Lab 9: Regression Discontinuity Design

PHS2000B, Spring 2023

Sudipta Saha

March 27, 2023

Overview

1. **Intuition behind RDD**
2. **When can we use RDD and why do folks like it?**
3. **Sharp vs Fuzzy RDD**
4. **Identifying Assumptions**
5. **Effect Types**
6. **Questions**
7. **How-to**
8. **Group work :)**

Intuition behind RDD

Discontinuities in a Continuous World

We generally expect things to be continuous or gradual in nature

But we often create various thresholds, rules and eligibility criteria

What are some reasons?

"... the experience of operating within a system of formalized rules and regulations ... actually does hold—for many of us much of the time, for all of us at least some of the time—a kind of covert appeal."

"Bureaucracy holds out at least the possibility of dealing with other human beings in ways that do not demand either party has to engage in all those complex and exhausting forms of interpretive labor ..."- David Graeber

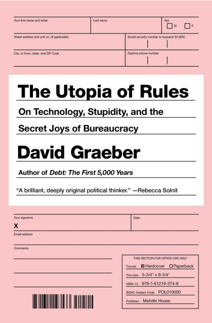

Discontinuities in a Continuous World

Some examples:

- Having thresholds make decisions easier for medical providers
  - Low birth weight
  - Advanced maternal age
  - Lab test ranges, thresholds for clinical risk scores
- Thresholds as a form of austerity
  - Means tested benefits
  - Allocation of scarce medications (age thresholds for COVID vaccines)

Sometimes, rules and thresholds can be arbitrarily cruel, preventing much needed resources from going to someone just above or below a threshold. Maybe a RDD could provide evidence of this to make the case for expanding access?

Discontinuities in a Continuous World

Lets think of an arbitrary threshold

Imagine we have a state flagship university that has a strict cut-off for admissions, based on SAT scores. All applicants with a score of 1250 and above are offered admissions, and no applicant with a score of 1240 or below are offered admissions (SAT scores change in increments of 10)

Consider two people randomly selected from the groups on either side of the threshold:

Maya scored 1240 | **Aaliyah scored 1250**

Other than their admission offer, do we expect people like Maya and Aaliyah to be very different from one another?

What might we assume about the group of individuals who scored 1240 vs the group that scored 1250?

Thinking in terms of potential outcomes or counterfactuals

RDD will answer questions like:

- What would happen if people at the threshold got admitted compared to if they did not get admitted?

We can't answer what would have happened if **everyone** got admitted, admissions are "as good as random" only for those with SAT scores near the threshold.

SAT Score and Enrollment

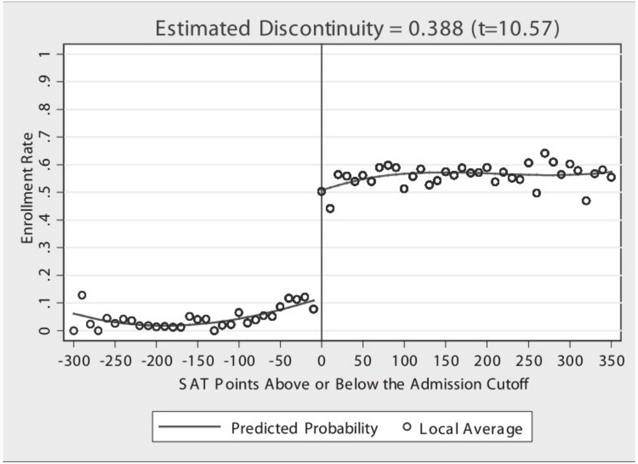

Hoekstra, Mark. 2009. "The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach." Review of Economics and Statistics 91 (4): 717–24.

SAT Score and Future Earnings

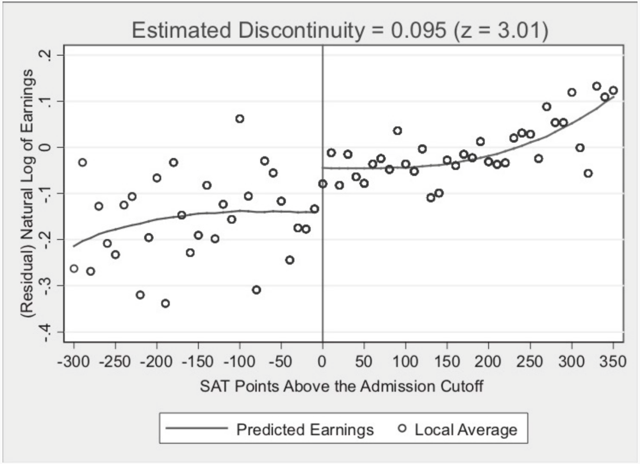

Hoekstra, Mark. 2009. "The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach." Review of Economics and Statistics 91 (4): 717–24.

