Identification in Regression Discontinuity Designs with Multiple Cutoffs[*]

Matias D. Cattaneo[†]   Luke Keele[‡]   Rocío Titiunik[§]

Gonzalo Vazquez-Bare[¶]

February 4, 2015

Abstract

We consider a regression discontinuity design where the treatment is received if a score is above a cutoff, but the cutoff may vary for each unit in the sample instead of being the same for all units. This multi-cutoff regression discontinuity design is very common in practice, and researchers often normalize the score variable and then use the zero cutoff on the normalized score for all observations to estimate a pooled regression discontinuity treatment effect. We formally derive the form that this pooled parameter takes, and discuss its interpretation under different assumptions. We show that this normalizing-and-pooling strategy so commonly employed in empirical work may not fully exploit all the information available in the multi-cutoff regression discontinuity design. We illustrate our results with two regression discontinuity examples based on vote shares, using data from Brazilian mayoral elections and U.S. Senate elections.

Word Count: $9,939$

1 Introduction

The regression discontinuity (RD) design has become one of the preferred quasi-experimental research designs in the social sciences, mostly as a result of the relatively weak assumptions that it requires in order to recover causal effects. In the "sharp" version of the RD design, every subject is assigned a score and a treatment is given to all units whose score is above the cutoff and withheld from all units whose score is below it. Under the assumption that all possible confounders vary smoothly at the cutoff as a function of the score (also known as "running variable"), a comparison of units barely above and barely below the cutoff can be used to recover the causal effect of the treatment—see Imbens and Lemieux (2008) and Lee and Lemieux (2010) for reviews. The RD design is particularly important in political science, where the discontinuous assignment of victory in close elections often provides a credible research design to make causal inferences about elite behavior. Although the RD design has been found to fail in U.S. House elections (Caughey and Sekhon, 2011), Eggers et al. (2015) show that RD designs based on elections seem to be generally valid as an identification strategy to recover causal effects in other electoral contexts.

In a standard RD design, the cutoff in the score that determines treatment assignment is known and the same for all units. For example, in the classic education example where a scholarship is awarded to students who score above a threshold on a standardized test (Thistlethwaite and Campbell, 1960), the cutoff for the scholarship is known (at least to the test administrators), and it is the same for every student. However, in many applications of the RD design, the value of the cutoff may vary by unit. The most common example of variable cutoffs occurs in political science applications where the score is a vote share, the unit is an electoral constituency, and the treatment is winning an election under plurality rules.

When there are only two options or candidates in an election, the victory cutoff is always 50% of the vote, and it suffices to know the vote share of one candidate to determine the

winner of the election and the margin by which the election was won. This occurs most naturally in political systems dominated by exactly two parties, or elections such as ballot initiatives where the vote is restricted to only two yes/no options (see, e.g., DiNardo and Lee, 2004). However, when there are three or more candidates, two races decided by the same margin might result in winners with very different vote shares. For example, in one district a party may barely win an election by 1 percentage point with 34% of the vote against two rivals that get 33% and 33%, while in another district a party may win by the same margin with 26% of the vote in a four-way race where the other parties obtain, respectively, 25%, 25% and 24% of the vote.

Standard practice for dealing with this heterogeneity in the value of the cutoff has been to normalize the score so that the cutoff is zero for all units. For example, researchers often use as running variable the margin of victory of the party of interest, defined as the vote share of the party minus the vote share of its strongest opponent. Using margin of victory as the score allows researchers to pool all observations together, regardless of the number of candidates in each particular district, and make inferences as in a standard RD design with a single cutoff. This normalizing-and-pooling approach is ubiquitous in political science and also in other disciplines. Table 1 lists twenty-three different examples in political science that use an RD design based on vote shares and adopt this approach. Table **??** in the Supplemental Appendix shows twenty-three additional examples from other disciplines, including education, economics and criminology, where this approach has been applied.

Despite the widespread use of the normalizing-and-pooling strategy in applications, the exact form and interpretation of the treatment effect recovered by this approach has not been formally explored. This is the motivation for our article. We generalize the conventional RD setup with a single fixed cutoff to an RD design where the cutoff is a random variable, and use this generalization to characterize the treatment effect parameter estimated by the pooling approach. We formally show that the pooled parameter can be interpreted as a double average: the weighted average across cutoffs of the local average treatment effects

across all units facing each particular cutoff value. This weighted average gives higher weights to those values of the cutoff that are most likely to occur and concentrate a high number of observations. Our derivations thus show that the pooled estimand is not equal to the overall average of the average treatment effects at every cutoff value, except under particular assumptions.

We also use our framework to characterize the heterogeneity that is aggregated in the pooled parameter and the assumptions under which this heterogeneity can be used to learn about the causal effect of the treatment at different values of the score. As we show, the probability of facing a particular value of the cutoff may vary with characteristics of the units. If these characteristics of the units also affect the outcome of interest, then differences between treatment effects at different values of the cutoff variable may be due to inherent differences in the types of units that happen to concentrate around every cutoff value. However, if the cutoff value does not directly affect the outcomes and units are placed as-if randomly at each cutoff value, then a treatment effect curve can be obtained.

We illustrate our results with two RD examples based on vote shares. The first example analyzes the effect of the Democratic party barely winning a U.S. Senate seat in the 1914-2010 period on the probability of winning the same seat in the following election. The second example uses mayoral elections in Brazil, and studies the effect of the Party of Brazilian Social Democracy (PSDB, Partido da Social Democracia Brasileira) winning the election in the 1996-2012 period on the probability that the party wins the mayor's office in the following election.[1] Our examples illustrate the different situations that researchers are likely to encounter in practice: one in which the normalizing-and-pooling approach hides very little heterogeneity because most races have an effective number of parties close to two, and one where the normalizing-and-pooling approach hides considerable heterogeneity because the average number of effective parties is larger.

Before concluding, we discuss recommendations for practice to guide researchers in the

---

[1]For details on the data sources for the U.S. and the Brazil examples see, respectively, Cattaneo, Frandsen, and Titiunik (2015) and Klašnja and Titiunik (2014).

Table 1: Empirical Examples in Political Science of RD Designs with Normalization and Pooling

Citation	Place	Running Variable	Outcome
Albouy (2013)	U.S.	Vote Share	Federal Spending
Boas and Hidalgo (2011)	Brazil	Vote Share	Incumbency
Boas, Hidalgo, and Richardson (2014)	Brazil	Vote Share	Govt Contracts
Brollo et al. (2013)	Brazil	Vote Share	Federal Transfers
Broockman (2009)	U.S.	Vote Share	Reverse Coattails
Butler (2009)	U.S.	Vote Share	Incumbency
Duraisamy, Lemennicier, and Khouri (2014)	India	Vote Share	Incumbency
Eggers and Hainmueller (2009)	UK	Vote Share	Wealth
Eggers et al. (2015)	Several	Vote Share	Incumbency
Ferreira and Gyourko (2009)	U.S.	Vote Share	Policy Outcomes
Folke and Snyder (2012)	U.S.	Vote Share	Gov. Vote Share
Gagliarducci and Paserman (2012)	Italy	Vote Share	Early Termination
Gerber and Hopkins (2011)	U.S.	Vote Share	Municipal Spending
Hainmueller and Kern (2008)	Germany	Vote Share	Incumbency
Kendall and Rekkas (2012)	Canada	Vote Share	Incumbency
Klašnja (2014)	Romania	Vote Share	Incumbency
Klašnja and Titiunik (2014)	Brazil	Vote Share	Incumbency
Lee, Moretti, and Butler (2004)	U.S.	Vote Share	Incumbency
Lee (2008)	U.S.	Vote Share	Incumbency
Pettersson-Lidbom (2008)	Sweden	Vote Share	Fiscal Policy
Trounstine (2011)	U.S.	Vote Share	Incumbency
Uppal (2009)	India	Vote Share	Incumbency
Uppal (2010)	U.S.	Vote Share	Incumbency

interpretation and analysis of RD designs with multiple cutoffs. We also outline several directions in which our framework can be extended, including the analysis of fuzzy RD designs, Regression Kink designs, RD designs with multiple running variables, different population parameters of interest, and endogenous selection of units into cutoffs. Several of these extensions are fully developed in our Supplemental Appendix, where we also discuss the connections between our framework and the framework developed by Lee (2008) where heterogeneity in the RD estimand arises from units having different unobservable types.

2 Motivating Examples

We now introduce two examples that we will use to motivate and illustrate concepts. Both examples use an RD design to study whether a party improves its future electoral outcomes by gaining access to office (i.e., by becoming the incumbent party). The treatment of interest is whether the party wins the election in year t, and the outcome of interest is the electoral victory or defeat of the party in the following election (for the same office), which we refer to as election at $t+1$.

In the first example, we analyze U.S. Senate elections between 1914 and 2010, pooling all election years and focusing on the effect of the Democratic party winning a Senate seat on the party's probability of victory in the following election for that same seat. In the second example, we analyze Brazilian mayoral elections for the PSDB between 1996 and 2012. As in the U.S. Senate example, we pool all election years and focus on the effect of the party's winning office at t on the party's probability of victory in the following election at $t+1$, which occurs four years later.

