Regression Discontinuity Design

Regression Discontinuity Design (RDD)

- RDD is a fairly old idea (Thistlethwaite and Campbell, 1960) but this design experienced a renaissance in recent years.
- Assignment to treatment and control is not random, but we know the assignment rule influencing how people are assigned or selected in to treatment
- Widely applicable in a rule based world (administrative programs, elections, etc.)
- High internal validity: In their validation study Cook and Wong (2008) identify RDD as one of the few observational study designs that can accurately reproduce experimental benchmarks

Outline

Outline

Sharp Regression Discontinuity Design

- Imagine a binary treatment D that is completely determined by the value of a predictor X_i being on either side of a fixed cutoff point c:

$$D_i = 1\{X_i > c\} \text{ so } D_i = \begin{cases} D_i = 1 & \text{if } X_i > c \\ D_i = 0 & \text{if } X_i < c \end{cases}$$

- X, called the forcing variable, may be correlated with the outcomes Y so comparing treated and untreated units does not provide causal estimates

- Design arises often from administrative decisions, where the allocation of units to a program is partly limited for reasons of resource constraints, and sharp rules rather than discretion by administrators is used for allocation.

Hypothetical Illustration Example: Sharp RDD

- Thistlethwaite and Campbell (1960) study the effects of college scholarships on later students' achievements
- Scholarships are given on the basis of whether or not a student's test score exceeds some threshold c
  - Treatment D is scholarship
  - Forcing variable X is SAT score with cutoff c
  - Outcome Y is subsequent earnings
  - Y_0 denotes potential earnings without the scholarship
  - Y_1 denotes potential earnings with the scholarship
- Y_1 and Y_0 are correlated with X: on average, students with higher SAT scores obtain higher earnings

Sharp RDD: Graphical Illustration

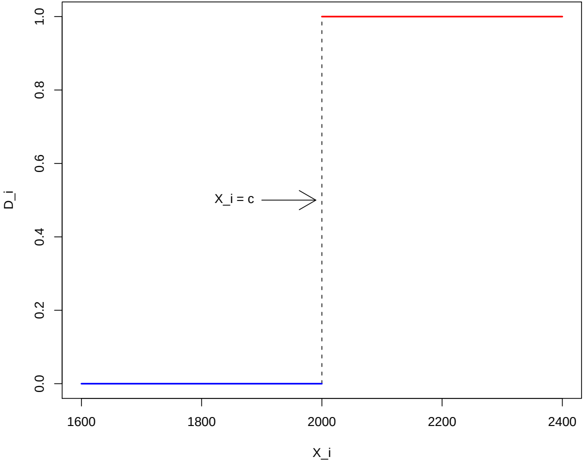

Sharp RDD: Graphical Illustration

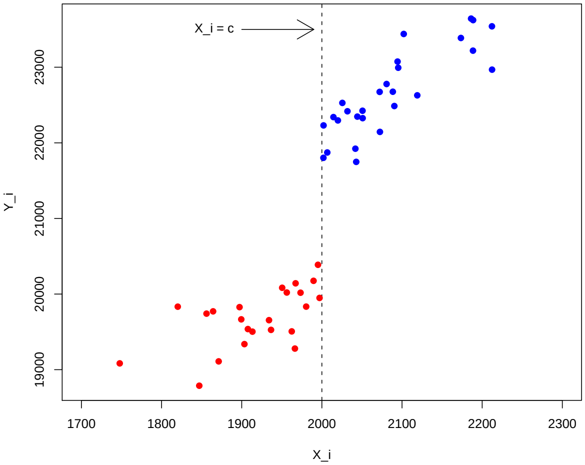

Sharp RDD: Graphical Illustration

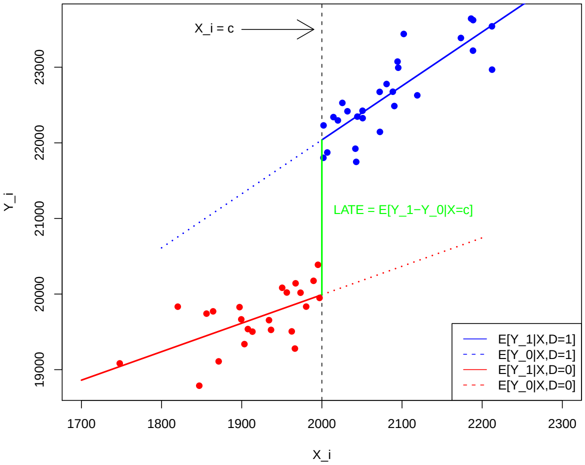

Sharp RDD: Identification

Identification Assumption

1. $Y_1, Y_0 \perp\!\!\!\perp D|X$ *(trivially met)*
2. $0 < P(D = 1|X = x) < 1$ *(always violated in Sharp RDD)*
3. $E[Y_1|X, D]$ *and* $E[Y_0|X, D]$ *are continuous in* X *around the threshold* $X = c$ *(to compensate for failure of common support)*

Identification Result

The treatment effect is identified at the threshold as:

$$\begin{aligned} \alpha_{SRDD} &= E[Y_1 - Y_0|X = c] \\ &= E[Y_1|X = c] - E[Y_0|X = c] \\ &= \lim_{x \downarrow c} E[Y_1|X = c] - \lim_{x \uparrow c} E[Y_0|X = c] \end{aligned}$$

Without further assumptions α_{SRDD} *is only valid at the threshold.*

Outline

Duverger's Law (Eggers, 2010)