When can we use RDD and why do folks like it?

When can we use regression discontinuity?

The examples we have discussed in class have three things in common:

- There is a continuous variable that is used to determine eligibility for treatment. In RDD we will call this the **running (or ”forcing”) variable**. Examples: test scores, GPA, birth weight, income, CD4 count, maternal age, vote share
- **There is a cut-off or threshold** in the running variable, at which point the probability of treatment or exposure changes discontinuously (it “jumps”). The location of this cut-off can be somewhat arbitrary, but above it, the probability of treatment is different than the probability below it.
- There is a **bandwidth** of observations around the cut-off that are considered comparable

Some questions to discuss

Imagine that there is a strictly enforced cut-off for a program, such that you are only eligible above a threshold and no one below the threshold is eligible. If this cut-off is enforced without any exceptions, what does this tell you about the discontinuity at the threshold?

Now imagine a threshold in a forcing variable that is entirely ignored? What does this imply about the probability of treatment at the threshold?

We keep saying that thresholds in RDD should be, for all intents and purposes, arbitrary – can you think of some thresholds where a threshold many not actually be as arbitrary as it seems?

Why do folks like regression discontinuity designs?

1. **They reflect real world treatment assignment mechanisms**

RDDs are one of the few study designs where an investigator knows the treatment assignment mechanism. Knowing the assignment mechanism (or "rule") allows us to develop models which more accurately capture the "selection mechanism", i.e., the mechanism that leads to certain people being treated while others not.

2. **They have high internal validity**

Studies have shown that RDD estimates have very high internal validity. For example, Chaplin et. al. (2018) test estimates from 15 RDD studies against comparable randomized controlled trials and find that the RDD studies very closely approximate the results from the RCTs.

RDD's popularity

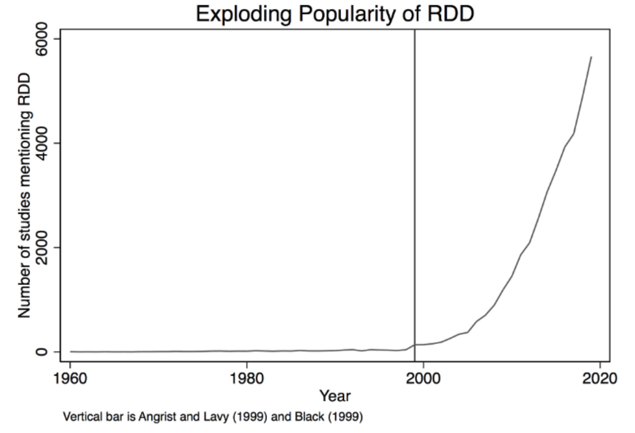

Cunningham, Scott. Causal inference: The Mixtape. Yale University Press, 2021.

At this point you are probably dying for a DAG

We love DAGs. But sometimes, our love for DAGs is unrequited 💔

> **Dr Ellie Murray, ScD**
> @EpiEllie
>
> Interesting discussion here. Anyone have ideas on a DAG or SWIG for regression discontinuity? 👇
>
> > **scott cunningham** @causalinf · Jun 11, 2019
> > In my Mixtape, I tried to convey an RDD via a DAG and the more I look at it, the more problematic it looks. First, the discontinuity (c0) is presented alongside the ray from X to D, but c0 is not a variable. It's a value.
> > Show this thread

Sharp vs Fuzzy RDD

Sharp vs Fuzzy

Sharp RDD

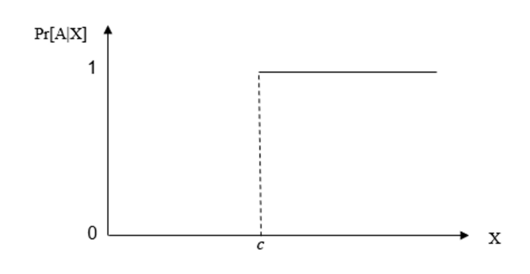

Fuzzy RDD

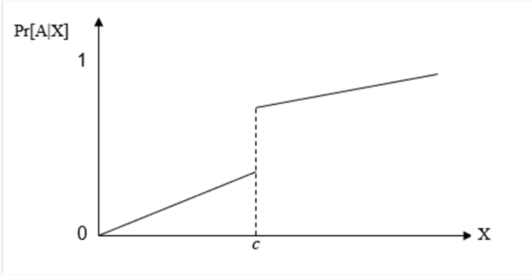

Note that the y-axis here is the probability of treatment

In the SAT score and enrollment example, what would be the difference between the sharp vs fuzzy RDD?

