right]-\mathbb{E}\left[y_{0i}|D_i=0\right]}_{\text{Selection bias}}$$

$\mathbb{E}\left[y_{1i}|D_i=1\right]$: exp. outcome with the treatment for those who got it (**observed**)
$\mathbb{E}\left[y_{0i}|D_i=1\right]$: ... without the treatment for those who got it (**unobserved**)

Note that $\mathbb{E}\left[y_i|D_i=1\right]=\mathbb{E}\left[y_{1i}|D_i=1\right]$ and
$\mathbb{E}\left[y_{1i}|D_i=1\right]-\mathbb{E}\left[y_{0i}|D_i=1\right]=\mathbb{E}\left[y_{i1}-y_{0i}|D_i=1\right]$

Selection Bias

What we see in the data (and what we don't)

$$\underbrace{\mathbb{E}\left[y_i|D_i=1\right]-\mathbb{E}\left[y_i|D=0\right]}_{\text{Observed difference}} = \underbrace{\mathbb{E}\left[y_{1i}|D_i=1\right]-\mathbb{E}\left[y_{0i}|D_i=1\right]}_{\text{ATT}}$$
$$+\underbrace{\mathbb{E}\left[y_{0i}|D_i=1\right]-\mathbb{E}\left[y_{0i}|D_i=0\right]}_{\text{Selection bias}}$$

$\mathbb{E}\left[y_{1i}|D_i=1\right]$: exp. outcome with the treatment for those who got it (**observed**)
$\mathbb{E}\left[y_{0i}|D_i=1\right]$: ... without the treatment for those who got it (**unobserved**)

Note that $\mathbb{E}\left[y_i|D_i=1\right]=\mathbb{E}\left[y_{1i}|D_i=1\right]$ and
$\mathbb{E}\left[y_{1i}|D_i=1\right]-\mathbb{E}\left[y_{0i}|D_i=1\right]=\mathbb{E}\left[y_{i1}-y_{0i}|D_i=1\right]$

- Selection bias: [*what would have happened* to the treatment group without the treatment] − [what happened to the control group]
- e.g. selection into education by "ability"

Empirical Strategies

Appoaches for solving the selection problem

1. Experiments
2. Statistical control
3. Differences-in-differences
4. Instrumental variables
5. Regression discontinuity designs

All attempt to *construct a plausible counterfactual*

- a control group that tells us of what would have happened to the treatment group in the absence of the treatment

The treatment effects is **identified** if such control group exists

Point Estimates and Standard Errors

- Most of the time, we are interested in two numbers
  - Point estimate
  - Standard error

Point Estimates and Standard Errors

- Most of the time, we are interested in two numbers
  - Point estimate
  - Standard error
- An estimate is **statistically significant** when it is at least two standard errors away from zero

Point Estimates and Standard Errors

- Most of the time, we are interested in two numbers
  - Point estimate
  - Standard error
- An estimate is **statistically significant** when it is at least two standard errors away from zero
- An estimate may be statistically insignificant because
  - precisely estimated point estimate is close to zero
  - standard errors are large (i.e. sample size is too small)
    $\rightarrow$ this does <u>not</u> prove that the true effect is zero/small.

Point Estimates and Standard Errors

- Most of the time, we are interested in two numbers
  - Point estimate
  - Standard error
- An estimate is **statistically significant** when it is at least two standard errors away from zero
- An estimate may be statistically insignificant because
  - precisely estimated point estimate is close to zero
  - standard errors are large (i.e. sample size is too small)
    $\rightarrow$ this does <u>not</u> prove that the true effect is zero/small.
- **Beware of underpowered studies!** Read Gelman, Weakliem (2009): Of Beauty, Sex and Power, *American Scientist 97*, available at www.stat.columbia.edu/~gelman/research/published/power4r.pdf

Ordinary Least Squares (OLS)

Estimation equation

$$y_{it} = \alpha D_{it} + X_{it}\beta + \epsilon_{it}$$

y_{jt} = outcome, D_{it} = treatment, X_{it} = a vector of controls (including constant), ϵ_{it} = unobservable factors.

Ordinary Least Squares (OLS)

Estimation equation

$$y_{it} = \alpha D_{it} + X_{it}\beta + \epsilon_{it}$$

y_{jt} = outcome, D_{it} = treatment, X_{it} = a vector of controls (including constant), ϵ_{it} = unobservable factors.