Figure 1 presents RD plots depicting the RD effect of the party barely winning an election on the probability of victory in the following election for both examples. These plots were constructed following the method in Calonico, Cattaneo, and Titiunik (2015a). Specifically, we plot the probability that the party wins election $t+1$ (y-axis) against the party's margin of the victory in the previous election (x-axis). The dots are binned means of district-level binary victory variables, and the solid blue line is a 4th order polynomial fit, estimated separately to the right and left of the cutoff (which is located at zero). In the plot, all observations to the right of the cutoff correspond to districts/municipalities where the party won election t, and all observations to the left correspond to locations where the party lost election t. Figure 1(a) shows that, in Brazilian mayoral elections, the PSDB's bare victory at t does not translate into a higher probability of victory at $t+1$. In contrast, as shown in Figure 1(b), a Democratic Party's victory in the Senate election at t considerably increases

the party's probability of winning the following election at the cutoff for the same Senate seat.
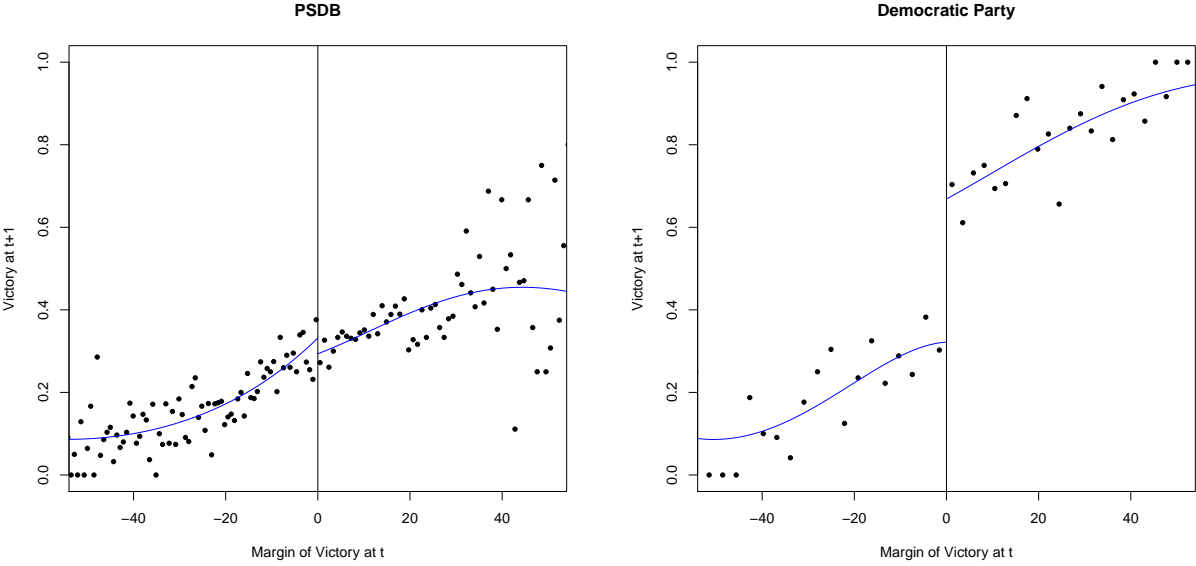
(a) Brazilian Mayoral Elections, 1996-2012   (b) U.S. Senate Elections, 1914-2010

Figure 1: RD Effect of Party Winning on Party's Future Victory: Brazil and the United States

For the statistical analysis of these two RD designs in Figure 1, we followed standard practice and used margin of victory as the score, thus normalizing the cutoff to zero for all elections. This score normalization is a practical strategy that allows researchers to analyze all elections simultaneously regardless of the number of parties contesting each electoral district. However, this approach pools together elections that are potentially highly heterogeneous. In fact, as we will demonstrate throughout this paper, the two RD designs in Figure 1 are quite different, with the U.S. Senate example containing much less heterogeneity than the Brazil example. In the next section, we show that the source of this heterogeneity is the number of effective parties.

2.1 Effective Number of Parties as RD Design Heterogeneity

In the examples introduced above, if there were exactly two parties contesting the election in each state or municipality, the running variable or score that determines treatment would be

the vote share obtained by the party at t, as this vote share alone would determine whether the party wins or loses election t. However, this is rarely the case in applications. In Table 2, we display the percentage of elections in our examples that are contested by more than two parties or candidates: roughly 68% of U.S. Senate elections and 50% of Brazilian mayoral elections are contested by three or more candidates in the periods for which we have data.[2] That this occurs in Brazil, a multi-party system where eight different parties win at least 5% of municipal races in the 1996-2012 period is not surprising, but it also occurs in the U.S., a two-party system where the Democratic and Republican parties dominate elections.[3]

Table 2: Number of Parties Contesting Elections

| | U.S. Senate | | Brazilian Municipalities | |
	Frequency	Percent	Frequency	Percent
One	40	2.88	748	2.72
Two	408	29.35	13,722	49.98
Three or more	942	67.77	12,985	47.30
Total	1,390	100.00	27,455	100.00

Note: *U.S. Senate* columns include state-level elections between 1914 and 2010, pooling across years, and *Brazilian Municipalities* columns include municipality-level mayoral elections between 1996 and 2012.

Based on the information in Table 2, mayoral elections in Brazil and Senate elections in the U.S. do not seem very different. However, as we show below, while the two examples differ little in terms of the number of parties, the number of effective parties is quite different. In a race with three or more parties, in order to know whether a party's vote share led the party to win the election, and by how much, we need to know the vote share obtained by the party's strongest opponent—the runner-up when the party wins and the winner when the party loses. Thus, the score or running variable that determines whether a party wins is the party's margin of victory—the vote share obtained by the party minus the vote share obtained by its strongest opponent. When this score is above zero, the party wins, otherwise

---

[2]We use the terms "parties" and "candidates" interchangeably throughout, but we note that in U.S. Senate elections some third candidates are unaffiliated with a political party.

[3]Combined, the Democratic and Republican parties win 98.6% of all Senate races in our sample.

it loses; and the closer this score is to zero, the more competitive the election is. In the above example, if the Democratic candidate obtains 33.4% of the vote against two candidates that obtain 33.3% and 33.33%, its margin of victory is $33.4 - 33.3 = 0.1\%$ and it barely wins the election. In contrast, when the other two parties obtain 60% and 6.6%, its margin of victory is 33.4-60=-26.6 and it loses the election by a large margin.

Figure 2 summarizes this information for our two examples. Figure 2(a) shows the histogram of the vote share obtained by the PSDB's strongest opponent at election t only for races where the PSDB won or lost by three percentage points—that is, for races where the absolute value of the PSDB's margin of victory at t is three percentage points or less. Figure 2(b) shows the analogous figure for the Democratic party in the Senate elections example.
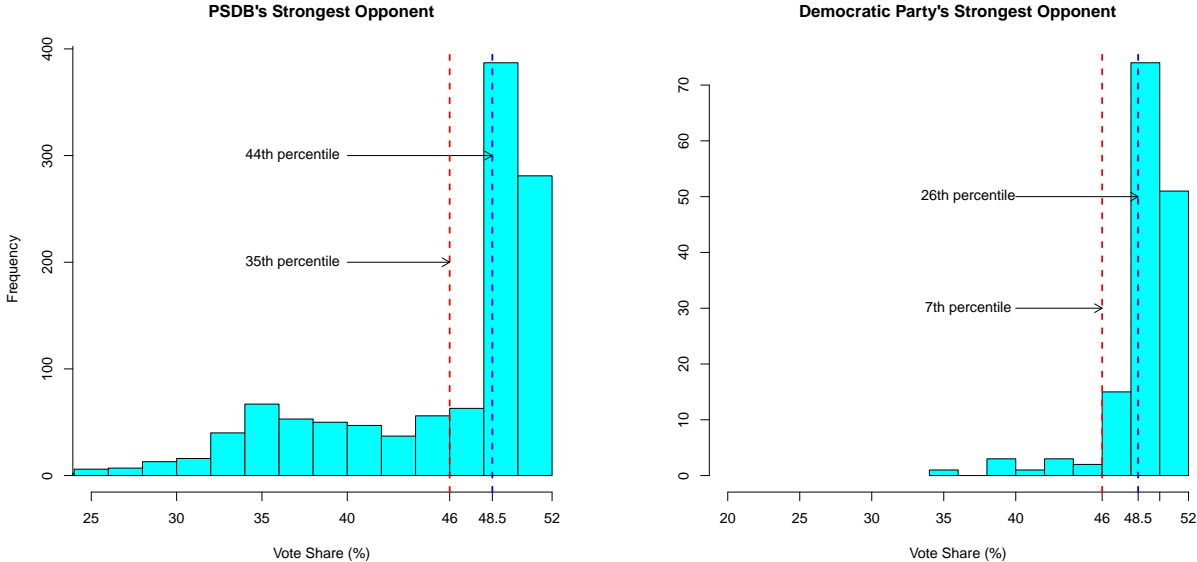
(a) Brazilian Mayoral Elections, 1996-2012   (b) U.S. Senate Elections, 1914-2010

Figure 2: Histogram of Vote Share of Strongest Opponent in Elections Where the PSDB and the Democratic Party Won or Lost by Less than 3 Percentage Points

Figure 2 reveals that the degree of heterogeneity differs greatly between the two examples. In a perfect two-party system, the vote share of the party's strongest opponent in races decided by 3 percentage points or less would range from 51.5% to 48.5%. That is, 48.5% is the minimum vote percentage that a party could get in a two-party race if it lost to the

Democratic Party by a margin no larger than 3 percentage points—and, similarly, 51.5% is the maximum possible value. As illustrated in Figure 2(b), in Senate elections where the Democratic party wins or loses by less than 3 percentage points, only 26% of the observations are below 48.5%, the minimum value that would occur in a perfect two-party system with a 50% cutoff for victory. Moreover, in 93% of the elections in the figure the Democratic party's strongest opponent gets 46% or more of the vote. In other words, despite most Senate elections having a third candidate (as shown in Table 2), in the overwhelming majority of these races the vote share obtained by such candidates is negligible, and there is little heterogeneity in the location of close races along the values of the strongest opponent's vote share.