Challenge of causal inference in a sharp RDD

Let $A = \{0, 1\}$ be a binary treatment. $A = 1$ when $X \geq c$. That is, c is the threshold.

$E[Y^a|X]$ is the expected value of the outcome Y when treatment is set to a, given a specific value of X.

Solid line = observed potential outcomes
Dashed line = unobserved potential outcomes

We can only observe potential outcomes under treatment when $X \geq c$ and potential outcomes under no treatment when $X < c$. In other words, sharp RDDs suffer from a violation of structural positivity.

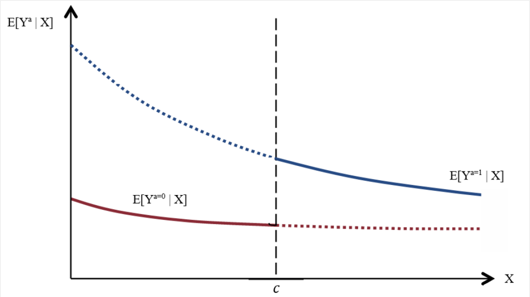

Adapted from Bor et al. (2014)

Challenge of causal inference in a fuzzy RDD

Fuzzy RDDs can be though of as sharp RDDs with **imperfect compliance or adherence**

Imperfect compliance might make us worry about why some individuals with $X < c$ are receiving treatment and why some individuals with $X \geq c$ are not

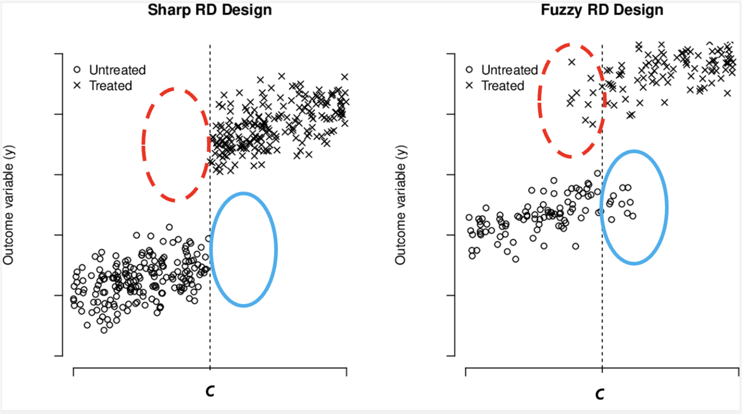

Challenge of causal inference in a fuzzy RDD

Fuzzy RDDs can be though of as sharp RDDs with **imperfect compliance or adherence**

Why might the folks in the dashed red circle in the fuzzy RD design be treated, even though they are below the threshold?

Why might folks in the solid blue circle be untreated, even though they are above the threshold?

Perhaps the things that cause people to choose or not choose treatment is related to the outcome (i.e. confounding)

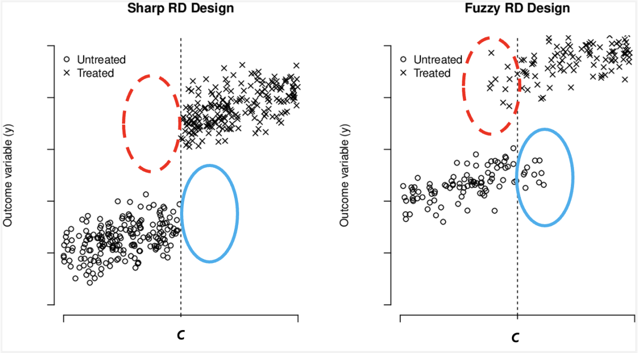

Fuzzy RDD

- We have an arbitrary threshold that is strongly (but not perfectly) determines treatment status
- We assume that the arbitrariness of the threshold means that there aren't other changes at the threshold that can affect the outcome
- This is exactly like an instrumental variable situation!

