JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
2021, VOL. 116, NO. 536, 1941–1952: Theory and Methods
https://doi.org/10.1080/01621459.2020.1751646

Extrapolating Treatment Effects in Multi-Cutoff Regression Discontinuity Designs

Matias D. Cattaneo[a], Luke Keele[b], Rocío Titiunik[c], and Gonzalo Vazquez-Bare[d]

[a]Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ; [b]Department of Surgery and Biostatistics, University of Pennsylvania, Philadelphia, PA; [c]Department of Politics, Princeton University, Princeton, NJ; [d]Department of Economics, University of California at Santa Barbara, Santa Barbara, CA

ABSTRACT

In nonexperimental settings, the regression discontinuity (RD) design is one of the most credible identification strategies for program evaluation and causal inference. However, RD treatment effect estimands are necessarily local, making statistical methods for the extrapolation of these effects a key area for development. We introduce a new method for extrapolation of RD effects that relies on the presence of multiple cutoffs, and is therefore design-based. Our approach employs an easy-to-interpret identifying assumption that mimics the idea of "common trends" in difference-in-differences designs. We illustrate our methods with data on a subsidized loan program on post-education attendance in Colombia, and offer new evidence on program effects for students with test scores away from the cutoff that determined program eligibility. Supplementary materials for this article are available online.

ARTICLE HISTORY
Received March 2019
Accepted March 2020

KEYWORDS
Causal inference; Extrapolation; Regression discontinuity

1. Introduction

The regression discontinuity (RD) design is one of the most credible strategies for estimating causal treatment effects in nonexperimental settings. In an RD design, units receive a score (or running variable), and a treatment is assigned based on whether the score exceeds a known cutoff value: units with scores above the cutoff are assigned to the treatment condition, and units with scores below the cutoff are assigned to the control condition. This treatment assignment rule creates a discontinuity in the probability of receiving treatment which, under the assumption that units' average characteristics do not change abruptly at the cutoff, offers a way to learn about the causal treatment effect by comparing units barely above and barely below the cutoff. Despite the popularity and widespread use of RD designs, the evidence they provide has an important limitation: the RD causal effect is only identified for the very specific subset of the population whose scores are "just" above or below the cutoff, and is not necessarily informative or representative of what the treatment effect would be for units whose scores are far from the RD cutoff. Thus, by its very nature, the RD parameter is local and has limited external validity.

We empirically illustrate the advantages and limitations of RD designs employing a recent study of the Acceso con Calidad a la Educación Superior (ACCES) program by Melguizo, Sanchez, and Velasco (2016). ACCES is a subsidized loan program in Colombia, administered by the Colombian Institute for Educational Loans and Studies Abroad (ICETEX), that provides tuition credits to underprivileged populations for various post-secondary education programs such as technical, technical-professional, and university degrees. To be eligible for an ACCES credit, students must be admitted to a qualifying higher education program, have good credit standing and, if soliciting the credit in the first or second semester of the higher education program, achieve a minimum score on a high school exit exam known as SABER 11. In other words, to obtain ACCES funding students must have an exam score above a known cutoff. Students who are just below the exam cutoff are deemed ineligible, and therefore are not offered financial assistance. This discontinuity in program eligibility based on the exam score leads to a RD design: Melguizo, Sanchez, and Velasco (2016) found that students just above the threshold in SABER 11 test scores were significantly more likely to enroll in a wide variety of post-secondary education programs. The evidence from the original study is limited to the population of students around the cutoff. This standard causal RD treatment effect is informative in its own right but, in the absence of additional assumptions, it cannot be used to understand the effects of the policy for students whose test scores are outside the immediate neighborhood of the cutoff. Treatment effects away from the cutoff are useful for a variety of purposes, ranging from answering purely substantive questions to addressing practically important policy making decisions such as whether to roll-out the program or not.

We propose a novel approach for estimating RD causal treatment effects away from the cutoff that determines treatment assignment. Our extrapolation approach is design-based as it exploits the presence of multiple RD cutoffs across different subpopulations to construct valid counterfactual extrapolations of the expected outcome of interest, given different scores levels, in the absence of treatment assignment. In sum, our approach

CONTACT Matias D. Cattaneo cattaneo@princeton.edu Department of Operations Research and Financial Engineering, Princeton University Sherrerd Hall, Charlton Street Princeton, NJ 08544

Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

imputes the average outcome in the absence of treatment of a treated subpopulation exposed to a given cutoff, using the average outcome of another subpopulation exposed to a higher cutoff. Assuming that the difference between these two average outcomes is constant as a function of the score, this imputation identifies causal treatment effects at score values higher than the lower cutoff.

The rest of the article is organized as follows. The next section presents further details on the operation of the ACCES program, discusses the particular program design features that we use for the extrapolation of RD effects, and presents the intuitive idea behind our approach. In that section, we also discuss related literature on RD extrapolation as well as on estimation and inference. Section 3 presents the main methodological framework and extrapolation results for the case of the "sharp" RD design, which assumes perfect compliance with treatment assignment (or a focus on an intention-to-treat parameter). Section 4 applies our results to extrapolate the effect of the ACCES program on educational outcomes, while Section 5 illustrates our methods using simulated data. Section 6 presents an extension to the "fuzzy" RD design, which allows for imperfect compliance. Section 7 concludes. The supplemental appendix contains additional results, including further extensions and generalizations of our extrapolation methods.