In contrast, Figure 2(a) shows that the PSDB exhibits much higher heterogeneity, with strongest opponent vote shares that fall below the two-party system minimum of 48.5% for 44% of the observations. Moreover, more than a third of the elections (35%) have strongest opponent vote shares below 46%. In other words, a non-negligible proportion of the elections where the PSDB wins or loses by 3 points are elections in which third parties obtain a significant proportion of the vote. The histogram shows that, below 46%, the density peaks at around 35%, showing that in the roughly third of races below 46% the third party obtains a minimum of about 20% of the vote.

The differences illustrated in Figure 2 suggest that we ought to interpret the RD results in Figure 1 differently. In the case of U.S. Senate elections, the average effect in Figure 1(b) can be interpreted as roughly the average effect of the Democratic barely reaching the 50% cutoff and thus winning a two-way race. Although the existence of third parties means that the real cutoff is not exactly 50%, in practice most close races are decided very close to this cutoff, so that the average RD effect can be roughly interpreted as the effect of winning at 50%. However, the average effect in Brazil includes a significant proportion of elections where the cutoff is very far from 50%. As a consequence, this overall effect cannot be interpreted simply as the effect of barely winning at the 50% cutoff. Rather, it is the average effect of

barely winning at different cutoffs that range roughly from 20% to 50% of the vote.

For example, consider two mayoral elections in Brazil where the PSDB wins with a 1-percentage-point margin of victory. This could be an election in a municipality where the PSDB obtains 51% of the vote against a single challenger that obtains 49%, or an election in a municipality where the PSDB obtains 36% of the vote against two challengers that obtain 35% and 29% of the vote. The former type of municipality is one where two parties dominate the election and the PSDB's victory translates into the support of the majority of the electorate, while the latter type of municipality is one in which despite the PSDB's victory, 64% of the electorate does not support the party. Governing may be different in these two political environments, and thus we might expect the effect of gaining access to office on the party's future electoral success to differ between the two. This is the heterogeneity that gets "hidden" or averaged in the normalizing-and-pooling strategy.

Importantly, the heterogeneity in the Brazil example is not unique or unusual. Many political systems around the world have third candidates that obtain a sizable proportion of the vote. Figure 3 shows the distribution of the vote share obtained by a reference party's strongest opponent in six different countries across different time periods and types of elections, using the data compiled by Eggers et al. (2015). These histograms show only the subset of races decided by less than 3 percentage points for legislative elections in Canada, the United Kingdom, Germany, India, New Zealand and mayoral elections in Mexico—the reference party is indicated in each case.[4] In all the elections illustrated in Figure 3, there is a non-negligible proportion of cases where the vote share of the party's strongest opponent falls below the range that would be observed in a perfect two-party system with 50% cutoff. This means that, in all these cases, a pooled RD design that normalizes all cutoffs to zero would potentially contain substantial heterogeneity. For example, Figure 3(a) shows that in the elections for the Canadian House of Commons where the Liberal Party of Canada won or lost by less than three percentage points, the party's strongest opponent obtained less than

---

[4]See the supplemental information to Eggers et al. (2015) for details about the data sources.

46% of the vote in almost half of the cases, with most of these observations concentrated between 30% and 46%.

In the sections below, we formally describe the heterogeneity within the RD treatment effect parameter that arises from normalizing and pooling election vote shares. We also discuss the interpretation of this parameter, and explore how to recover different quantities of interest. Later we introduce some additional assumptions that researchers may wish to invoke in order to use this heterogeneity to learn about treatment effects of a more global nature, and also discuss recommendations for practice and possible extensions.

## 3 Identification for an RD Design with Multiple Cut-offs

We now present formal results for the sharp RD design, assuming that the cutoff has finite support–i.e., that the cutoff can only take a finite number of different values. We adopt these simplifications to ease the exposition, but our results extend to fuzzy RD designs, Regression Kink (RK) designs and cutoffs with continuous support among other possibilities—see Section 7 below and the Supplemental Appendix. Our assumptions and identification results reduce to those in Hahn, Todd, and van der Klaauw (2001), Lee (2008) and Card et al. (2014) for the special case of single-cutoff RD designs.

We begin by defining the pooled estimand in a multi-cutoff RD design. We adopt the standard RD framework with one additional modification to account for multiple cutoffs. We let X_i denote the running variable or score for unit i, which we assume continuous with a continuous density $f_X(x)$. We introduce the random variable C_i to denote the cutoff that unit i faces, which we assume has support $\mathcal{C} = \{c_1, c_2, ..., c_J\}$ with $\mathbb{P}[C_i = c] = p_c \in [0, 1]$ for $c \in \mathcal{C}$. We let $f_{X|C}(x|c)$ denote a conditional density of $X_i|C_i = c$ [5]. Note that in a standard RD, C_i would be a fixed value, but in our framework it is a random variable. As
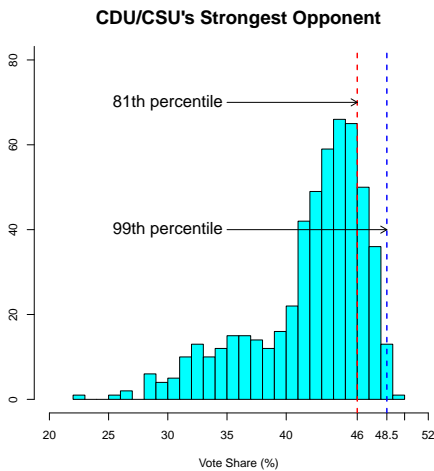
---

[5]Throughout the paper, we assume that all densities and conditional densities are positive and that all (conditional) expectations exist.
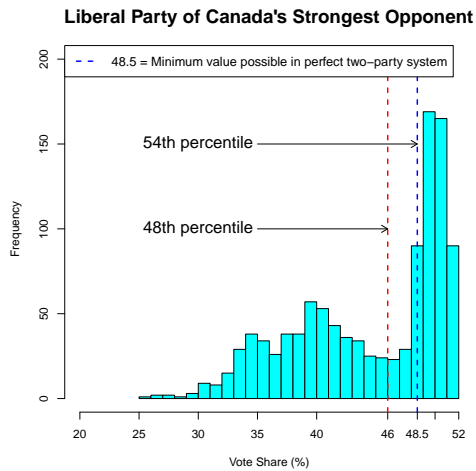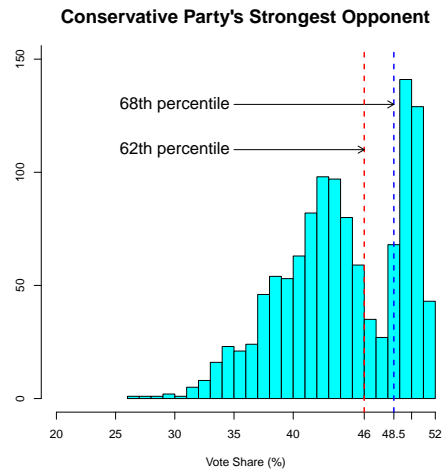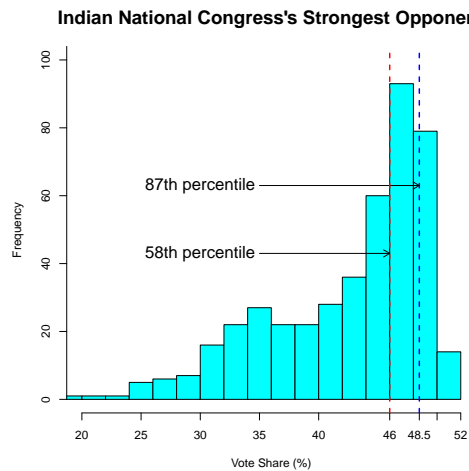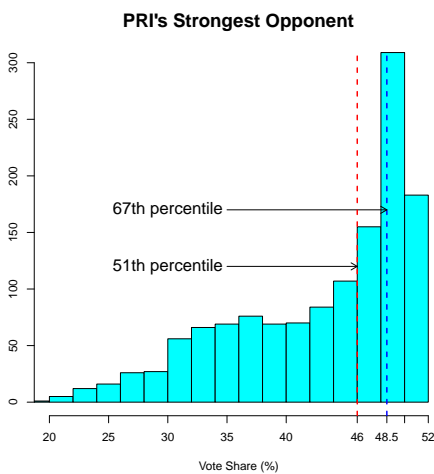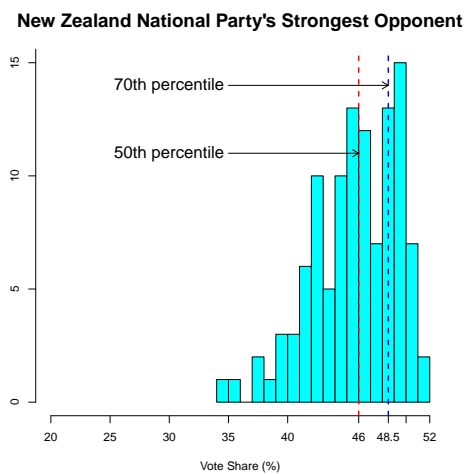
Figure 3: Strongest Opponent's Vote Share In Elections Decided by Less than 3 Percentage Points