- Duverger (1972): "a majority vote on one ballot is conducive to a two-party system; proportional representation is conducive to a multiparty system"
- Therefore, we expect the number of parties to increase when going from a majority to a proportional electoral system
- In French municipalities, the electoral rule used to elect the municipal council depends on the city's population:
  - cities with fewer than 3,500 people elect their councils by a form of plurality rule
  - cities with a population of 3,500 or more use a form of PR rule

Sharp RDD: Duverger's Law

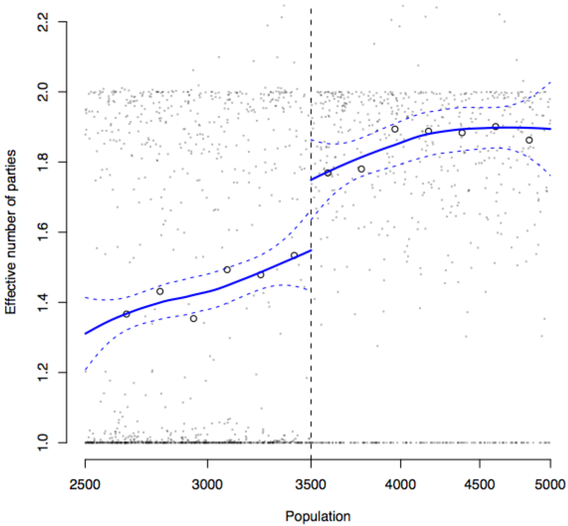

Party Incumbency Advantage (Lee, 2006)

- What is the effect of incumbency status on vote shares?
- Let i indicate congressional districts, j indicate parties, and t indicate elections, d indicate incumbency status
- V_{ditj} is the vote share of j in i at t as incumbent $d = 1$ or non-incumbent $d = 0$
- Party Incumbency Effect: $V_{1itj} - V_{0itj}$
- Margin of Victory for party j: $Z_{itj} = V_{itj} - V_{itk}$ where k indicates the strongest opposition party.
- Party Incumbency status is then assigned as:

$$D_{ij,t+1} = 1\{Z_{itj} > 0\} \text{ so } D_i = \begin{cases} D_{ij,t+1} = 1 & \text{if } Z_{itj} > 0 \\ D_{ij,t+1} = 0 & \text{if } Z_{itj} < 0 \end{cases}$$

- With only two parties we can also use $Z = V - c$ with $c = .5$

Sharp RDD: Incumbency Advantage

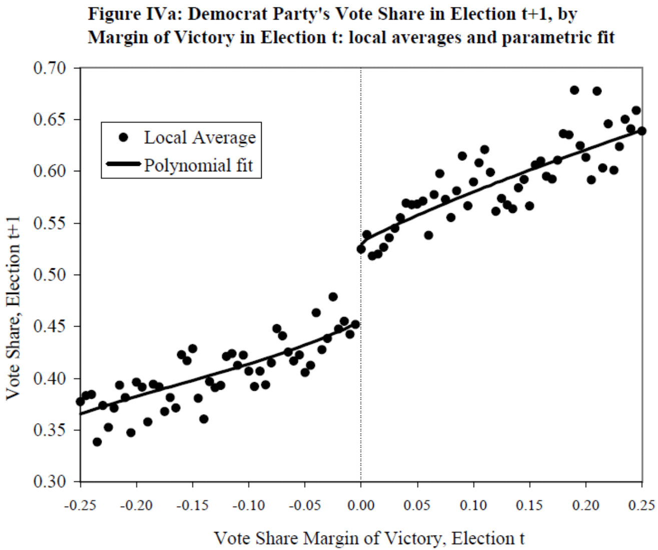

Other Recent Examples

- Effect of class size on student achievement (class size is determined by a cutoff in class size)
- Effect of access to credit on development outcomes (loan offer is determined by credit score threshold)
- Effect of party democratic versus republican mayor
- Effect of wages increase for mayors on policy performance (wage jumps at population cutoffs)
- Effect of an additional night in the hospital, a newborn delivered at 12:05 a.m. will have an extra night of reimbursable care
- Effect of school district boundaries on home values
- Effect of colonial borders on development outcomes

Outline

Estimate $\alpha_{SRDD} = E[Y_1|X = c] - E[Y_0|X = c]$

1. Trim the sample to a reasonable window around the threshold c (discontinuity sample):
   - $c - h \leq X_i \leq c + h$, were h is some positive value that determines the size of the window
   - h may be determined by cross-validation
2. Recode running variable to deviations from threshold: $\tilde{X} = X - c$.
   - $\tilde{X} = 0$ if $X = c$
   - $\tilde{X} > 0$ if $X > c$ and thus $D = 1$
   - $\tilde{X} < 0$ if $X < c$ and thus $D = 0$
3. Decide on a model for $E[Y|X]$
   - linear, same slope for $E[Y_0|X]$ and $E[Y_1|X]$
   - linear, different slopes
   - non-linear
   - always start with visual inspection (scatter plot with kernel/lowess) to check which model is appropriate

SRRD Estimation: Linear with Same Slope

- $E[Y_0|X]$ is linear and treatment effect, α, does not depend on X:

$$E[Y_0|X] = \mu + \beta X\,, \qquad E[Y_1 - Y_0|X] = \alpha$$

- Therefore $E[Y_1|X] = \alpha + E[Y_0|X] = \alpha + \mu + \beta X$
- Since D is determined given X, we have that:

$$\begin{aligned} E[Y|X,D] &= D \cdot E[Y_1|X] + (1-D) \cdot E[Y_0|X] \\ &= \mu + \alpha D + \beta X \\ &= (\mu - \beta c) + \alpha D + \beta(X - c) \\ &= \gamma + \alpha D + \beta \tilde{X} \end{aligned}$$