2. The RD Design in the ACCES Program

The SABER 11 exam that serves as the basis for eligibility to the ACCES program is a national exam administered by the Colombian Institute for the Promotion of Postsecondary Education (ICFES), an institute within Colombia's National Ministry of Education. This exam may be taken in the fall or spring semester each year, and has a common core of mandatory questions in seven subjects—chemistry, physics, biology, social sciences, philosophy, mathematics, and language. To sort students according to their performance in the exam, ICFES creates an index based on the difference between (i) a weighted average of the standardized grades obtained by the student in each common core subject, and (ii) the within-student standard deviation across the standardized grades in the common core subjects. This index is commonly referred to as the SABER 11 score.

Each semester of every year, ICFES calculates the 1000 quantiles of the SABER 11 score among all students who took the exam that semester, and assigns a score between 1 and 1000 to each student according to their position in the distribution—we refer to these scores as the SABER 11 position scores. Thus, the students in that year and semester whose scores are in the top 0.1% are assigned a value of 1 (first position), the students whose scores are between the top 0.1% and 0.2% are assigned a value of 2 (second position), etc., and the students whose scores are in the bottom 0.1% are assigned a value of 1000 (the last position). Every year, the position scores are created separately for each semester, and then pooled. Melguizo, Sanchez, and Velasco (2016) provided further details on the Colombian education system and the ACCES program.

In this sharp RD design, the running variable is the SABER 11 position score, denoted by X_i for each unit i in the sample, and the treatment of interest is receiving approval of the ACCES credit. Between 2000 and 2008, the cutoff to qualify for an ACCES credit was 850 in all Colombian departments (the largest subnational administrative unit in Colombia, equivalent to U.S. states). To be eligible for the program, a student must have a SABER 11 position score at or below the 850 cutoff.

2.1. The Multi-Cutoff RD Design

In the canonical RD design, a single cutoff is used to decide which units are treated. As we noted above, eligibility for the ACCES program between 2000 and 2008 followed this template, since the cutoff was 850 for all students. However, in many RD designs, the same treatment is given to all units based on whether the RD score exceeds a cutoff, but different units are exposed to different cutoffs. This contrasts with the assignment rule in the standard RD design, in which all units face the same cutoff value. RD designs with multiple cutoffs, which we call multi-cutoff RD designs, are fairly common and have specific properties (Cattaneo et al. 2016).

In 2009, ICFES changed the program eligibility rule, and started employing different cutoffs across years and departments. Consequently, after 2009, ACCES eligibility follows a multi-cutoff RD design: the treatment is the same throughout Colombia—all students above the cutoff receive the same financial credits for educational spending—but the cutoff that determines treatment assignment varies widely by department and changes each year, so that different sets of students face different cutoffs. This design feature is at the core of our approach for extrapolation of RD treatment effects.

2.2. The Pooled RD Effect of the ACCES Program

Multi-cutoff RD designs are often analyzed as if they had a single cutoff. For example, in the original analysis, Melguizo, Sanchez, and Velasco (2016) redefined the RD running variable as distance to the cutoff, and analyzed all observations together using a common cutoff equal to zero. In fact, this normalizing-and-pooling approach (Cattaneo et al. 2016), which essentially ignores or "averages over" the multi-cutoff features of the design, is widespread in empirical work employing RD designs. See the supplemental appendix for a sample of recent papers that analyze RD designs with multiple cutoffs across various disciplines.

We first present some initial empirical results using the normalizing-and-pooling approach as a benchmark for later analyses. The outcome we analyze is an indicator for whether the student enrolls in a higher education program, one of several outcomes considered in the original study by Melguizo, Sanchez, and Velasco (2016). To maintain the standard definition of RD assignment as having a score above the cutoff, we multiply the SABER 11 position score by -1. We focus on the intention-to-treat effect of program eligibility on higher education enrollment, which gives a sharp RD design. We discuss an extension to fuzzy RD designs in Section 6. We focus our analysis on the population of students exposed to two different cutoffs, -850 and -571.

For our main analysis, we employ statistical methods for RD designs based on recent methodological developments in

Calonico, Cattaneo, and Titiunik (2014, 2015), Calonico, Cattaneo, and Farrell (2018, 2020a, 2020b), Calonico et al. (2019), and references therein. In particular, point estimators are constructed using mean squared error (MSE) optimal bandwidth selection, and confidence intervals are formed using robust bias correction (RBC). We provide details on estimation and inference in Section 3.2.

Figure 1 reports two RD plots of the data, reporting linear and quadratic global polynomial approximations. Using RBC local polynomial inference, Table 1 reports that the pooled RD estimate of the ACCES program treatment effect on expected higher education enrollment is 0.125, with corresponding 95% RBC confidence interval [0.012, 0.219]. These results indicate that, in our sample, students who barely qualify for the ACCES program based on their SABER 11 score are 12.5 percentage points more likely to enroll in a higher education program than students who are barely ineligible for the program. These results are consistent with the original positive effects of ACCES eligibility on higher education enrollment rates reported in Melguizo, Sanchez, and Velasco (2016). However, the pooled RD estimate only pertains to a limited set of ACCES applicants: those whose scores are barely above or below one of the cutoffs.