a result, it is possible for different units to face different cutoff values. In the examples, X_i is the vote share obtained by the party of interest (PSDB or Democratic Party), C_i is the vote share of the party's strongest opponent, and the units of analysis indexed by i are U.S. states or Brazilian municipalities. We let $D_i \in \{0, 1\}$ be the treatment indicator. In our examples, $D_i = 1$ when the party wins the t election in state/municipality i, and $D_i = 0$ if it loses. Like in the usual sharp RD design, assignment to treatment depends on both the running variable X_i and the cutoff C_i. The unit receives treatment if the value of X_i exceeds the value of the cutoff C_i and receives the control condition otherwise, leading to $D_i = D_i(X_i, C_i) = \mathbb{1}(X_i \geq C_i)$, where $\mathbb{1}(\cdot)$ is the indicator function.

As discussed above, a common practice in the context of multiple cutoffs is to define a normalized score $\tilde{X}_i := X_i - C_i$, pool all the observations as if there was only one cutoff at $\tilde{X}_i = 0$, and use standard RD techniques. In our examples, $\tilde{X}_i$ is the party's margin of victory at election t—i.e., the party's vote share (X_i) minus the vote share of its strongest opponent (C_i)—and the party wins the election when this margin is above zero. That is, we can write $D_i = \mathbb{1}(\tilde{X}_i \geq 0)$. It follows that the limit of D_i as X_i approaches $C_i = c$ from the left (i.e., from the region where $X_i \leq C_i$) is equal to zero, and it is equal to one when X_i approaches $C_i = c$ from the right. We formalize this in the assumption below.

Assumption 1 (Sharp RD)

$\lim_{\varepsilon \to 0^+} \mathbb{E}[D_i \mid X_i = c + \varepsilon, C_i = c] = 1$ *and* $\lim_{\varepsilon \to 0^+} \mathbb{E}[D_i \mid X_i = c - \varepsilon, C_i = c] = 0$, *for all $c \in \mathcal{C}$.*

To complete the RD model, we assume the observed outcome is $Y_i = Y_{1i}(C_i)D_i + Y_{0i}(C_i)(1 - D_i)$, where $Y_{1i}(c)$ and $Y_{0i}(c)$ are, respectively, the potential outcomes under treatment and control at each level $c \in \mathcal{C}$.[6] Finally, as it is common in the literature, we also assume that we observe a random sample (across i) from a well-defined population. In our examples, $Y_{1i}(c)$ is the party's victory or defeat that would be observed at election $t + 1$ if the party won the previous election at t, and $Y_{0i}(c)$ is the party's victory or defeat that would be observed at election $t+1$ if the party lost the previous election. Note that, for each

---

[6]We employ the usual notation $Y_{di}(C_i) = \sum_{c \in \mathcal{C}} \mathbb{1}(C_i = c)Y_{di}(c)$ for $d = 0, 1$.

state or municipality, we only get to observe $Y_{0i}(c)$ or $Y_{1i}(c)$, but not both, since the party cannot simultaneously lose and win election t. Instead, we observe Y_i a (binary) variable equal to one if the party wins election $t+1$.

Our notation allows the cutoff for winning an election to affect the potential outcomes directly. More generally, the potential outcomes may be related to several variables: the running variable X_i, the cutoff C_i, and other unit-specific (unobserved) characteristics. The latter variables are usually referred to as the unit's "type"—see Lee (2008) and Section **??** in the Supplemental Appendix for further discussion. Thus, in our examples, we not only let the party's potential electoral success in election $t+1$ be related to its vote share and the vote share of its strongest opponent at t, but also to other (potentially unobservable) characteristics of the state or municipality where the elections occur, such as its geographic location, the underlying partisan preferences of the electorate and its demographic makeup.

The RD pooled estimand, τ^p, is defined as follows:

$$\tau^p = \lim_{\varepsilon \to 0^+} \mathbb{E}[Y_i \mid \tilde{X}_i = \varepsilon] - \lim_{\varepsilon \to 0^+} \mathbb{E}[Y_i \mid \tilde{X}_i = -\varepsilon] \tag{1}$$

Equation 1 is the general form of the causal estimand in a multi-cutoff RD where the score has been normalized and all observations have been pooled. Estimation of this pooled estimand is straightforward and, as discussed above, is done routinely by applied researchers. After normalization of the running variable, estimation just proceeds as in a standard RD design with a single cutoff—for example, using local non-parametric regression methods, as is now standard practice.[7] Although estimation of τ^p is straightforward, the interpretation of this estimand differs in a number of important ways from the interpretation of the causal estimand in a standard single-cutoff RD design. We turn to this issue in the following section.

---

[7]See Imbens and Lemieux (2008) and Lee and Lemieux (2010) for reviews on estimation methods, and Calonico, Cattaneo, and Titiunik (2014b) for recent results on local polynomial regression methods.

3.1 General case: heterogeneity within and across cutoffs

We start by considering the most general form of treatment effect heterogeneity where the treatment effect varies both across and within cutoffs. In this general case, individuals may respond to treatment differently if they face different cutoffs but also if they face the same one. Formally, this individual-level treatment effect is $\tau_i(c) = Y_{1i}(c) - Y_{0i}(c)$. In our empirical example, this implies that the incumbency effect may vary in districts with different vote shares of the party's strongest opponent, but it may also vary across districts with the same value of this variable. In order to derive the expression for τ^p we invoke the following two assumptions.

Assumption 2 (Continuity of regression functions)
$\mathbb{E}[Y_{0i}(c) \mid X_i = x, C_i = c]$ and $\mathbb{E}[Y_{1i}(c) \mid X_i = x, C_i = c]$ are continuous in x at $x = c$ for all $c \in \mathcal{C}$.

Assumption 3 (Continuity of the density)
$f_{X|C}(x|c)$ is continuous in x at $x = c$ for all $c \in \mathcal{C}$.

Assumption 2 says that expected outcome under control is a continuous function of the running variable for all values of the score, implying that units barely below the cutoff are valid counterfactuals of units barely above it. This is the fundamental identifying assumption in all RD designs. Assumption 3 rules out discontinuous changes in the density of the running variable. Lemma 1 below characterizes the pooled estimand under complete heterogeneity.

Lemma 1 (Heterogenous treatement effects)
If Assumptions 1, 2 and 3 hold, the pooled RD causal estimand is

$$\tau^p = \sum_{c \in \mathcal{C}} \mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid X_i = c, C_i = c]\, \omega(c), \qquad \omega(c) = \frac{f_{X|C}(c|c)\mathbb{P}[C_i = c]}{\sum\limits_{c \in \mathcal{C}} f_{X|C}(c|c)\mathbb{P}[C_i = c]}.$$

The proof is in Section **??** of the Supplemental Appendix. Lemma 1 says that whenever heterogeneity within and across cutoffs is allowed, the pooled RD estimand recovers a double

16

average: the weighted average across cutoffs of the average treatment effects $\mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid X_i = c, C_i = c]$ across all units facing each particular cutoff value. Importantly, this derivation shows that the pooled estimand is not the equal to the overall average of the (average) treatment effects at every cutoff value. In Section **??** of the Supplemental Appendix we discuss this point formally and show the differences between the average of the cutoff-specific effects and τ^p. In that section, we also discuss how the pooled estimand can be written as an average across individuals of different types as in Lee (2008).

Two things should be noted in order interpret the estimand in Lemma 1. First, the weight $\omega(c)$ determines the effects that are included in the pooled parameter τ^p, and how much each effect contributes to this parameter. The term $\mathbb{P}[C_i = c]$ is simply the probability of observing the particular realization of each cutoff, and implies that $\omega(c)$ will be higher for those values of c that are more likely to occur. The term $f_{X|C}(c|c)$ increases the weight of effects that occur at values of c where the density of the running variable is high.