- So we just run a regression of Y on D and $\tilde{X}$

SRDD: Linear with Same Slope

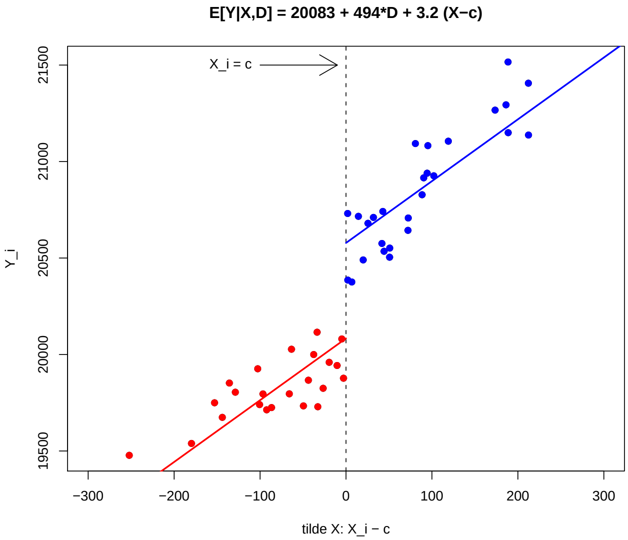

Sharp RDD Estimation: Differential Slopes

- $E[Y_0|X]$ and $E[Y_1|X]$ are distinct linear functions of X, so the average effect of the treatment $E[Y_1 - Y_0|X]$ varies with X:

$$E[Y_0|X] = \mu_0 + \beta_0 X, \qquad E[Y_1|X] = \mu_1 + \beta_1 X$$

- So $\alpha(X) = E[Y_1 - Y_0|X] = (\mu_1 - \mu_0) + (\beta_1 - \beta_0)X$ we have

$$\begin{aligned} E[Y|X, D] =& D \cdot E[Y_1|X] + (1 - D) \cdot E[Y_0|X] \\ =& \mu_1 D + \beta_1 (X \cdot D) + \mu_0 (1 - D) + \beta_0 (X \cdot (1 - D)) \\ =& \gamma + \beta_0 (X - c) + \alpha D + \beta_1 ((X - c) \cdot D) \end{aligned}$$

- Regress Y on $(X - c)$, D and the interaction $(X - c) \cdot D$, the coefficient of D reflects the average effect of the treatment at $X = c$

SRDD: Linear with Differential Slope

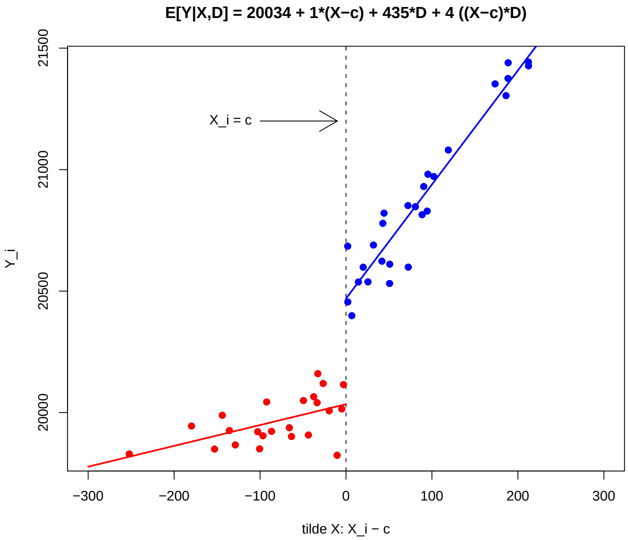

Sharp RDD Estimation: Non-Linear Case

- $E[Y_0|X]$ and $E[Y_1|X]$ are distinct non-linear functions of X and the average effect of the treatment $E[Y_1 - Y_0|X]$ varies with X
- Include quadratic and cubic terms in $(X - c)$ and their interactions with D in the equation
- The specification with quadratic terms is

$$\begin{aligned} E[Y|X, D] =& \gamma_0 + \gamma_1(X - c) + \gamma_2(X - c)^2 \\ &+ \alpha_0 D + \alpha_1((X - c) \cdot D) + \alpha_2((X - c)^2 \cdot D) \end{aligned}$$

The specification with cubic terms is

$$\begin{aligned} E[Y|X, D] =& \gamma_0 + \gamma_1(X - c) + \gamma_2(X - c)^2 + \gamma_3(X - c)^3 + \alpha_0 D \\ &+ \alpha_1((X - c) \cdot D) + \alpha_2((X - c)^2 \cdot D) + \alpha_3((X - c)^3 \cdot D) \end{aligned}$$