Table 1. Main empirical results for ACCES loan eligibility on post-education enrollment.

				Robust BC inference	
	Estimate	Bw	Eff. N	p-Value	95% CI
RD effects					
$C = -850$	0.137	71.7	71	0.007	[0.036, 0.232]
$C = -571$	0.169	136.3	133	0.103	[−0.039, 0.428]
Weighted	0.156		204	0.021	[0.025, 0.314]
Pooled	0.125	147.6	291	0.029	[0.012, 0.219]
Naive difference					
$\mu_\ell(-650)$	0.756	240.1	441		
$\mu_h(-650)$	0.706	131.2	202		
Difference	0.050			0.179	[−0.020, 0.107]
Bias					
$\mu_\ell(-850)$	0.525	54.9	54		
$\mu_h(-850)$	0.667	144.2	230		
Difference	−0.142			0.004	[−0.274, −0.054]
Extrapolation					
$\tau_\ell(-650)$	0.191			0.001	[0.080, 0.336]

NOTES: Local polynomial regression estimation with MSE-optimal bandwidth selectors and robust bias corrected inference. See Calonico, Cattaneo, and Titiunik (2014) and Calonico, Cattaneo, and Farrell (2018) for methodological details, and Calonico et al. (2017) and Cattaneo, Titiunik, and Vazquez-Bare (2020a) for implementation. "Eff. N" indicates the effective sample size, that is, the sample size within the MSE-optimal bandwidth. "Bw" indicates the MSE-optimal bandwidth.

Cattaneo et al. (2016) showed that the pooled RD estimand is a weighted average of cutoff-specific RD treatment effects for subpopulations facing different cutoffs. The empirical results for the pooled and cutoff-specific estimates can be seen in the upper panel of Table 1. In our sample, the pooled estimate of 0.125 is a (linear in large samples) combination of two cutoff-specific RD estimates, one for units facing the low cutoff −850 and one for units facing the high cutoff −571. We provide a detailed analysis of these estimates in Section 4. These cutoff-specific estimates are not directly comparable, as these magnitudes correspond not only to different values of the running variable but also to different subpopulations. We discuss next how the availability of multiple cutoffs can be exploited to learn about treatment effects far from the cutoff in the context of the ACCES policy intervention.

2.3. Using the Multi-Cutoff RD Design for Extrapolation

Our key contribution is to exploit the presence of multiple RD cutoffs to extrapolate the standard RD average treatment effects (at each cutoff) to students whose SABER 11 scores are away from the cutoff actually used to determine program eligibility. Our method relies on a simple idea: when different units are exposed to different cutoffs, different units with the same value of the score may be assigned to different treatment conditions, relaxing the strict lack of overlap between treated and control scores that is characteristic of the single-cutoff RD design.

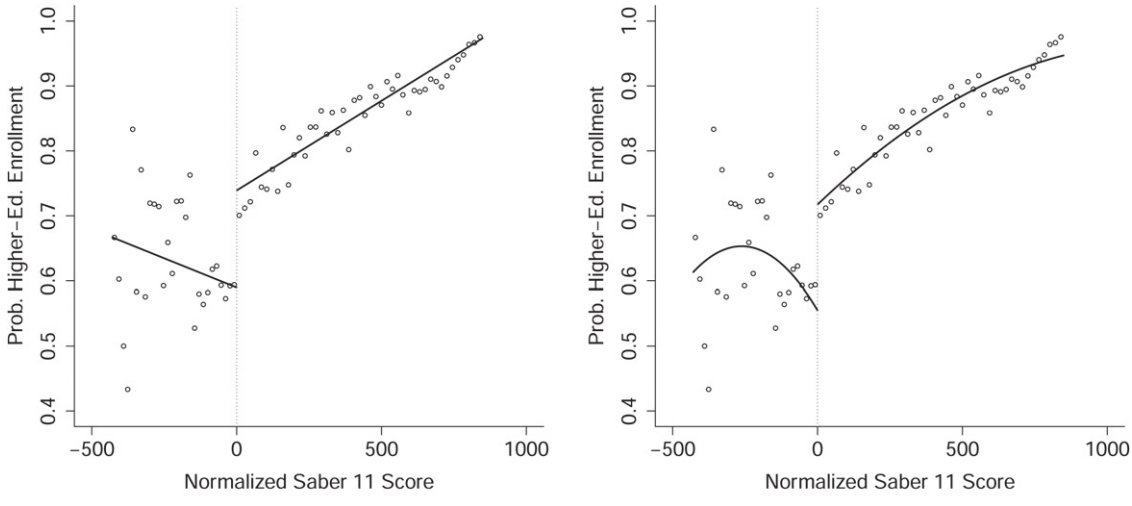

(a) Global Linear Fit (b) Global Quadratic Fit

Figure 1. Normalizing-and-pooling RD Plot of ACCES loan eligibility on post-education enrollment. NOTES: RD plot constructed using evenly spaced binning and global linear (left) and quadratic (right) polynomial fits for normalized (to zero) and pooled (across cutoffs) score variable. See Calonico, Cattaneo, and Titiunik (2015) and Cattaneo et al. (2016) for methodological details, and Calonico et al. (2017) and Cattaneo, Titiunik, and Vazquez-Bare (2020a) for implementation.

