

ECON 594: Applied Economics

Regression Discontinuity

Sam Norris
University of British Columbia

Where we are

Last classes: panel data, difference-in-differences, event studies

All of them leaned on parallel trends, which we can't fully test

Today: regression discontinuity

- Treatment switches at a known cutoff → something close to random assignment
- You have seen the basics, so today is about doing it well

The big idea

Some cutoffs are arbitrary, and that makes them experiments

- Retirement possible at 65, loan denied below a credit score, admitted above a GPA, legal to drink at 21
- Compare units just above the cutoff to units just below
- Where the cutoff is arbitrary, which side you land on is as good as random
- (Works even better when people don't know the cutoff—why?)

Sharp RD: treatment is a deterministic function of the running variable X

$$D = \mathbb{1}[X \geq c]$$

- Once you know X , you know D
- Identification comes from a thin slice of data around c

Running example: the drinking age


The minimum legal drinking age switches on at 21

- A few days of age means a big change in legal access to alcohol
- Nothing else about a person changes overnight on their 21st birthday

If legal access to alcohol kills, deaths should jump right at the 21st birthday

- And only at the 21st, not at the 20th or 22nd

Deaths spike right at 21



From idea to regression

Aggregate to monthly age cells and fit a line on each side of 21

$$\bar{m}_a = \alpha + \rho D_a + \gamma a + e_a$$

- ρ is the jump at the cutoff: the MLDA effect on mortality
- γa absorbs the smooth trend in mortality with age
- Estimate: about 7.7 more deaths per 100,000

The estimate is just the gap between the two fitted lines at the cutoff

Always be plotting

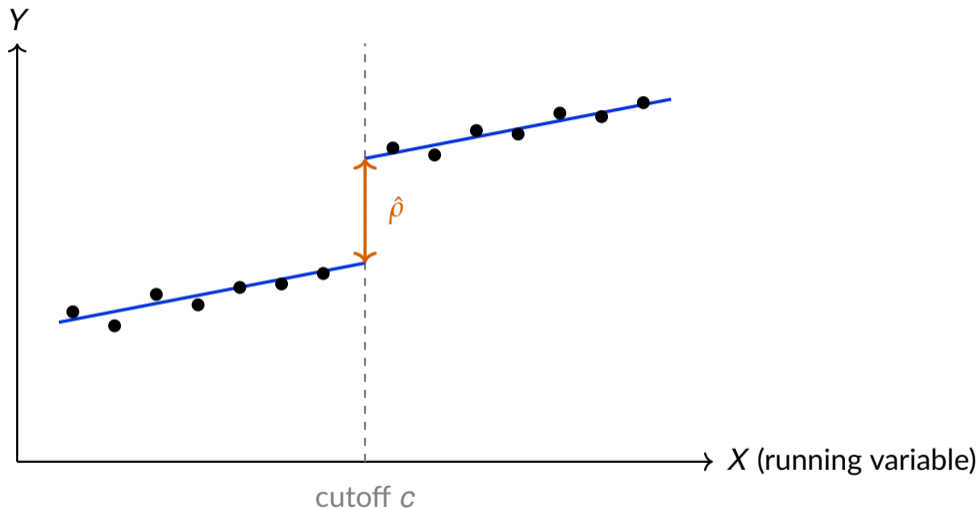
The identifying variation is visible to the naked eye!

- Bin the running variable, plot the mean outcome in each bin
- Fit a low-order curve on each side
- This is a binscatter with a cutoff: the RD plot
- A real effect is a jump you can see before running any regression

Readers trust the picture more than the coefficient

- If you cannot see it in the binned means, be skeptical of the regression

The RD plot



Binned means on each side, a low-order fit, and the jump $\hat{\rho}$ read straight off the gap.

The two regressions behind the picture

The dots are bin means: a saturated regression on bin indicators

$$Y_i = \sum_b \mu_b \mathbb{1}[i \in \text{bin } b] + u_i$$

- Plot each $\hat{\mu}_b$ against its bin's midpoint
- No functional form imposed, except bin choice

The line is a local linear fit with the running variable recentered at c

$$Y_i = \alpha + \rho D_i + \beta(X_i - c) + \delta D_i(X_i - c) + \varepsilon_i, \quad D_i = \mathbb{1}[X_i \geq c]$$

- Separate slopes: β below the cutoff, $\beta + \delta$ above
- Rescaling X_i lets you read ρ off the regression
- Fit it within the bandwidth, $|X_i - c| \leq h$

RD plots in Stata

```
ssc install rdrobust    // rdplot, rdrobust, rdbwselect
ssc install rddensity  // manipulation test
```