Solved as

$$\arg\min_{\alpha,\beta} \sum_{i=1}^{N} [y_{it} - \alpha D_{it} - X_{it}\beta]^2$$

i.e. the parameter estimates minimize the sum of squared residuals. Hence it corresponds to **expectations**

Ordinary Least Squares (OLS)

Estimation equation

$$y_{it} = \alpha D_{it} + X_{it}\beta + \epsilon_{it}$$

y_{jt} = outcome, D_{it} = treatment, X_{it} = a vector of controls (including constant), ϵ_{it} = unobservable factors.

Solved as

$$\arg\min_{\alpha,\beta} \sum_{i=1}^{N} [y_{it} - \alpha D_{it} - X_{it}\beta]^2$$

i.e. the parameter estimates minimize the sum of squared residuals. Hence it corresponds to **expectations**

Conditional Expectation Function (CEF)

$$\mathbb{E}[y_i | X_i = x] = \int t f_y (t | X_i = x)\, dt$$

where $f_y(\mathbf{.}|X_i = x)$ is the conditional density for y. That is, CEF gives the expectation—the population average—for y given that the vector X_i equals x

CEF: an example

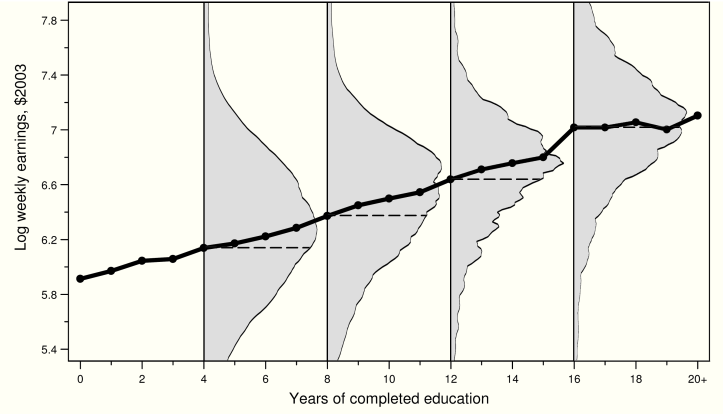

Figure 3.1.1: Raw data and the CEF of average log weekly wages given schooling. The sample includes white men aged 40-49 in the 1980 IPUMS 5 percent file.

Source: Angrist and Pischke (2009). The estimates correspond to OLS regression $\log(W_i) = \sum_{y=0}^{20} \eta_y d_{iy} + \epsilon_i$ where d_{iy} takes value one if person i has y years of education and zero otherwise.

OLS and Causality

- OLS approximates CEF (regardless of whether the true CEF is linear)
- Thus α gives us
  - (weighted) average difference in y
  - ... between the treated and untreated
  - ... who have identical observed characteristics

OLS and Causality

- OLS approximates CEF (regardless of whether the true CEF is linear)
- Thus α gives us
  - (weighted) average difference in y
  - ... between the treated and untreated
  - ... who have identical observed characteristics
- Causal interpretation of OLS estimates relies on the **Conditional Independence Assumption** (CIA)

$$y_0, y_1 \perp D | X$$

  i.e. potential outcomes are independent of the treatment status (conditional on X)

  - This implies that, on expectation, the treatment and control group do not differ in ϵ (cond. X)

OLS and Causality

- OLS approximates CEF (regardless of whether the true CEF is linear)
- Thus α gives us
  - (weighted) average difference in y
  - ... between the treated and untreated
  - ... who have identical observed characteristics
- Causal interpretation of OLS estimates relies on the **Conditional Independence Assumption** (CIA)

$$y_0, y_1 \perp D | X$$

  i.e. potential outcomes are independent of the treatment status (conditional on X)

  - This implies that, on expectation, the treatment and control group do not differ in ϵ (cond. X)
- Always ask: If we control for everything relevant, why is there variation in the treatment status?

OLS: Miscellaneous

- The intuition of binary treatment applies to continuous treatments
- OLS and various matching methods are closely related
  - all based on CIA
  - OLS does not work well when the distribution of X is very different among the treatment and control groups (see Imbens, Wooldridge (2009): "Recent Developments in the Econometrics of Program Evaluation", *JEL* 47(1) for details)
- Linear methods usually fine also for limited dependent variables (e.g. binary employment status)

OLS: Summary

- OLS is always informative
  - average differences between the treatment and control group after taking into account differences in observables
  - this is a causal effect if CIA holds
- However, CIA often violated
  - unobservable factors—motivation, social skills, ability and so forth—typically affect selection into the treatment and directly the outcomes
- We will next consider approaches that sometimes allow us to recover the causal effect even when in the selection-on-unobservables case