For example, consider the simplest multi-cutoff RD design with two cutoffs, ℓ and $\hbar$, with $\ell < \hbar$, where we wish to estimate the average treatment effect at a point $\bar{x} \in (\ell, \hbar)$. Units exposed to ℓ receive the treatment according to $\mathbb{1}(X_i \geq \ell)$, where X_i is unit's i score and $\mathbb{1}(.)$ is the indicator function, so they are all treated at $X = \bar{x}$. However, the same design contains units who receive the treatment according to $\mathbb{1}(X_i \geq \hbar)$, so they are controls at both $X = \bar{x}$ and $X = \ell$. Our idea is to compare the observable difference in the control groups at the low cutoff ℓ, and assume that the same difference in control groups occurs at the interior point $\bar{x}$. This allows us to identify the average treatment effect for all score values between the cutoffs ℓ and $\hbar$.

Our identifying idea is analogous to the "parallel trends" assumption in difference-in-difference designs (see, e.g., Abadie 2005, and references therein), but over a continuous dimension—that is, over the values of the continuous score variable X_i.

2.4. Related Literature

We contribute to the causal inference and program evaluation literatures (Imbens and Rubin 2015; Abadie and Cattaneo 2018) and, more specifically, to the methodological literature on RD designs. See Imbens and Lemieux (2008), Cattaneo, Titiunik, and Vazquez-Bare (2017), Cattaneo, Titiunik, and Vazquez-Bare (2020b), and Cattaneo, Idrobo, and Titiunik (2019, 2020), for literature reviews, background references, and practical introductions.

Our article adds to the recent literature on RD treatment effect extrapolation methods, a nonparametric identification problem in causal inference. This strand of the literature can be classified into two groups: strategies assuming the availability of external information, and strategies based only on information from within the research design. Approaches based on external information include Mealli and Rampichini (2012), Wing and Cook (2013), Rokkanen (2015), and Angrist and Rokkanen (2015). The first two papers rely on a pre-intervention measure of the outcome variable, which they use to impute the treated-control differences of the post-intervention outcome above the cutoff. Rokkanen (2015) assumed that multiple measures of the running variable are available, and all measures capture the same latent factor; identification relies on the assumption that the potential outcomes are conditionally independent of the available measurements given the latent factor. Angrist and Rokkanen (2015) relied on pre-intervention covariates, assuming that the running variable is ignorable conditional on the covariates over the whole range of extrapolation. All these approaches assume the availability of external information that is not part of the original RD design.

In contrast, the extrapolation approaches in Dong and Lewbel (2015) and Bertanha and Imbens (2020) require only the score and outcome in the standard (single-cutoff) RD design. Dong and Lewbel (2015) assume mild smoothness conditions to identify the derivatives of the average treatment effect with respect to the score, which allows for a local extrapolation of the standard RD treatment effect to score values marginally above the cutoff. Bertanha and Imbens (2020) exploit variation in treatment assignment generated by imperfect treatment compliance imposing independence between potential outcomes and compliance types to extrapolate a single-cutoff fuzzy RD treatment effect (i.e., a local average treatment effect at the cutoff) away from the cutoff. Our paper also belongs to this second type, as it relies on within-design information, using only the score and outcome in the multi-cutoff RD design.

Cattaneo et al. (2016) introduced the causal Multi-cutoff RD framework, which we employ herein, and studied the properties of normalizing-and-pooling estimation and inference in that setting. Building on that paper, Bertanha (2020) discusses estimation and inference of an average treatment effect across multi-cutoffs, assuming away cutoff-specific treatment effect heterogeneity. Neither of these papers addressed the topic of RD treatment effect extrapolation across different levels of the score variable, which is the main goal and innovation of the present article.

All the papers mentioned above focus on extrapolation of RD treatment effects away from the cutoff by relying on continuity-based methods for identification, estimation and inference, which are implemented using local polynomial regression (Fan and Gijbels 1996). As pointed out by a reviewer, an alternative approach to analyzing RD designs is to employ the local randomization framework introduced by Cattaneo, Frandsen, and Titiunik (2015). This framework has been later used in the context of geographic RD designs (Keele, Titiunik, and Zubizarreta 2015), principal stratification (Li, Mattei, and Mealli 2015), and kink RD designs (Ganong and Jäger 2018), among other settings. More recently, this alternative RD framework was expanded to allow for finite-sample falsification testing and for local regression adjustments (Cattaneo, Titiunik, and Vazquez-Bare 2017). See also Sekhon and Titiunik (2016, 2017) for further conceptual discussions, Cattaneo, Titiunik, and Vazquez-Bare (2020b) for another review, and Cattaneo, Idrobo, and Titiunik (2020) for a practical introduction.