* The picture first: binned means + a low-order fit on each side

```
rdplot y x, c(0)
```

* More bins to show the raw variation; fewer to smooth it

```
rdplot y x, c(0) nbins(20 20)
```

rdplot picks the number of bins for you, but always eyeball a few choices

- The jump should be there whether you use 10 bins or 40

Two ways to control for the running variable

You need a model for the trend on each side, and there are two philosophies

- Global polynomial: fit one flexible curve to all the data on each side
- Local linear: fit a straight line in a narrow window around the cutoff

Modern practice favors local linear, or local quadratic

- Gelman & Imbens (2019): high-order global polynomials put strange weight on far-away points and give noisy, knife-edge estimates
- If in doubt, go linear

The bandwidth tradeoff

Local linear means picking a bandwidth h , then keeping only $|X - c| \leq h$

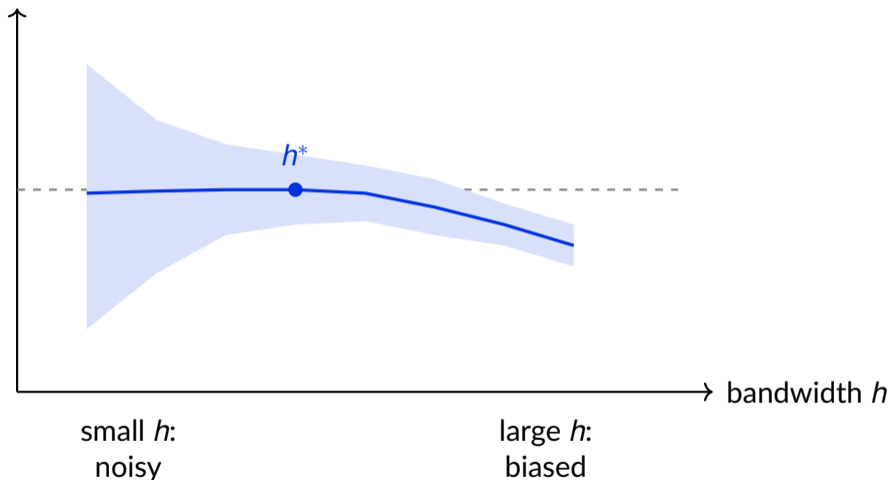
- Wider h : more data and less variance, but more bias if the trend is curved
- Narrower h : less bias, but noisier
- The optimal h trades the two off (MSE-optimal)

Don't hand-pick the bandwidth

- Use a data-driven selector (Calonico, Cattaneo & Titiunik 2014)
- Robust confidence intervals correct for the bias the optimal bandwidth leaves behind

Show that it doesn't hinge on the bandwidth

RD estimate $\hat{\rho}$



Re-estimate across a range of bandwidths and plot the coefficient. Flat across sensible h is reassuring; flipping sign means the result is a bandwidth artifact.

Bandwidth and estimation in Stata

```
* Local linear, triangular kernel, MSE-optimal bandwidth, robust CIs  
rdrobust y x, c(0) p(1) kernel(triangular)
```

```
* See the bandwidth it chose  
rdbwselect y x, c(0)
```

```
* Sensitivity: loop over bandwidths and store the estimate  
foreach h in 2 4 6 8 10 {  
    rdrobust y x, c(0) h('h')  
    * collect e(tau_cl), e(ci_l), e(ci_r) and plot vs. h  
}
```

The default `rdrobust` call is a sensible starting point for a thesis

- Can show the bandwidth-sensitivity plot in the appendix
- As always, institutional knowledge should guide bandwidth choice

Four things to investigate

Checklist for RD

1. Manipulation and sorting
2. Do other parts of the treatment bundle change at the threshold?
3. Is there an IV interpretation?
4. RD complications (things worth knowing)

The rest of today works through each in turn

Placebo outcomes: things that shouldn't move

Table 4.1: Sharp RD estimates of MLDA effects on mortality

Dependent Variable	Ages 19-23		Ages 20-22	
	(1)	(2)	(3)	(4)
All Deaths	7.66 (1.51)	9.55 (1.83)	9.75 (2.06)	9.61 (2.29)
Motor Vehicle Accidents	4.53 (0.72)	4.66 (1.09)	4.76 (1.08)	5.89 (1.33)
Suicide	1.79 (0.50)	1.81 (0.78)	1.72 (0.73)	1.30 (1.14)
Homicide	0.10 (0.45)	0.20 (0.50)	0.16 (0.59)	-0.45 (0.93)
Other External Causes	0.84 (0.42)	1.80 (0.56)	1.41 (0.59)	1.63 (0.75)
All Internal Causes	0.39 (0.54)	1.07 (0.80)	1.69 (0.74)	1.25 (1.01)
Alcohol Related Causes	0.44 (0.21)	0.80 (0.32)	0.74 (0.33)	1.03 (0.41)
Controls	age	age, age ² , both interacted with <i>D</i>	age	age, age ² , both interacted with <i>D</i>

Notes: This table reports coefficients on an over-21 dummy from regressions of month-of-age-specific death rates by cause on an over-21 dummy and linear or interacted quadratic age controls. Standard errors are reported in parentheses.