Second, each of the conditional effects being averaged, $\mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid X_i = c, C_i = c]$, is the average effect of treatment given that both the running variable X and the cutoff C are equal to a particular value c. In the standard single-cutoff RD design, the effect recovered is the average effect of treatment at the point $X = c$, an effect that is typically characterized as local because it reflects the average effect of a treatment at a particular value of the running variable and is not necessarily generalizable to other values of X. Therefore, the conditional effects in the pooled RD case intensify the local nature of the effect, because they represent the average effect of treatment when both the running variable and the cutoff take the same particular value.

For example, in a perfect two-party system, the RD effect of a party winning election t on the party's future victory at $t+1$ recovers a single effect—the effect of this party winning with a vote share just above 50%, not the effect of winning in general. In contrast, in the pooled RD design, this is just one of the effects that are included in τ^p. The pooled RD estimand τ^p includes other effects such as the average of the party winning with 40% of the

vote against a strongest opponent that gets just below 40%, the average effect of the party winning with 30% of the vote against a strongest opponent that gets just below 30%, etc. This heterogeneity in τ^p makes it a richer estimand, but it also makes each of its component effects more local or specific, because each reflects only one of the multiple ways in which "barely winning" can occur.

Moreover, τ^p is subtle in other ways. In the pooled multi-cutoff RD design, just like in the standard single-cutoff RD design, units whose score X is close to a cutoff may be systematically different from the units whose score is far from it. In the pooled RD design, however, units can also differ systematically in their probabilities of facing a particular value of the cutoff. For example, in the Brazilian mayoral context, municipalities where the PSDB gets 50% of the vote might be different in relevant ways from municipalities where the PSDB gets 35% of the vote. In addition, even within those municipalities where the PSDB gets 35% of the vote, municipalities where the strongest opponent also gets roughly 35% may be very different from those where the strongest opponent gets 10% or 15% and the election is uncompetitive. In terms of our example, this means that, at every value c, the effects that contribute to τ^p are the average effect of the party barely defeating an opponent that obtained a vote share equal to c. While this effect is uninformative about the effects at other values of c, it does imply that when there are many values of c the pooled RD estimand contains information about the causal effect of barely winning in a number of different contexts. This aspect of the pooled RD estimand, by which many different local effects are combined when many different values of C may occur, shows that multi-cutoff RD designs contain a richer set of information relative to single-cutoff settings.

This means that the causal estimand in a multi-cutoff RD design is something of a paradox. On the one hand, τ^p is a more local parameter in the sense that it is the effect of the treatment for those units for which X_i barely exceeds C_i in only one of the multiple ways in which X_i could barely exceed C_i. On the other hand, when C_i takes a wide range of values, the average effect of treatment is recovered for many different ways in which X_i can

barely exceed C_i, potentially leading to a more global interpretation of the RD effect. Later, we use our two empirical examples to illustrate how researchers can explore the richness in τ^p.

4 Multi-Cutoff RD Design Under Additional Assumptions

A usual concern with single-cutoff RD designs is that they only offer estimates of the treatment effect at the cutoff and are thus uninformative about the magnitude of the treatment effect at other values of the running variable. In our examples, the multi-cutoff RD gives us the effect of barely defeating the opponent party with a range of different values—in Brazil mayoral elections this range is roughly 20% to 50%. Can we use this wider range of values to learn about a more global effect? We now consider assumptions under which the diverse information contained in the pooled estimand can be used to disaggregate the information in τ^p and learn about treatment effects of a more global nature.

4.1 Constant treatment effects

We first consider a simplification of the general case, where the treatment effect is different across cutoffs but constant for all individuals who face the same cutoff, i.e. $Y_{1i}(c) - Y_{0i}(c) = \tau(c)$ with $\tau(c)$ a fixed constant for all i. Note that $\tau(c)$ varies by unit only insofar as c varies by unit, but there is no i subindex in $\tau(c)$, indicating that two units facing the same given cutoff c will have the same treatment effect $\tau(c)$. In terms of our example, this assumption implies that the effect of the party winning an election on its future electoral success is the same in all municipalities/districts where its strongest opponent obtains the same proportion of the vote. This is undoubtedly a very strong assumption. We include it here to illustrate one possible way in which the treatment effects recovered by the multi-cutoff RD design can be given a more global interpretation, but we discuss weaker assumptions in the subsequent sections.

The proposition below shows that when there is no heterogeneity within cutoffs, the relationship between the pooled RD estimand and the cutoff-specific effects simplifies considerably.

Proposition 1 (Constant treatment effects)

If Assumptions 1 and 3 hold, and $\tau_i(c) = \tau(c)$ for all i and $\tau(c)$ fixed for each c, then the pooled RD estimand is

$$\tau^p = \sum_{c \in \mathcal{C}} \tau(c)\,\omega(c),$$

where the weights are the same as in Lemma 1.

Thus, when effects are constant within cutoffs, $\tau(c)$ captures the effect of treatment for all values individuals facing cutoff c. Naturally, Proposition 1 simplifies considerably when the treatment effect is the same for all individuals at all cutoffs (i.e., $Y_{1i}(c) - Y_{0i}(c) = \tau$ for all i and all c, and thus $\tau(c) = \tau$ for all c). In this case, the pooled estimand becomes $\tau^p = \sum_{c \in \mathcal{C}} \tau(c)\omega(c) = \tau \sum_{c \in \mathcal{C}} \omega(c) = \tau$, recovering the single (and therefore global) constant treatment effect. This global interpretation of the multi-cutoff RD estimand under constant treatment effects is analogous to the interpretation in a single-cutoff RD design, where the assumption of homogeneous treatment effects leads to the identification of the overall constant effect of treatment.

4.2 Ignorable running variable

The case introduced above is very restrictive, as it is natural to expect some heterogeneity in treatment affects among units facing the same value of the cutoff. We now consider the less restrictive case of unit-heterogeneity within cutoffs, but with an average treatment effect at every value of the cutoff that does not depend on the particular value taken by the score. We summarize this in the following assumption.

Assumption 4 (Score Ignorability)
$\mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid X_i, C_i = c] = \mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid C_i = c]$ *for all $c \in \mathcal{C}$.*

Under Assumption 4, the running variable is ignorable once we condition on the value of the cutoff—that is, once the value of the cutoff is fixed, we assume that the average effect of treatment is the same regardless of the value taken by the score. The proposition below shows the form of the pooled RD estimand in this case.

Proposition 2 (Score-ignorable treatment effects)

If Assumptions 1, 3 and 4 hold, then the pooled RD estimand is

$$\tau^p = \sum_{c \in \mathcal{C}} \mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid C_i = c]\, \omega(c),$$

where the weights are the same as in Lemma 1.

Thus, when the average effect of treatment does not vary with the running variable X, $\mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid C_i = c]$ captures the effect of treatment for all values of X, not necessarily those that are close to the cutoff c. For example, $\mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid C_i = c]$ may reflect the average effect of the Democratic party winning election t on its future electoral success for a given value of its strongest opponent's vote share, regardless of whether the party defeated its opponent barely or by a large margin.

In this sense, the effects in Proposition 2 are global in nature. Note, however, that the treatment effects are allowed to vary with the value of C, and therefore the expression for τ^p in Proposition 2, though not necessarily local, is only averaging over the set of values that C can take, and the values of C that will be given positive weight are only those values where the density of X given C at $X = C = c$, $f_{X|C}(c|c)$, is positive. As such, τ^p still retains a local aspect.

4.3   Ignorable cutoffs

We now consider the case where the running variable is not ignorable, but where the heterogeneity brought about by the multiple cutoffs to can be restricted in ways that allow extrapolation. It is useful to introduce the analogy between the RD design with multiple

cutoffs and an experiment that is performed in different sites or locations. In the latter case, internally valid treatment effect estimates from experiments in multiple sites are not necessarily informative about the effect that the treatment would have in a different site where the experiment has not been run. This means that the results from multi-site experiments may not allow researchers to extrapolate to the overall population, a concern that is not necessarily eliminated if the number of sites is large (Allcott, 2014). The problem arises because the sites that are selected to run an experimental trial may differ from the overall population of sites in ways that are correlated with the treatment effect. For example, sites where the treatment is expected to have large effects may be more likely to run experimental trials, leading to a positive "site selection bias" that would overestimate the effects that the treatment would have if it were implemented in the overall population. Alternatively, the population may differ across sites in a characteristic that is associated with treatment effectiveness (Hotz, Imbens, and Mortimer, 2005). Of course, generalizing the treatment effect from one particular site to other locations can be done under additional assumptions.

Like in a multi-site experiment, we have a series of internally valid estimates in the multi-cutoff RD that we might wish to think of as more general. In the multi-site experiments literature, the strongest and simplest assumption under which the generalization of effects is possible is independence of locations with respect to potential outcomes. This condition is guaranteed by design when the units in the population are randomly assigned to different sites. In our context, we can make the analogous assumption that, conditional on the value of the running variable, the cutoff faced by a unit is unrelated to the potential outcomes. Formally, we can write this assumption as follows.

Assumption 5 (Cutoffs Ignorability)

(a) $\mathbb{E}[Y_{1i}(c) \mid X_i, C_i = c] = \mathbb{E}[Y_{1i}(c) \mid X_i]$ *and* $\mathbb{E}[Y_{0i}(c) \mid X_i, C_i = c] = \mathbb{E}[Y_{0i}(c) \mid X_i]$ *for all* $c \in \mathcal{C}$.