Local randomization RD methods implicitly give extrapolation within the neighborhood where local randomization is assumed to hold because they assume a parametric (usually constant) treatment effect model as a function of the score. However, those methods cannot aid in extrapolating RD treatment effects beyond such neighborhood without additional assumptions, which is precisely our goal. Since local randomization methods explicitly view RD designs as local randomized experiments, we can summarize the key conceptual distinction between that literature and our paper as follows: available local randomization methods for RD designs have only internal validity (i.e., within the local randomization neighborhood), while our proposed method seeks to achieve external validity (i.e., outside the local randomization neighborhood), which we achieve by exploiting the presence of multiple cutoffs (akin to multiple local experiments) together with an additional identifying assumption within the continuity-based approach to RD designs (i.e., parallel control regression functions across cutoffs).

Our core extrapolation idea can be developed within the local randomization framework, albeit under considerably stronger assumptions. To conserve space, in the supplemental appendix, we discuss multi-cutoff extrapolation of RD treatment effects using local randomization ideas, and develop randomization-based estimation and inference methods (Rosenbaum 2010; Imbens and Rubin 2015). We also empirically illustrate these methods in the supplemental appendix.

3. Extrapolation in Multi-Cutoff RD Designs

We assume (Y_i, X_i, C_i, D_i), $i = 1, 2, \ldots, n$, is an observed random sample, where Y_i is the outcome of interest, X_i is the score (or running variable), C_i is the cutoff indicator, and D_i is a treatment status indicator. We assume the score has a continuous positive density $f_X(x)$ on the support $\mathcal{X}$. Unlike the canonical RD design where the cutoff is a fixed scalar, in the multi-cutoff RD design the cutoff faced by unit i is the random variable C_i taking values in a set $\mathcal{C} \subset \mathcal{X}$. For simplicity, we consider two cutoffs: $\mathcal{C} = \{\ell, \hbar\}$, with $\ell < \hbar$ and $\ell, \hbar \in \mathcal{X}$. Extensions to more than two cutoffs and to geographic and multi-score RD designs are conceptually straightforward, and hence discussed in the supplemental appendix.

The conditional density of the score at each cutoff is $f_{X|C}(x|c)$, $c \in \mathcal{C}$. In sharp RD designs treatment assignment and status are identical, and hence $D_i = \mathbb{1}(X_i \geq C_i)$. Section 6 discusses an extension to fuzzy RD designs. Finally, we let $Y_i(1)$ and $Y_i(0)$ denote the potential outcomes of unit i under treatment and control, respectively, and $Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0)$ is the observed outcome.

The potential outcome regression functions are $\mu_{d,c}(x) = \mathbb{E}[Y_i(d)|X_i = x, C_i = c]$, for $d = 0, 1$. We express all parameters of interest in terms of the "response" function

$$\tau_c(x) = \mathbb{E}[Y_i(1) - Y_i(0) \mid X_i = x, C_i = c]. \tag{1}$$

This function measures the treatment effect for the subpopulation exposed to cutoff c when the running variable takes the value x. For a fixed cutoff c, it records how the treatment effect for the subpopulation exposed to this cutoff varies with the running variable. As such, it captures a key quantity of interest when extrapolating the RD treatment effect. The usual parameter of interest in the standard (single-cutoff) RD design is a particular case of $\tau_c(x)$ when cutoff and score coincide

$$\tau_c(c) = \mathbb{E}[Y_i(1) - Y_i(0) \mid X_i = c, C_i = c] = \mu_{1,c}(c) - \mu_{0,c}(c).$$

It is well known that, via continuity assumptions, the function $\tau_c(x)$ is nonparametrically identifiable at the single point $x = c$. Our approach exploits the presence of multiple cutoffs to identify this function at other points on a portion of the support of the score variable.

Figure 2 contains a graphical representation of our extrapolation approach for multi-cutoff RD designs. In the plot, there are two populations, one exposed to a low cutoff ℓ, and another exposed to a high cutoff $\hbar$. The RD effects for each subpopulation are, respectively, $\tau_\ell(\ell)$ and $\tau_\hbar(\hbar)$. We seek to learn about the effects of the treatment at points other than the particular cutoff to which units were exposed, such as the point $\bar{x}$ in Figure 2. Below, we develop a framework for the identification of $\tau_\ell(x)$ for $\ell < x \leq \hbar$ so that we can assess what would have been the average treatment effect for the subpopulation exposed to the cutoff ℓ at score values above ℓ (illustrated by the effect $\tau_\ell(\bar{x})$ in Figure 2 for the intermediate point $X_i = \bar{x}$).

In our framework, the multiple cutoffs define different subpopulations. In some cases, the cutoff to which a unit is exposed depends only on characteristics of the units, such as when the cutoffs are cumulative and increase as the score falls in increasingly higher ranges. In other cases, the cutoff depends on external features, such as when different cutoffs are used in different geographic regions or time periods. This means that, in our framework, the cutoff C_i acts as an index for different subpopulation "types," capturing both observed and unobserved characteristics of the units.