Differences-in-Differences (Dif-in-Dif or DD)

- Dif-in-dif is based on the assumption that **the counterfactual *trend* in the treatment and control groups are the same**
  - i.e. there may be permanent unobserved differences between the treated and untreated

Differences-in-Differences (Dif-in-Dif or DD)

- Dif-in-dif is based on the assumption that **the counterfactual *trend* in the treatment and control groups are the same**
  - i.e. there may be permanent unobserved differences between the treated and untreated
- The most basic dif-in-dif setup looks like this

	Pre-Period	Post-Period	Dif
Treatment Group	a	b	b-a
Control Group	c	d	d-c
		Dif-in-Dif	(b-a)-(d-c)

where a,b,c and d are means of the outcome

How does Dif-in-Dif work?

Suppose that the true data generating process is

$$y_{it} = \alpha D_{it} + \delta_t + \eta_i + \epsilon_{it}$$

y_{it} is the outcome, D_{it} treatment status (0/1), δ_t is time-variant unobserved factor that affects everyone, and η_i and ϵ_{it} are person specific time-invariant and time-variant unobserved factors. Note that including η_i implies that $\mathbb{E}[\epsilon_{it}] = 0$.

How does Dif-in-Dif work?

Suppose that the true data generating process is

$$y_{it} = \alpha D_{it} + \delta_t + \eta_i + \epsilon_{it}$$

y_{it} is the outcome, D_{it} treatment status (0/1), δ_t is time-variant unobserved factor that affects everyone, and η_i and ϵ_{it} are person specific time-invariant and time-variant unobserved factors. Note that including η_i implies that $\mathbb{E}[\epsilon_{it}] = 0$.

Dif-in-Dif

Period 0: no-one treated. Period 1: treatment group (T=1) treated

$$\begin{aligned}
\mathbb{E}\left[y_1^T\right] - \mathbb{E}\left[y_0^T\right] &= (\alpha + \delta_1 + \mathbb{E}[\eta_i | T = 1]) - (\delta_0 + \mathbb{E}[\eta_i | T = 1]) \\
&= \alpha + \delta_1 - \delta_0 \\
\mathbb{E}\left[y_1^C\right] - \mathbb{E}\left[y_0^C\right] &= (\delta_1 + \mathbb{E}[\eta_i | T = 0]) - (\delta_0 + \mathbb{E}[\eta_i | T = 0]) \\
&= \delta_1 - \delta_0
\end{aligned}$$

Even if $\mathbb{E}[\eta_i | T = 1] \neq \mathbb{E}[\eta_i | T = 0]$, we have $\mathbb{E}\left[\left(y_1^T - y_0^T\right) - \left(y_1^C - y_0^C\right)\right] = \alpha$

Dif-in-Dif Example

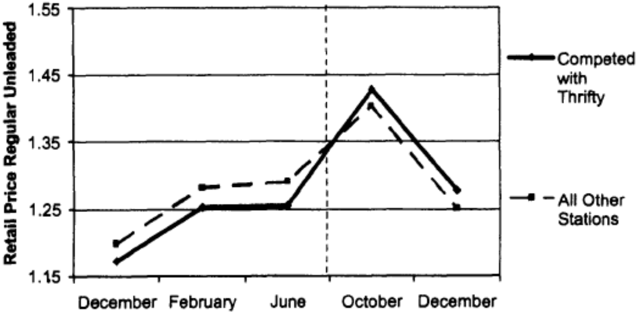

Note: Comparison of gas prices in stations close to and far away to an independent competitor that ceased to exist between June and October. Source: Hastings, 2004, "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California," *American Economic Review* (94):1 317-328.

Regression Dif-in-Dif

Dif-in-dif is typically implemented as

$$y_{it} = \alpha + \gamma_o D_i + \sum \gamma_{1t}\delta_t + \gamma_2 I_{it} + X_{it}\beta + \epsilon_{it}$$

where D_i is dummy for individual i belonging to the treament group (pre- *or* post), δ_t is dummy for time period t, I_{it} is a dummy for i belonging to the treament group *and* time period being post-treatment, X_{it} is a vector of control variables (inc. constant)

Regression Dif-in-Dif

Dif-in-dif is typically implemented as

$$y_{it} = \alpha + \gamma_o D_i + \sum \gamma_{1t}\delta_t + \gamma_2 I_{it} + X_{it}\beta + \epsilon_{it}$$

where D_i is dummy for individual i belonging to the treament group (pre- *or* post), δ_t is dummy for time period t, I_{it} is a dummy for i belonging to the treament group *and* time period being post-treatment, X_{it} is a vector of control variables (inc. constant)