- In both cases $\alpha_0 = E[Y_1 - Y_0|X = c]$

SRDD: Non-Linear Case

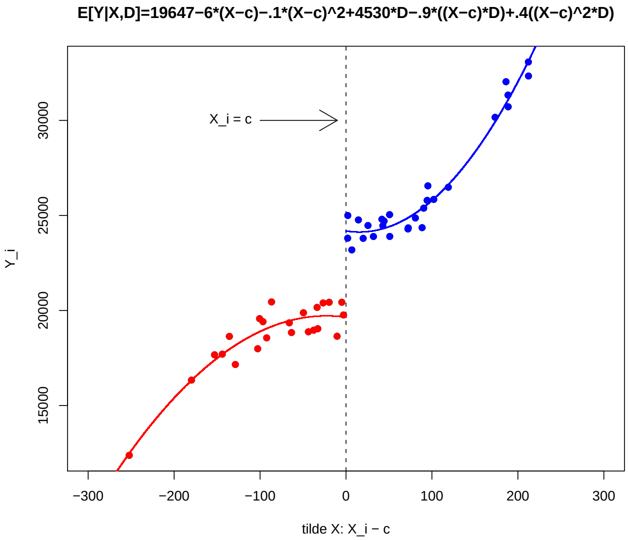

Outline

Sharp RDD: Falsification Checks

1. Sensitivity: Are results sensitive to alternative specifications?
2. Balance Checks: Do covariates Z jump at the threshold?
3. Check if jumps occur at placebo thresholds $c^{\star}$?
4. Sorting: Do units sort around the threshold?

Sharp RDD: Falsification Checks

1. Sensitivity: Are results sensitive to alternative specifications?
2. Balance Checks: Do covariates Z jump at the threshold?
3. Check if jumps occur at placebo thresholds $c^{\star}$?
4. Sorting: Do units sort around the threshold?

Sharp RDD: Sensitivity to Specification

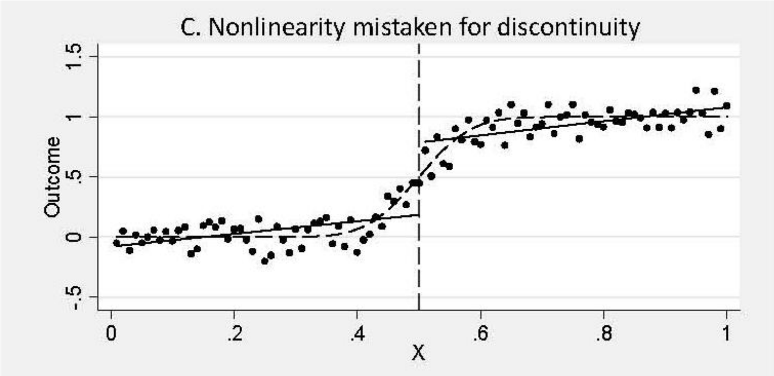

- $Y = f(X) + \alpha D + \varepsilon$: A miss-specified control function $f(X)$ can lead to a spurious jump: Take care not to confuse a nonlinear relation with a discontinuity
- More flexibility (e.g. adding polynomials) reduces bias, decreases efficiency
- Check sensitivity to size of bandwidth (i.e. estimation window)

SRDD: Balance Checks Test

- Test for comparability of agents around the cut-off:
  - Visual tests: Plot $E[Z|X, D]$ and look for jumps, ideally the relation between covariates and treatment should be smooth around threshold
  - Run the RDD regression using Z as the outcome:
    $E[Z|X, D] = \beta_0 + \beta_1(X - c) + \alpha_z D + \beta_3((X - c) \cdot D)$
    ideally should yield $\alpha_z = 0$ if Z is balanced at the threshold.
- Finding a discontinuity in Z does not necessarily invalidate the RDD
  - Can incorporate Z as additional controls into our main RDD regression. Ideally, this should only impact statistical significance, not magnitude of treatment effect.
  - Alternatively, regress the outcome variable on a vector of controls and use the residuals in the RDD, instead of the outcome itself
- Balance checks address only observables, not unobservables