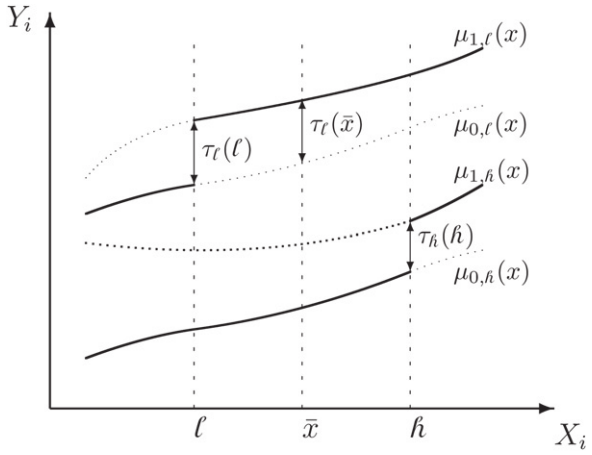

Figure 2. Estimands of interest with two cutoffs.

Given the subpopulations defined by the cutoff values actually used in the multi-cutoff RD design, we consider the effect that the treatment would have had for those subpopulations had the units had a higher score value than observed. This is why, in our notation, the index for the cutoff value is fixed, and the index for the score is allowed to vary and is the main argument of the regression functions. This conveys the idea that the subpopulations are defined by the multiple cutoffs actually employed, and our exercise focuses on studying the treatment effect at different score values for those predefined subpopulations. For example, this setting covers RD designs with a common running variable but with cutoffs varying by regions, schools, firms, or some other group-type variable. Our method is not appropriate to extrapolate to populations outside those defined by the multi-cutoff RD design.

3.1. Identification Result

The main challenge to the identification of extrapolated treatment effects in the single-cutoff (sharp) RD design is the lack of observed control outcomes for score values above the cutoff. In the multi-cutoff RD design, we still face this challenge for a given subpopulation, but we have other subpopulations exposed to higher cutoff values that, under some assumptions, can aid in solving the missing data problem and identify average treatment effects. Before turning to the formal derivations, we illustrate the idea graphically.

Figure 3 illustrates the regression functions for the populations exposed to cutoffs ℓ and $\hbar$, with the function $\mu_{1,\hbar}(x)$ omitted for simplicity. We seek an estimate of $\tau_\ell(\bar{x})$, the average effect of the treatment at the point $\bar{x} \in (\ell, \hbar)$ for the subpopulation exposed to the lower cutoff ℓ. In the figure, this parameter is represented by the segment $\overline{ab}$. The main identification challenge is that we only observe the point a, which corresponds to $\mu_{1,\ell}(\bar{x})$, the treated regression function for the population exposed to ℓ, but we fail to observe its control counterpart $\mu_{0,\ell}(\bar{x})$ (point b), because all units exposed to cutoff ℓ are treated at any $x > \ell$. We

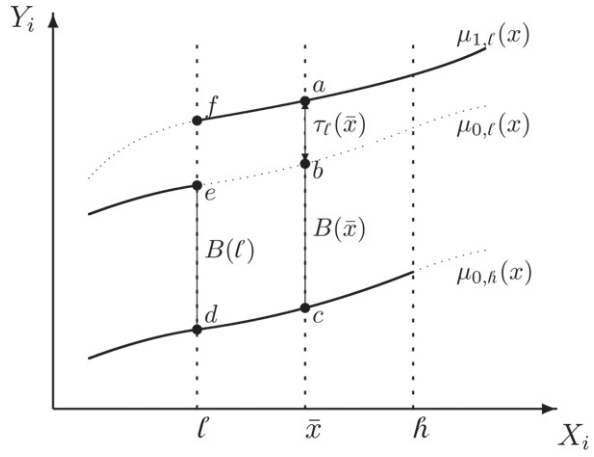

Figure 3. RD extrapolation with constant bias ($B(\ell) = B(\bar{x})$).

Table 2. Segments and corresponding parameters in Figure 2.

Segment	Parameter	Description
$\overline{ab}$	$\tau_\ell(\bar{x}) = \underbrace{\mu_{1,\ell}(\bar{x})}_{\text{Observable}} - \underbrace{\mu_{0,\ell}(\bar{x})}_{\text{Unobservable}}$	Extrapolation parameter of interest
$\overline{bc}$	$B(\bar{x}) = \underbrace{\mu_{0,\ell}(\bar{x})}_{\text{Unobservable}} - \underbrace{\mu_{0,\bar{\ell}}(\bar{x})}_{\text{Observable}}$	Control facing $\bar{\ell}$ vs. control facing $\bar{\ell}$, at $X_i = \bar{x}$
$\overline{ac}$	$\tau_\ell(\bar{x}) + B(\bar{x}) = \underbrace{\mu_{1,\ell}(\bar{x})}_{\text{Observable}} - \underbrace{\mu_{0,\bar{\ell}}(\bar{x})}_{\text{Observable}}$	Treated facing ℓ vs. control facing $\bar{\ell}$, at $X_i = \bar{x}$
$\overline{ed}$	$B(\ell) = \underbrace{\mu_{0,\ell}(\ell)}_{\text{Observable}} - \underbrace{\mu_{0,\bar{\ell}}(\ell)}_{\text{Observable}}$	Control facing ℓ vs. control facing $\bar{\ell}$, at $X_i = \ell$

use the control group of the population exposed to the higher cutoff, $\bar{\ell}$, to infer what would have been the control response at $\bar{x}$ of units exposed to the lower cutoff ℓ. At the point $X_i = \bar{x}$, the control response of the population exposed to $\bar{\ell}$ is $\mu_{0,\bar{\ell}}(\bar{x})$, which is represented by the point c in Figure 3. Since all units in this subpopulation are untreated at $\bar{x}$, the point c is identified by the average observed outcomes of the control units in the subpopulation $\bar{\ell}$ at $\bar{x}$.