Estimating causal effects in sharp vs fuzzy RDD

Sharp RDD

We are usually interested in estimating the **Average Causal Effect (ACE)** at the threshold

$$ACE_{SRD} = E[Y_i^{a=1} - Y_i^{a=0} | X = c]$$

Using our data:

$$ACE_{SRD} = \lim_{x \downarrow c} E[Y_i | X_i = x] - \lim_{x \uparrow c} E[Y_i | X_i = x]$$

Fuzzy RDD

We are usually interested in estimating the **Complier Average Causal Effect (CACE)** at the threshold

$$CACE_{FRD} = E[Y_i^{a=1} - Y_i^{a=0} | X = c, \text{compliers}]$$

Using our data:

$$CACE_{FRD} = \frac{\lim_{x \downarrow c} E[Y_i | X_i = x] - \lim_{x \uparrow c} E[Y_i | X_i = x]}{\lim_{x \downarrow c} E[A_i | X_i = x] - \lim_{x \uparrow c} E[A_i | X_i = x]}$$

Estimating causal effects in sharp vs fuzzy RDD

Sharp RDD Model

$$Y_i = f(X_i) + \beta A_i + \epsilon_i$$

Fuzzy RDD Model

Same as IV, estimate two-stage least squares (2SLS) model

Identifying Assumptions

Key identifying assumption

The key identifying assumption in RDD is about the potential outcomes at the threshold

We assume that **the potential outcomes are continuous at the threshold**.

That is, the outcome that we would have observed were everyone treated is continuous at the threshold. And similarly the outcome that we would have observed were everyone not treated is continuous at the threshold

Let's try to draw this

Identifying Assumptions

		Sharp RD design	Fuzzy RD design
	Relevance	By definition	Needed (check that threshold is actually causing a discontinuity in treatment)
Expected potential outcomes are continuous in the running variable around the threshold	Independence/ Exchangeability	People cannot manipulate the running variable The threshold cannot *coincide with* a jump in another variable which affects the potential outcome	
	Exclusion	The threshold cannot *cause* a jump in another variable which affects the potential outcome	
	Monotonicity	Not needed (everyone is a complier)	Needed (IV framework)

Effect Types

Keeping up with the effects

Average Treatment Effect (ATE) is the effect of treatment in the whole population. Sometimes used interchangeably with ACE below, but in this course we will stick to this definition of ATE. $E[Y_i^{a=1} - Y_i^{a=0}]$

Average Causal Effect (ACE) is the effect of treatment at the threshold. $E[Y_i^{a=1} - Y_i^{a=0} | X = c]$. ACE = ATE if we assume that the effect at the threshold applies to everyone (very very unlikely)

Local Average Treatment Effect (LATE) is the effect of treatment among compliers. $E[Y_i^{a=1} - Y_i^{a=0} | \text{compliers}]$

Complier Average Causal Effect (CACE) is the effect of treatment among compliers at the threshold. $E[Y_i^{a=1} - Y_i^{a=0} | X = c, \text{compliers}]$

Effect of the threshold / ITT at the threshold is the effect of the threshold mechanism. $E[Y_i^{z=1} - Y_i^{z=0} | X = c]$, where z is an indicator for being on one side of the threshold.

Questions (pollev.com/shu)

True, False or Uncertain? (pollev.com/shu)

You are trying to identify the causal effect of a new antimalarial drug on child survival.

You find that this drug is recommended for children who have malaria and a body temperature above 104F.

You examine the data and find that utilization increases from 20% for children with body temperature $<$104F to 40% for children with body temperature $>$104F.

You make no assumption regarding the distribution of compliers at the threshold of 104F.

A RDD analysis of these data will allow us to identify the Average Causal Effect (ACE) as well as the Average Treatment Effect (ATE).

True, False or Uncertain? (pollev.com/shu)

You are trying to identify the causal effect of a new antimalarial drug on child survival.

You find that this drug is recommended for children who have malaria and a body temperature above 104F.

You examine the data and find that utilization increases from 20% for children with body temperature $<$104F to 40% for children with body temperature $>$104F.

You make no assumption regarding the distribution of compliers at the threshold of 104F.

A RDD analysis of these data will allow us to identify the Average Causal Effect (ACE) as well as the Average Treatment Effect (ATE).

Generally false. This is a fuzzy RDD so what we can identify using the data is the Complier Average Causal Effect (CACE). Under the practically implausible assumption that the compliers are a random sample of individuals at the cut-off, then we might argue that the CACE equals the ACE.

True, False or Uncertain? (pollev.com/shu)

In a sharp RDD, the Average Causal Effect (ACE) will equal the Complier Average Causal Effect (CACE).

True, False or Uncertain? (pollev.com/shu)

In a sharp RDD, the Average Causal Effect (ACE) will equal the Complier Average Causal Effect (CACE).

True. Everyone is a complier.

True, False or Uncertain? (pollev.com/shu)

In a sharp RDD, the Average Causal Effect (ACE) will equal the Average Treatment Effect (ATE).

True, False or Uncertain? (pollev.com/shu)

In a sharp RDD, the Average Causal Effect (ACE) will equal the Average Treatment Effect (ATE).