Interpretation of the estimates

- γ_0: permanent difference between treatment/control groups
- γ_{1t}: time-effects (e.g. business cycle)
- γ_2: the treatment effect

Dif-in-Dif: Miscellaneous

- Extends to many time-periods, treatment/control groups
- Getting the standard errors right sometimes tricky
  - The issue is now well understood, but many of the early dif-in-dif papers report too small standard errors
- Group specific time trends can be added to strenghten the identification

Instrumental Variables (IV)

- The IV approach is based on **factors that affect the likelihood of receiving a treatment but not the outcomes** (except through the treatment)
  - the first part is know as "having a first-stage"
  - the second part is often refered to as **exclusion restriction** (the instrument does not appear in the main structural equation)

Instrumental Variables (IV)

- The IV approach is based on **factors that affect the likelihood of receiving a treatment but not the outcomes** (except through the treatment)
  - the first part is know as "having a first-stage"
  - the second part is often refered to as **exclusion restriction** (the instrument does not appear in the main structural equation)
- Formally, we need an instrument Z for which
  1. $\mathbb{E}\left(D=1|X,Z=z\right)\neq\mathbb{E}\left(D=1|X,Z=z'\right)$
  2. $\mathbb{E}\left[\epsilon|X,Z\right]=\mathbb{E}\left[\epsilon|X\right]$

Instrumental Variables (IV)

- The IV approach is based on **factors that affect the likelihood of receiving a treatment but not the outcomes** (except through the treatment)
  - the first part is know as "having a first-stage"
  - the second part is often refered to as **exclusion restriction** (the instrument does not appear in the main structural equation)
- Formally, we need an instrument Z for which
  1. $\mathbb{E}(D = 1|X, Z = z) \neq \mathbb{E}(D = 1|X, Z = z')$
  2. $\mathbb{E}[\epsilon|X, Z] = \mathbb{E}[\epsilon|X]$
- Finding instruments hard, but potential sources include
  - randomize treatment/outreach (here IV deals with imprefect compliance)
  - policy details, how things are implemented
  - "natural experiments"

How does IV work? (binary instrument)

Wald Estimator

Expected values of the outcome conditional on X and Z

$$\begin{aligned}\mathbb{E}\left[Y|X,Z=1\right] &= \alpha\mathbb{E}\left(D_{it}|X,Z=1\right)+X\beta+\mathbb{E}\left(\epsilon_{it}|X,Z=1\right)\\ \mathbb{E}\left[Y|X,Z=0\right] &= \alpha\mathbb{E}\left(D_{it}|X,Z=0\right)+X\beta+\mathbb{E}\left(\epsilon_{it}|X,Z=0\right)\end{aligned}$$

Substracting the second from the first yields

$\mathbb{E}\left[Y|X,Z=1\right]-\mathbb{E}\left[Y|X,Z=0\right]=\alpha\left[\mathbb{E}\left(D_{it}|X,Z=1\right)-\mathbb{E}\left(D_{it}|X,Z=0\right)\right]$ and rearranging

$$\alpha=\frac{\mathbb{E}\left[Y|X,Z=1\right]-\mathbb{E}\left[Y|X,Z=0\right]}{\mathbb{E}\left[D|X,Z=1\right]-\mathbb{E}\left[D_{it}|X,Z=0\right]}$$

- The numerator is the “reduced form” or “intention to treat”
- The denominator is the “first-stage”

IV interpretation

- With heterogeneous treatment effects, IV yieds a local average treament effect (**LATE**)
  - (weighted) average of the impact for '**compliers**', i.e. those who got the treatment because of the instrument and would not have gotten it without the instrument
  - we do not learn anything about 'never-takes' or 'always-takers'
  - e.g. Minneapolis domestic violence experiment see Angrist (2006, *Journal of Experimental Criminology*)
- External validity: How representative are the compliers?
- We also need the monotonicity assumption
  - the instrument inreases (or has zero effect) the likelihood of being treated for *everyone*
  - note that this is always implicit in 2SLS

IV implementation

Two-Stage Least-Squares

First-stage

$$D_{it} = \pi_0 + X_{it}\pi_1 + Z_{it}\pi_2 + \nu_{it}$$

Second-stage

$$y_{it} = \alpha\hat{D}_{it} + X_{it}\beta + \epsilon_{it}$$

where $\hat{D}_{it}$ is the predicted values from the first-stage.