(b) $Y_{1i}(c) = Y_{1i}$ *and* $Y_{0i}(c) = Y_{0i}$.

Assumption 5(a) says that, conditional on the running variable, the potential outcomes

are mean independent of the cutoff variable C. In addition, we need to ensure that the value of the cutoff does not affect the potential outcomes.[8] Assumption 5(b) above formalizes this idea as an exclusion restriction, requiring that the cutoff level does not affect the potential outcomes directly.

Now define a point x_0 such that $c_0 \leq x_0 < c_1$. Assumption 5 leads to the following result for observed random variables:

$$\mathbb{E}[Y_i \mid X_i = x_0, C_i = c_0] - \mathbb{E}[Y_i \mid X_i = x_0, C_i = c_1]$$
$$= \mathbb{E}[Y_{1i}(c_0) \mid X_i = x_0, C_i = c_0] - \mathbb{E}[Y_{0i}(c_1) \mid X_i = x_0, C_i = c_1]$$
$$= \mathbb{E}[Y_{1i}(c_1) - Y_{0i}(c_0) \mid X_i = x_0] = \mathbb{E}[Y_{1i} - Y_{0i} \mid X_i = x_0],$$

which is the average treatment effect at x_0. Since x_0 is not a value in the support of the random cutoff variable, this shows that under these assumptions we can estimate the average treatment effect away from the cutoff, and thus obtain a more global effect.

However, the following lemma shows that, as before, the ability to recover a global effect from the pooled multi-cutoff RD design even under Assumption 5 is limited by the fact that τ^p weighs these average effects by the probability of observing a realization of the cutoff variable C_i at the particular value c.

Proposition 3 (Cutoff-ignorable treatment effects) *If Assumptions 1, 2, 3 and 5 hold, the pooled RD estimand becomes*

$$\tau^p = \sum_{c \in \mathcal{C}} \mathbb{E}[Y_{1i} - Y_{i0} \mid X_i = c]\,\omega(c),$$

where the weights are the same as in Lemma 1.

Thus, under these assumptions, τ^p averages the average treatment effects $\mathbb{E}[Y_{1i} - Y_{i0} \mid X_i = c]$, each of which is the average effect of receiving treatment conditional on the running

---

[8]Note that this is equivalent to the "no macro-level variables" assumption in Hotz, Imbens, and Mortimer (2005).

variable X_i being at the value c, regardless of the value taken by C_i. In our example, this represents the average effect of a party winning the t election given that the party's vote share is c and regardless of the vote share obtained by its strongest opponent, i.e. regardless of whether it won barely or by a large margin. However, these averages are still evaluated only at values of c that are in the support of the random cutoff variable C_i. So, although they are more global effects, they can only be recovered at feasible values of C_i. Moreover, the weights entering τ^p still depend on $\mathbb{P}[C_i = c]$ through the weights $\omega(c)$.

Note that if, in addition to the assumptions imposed in Proposition 3, we imposed the assumption that the conditional density of the score X_i given C_i is constant in the support of C_i,[9] the pooled RD parameter τ^p simplifies to:

$$\tau^p = \sum_{c \in \mathcal{C}} \mathbb{E}[Y_{1i} - Y_{0i} \mid X_i = c]\mathbb{P}[C_i = c]$$

and now, if the support of C_i is equal to the support of X_i (which will only be possible if both are discrete or both are continuous), we can recover the average of the average treatment effect at all values of X_i—see Section ?? in the Supplemental Appendix for further discussion. All these assumptions combined would thus make τ^p a truly global estimand, without the assumption of constant effects as we did before.

Assumption 5 also has another application. Under the conditions imposed in that assumption, $\mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid X_i = c, C_i = c] = \mathbb{E}[Y_{1i} - Y_{0i} \mid X_i = c]$. This shows that when these assumptions hold, estimating the RD effects separately for each value c will provide a treatment-effect curve that will summarize the effects of the treatment at different values of the running variable (independently of the value taken by the cutoff). In other words, under these assumptions, we can estimate multiple RD treatment effects for different values of the running variable.

Of course, Assumption 5 is generally strong and may be too restrictive in some empirical applications. In Section 7 below, we discuss how our framework can be used to relax this

---

[9]Formally, this assumption can be written as $f_{X|C}(c|c) = k$ for all $c \in \mathcal{C}$ with k a constant.

assumption and allow for endogenous sorting of different unit types across different values of the cutoff.

5   Empirical Examples

We now illustrate how the formal results derived in the paper can have empirical implications using our two examples. In particular, we show how in the Brazil example we can estimate different effects for different values of the running variable. First, we use an exploratory statistical analysis to show that extrapolating the pooled treatment effect based on the information contained in the multiple cutoffs is possible in the Brazil example but not in the Senate example, since in the latter there is simply not enough observations for lower values of the cutoff variable.

As we highlighted earlier, the two examples differ sharply in the density of observations at different cutoff values. There are very few U.S. Senate elections where a third party obtains anything more than a very small fraction of the vote. In the Brazilian mayoral elections, however, about a third of races occurs in municipalities where the two top-getters combined obtain less than 70% of the vote. Table 3 presents the frequency of races in our sample by different levels of strongest opponent's vote shares at t for the Democratic Party and the PSDB. Since this variable is continuous, we divide its support in four exclusive intervals: $[0, 35)$, $[35, 40)$, $[40, 45)$, and $[45, 50)$. Within each of these intervals of strongest opponent's vote share at t, Table 3 reports the number of elections that each party won and lost at t. Note that in a perfect two-party system, knowing the value of a party's strongest opponent's vote share is *equivalent* to knowing whether the party won or lost the election, but this equivalency is broken in a multi-party RD design.

For example, the columns corresponding to the PSDB show that, of the 1346 races in our sample where the PSBD's strongest opponent obtained between 35% and 40% of the t vote, the PSDB won roughly 85% and lost the rest. The proportion of victories decreases for higher values of this variable, with the PSDB winning no more than 64% of the races in all

cells where vote share of its strongest opponent is 35% or higher. A very different situation occurs in U.S. Senate elections where, for example, the Democratic Party won all 264 races where the strongest opponent obtained less than 35% of the vote, as would occur in a perfect two-party system. Similarly, of the 118 races in our sample where the Democratic Party's strongest opponent obtained between 35% and 40% of the t vote, the party won 111 and lost only 7. It is only in the $[45, 50)$ range where the party loses 20% of races—a non-neligible but still small proportion.

Table 3: Frequency of Observations for Different Levels of Strongest Opponent's Vote Shares at t

	Democratic Party U.S. Senate Elections			PSDB Brazil Mayoral Elections		
Opponent Vote (%)	Total	Victories (%)	Defeats (%)	Total	Victories (%)	Defeats (%)
$[0, 35)$	264	100.0	0.0	1346	84.9	15.1
$[35, 40)$	118	94.1	5.9	986	63.9	36.1
$[40, 45)$	161	96.3	3.7	1251	62.3	37.7
$[45, 50)$	221	77.8	22.2	1490	61.5	38.5

Note: Columns corresponding to Democratic Party report number of U.S. Senate elections in 1910–2010. Columns corresponding to PSDB report number of mayoral elections in Brazil in 1996-2012.

We explore the heterogeneity in the Brazil example by separately estimating the RD effects at different levels of strongest opponent's vote share. We choose a grid of values in the support of the vote share of the PSDB's strongest opponent and, for each value in this grid, we separately estimate the RD effect of the PSDB's winning at t on the PSDB's future success using only the 600 treated observations closest to the grid value and the 600 control observations closest to the grid value.

Figure 4 summarizes the results, showing the treatment effects at six different, equidistant values of strongest opponent vote shares between 34% and 50%. The dots are the treatment effect and bars are the 95% confidence intervals. We estimated these effects using local-linear regression with a mean-squared-error (MSE) optimal bandwidth and confidence intervals based on procedures developed by Calonico, Cattaneo, and Titiunik (2014b). Their method
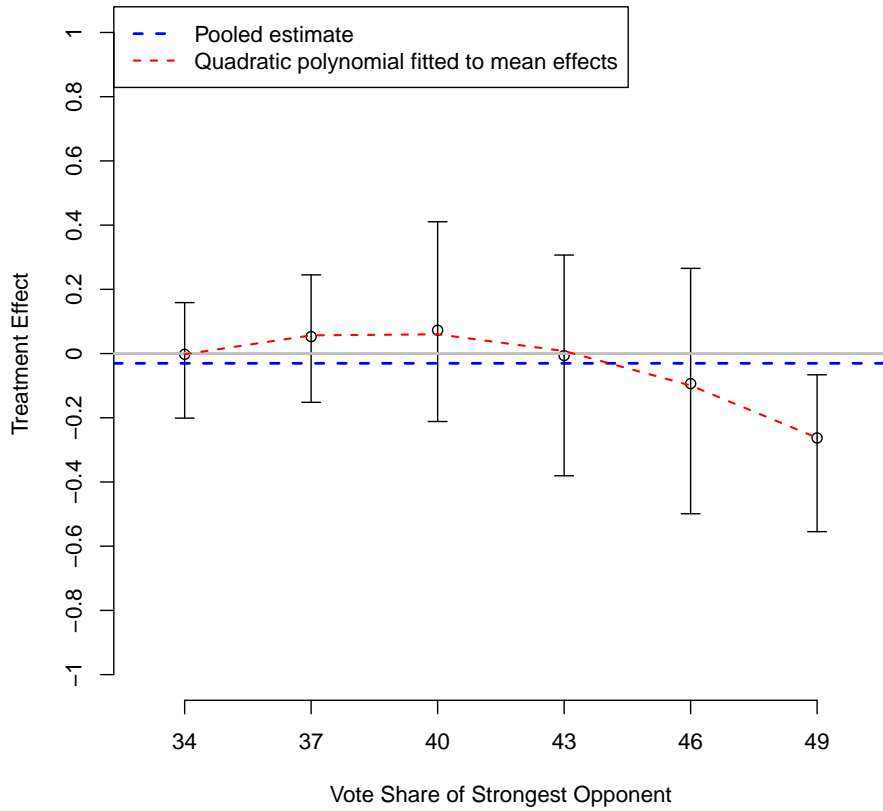
Figure 4: RD Effects of PSDB's Victory on Future Vote Share at Different Levels of Strongest Opponent's Vote Share