SRDD: Falsification Test

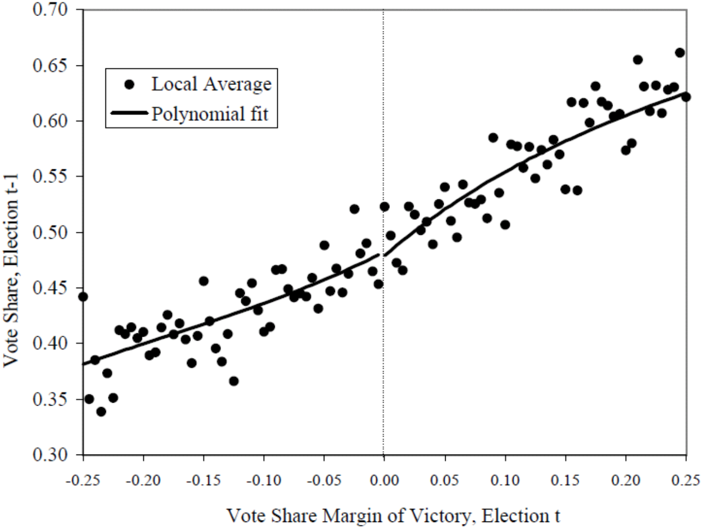

SRDD: Falsification Test

TABLE 6. Effect of Serving on Placebo Outcomes

	Conservative Party			Labour Party		
Placebo Outcome	Placebo Effect	95. UB	95 LB	Placebo Effect	95. UB	95 LB
Year of birth	2.79	8.10	−2.62	2.50	8.62	−3.77
Year of death	2.08	5.97	−1.89	2.23	6.23	−1.91
Age at death	0.12	−6.32	6.56	1.41	−5.78	8.60
Female	−0.01	0.14	−0.16	−0.03	0.06	−0.12
Teacher	−0.09	0.06	−0.23	−0.23	0.01	−0.47
Barrister	0.09	0.25	−0.09	−0.07	0.05	−0.18
Solicitor	−0.13	0.07	−0.33	0.03	0.15	−0.10
Doctor	−0.00	0.12	−0.13	0.03	0.14	−0.09
Civil servant	0.04	0.10	−0.02	−0.03	0.03	−0.10
Local politician	−0.01	0.23	−0.25	0.10	0.40	−0.21
Business	−0.05	0.21	−0.31	0.00	0.13	−0.13
White collar	−0.00	0.19	−0.19	−0.00	0.15	−0.16
Union official	0.00	NA	NA	−0.04	0.12	−0.20
Journalist	−0.08	0.07	−0.22	0.05	0.29	−0.20
Miner	0.00	NA	NA	−0.02	0.02	−0.07
Schooling: Eton	0.12	0.28	−0.04	−0.04	0.02	−0.11
Schooling: public	−0.22	0.07	−0.52	0.03	0.23	−0.17
Schooling: regular	−0.15	0.12	−0.42	−0.01	0.32	−0.35
Schooling: not reported	0.25	0.46	0.03	0.02	0.33	−0.30
University: Oxbridge	0.10	0.36	−0.17	−0.04	0.21	−0.30
University: degree	−0.02	0.25	−0.30	0.10	0.42	−0.23
University: not reported	−0.08	0.21	−0.37	−0.06	0.25	−0.37
Aristocrat	0.05	0.19	−0.09	0.06	0.17	−0.06
Previous races	0.22	0.59	−0.16	0.24	0.76	−0.29
Vote margin in previous race	−0.00	0.04	−0.05	−0.05	0.01	−0.11
Size of electorate	−622	−8056	6812	−545	−7488	6397
Turnout	−0.01	−0.04	0.03	0.02	−0.02	0.05

SRDD: Adding Covariates

TABLE 4. Regression Discontinuity Design Results: Effect of Serving in House of Commons on (Log) Wealth at Death

	Conservative Party		Labour Party	
Effect of serving	0.61	0.66	−0.20	−0.25
Standard error	(0.27)	(0.37)	(0.26)	(.26)
Covariates		x		x
Percent wealth increase	83	94	−18	−23
95% Lower bound	8	−7	−52	−65
95% Upper bound	212	306	31	71

Note: Effect estimates at the threshold of winning $\tau_{RDD} = E[Y(1) - Y(0) \mid Z = 0]$. Estimates without covariates from local polynomial regression fit to both sides of the threshold with bootstrapped standard errors. Estimates with covariates from local linear regression with rectangular kernel (equation 2); bandwidth is 15 percentage point of vote share margin with

SRDD: Placebo Threshold

- Test whether the treatment effect is zero when it should be
- Let $c^\star$ be a placebo thershold value. Run the regression of: $E[Y|X,D] = \beta_0 + \beta_1(X - c^\star) + \alpha D + \beta_3((X - c^\star) \cdot D)$ and check if α large and significant?
  - Usually we split the sample to the left and the right of the actual threshold c in order to avoid miss-specification by imposing a zero jump at c
- The existence of large placebo jumps does not invalidate the RDD, but does require an explanation
- Concern is that the relation is fundamentally discontinuous and jump at cut-off is contaminated by other factors.
- Maybe data exists in a period where there was no program

SRDD: Sorting Around the Threshold

- Can subjects behavior invalidate the local continuity assumption?
  - Can they exercise control over their values of the assignment variable?
  - Can administrators strategically choose what assignment variable to use or which cut-off point to pick?
  - Either can invalidate the comparability of subjects near the threshold because of sorting of agents around the cut-off, where those below may differ on average form those just above
  - Does not necessarily invalidate the design unless sorting is very widespread and very precise
- What else changes at c? Continuity violated in the presence of other programs that use a discontinuous assignment rule with the exact same assignment variable and cut-off

SRDD: Sorting Around the Threshold

Example: Beneficial job training program offered to agents with income $< c$. Concern, people will withhold labor to lower their income below the cut-off to gain access to the program.

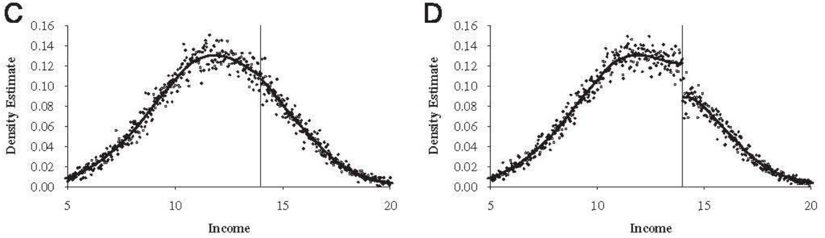

SRDD: Sorting Around the Threshold

- Test for discontinuity in density of forcing variable:
  - Visual Histogram Inspection:
    - Construct equal-sized non-overlapping bins of the forcing variable such that no bin includes points to both the left and right of the cut-off
    - For each bin, compute the number of observations and plot the bins to see if there is a discontinuity at the cut-off
  - Formal tests (e.g. McCrary, 2008)

Sorting Around the Threshold (Eggers, 2010)