- Most of the time 2SLS is fine
- When there are many 'weak' instruments, 2SLS is biased to the same direction as OLS. Then LIML works better (but at this point, you don't need to understand what that means)

Regression Discontinuity Designs (RDD)

- A RDD design is present when **selection into the treatment depends on whether an observed continuous variable is below/above a known threshold** and potential outcomes are continuous at the threshold
  - Shard RDD: everyone above the threshold are treated
  - Fuzzy RDD: a discontinuous increase in the likelihood of being treated at the threshold

Regression Discontinuity Designs (RDD)

- A RDD design is present when **selection into the treatment depends on whether an observed continuous variable is below/above a known threshold** and potential outcomes are continuous at the threshold
  - Shard RDD: everyone above the threshold are treated
  - Fuzzy RDD: a discontinuous increase in the likelihood of being treated at the threshold
- Typical examples of RDD design include
  - administrative rules (e.g. school admission)
  - elections, geographical borders

Regression Discontinuity Designs (RDD)

- A RDD design is present when **selection into the treatment depends on whether an observed continuous variable is below/above a known threshold** and potential outcomes are continuous at the threshold
  - Shard RDD: everyone above the threshold are treated
  - Fuzzy RDD: a discontinuous increase in the likelihood of being treated at the threshold
- Typical examples of RDD design include
  - administrative rules (e.g. school admission)
  - elections, geographical borders
- Interpretation
  - Sharp: impact for those at the threshold
  - Fuzzy: LATE for those at the threshold

Sharp RDD Illustration

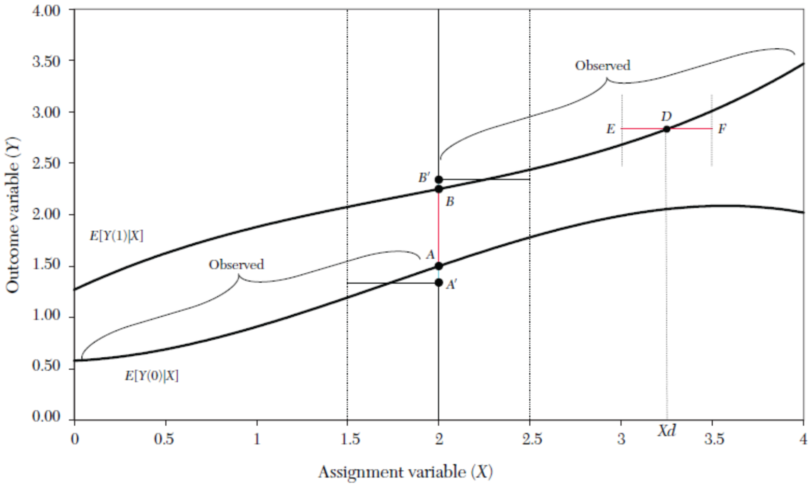

RD estimate $(B - A)$ is the impact of the treatment **for those with** $X = 2$.
Lee, Lemieux (2010): Regresion Discontinuity Designs in Economics. *Journal of Economic Literature*

Fuzzy RD is (almost) IV

Fuzzy RD

$$\alpha = \frac{y^+ - y^-}{p^+ - p^-}$$

where the numerator corresponds to $(B - A)$ in the previous figure and the denominator to similar figure when looking at the fraction of treated. Formally, $y^+ = \lim_{z \downarrow z_0} \mathbb{E}[y|z_i = z]$, $y^- = \lim_{z \uparrow z_0} \mathbb{E}[y|z_i = z]$, $p^+ = \lim_{z \downarrow z_0} \mathbb{E}[D|z_i = z]$, $p^- = \lim_{z \uparrow z_0} \mathbb{E}[D|z_i = z]$, i.e. expectations at the thresholds when approaching it from below/above.

Fuzzy RD is (almost) IV

Fuzzy RD

$$\alpha = \frac{y^+ - y^-}{p^+ - p^-}$$

where the numerator corresponds to $(B - A)$ in the previous figure and the denominator to similar figure when looking at the fraction of treated. Formally, $y^+ = \lim_{z \downarrow z_0} \mathbb{E}[y|z_i = z]$, $y^- = \lim_{z \uparrow z_0} \mathbb{E}[y|z_i = z]$, $p^+ = \lim_{z \downarrow z_0} \mathbb{E}[D|z_i = z]$, $p^- = \lim_{z \uparrow z_0} \mathbb{E}[D|z_i = z]$, i.e. expectations at the thresholds when approaching it from below/above.