ensures that the distributional approximation used is valid for this bandwidth. Specifically we used the `rdrobust` package (Calonico, Cattaneo, and Titiunik, 2014a, 2015b). Note that, for every value of the PSDB's strongest opponent vote share that is displayed in the figure, we are estimating the effect of the PSDB's *barely* defeating its strongest opponent, so that all the effects in this figure are local RD effects.

We begin by estimating τ^p, which is the pooled RD estimand that uses margin of victory as the score and normalizes all cutoffs to zero, by local linear regression and MSE-optimal bandwidth. The pooled RD point estimate is -0.03, an effect that cannot be statistically distinguished from zero (p-value = 0.44). The robust 95% confidence interval is $[-0.11, 0.05]$. The blue dotted line indicates this point estimate in Figure 4. As we noted in Section 3, τ^p has a different interpretation than the estimand in a standard single-cutoff RD design. Recall that the interpretation of τ^p is both more local and more global than would be the case in

an RD design with a single cutoff, because this estimand is the effect of barely winning in a prior election averaging across the multiple ways in which the party can barely win the election.

As a result, the pooled effect may contain or "hide" significant heterogeneity. Figure 4 reveals that this is the case in the Brazilian mayoralty data. For values of strongest opponent vote shares that fall near 46% or below, the effect of barely winning is relatively small and cannot be distinguished from zero. This estimate is also consistent with the results from the pooled analysis. That is, for most of the range of the running variable, the disaggregated estimates are not statistically distinguishable from the pooled estimate. However, for those elections where the PSDB's strongest opponent obtains a vote share near 49%, the effect is negative, large in absolute value, and significantly different from zero.

The heterogeneity illustrated in Figure 4 must be interpreted with care for two reasons. The first reason is practical. As shown in Table 3, the number of observations at every level of strongest opponent's vote share is moderate, which may lead to noisy estimates of the conditional expectations. The width of the confidence intervals in Figure 4 vary significantly across the range of the running variable. Without sufficient data density, it may be difficult to separate heterogeneity from noise.

Second, following our discussion in Section 4, the interpretation of the treatment-effect curve in Figure 4 depends crucially on the assumptions surrounding the factors that affect the strongest opponent's vote shares. If we were willing to assume that, at every level of vote share obtained by the PSDB at t, the vote share obtained by its strongest opponent is mean independent of the PSDB's potential victory at $t + 1$ (Assumption 5a) and the strongest opponent's vote shares affect the potential future performance of the PSDB only through the PSDB's winning or losing the election but not directly (Assumption 5b), then each of these effects would be the effect of the PSDB winning election t with a vote share in each interval, regardless of whether it won barely or by a large margin.

If however, we believe that the more plausible scenario is one in which elections that

differ in the strongest opponent's vote share also differ systematically in observed and unobserved factors that affect the PSDB future vote shares (e.g., municipalities with strong third parties may be systematically different from municipalities where only two parties contest the election), then the interpretation of Figure 4 changes considerably. Under this scenario, the potential differences between the effects also reflect the different electoral environments that occur at different levels of strongest opponent's vote shares, and cannot be simply interpreted as the effect of treatment at those levels of the PSDB's t vote share (the running variable).

6   Recommendations for Practice

We now outline a few basic recommendations for applied researchers. These recommendations are specifically tailored to multi-cutoff RD designs that arise from multiple candidates or parties, but the general ideas apply more broadly to other settings in social science involving multi-cutoff RD designs.

As a natural starting point, we suggest some visual diagnostics. When there are more than 2 parties, the analyst should create a histogram of the strongest opponent's vote shares, as we did in Figure 2. If most of the mass in the distribution is near the same cutoff value, then the analyst can treat the design as equivalent to a single-cutoff RD design, since the heterogeneity is minimal. If the density of the vote share of the strongest opponent is more dispersed as in Figure 2(a), then the pooled estimand is potentially heterogeneous. When heterogeneous effects are present, the analyst has several options.

First, one could simply pool the estimates and ignore (i.e., average) the heterogeneity or assume constant treatment effects. Second, one could acknowledge the presence of heterogeneity, but leave it unexplored claiming that the main object of interest is the pooled estimand. Third, one could explore whether the pooled estimate is robust to excluding some of the observations. For example, in a case that looks like our Brazil example, one could split the sample into two subsets: races where the strongest opponent gets 45% or more of

29

the vote, and the rest. If most of the mass is in the first subset, an interesting question is whether the pooled estimate is actually close to the estimate that uses only this subset. Since the pooled estimand is a weighted average, a low mass of observations below the 45% cut point would receive little weight but an aberrant treatment effect in this range could lead to an "uninformative" pooled effect.

Finally, one could develop substantive hypotheses about how the heterogeneity is expected to change from one cutoff to the next, and explore these hypotheses and heterogeneity fully, estimating several treatment effects along the cutoff variable. For example, one could formally investigate the presence of monotonic treatment effects along the running variable.

7 Extensions

The results developed above can be extended in many different directions. Some of these extensions are undertaken in our Supplemental Appendix, while others are suggested as topics for future research.

1. **Fuzzy RD Designs.** All the results presented above can be derived for the more general case of a multi-cutoff *fuzzy* RD design, where compliance with treatment is imperfect. In this case, some units below the cutoff may receive the treatment and some units above it may refuse it, leading to a jump in the probability of receiving treatment at the cutoff that is less than one. This extension is straightforward and is given in Section **??** of the Supplemental Appendix. Despite the necessary technical modifications, all the conceptual issues discussed above apply directly to the fuzzy RD case.

2. **Kink RD Designs.** Our results also apply to the quite recent literature on Regression Kink (RK) designs, in which a treatment or policy is assigned on the basis of a score via a formula exhibiting "kink" points at which the specific formula that relates the assignment variable to the treatment changes discontinuously. See Card et al. (2014) for more details. To our knowledge, this design has not yet been used in Political

Science. Nonetheless, in Section **??** of the Supplemental Appendix we show how our main identification results extend to this case, with appropriate modifications.

3. **RD Designs with Multiple Scores.** There are strong connections between our multi-cutoff framework and RD designs with multiple scores or running variables, which have received recent attention (see, e.g., Imbens and Zajonc, 2011; Keele and Titiunik, 2015; Papay, Willett, and Murnane, 2011; Wong, Steiner, and Cook, 2013). In Section **??** of the Supplemental Appendix, we discuss these connections and how our main results apply, with proper modifications, to this case. In particular, we show how an RD design with one score and multiple cutoffs can be recast as an RD design with two running variables.

4. **Population Parameters of Interest.** The main goal of this paper is to present a framework to analyze multi-cutoff RD designs, and to clarify the interpretation of the commonly targeted pooled RD estimand. Our results, nonetheless, also allow to investigate the construction of many other estimands by altering the weighting scheme given to the different local RD treatment effects for each value in the support of C_i. In other words, instead of settling for the weights $\omega(c)$ in the lemma and propositions above, researchers could consider employing appropriate weighting schemes to modify the final targeted RD estimand. Due to space limitations, we did not pursue this line of research here but we plan to investigate this issue further in upcoming work.

5. **Endogenous Cutoffs.** This is perhaps the most important extension of our work. Some of our results impose simple and easy-to-interpret assumptions, effectively treating the variable C_i as exogenous. However, as we discussed above, the average treatment effects $\mathbb{E}[Y_{1i}(c) - Y_{0i}(c) \mid X_i = c, C_i = c]$ may differ for different values of c because the units close to one value of c may be systematically different from the units who are close to a different value, and this difference may be correlated, for example, with the effectiveness of the treatment. This implies that differences in the average

treatment effect at different values of the cutoff may arise due to "selection" of different unit types into cutoffs. Under Assumption 5 above, the potential outcomes are mean independent of the cutoff variable C_i, and we can rule out the phenomenon of self-selection. This is undoubtedly a strong assumption. In ongoing research, we are exploring how to relax this assumption to allow for self-selection into cutoffs based on observable characteristics, as in the framework of Hotz, Imbens, and Mortimer (2005) for multi-site experiments.