False. ACE=ATE is only true if the treatment effect is assumed to be homogenous along the running variable.

True, False or Uncertain? (pollev.com/shu)

We can only conduct a regression discontinuity analysis when individuals are randomly sorted around the cut-off. In other words, we can only conduct a regression discontinuity analysis when we have a ”locally randomized experiment” (i.e., local randomization is a necessary condition for identification).

True, False or Uncertain? (pollev.com/shu)

We can only conduct a regression discontinuity analysis when individuals are randomly sorted around the cut-off.

Local Randomization Assumption

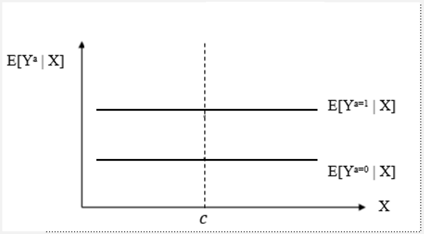

RDD Identifying Assumption

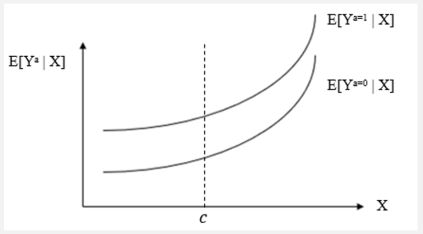

False. Treatment effects in a RDD are identified assuming that the potential outcomes are continuous at the cut-off. Although local randomization is a sufficient condition for continuity of potential outcomes at the cut-off, it is not necessary.

True, False or Uncertain? (pollev.com/shu)

Covariate imbalance on either side of the cut-off necessarily implies that the continuity of potential outcomes is invalid and therefore a RDD analysis is not feasible.

True, False or Uncertain? (pollev.com/shu)

Covariate imbalance on either side of the cut-off necessarily implies that the continuity of potential outcomes is invalid and therefore a RDD analysis is not feasible.

False. Covariate imbalance does not necessarily imply that potential outcomes are not continuous at the cut-off. If covariates jump discontinuously at the cut-off, then we would be worried about there being a violation of the continuity assumption. The fact that the distribution of covariates are different on either side of the cut-off is not necessarily problematic.

True, False or Uncertain? (pollev.com/shu)

While variables like self-reported income can be manipulated, data on biological variables like age or birth weight are always free from manipulation.

How-to

Steps in your RDD analysis

1. **Center the running variable around the cut-off**

The value 0 in the centered data will correspond to the threshold.

2. **Create a "discontinuity sample" by trimming your sample to include observations within a certain bandwidth h around the cut-off**

This focuses our analysis to the threshold. The selection of the size of bandwidth h can be done by the researcher or by using algorithms, which try to optimize the bias-efficiency trade-off.

3. **Fit the model**

Sharp RDD: $Y_i = f(X_i) + \beta A_i + \epsilon_i$

Fuzzy RDD: Estimate two-stage least squares (2SLS) model. Or just regress

Diagnostics

1. **Balance checks** Visually explore relationship of covariates with running variable. Fit your RDD model, except have the covariate as the outcome, to empirically check for balance of measured covariates.

2. **Visually check for sorting around the threshold for the running variable**

Sorting or "bunching" around the cut-off may indicate manipulation of the running variable. This violates the continuity assumption

We can visually inspect for bunching by creating a histogram/density plot of the running variable and assessing if there are discontinuous changes in the density at the cut-off

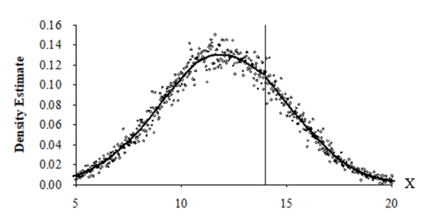

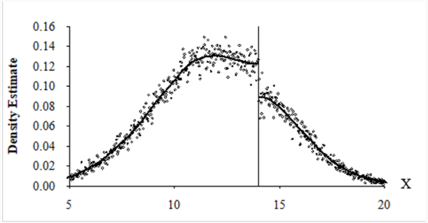

Diagnostics

3. **Placebo Thresholds**

Empirically test for jumps in the outcome along the running variable by splitting the dataset at the cut-off and, in each dataset, estimating the RDD model at "placebo thresholds"

4. **Sensitivity analysis (bin size, functional form specification)**

A good RDD analysis will present results from regressions with different functional form specifications and different bandwidth sizes. These are often referred to as "robustness" or "specification tests."

Group Work :) (https://tinyurl.com/nhbts9yv)