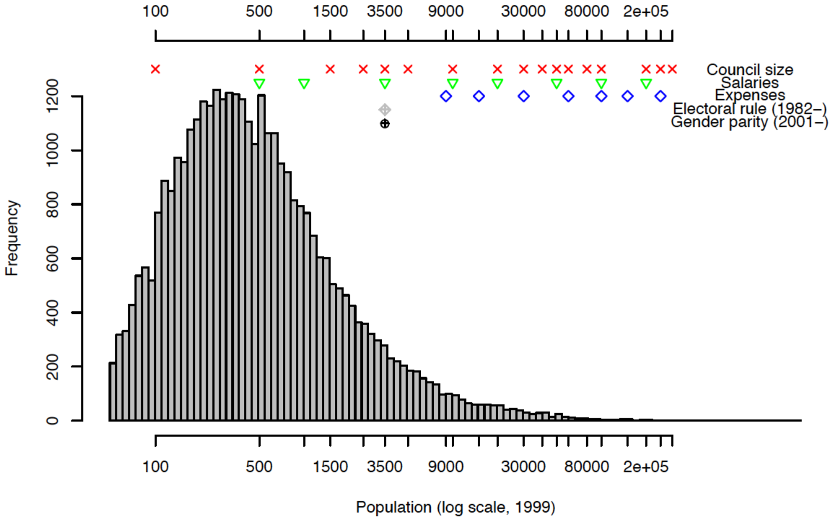

Outline

Outline

Fuzzy Regression Discontinuity Design

- Threshold may not perfectly determine treatment exposure, but it creates a discontinuity in the probability of treatment exposure
- Incentives to participate in a program may change discontinuously at a threshold, but the incentives are not powerful enough to move all units from nonparticipation to participation
- We can use such discontinuities to produce instrumental variable estimators of the effect of the treatment (close to the discontinuity)

Fuzzy Regression Discontinuity Design

- Probability of being offered a scholarship may jump at a certain SAT score threshold (when applicants are given "special consideration")
- We shouldn't compare recipients with non-recipients (even close to threshold) since they are likely differ along unobservables related to outcome (e.g., letters of rec)
- But for applicants with scores close to the threshold we can exploit the discontinuity as an instrument to estimate the LATE for the subgroup of applicants for whom scholarship receipt depends on the difference between their score and the threshold
    - A complier in the framework is a student who switches from non-recipient to recipient if her scores crosses the threshold

Fuzzy RDD: Discontinuity in $E[D|X]$

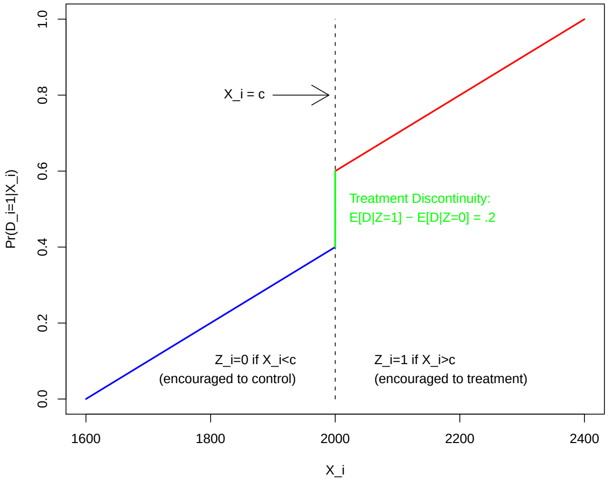

Fuzzy RDD: Discontinuity in $E[Y|X]$

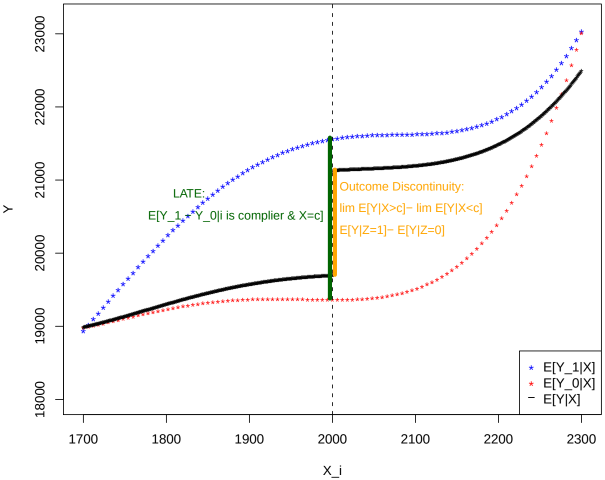

Fuzzy RDD: Identification

Identification Assumption

- *Binary instrument Z with $Z = 1\{X > c\}$*
- *Restrict sample to observations close to discontinuity where $E[Y|D, X]$ jumps so that $X \approx c$ and thus $E[X|Z = 1] - E[X|Z = 0] \approx 0$.*
- *Usual IV assumptions hold (ignorability, first stage, montonicity)*

Identification Result

$$
\begin{aligned}
\alpha_{FRDD} &= E[Y_1 - Y_0 | X = c \text{ and } \; i \; \text{ is a complier}] \\
&= \frac{\lim_{x \downarrow c} E[Y|X = c] - \lim_{x \uparrow c} E[Y|X = c]}{\lim_{x \downarrow c} E[D|X = c] - \lim_{x \uparrow c} E[D|X = c]} \\
&= \frac{\text{outcome discontinuity}}{\text{treatment discontinuity}} \\
&\approx \frac{E[Y|Z = 1] - E[Y|Z = 0]}{E[D|Z = 1] - E[D|Z = 0]}
\end{aligned}
$$

Fuzzy RDD: Identification

- Suppose $E[Y_0|X]$ is linear and treatment effect is constant:

$$E[Y_0|X] = \mu + \beta X,\ E[Y_1 - Y_0|X] = \alpha,\ E[Y_1|X] = \alpha + \mu + \beta X.$$