8   Concluding Remarks

The standard RD design assumes that a treatment is assigned on the basis of whether a score exceeds a single cutoff. However, in many empirical RD applications the cutoff varies by units, and researchers normalize the running variable so that all units face the same cutoff value and a single estimate can be obtained by pooling all observations. This is a useful approach to summarize the average effect across cutoffs, but in some cases it is possible to disaggregate the information contained in the pooled effect and provide a richer description of the underlying heterogeneity in the treatment effects.

As we showed above, when there are multiple cutoffs, the pooled RD estimand is the weighted average of the average effect of treatment at every cutoff value, with higher weight given to a particular cutoff value c when there is a high number of units whose scores are close to c. Our formalization of the pooled estimand as a weighted average thus shows that the degree of heterogeneity captured by this estimand will vary on a case-by-case basis depending on the density of the data used in each application.

Our two examples illustrated this issue. In the case of U.S. Senate elections, a pooled multi-cutoff RD design arises because races with more than two candidates are very common. However, the pooled estimand that uses margin of victory as the normalized running variable hides little heterogeneity, because the third candidate typically obtains a very low share of the vote and as a result the number of elections where the Democratic party wins with a vote

percentage that is far from 50% is very low. In contrast, a substantial proportion of Brazil mayoral elections are decided far from the 50% cutoff because it is common for the two top parties combined to obtain less than 80% or 70% of the vote. In this type of scenario, the heterogeneity underlying the pooled estimand can be substantial. As we show in the case of Brazil for the effect of the PSDB winning on its future electoral victory, the pooled estimate is statistically indistinguishable from zero but using only observations where the vote share obtained by the PSDB's strongest opponent is near 49%, this effect becomes negative, large and statistically different from zero.

In showing that the weights in the pooled approach combine the effects at different cutoffs in a particular way, our framework also suggests that researchers may want to choose different weights relevant to their application. But, as we discussed above, the interpretation of this heterogeneity depends on whether the probability that a unit faces a particular cutoff is related to characteristics that correlate with the potential effects of the treatment.

References

Albouy, David. 2013. "Partisan representation in Congress and the geographic distribution of Federal funds." *Review of Economics and Statistics* 95 (1): 127–141.

Allcott, Hunt. 2014. "Site Selection Bias in Program Evaluation." Working Paper.

Boas, Taylor C, and F Daniel Hidalgo. 2011. "Controlling the airwaves: Incumbency advantage and community radio in Brazil." *American Journal of Political Science* 55 (4): 869–885.

Boas, Taylor C, F Daniel Hidalgo, and Neal P Richardson. 2014. "The Spoils of Victory: Campaign Donations and Government Contracts in Brazil." *The Journal of Politics* 76 (2): 415–429.

Brollo, Fernanda, Tommaso Nannicini, Roberto Perotti, and Guido Tabellini. 2013. "The Political Resource Curse." *American Economic Review* 103 (5): 1759-96.

Broockman, David E. 2009. "Do congressional candidates have reverse coattails? Evidence from a regression discontinuity design." *Political Analysis* 17 (4): 418–434.

Butler, Daniel Mark. 2009. "A regression discontinuity design analysis of the incumbency advantage and tenure in the US House." *Electoral Studies* 28 (1): 123–128.

Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik. 2014a. "Robust Data-Driven Inference in the Regression-Discontinuity Design." *Stata Journal* 14 (4): 909-946.

Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik. 2014b. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82 (6): 2295–2326.

Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik. 2015a. "Optimal Data-Driven

Regression Discontinuity Plots." *Journal of the American Statistical Association*, forthcoming.

Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik. 2015b. "`rdrobust`: An R Package for Robust Inference in Regression-Discontinuity Designs." *R Journal,* forthcoming.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber. 2014. "Inference on Causal Effects in a Generalized Regression Kink Design." Working paper, UC-Berkeley.

Cattaneo, Matias D., Brigham Frandsen, and Rocío Titiunik. 2015. "Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate." *Journal of Causal Inference*, forthcoming.

Caughey, Devin, and Jasjeet S Sekhon. 2011. "Elections and the Regression Discontinuity Design: Lessons from Close US House Races, 1942–2008." *Political Analysis* 19: 385–408.

DiNardo, John, and David S Lee. 2004. "Economic Impacts of New Unionization on Private Sector Employers: 1984–2001." *Quarterly Journal of Economics* 119 (4): 1383–1441.

Duraisamy, P., Bertrand Lemennicier, and Michele Khouri. 2014. "The Incumbency Effect In The Indian Parliamentary Elections, 2004 and 2009: A Regression Discontinuity Approach." *Journal of Quantitative Economics* 12 (1): 12–30.

Eggers, Andrew, Anthony Fowler, Jens Hainmueller, Andrew B Hall, and James M Snyder. 2015. "On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races." *American Journal of Political Science* 59 (1): 259–274.

Eggers, Andrew C., and Jens Hainmueller. 2009. "MPs for Sale? Returns to Office in Postwar British Politics." *American Political Science Review* 103 (4): 513-533.

Ferreira, Fernando, and Joseph Gyourko. 2009. "Do Political Parties Matter? Evidence from US Cities." *Quarterly Journal of Economics* 124 (1): 399–422.

Folke, Olle, and James M Snyder. 2012. "Gubernatorial midterm slumps." *American Journal of Political Science* 56 (4): 931–948.

Gagliarducci, Stefano, and M. Daniele Paserman. 2012. "Gender Interactions within Hierarchies: Evidence from the Political Arena." *The Review of Economic Studies* 79 (3): 1021-1052.

Gerber, Elisabeth R, and Daniel J Hopkins. 2011. "When mayors matter: estimating the impact of mayoral partisanship on city policy." *American Journal of Political Science* 55 (2): 326–339.

Hahn, Jinyong, Petra Todd, and Wilbert van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica* 69 (1): 201–209.

Hainmueller, Jens, and Holger Lutz Kern. 2008. "Incumbency as a source of spillover effects in mixed electoral systems: Evidence from a regression-discontinuity design." *Electoral Studies* 27 (2): 213–227.

Hotz, V. Joseph, Guido Imbens, and Julie H. Mortimer. 2005. "Predicting the efficacy of future training programs using past experiences at other locations." *Journal of Econometrics* 125 (1-2): 241–70.

Imbens, Guido, and Thomas Lemieux. 2008. "Regression Discontinuity Designs: A Guide to Practice." *Journal of Econometrics* 142 (2): 615–635.

Imbens, Guido, and Tristan Zajonc. 2011. "Regression discontinuity design with multiple forcing variables." *Working paper.*

Keele, Luke J., and Rocío Titiunik. 2015. "Geographic Boundaries as Regression Disconti-
nuities." *Political Analysis* 23 (1): 127–155.

Kendall, Chad, and Marie Rekkas. 2012. "Incumbency advantages in the Canadian Parlia-
ment." *Canadian Journal of Economics/Revue canadienne d'économique* 45 (4): 1560–
1585.

Klašnja, Marko. 2014. "Corruption and the Incumbency Disadvantage: Theory and Evi-
dence." *Working paper*.

Klašnja, Marko, and Rocío Titiunik. 2014. "The Incumbency Curse: Weak Parties, Term
Limits, and Unfulfilled Accountability." *Working paper*.

Lee, David S. 2008. "Randomized Experiments from Non-random Selection in U.S. House
Elections." *Journal of Econometrics* 142 (2): 675–697.

Lee, David S, Enrico Moretti, and Matthew J Butler. 2004. "Do voters affect or elect policies?
Evidence from the US House." *Quarterly Journal of Economics*: 807–859.

Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics."
Journal of Economic Literature 48 (2): 281–355.

Papay, John P, John B Willett, and Richard J Murnane. 2011. "Extending the regression-
discontinuity approach to multiple assignment variables." *Journal of Econometrics* 161 (2):
203–207.

Pettersson-Lidbom, Per. 2008. "Do Parties Matter for Economic Outcomes? A Regression-
Discontinuity Approach." *Journal of the European Economic Association* 6 (5): 1037–1056.

Thistlethwaite, Donald L., and Donald T. Campbell. 1960. "Regression-discontinuity Analy-
sis: An Alternative to the Ex-Post Facto Experiment." *Journal of Educational Psychology*
51 (6): 309–317.

Trounstine, Jessica. 2011. "Evidence of a local incumbency advantage." *Legislative Studies Quarterly* 36 (2): 255–280.

Uppal, Yogesh. 2009. "The disadvantaged incumbents: estimating incumbency effects in Indian state legislatures." *Public Choice* 138 (1-2): 9–27.

Uppal, Yogesh. 2010. "Estimating Incumbency Effects in U.S. State Legislatures: A Quasi-Experimental Study." *Economics & Politics* 22 (2): 180–199.

Wong, Vivian C., Peter M. Steiner, and Thomas D. Cook. 2013. "Analyzing Regression-Discontinuity Designs With Multiple Assignment Variables A Comparative Study of Four Estimation Methods." *Journal of Educational and Behavioral Statistics* 38 (2): 107–141.