- The logic is similar to IV
- The interpretation is similar to LATE, but now we identify the effect for 'compliers at the threshold'
- Many ways to estimate RDD (2SLS the simplest option)

Example: First-Stage

FIGURE 1.—FRACTION ENROLLED AT THE FLAGSHIP STATE UNIVERSITY

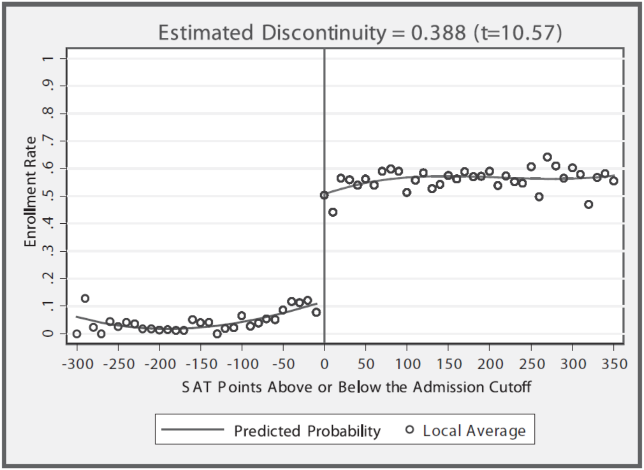

Hoekstra (2009): The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach, *Review of Economics and Statistics* 91(4): 717–724

Example: Second-Stage

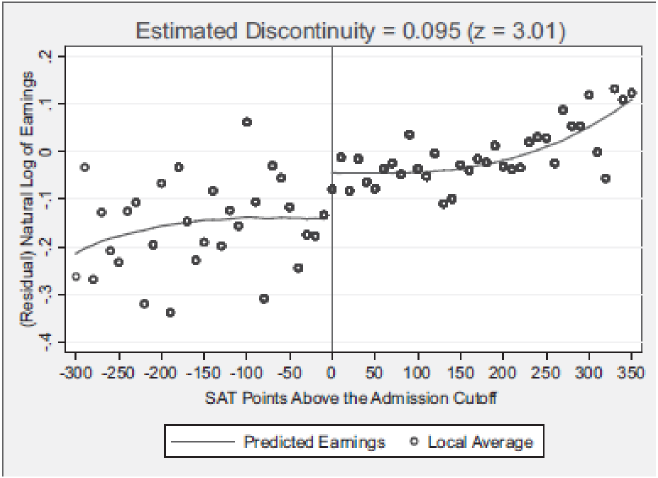

Example: Second-Stage

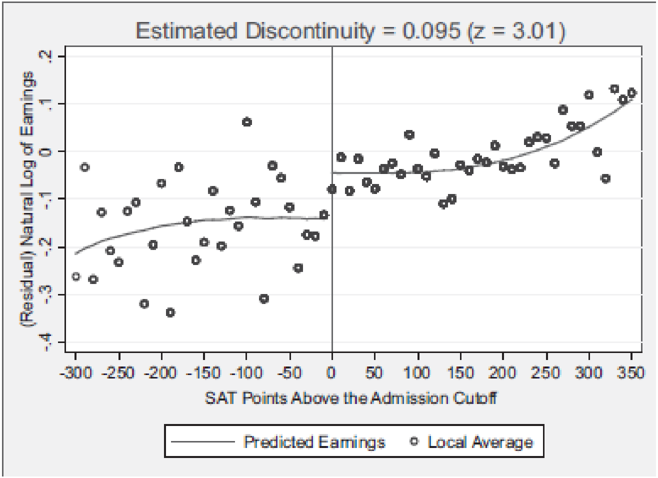

Earnings jump roughly 9.5% at the threshold; enrolment jumps 38.8 percentage points $\rightarrow$ RD estimate $0.095/0.388 = 0.245$

Empirical Strategies: Summary

- All approaches aim at constructing a valid control group
  - it is vital to understand who are we comparing to whom
- There are no 'magic bullets'
  - identification always case specific
  - finding/creating plausible research designs is hard $\rightarrow$ we often have to look at strange cases to gain internal validity (even though we would prefer representative populations)
- Key identifying assumptions
  - OLS and matching: CIA
  - Dif-in-dif: common counterfactual trends
  - IV: exclusion restriction
  - RDD: no jumps in potential outcomes at the threshold