- Suppose also that $E[D|X]$ has a discontinuity at c. For those who are close to c, the average outcomes are:

$$\begin{aligned} E[Y|Z=0] &= \mu + \alpha E[D|Z=0] + \beta E[X|Z=0] \\ E[Y|Z=1] &= \mu + \alpha E[D|Z=1] + \beta E[X|Z=1] \end{aligned}$$

- For those with $X \approx c$ $E[X|Z=1] - E[X|Z=0] \approx 0$. However, $E[D|Z=1] - E[D|Z=0] \neq 0$ because of discontinuity in the assignment probability

Outline

Fuzzy RDD: Estimation

- Cut the sample to a small window above and below the threshold (discontinuity sample)
- Code instrument $Z = 1\{X > c\}$
- Fit 2SLS: $Y = \beta_0 + \beta_1(X - c) + \beta_2(Z \cdot (X - c)) + \alpha D$ where D is instrumented with Z
- Specification can be more flexible by adding polynomials
- Using a larger window we may also fit 2SLS: $Y = \beta_0 + \beta_1(X - c) + \alpha D + \beta_2(D * (X - c))$ where D and $D * (X - c)$ are instrumented with Z and $Z \cdot (X - c)$
- Also helpful to separately plot (and estimate) the outcome discontinuity and treatment discontinuity

Outline

Early Release Program (HDC)

- Prison system in many countries is faced with overcrowding and high recidivism rates after release.
- Early discharge of prisoners on electronic monitoring or tag has become a popular policy
- Difficult to estimate impact of early release program on future criminal behaviour: best behaved inmates are usually the ones to be released early.
- Marie (2008) considers Home Detention Curfew (HDC) scheme in England and Wales:
  - Fuzzy RDD: Only offenders sentenced to more than three months (88 days) in prison are eligible for HDC, but obviously, not all those with longer sentences are offered HDC

Table 1: Descriptive Statistics for Prisoners Released by Length of Sentence and HDC and Non HDC Discharges

Panel A - Released Before 3 Months:			
Discharge Type	**Non HDC**	**HDC**	**Total**
Percentage Female	12.2	-	12.2
Mean Age	29.5	-	29.5
Percentage Incarcerated for Violence	17.6	-	17.6
Mean Number Previous Offences	8.8	-	8.8
Recidivism within 12 Months	52.4	-	52.4
Sample Size	42,987	0	42,987
Panel B - Released Between 3 and 6 Months:			
Discharge Type	**Non HDC**	**HDC**	**Total**
Percentage Female	8.8	8.8	8.8
Mean Age at Release	27.6	30.8	28.4
Percentage Incarcerated for Violence	20.3	18.3	19.8
Mean Number Previous Offences	10	6.5	9.1
Recidivism within 12 Months	60	30.2	52.6
Sample Size	52,091	17,222	69,313

Table 2: Descriptive Statistics for Prisoners Released by Length of Sentence and HDC and Non HDC Discharges and +/-7 Days Around Discontinuity Threshold

Panel A - Released +/- 7 Days of 3 Months (88 Days) Cut-off:			
Discharge Type	**Non HDC**	**HDC**	**Total**
Percentage Female	10.5	9.7	10.3
Mean Age at Release	28.9	30.7	29.3
Percentage Incarcerated for Violence	19.8	18.2	19.4
Mean Number Previous Offences	9.5	5.7	8.7
Recidivism within 12 Months	54.6	28.1	48.8
Sample Size	18,928	5,351	24,279
Panel B - Released +/- 7 Days of 3 Months (88 Days) Cu-off:			
Day of Release around Cut-off	**- 7 Days**	**+ 7 Days**	**Total**
Percentage Female	11	10.2	10.3
Mean Age at Release	28.8	29.4	29.3
Percentage Incarcerated for Violence	17.1	19.7	19.4
Mean Number Previous Offences	9.1	8.6	8.7
Recidivism within 12 Months	56.8	47.9	48.8
Percentage Released on HDC	0	24.4	22
Sample Size	2,333	21,946	24,279