Of course, units facing different cutoffs may differ in both observed and unobserved ways. Thus, there is generally no reason to expect that the average control outcome of the population facing cutoff $\bar{\ell}$ will be a good approximation to the average control outcome of the population facing cutoff ℓ. This is captured in Figure 3 by the fact that $\mu_{0,\ell}(\bar{x}) \equiv b \neq c \equiv \mu_{0,\bar{\ell}}(\bar{x})$. This difference in untreated potential outcomes for units facing different cutoffs can be interpreted as a bias driven by differences in observed and unobserved characteristics of the different subpopulations, analogous to "site selection" bias in multiple randomized experiments. We formalize this idea with the following definition.

Definition 1 (Cutoff bias). $B(x, c, c') = \mu_{0,c}(x) - \mu_{0,c'}(x)$, for $c, c' \in \mathcal{C}$. There is bias from exposure to different cutoffs if $B(x, c, c') \neq 0$ for some $c, c' \in \mathcal{C}$, $c \neq c'$ and for some $x \in \mathcal{X}$.

Table 2 defines the parameters associated with the corresponding segments in Figure 3. The parameter of interest, $\tau_\ell(\bar{x})$, is unobservable because we fail to observe $\mu_{0,\ell}(\bar{x})$. If we replaced $\mu_{0,\ell}(\bar{x})$ with $\mu_{0,\bar{\ell}}(\bar{x})$, we would be able to estimate the distance $\overline{ac}$. This distance, which is observable, is the sum of the parameter of interest, $\tau_\ell(\bar{x})$, plus the bias $B(\bar{x}, c, c')$ that arises from using the control group in the $\bar{\ell}$ subpopulation instead of the control group in the ℓ subpopulation. Graphically, $\overline{ac} = \overline{ab} + \overline{bc}$. Since we focus on the two-cutoff case, we denote the bias by $B(\bar{x})$ to simplify the notation.

We use the distance between the control groups facing the two different cutoffs at a point where both are observable, to approximate the unobservable distance between them at $\bar{x}$—that is, to approximate the bias $B(\bar{x})$. As shown in the figure, at ℓ, all units facing cutoff $\bar{\ell}$ are controls and all units facing cutoff ℓ are treated. But under standard RD assumptions, we can identify $\mu_{0,\ell}(\ell)$ using the observations in the ℓ subpopulation whose scores are just below ℓ. Thus, the bias term $B(\ell)$, captured in the distance $\overline{ed}$, is estimable from the data.

Graphically, we can identify the extrapolation parameter $\tau_\ell(\bar{x})$ assuming that the observed difference between the control functions $\mu_{0,\ell}(\cdot)$ and $\mu_{0,\bar{\ell}}(\cdot)$ at ℓ is constant for all values of the score:

$$\begin{aligned}\overline{ac} - \overline{ed} &= \left\{\mu_{1,\ell}(\bar{x}) - \mu_{0,\bar{\ell}}(\bar{x})\right\} - \left\{\mu_{0,\ell}(\ell) - \mu_{0,\bar{\ell}}(\ell)\right\} \\ &= \{\tau_\ell(\bar{x}) + B(\bar{x})\} - \{B(\ell)\} \\ &= \tau_\ell(\bar{x}).\end{aligned}$$

We now formalize this intuitive result employing standard continuity assumptions on the relevant regression functions. We make the following assumptions.

Assumption 1 (Continuity). $\mu_{d,c}(x)$ is continuous in $x \in [\ell, \bar{\ell}]$ for $d = 0, 1$ and for all c.

The observed outcome regression functions are $\mu_c(x) = \mathbb{E}[Y_i | X_i = x, C_i = c]$, for $c \in \mathcal{C} = \{\ell, \bar{\ell}\}$, and note that by standard RD arguments $\mu_{0,c}(c) = \lim_{\varepsilon \uparrow 0} \mu_c(c + \varepsilon)$ and $\mu_{1,c}(c) = \lim_{\varepsilon \downarrow 0} \mu_c(c + \varepsilon)$. Furthermore, $\mu_{0,\bar{\ell}}(x) = \mu_{\bar{\ell}}(x)$ and $\mu_{1,\ell}(x) = \mu_\ell(x)$ for all $x \in (\ell, \bar{\ell})$.

Our main extrapolation assumption requires that the bias not be a function of the score, which is analogous to the parallel trends assumption in the difference-in-differences design.

Assumption 2 (Constant bias). $B(\ell) = B(x)$ for all $x \in (\ell, \bar{\ell})$.