Pick outcomes the treatment cannot plausibly affect. The MLDA moves motor-vehicle and alcohol deaths; it should not move deaths from internal causes (cancer, disease) — and it doesn't.

Placebo cutoffs: nothing there to find

Re-run the design at fake cutoffs where nothing should happen

- Pretend the drinking age is at $c = 20$ or $c = 22$ and estimate the “effect”
- Keep the spec otherwise the same
- A clean design finds nothing at the placebo cutoffs
- A design that finds jumps everywhere is picking up curvature or noise, not treatment

Manipulation: can people pick their side?

McCrary's idea: people shouldn't be able to choose the running variable

- If they can, those who sort across the cutoff are selected, and continuity fails
- Check the density of X for a spike or a hole at c
- Formal test: `rddensity`


Watch for heaping and round numbers

- Mass points at salient values (age 85, 40 hours worked) signal manipulation/rounding
 - Not always the end of the world, but can create subtle issues (e.g., age heapers tend to be less educated)
- The donut RD drops observations right at the cutoff and re-estimates


What the density test is looking for

Plot the density of the running variable and look right at the cutoff

density no manipulation



density manipulation



Raw exam scores with an A cutoff: if graders nudge near-misses over the line, scores just below an A pile up at it—a spike at the cutoff and a hole beneath. `rdensity` tests whether that jump in the density is real.

Validity checks in Stata

* Manipulation: is the density of the running variable smooth at c ?

```
rddensity x, c(0)
```

* Placebo cutoffs: nothing should jump where there is no policy

```
rdrobust y x if x < 0, c(-2)
```

```
rdrobust y x if x >= 0, c(2)
```

Neither the density nor the placebo cutoffs should show a jump

Is the focal treatment the only thing that changes at the cutoff?

Bundling is a different worry from placebo tests

- Placebo tests ask whether the units on each side are comparable
- Bundling asks whether crossing c changes only the focal treatment, or other things too
- Can think of as exogeneity versus exclusion

Even with perfectly comparable units, the jump can bundle other changes

- Canada-US border RD — many things change at the border besides health insurance
- Close election RD — the winner's gender correlates with party, incumbency, etc

RD on a characteristic of the winner

RD on who wins a close election, where the winner has some characteristic

- Running variable: the female candidate's margin of victory
- Treatment: a woman holds the seat
- Outcome: corruption, public health, spending, child mortality

Close races make which candidate wins as good as random

- Lee (2008) is the canonical close-election RD
- Examples on gender: Brollo & Troiano (2016, Brazil mayors); Bhalotra & Clots-Figueras (2014, India assemblies)

But “a woman wins” is a bundle

A close win by a woman over a man is not the same as randomly assigning gender

- To stay competitive against men, the women who win close races differ in other ways
- Party, incumbency, experience, money – all can vary with gender

Marshall (2024): the design identifies gender plus its compensating differentials

- District covariates are still continuous at the cutoff, so design is valid
- It estimates is the effect of the whole bundle, not the one focal characteristic

Benchmark the bundle: put candidate characteristics on the LHS

Run the same RD with each of the winner's other traits as the outcome

- E.g. party, incumbency, experience, age, campaign spending
- Whatever jumps at the cutoff is bundled into your “effect of a woman”

$$\text{characteristic}_i = \alpha + \rho D_i + \gamma(X_i - c) + e_i, \quad \text{want } \rho = 0$$

```
foreach v in party incumbent experience age spend {  
  rdrobust 'v' female_margin, c(0)  
}
```

Here the goal is to size the bundle: report the jumps alongside the main estimate so the reader can weigh them

When treatment doesn't switch cleanly: fuzzy RD

Sometimes crossing the cutoff changes the probability of treatment, not treatment itself