Figure 1: Proportion Discharged on HDC by Sentence Length

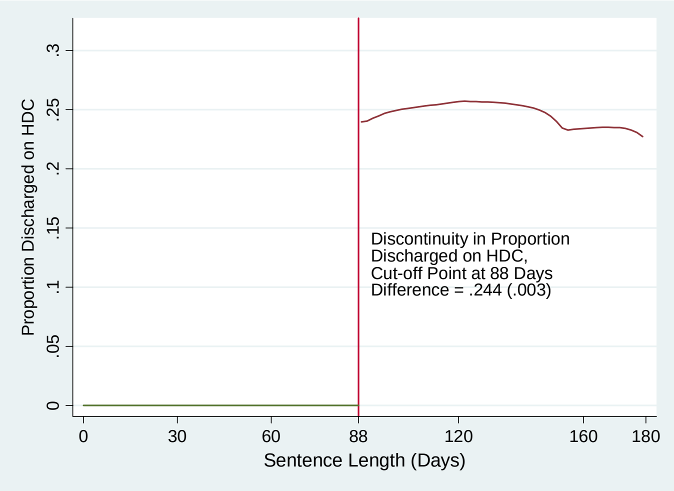

Figure 2: Mean Number of Previous Offence by Sentence Length

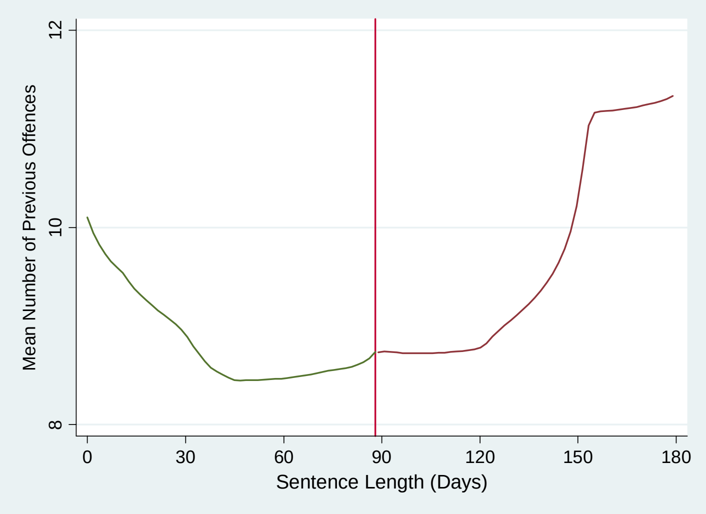

Figure 4: Recidivism within 1 Year by Sentence Length

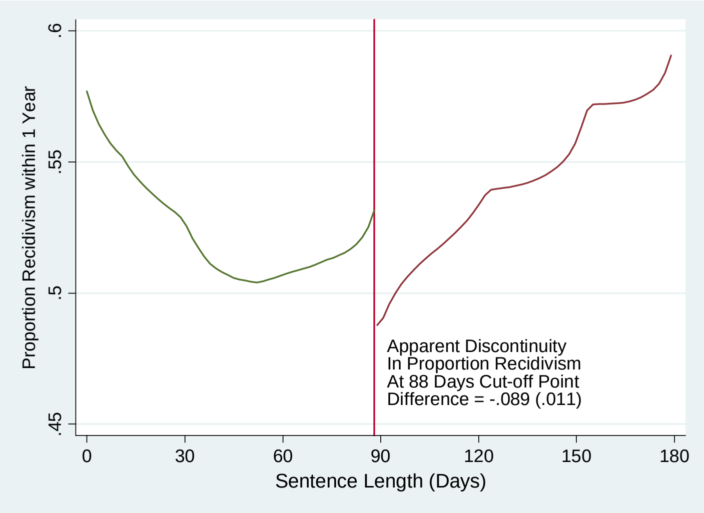

Table 4: RDD Estimates of HDC Impact on Recidivism – Around Threshold

	Dependent Variable = Recidivism Within 12 Months Estimation on Individuals Discharged +/- 7 Days of 88 Days Threshold		
	(1)	(2)	(3)
Estimated Discontinuity of HDC Participation at Threshold ($HDC^{+} - HDC^{-}$)	.243 (.009)	.223 (.009)	.243 (.003)
Estimated Difference in Recidivism Around Threshold ($Rec^{+} - Rec^{-}$)	-.089 (.011)	-.059 (.009)	-.044 (.014)
Estimated Effect of HDC on Recidivism Participation ($Rec^{+} - Rec^{-}$)/ ($HDC^{+} - HDC^{-}$)	-.366 (.044)	-.268 (.044)	-.181 (n.a.)
Controls	No	Yes	No
PSM	No	No	Yes
Sample Size	24,279	24,279	24,279

Outline

Internal and External Validity

- At best, Sharp and Fuzzy RDD estimate the average effect of the sub-population with X_i close to c
- Fuzzy RDD restricts this subpopulation even further to that of the compliers with X_i close to c
- Only with strong assumptions (e.g., homogenous treatment effects) can we estimate the overall average treatment effect
- So, RDD have strong internal validity but may have weak external validity (although it depends...)

Multiple Thresholds: Effect of Class Size

- Old Jewish law (Mainomides) says class size should not be over 40. Angrist and Lavy (1999) look at schools where cohorts are close to multiples of 40.
- If just over a multiple, the actual class size is much smaller than if the cohort size is just under 40.
- Within that set of classes they look at correlation between educational outcomes test score and class size.

Class Size Effect

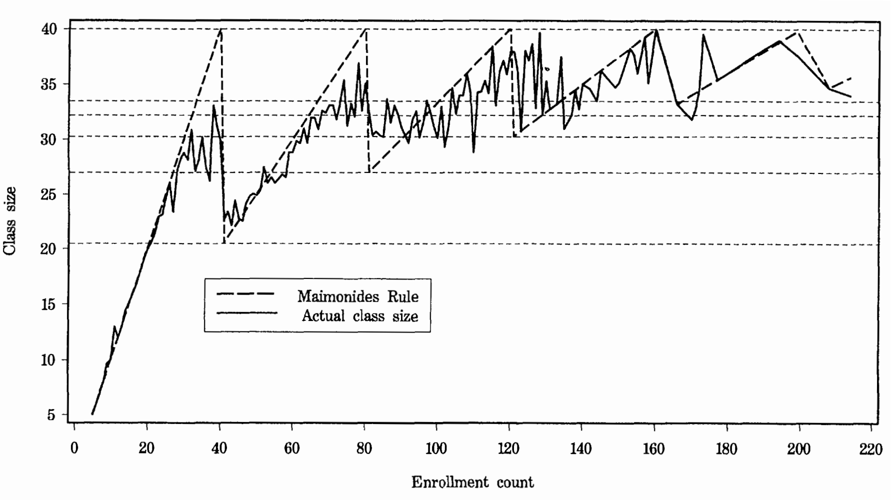

Class Size Effect