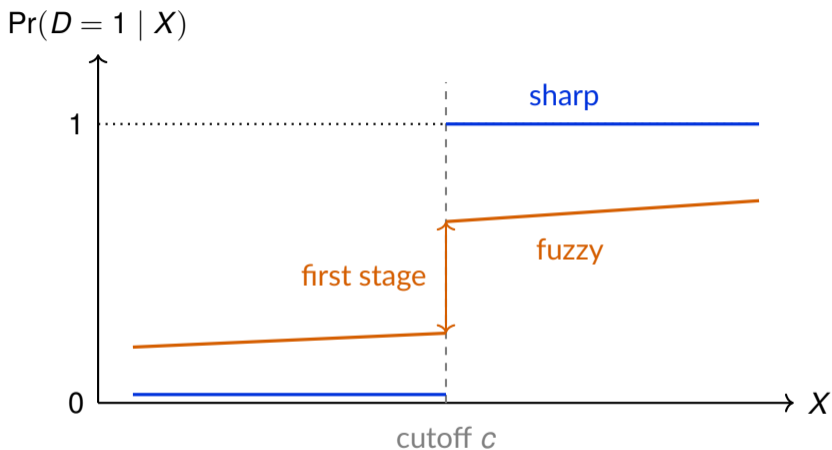
- Charter school admission raises the chance of attending, but not all the way to one
- Can use the jump in treatment at c as the first stage

The estimate becomes the Wald ratio:

$$\rho_{\text{fuzzy}} = \frac{\text{jump in } Y \text{ at } c}{\text{jump in } D \text{ at } c} = \frac{\text{reduced form}}{\text{first stage}}$$

- Clearing the cutoff is the instrument for treatment
- Exclusion restriction: clearing the cutoff affects Y only through D

Sharp vs fuzzy



Sharp: treatment goes from 0 to 1 at the cutoff. Fuzzy: the probability jumps, and that jump is the first stage.

IV light: rescaling to ballpark a magnitude

Sometimes you want an effect per unit of a channel, not as a reduced form

- Law change increases likelihood of parole for defendants sentenced after date c , but also slightly reduces average sentence
- First stage: the jump in parole likelihood at the cutoff
- Wald ratio: the reduced-form jump in Y over the jump in the channel

$$\frac{\text{jump in recidivism at } c}{\text{jump in parole at } c} \approx \text{“effect per parole”}$$

Be honest that the exclusion restriction almost surely fails here

- Policy changes many things at once, not just the parole channel
- So report the ratio as a back-of-envelope magnitude, not a clean causal parameter
- The more the other characteristics jump, the less interpretable the “IV light” number

Extra tools for specific situations

Discrete running variable, few distinct values

- Old advice was to cluster on the running variable (Lee & Card 2008)
- That undercovers; use honest CIs instead (Kolesár & Rothe 2018)

The cutoff is not known in advance

- Test for a discontinuity and estimate where it is (Porter & Yu 2015)
- Classic application: estimating tipping points (Card, Mas & Rothstein 2008)

Another policy changes at the same cutoff

- Difference out the second jump: difference-in-discontinuities (Grembi, Nannicini & Troiano 2016)
- For example, Deshpande (2016) on SSI removal at age 18

Treatment kinks rather than jumps at the cutoff

- The slope changes, not the level: regression kink design (Card, Lee, Pei & Weber 2016)

Where to read more

The practical how-to:

- Cattaneo, Idrobo & Titiunik – *A Practical Introduction to Regression Discontinuity Designs* (two short volumes, with code)
- Lee & Lemieux (2010, J. Econ. Literature) – the standard survey

On the methods choices:

- Calonico, Cattaneo & Titiunik (2014) on bandwidth and robust inference
- Gelman & Imbens (2019) on polynomials; McCrary (2008) on manipulation

On close elections:

- Caughey & Sekhon (2011) and Eggers et al. (2015) on when the design is valid
- Marshall (2024) on what a characteristic RD can and cannot identify

Checklist for your thesis

Always:

- Show the RD plot (binned means + low-order fit) before any regression
- Local linear, data-driven bandwidth, robust CIs (`rdrobust`)
- Plot the estimate across a range of bandwidths

Validity:

- Covariate balance, placebo outcomes, placebo cutoffs
- Density test for manipulation, and consider a donut

If it's fuzzy or a characteristic RD:

- Report the first stage and spell out the exclusion restriction
- For characteristic RDs, put the rival characteristics on the LHS

Recap

RD turns an arbitrary cutoff into an experiment, identified just around the cutoff

- The picture is the evidence: a jump you can see in the binned means
- Bandwidth and functional form are choices, so show the result survives them
- Placebos test continuity, and the density test catches sorting

Fuzzy RD is IV, and characteristic RDs estimate a bundle

- Rescale only with the exclusion restriction stated out loud, and ballpark honestly

Coming up

Next: Instrumental Variables

- The general version of the fuzzy-RD ratio we just saw

Then: Measurement Error and Other Empirical Concerns

Reminder: progress report due Wednesday June 3rd