While technically our identification result only needs this condition to hold at $x = \bar{x}$, in practice it may be hard to argue that the equality between biases holds at a single point. Combining the constant bias assumption with the continuity-based identification of the conditional expectation functions allows us to express the unobservable bias for an interior point, $\bar{x} \in (\ell, \bar{\ell})$, as a function of estimable quantities. The bias at the low cutoff ℓ can be written as

$$B(\ell) = \lim_{\varepsilon \uparrow 0} \mu_\ell(\ell + \varepsilon) - \mu_{\bar{\ell}}(\ell).$$

Under Assumption 2, we have

$$\mu_{0,\ell}(\bar{x}) = \mu_{\bar{\ell}}(\bar{x}) + B(\ell), \quad \bar{x} \in (\ell, \bar{\ell}),$$

that is, the average control response for the ℓ subpopulation at the interior point $\bar{x}$ is equal to the average observed response for the $\hbar$ subpopulation at the same point, plus the difference in the average control responses between both subpopulations at the low cutoff ℓ. This leads to our main identification result.

Theorem 1 (Extrapolation). Under Assumptions 1 and 2, for any point $\bar{x} \in (\ell, \hbar)$,

$$\tau_\ell(\bar{x}) = \mu_\ell(\bar{x}) - [\mu_\hbar(\bar{x}) + B(\ell)].$$

This result can be extended to hold for $\bar{x} \in (\ell, \hbar]$ by using side limits appropriately. In Section 3.3, we discuss two approaches to provide empirical support for the constant bias assumption. We extend our result to fuzzy RD designs in Section 6, and allow for nonparallel control regression functions and pre-intervention covariate-adjustment in the supplemental appendix.

While we develop our core idea for extrapolation from "left to right," that is, from a low cutoff to higher values of the score, it follows from the discussion above that the same ideas could be developed for extrapolation from "right to left." Mathematically, the problem is symmetric and hence both extrapolations are equally viable. However, conceptually, there is an important asymmetry. Theorem 1 requires the regression functions for *control* units to be parallel over the extrapolation region (Assumption 2), while a version of this theorem for "right to left" extrapolation would require that the regression functions for *treated* units be parallel. These two identifying assumptions are not symmetric because the latter effectively imposes a constant treatment effect assumption across cutoffs (for different values of the score), while the former does not because it pertains to control units only.

3.2. Estimation and Inference

We estimate all (identifiable) conditional expectations $\mu_{d,c}(x) = \mathbb{E}[Y_i(d)|X_i = x, C_i = c]$ using nonparametric local polynomial methods, employing second-generation MSE-optimal bandwidth selectors and robust bias correction inference methods. See Calonico, Cattaneo, and Titiunik (2014), Calonico, Cattaneo, and Farrell (2018, 2020a, 2020b), and Calonico et al. (2019) for more methodological details, and Calonico et al. (2017) and Calonico, Cattaneo, and Farrell (2019) for software implementation. See also Hyytinen et al. (2018), Ganong and Jäger (2018), and Dong, Lee, and Gou (2020) for some recent applications and empirical testing of those methods.

To be more precise, a generic local polynomial estimator is $\hat{\mu}_{d,c}(x) = \mathbf{e}_0'\hat{\boldsymbol{\beta}}_{d,c}(x)$, where

$$\hat{\boldsymbol{\beta}}_{d,c}(x) = \underset{\mathbf{b}\in\mathbb{R}^{p+1}}{\operatorname{argmin}} \sum_{i=1}^{n} (Y_i - \mathbf{r}_p(X_i - x)'\mathbf{b})^2 K\left(\frac{X_i - x}{h}\right) \mathbb{1}(C_i = c)\mathbb{1}(D_i = d),$$

$\mathbf{e}_0$ is a vector with a one in the first position and zeros in the rest, $\mathbf{r}_p(\cdot)$ is a polynomial basis of order p, $K(\cdot)$ is a kernel function, and h a bandwidth. For implementation, we set $p = 1$ (local-linear), K to be the triangular kernel, h to be a MSE-optimal bandwidth selector, unless otherwise noted. Then, given the two cutoffs ℓ and $\hbar$ and an extrapolation point $\bar{x} \in (\ell, \hbar]$, the extrapolated treatment effect at $\bar{x}$ for the subpopulation facing cutoff ℓ is estimated as

$$\hat{\tau}_\ell(\bar{x}) = \hat{\mu}_{1,\ell}(\bar{x}) - \hat{\mu}_{0,\hbar}(\bar{x}) - \hat{\mu}_{0,\ell}(\ell) + \hat{\mu}_{0,\hbar}(\ell).$$

The estimator $\hat{\tau}_\ell(\bar{x})$ is a linear combination of nonparametric local polynomial estimators at boundary and at interior points depending on the choice of $\bar{x}$ and data availability. Hence, optimal bandwidth selection and robust bias-corrected inference can be implemented using the methods and software mentioned above. By construction, $\hat{\mu}_{d,\ell}(\cdot)$ and $\hat{\mu}_{0,\hbar}(\cdot)$ are independent because the observations used for estimation come from different subpopulations. Similarly, $\hat{\mu}_{0,\ell}(\cdot)$ and $\hat{\mu}_{1,\ell}(\cdot)$ are independent since the first term is estimated using control units whereas the second term uses treated units. On the other hand, in finite samples, $\hat{\mu}_{0,\hbar}(\ell)$ and $\hat{\mu}_{0,\hbar}(\bar{x})$ can be correlated if the bandwidths used for estimation overlap (or, alternatively, if ℓ and $\bar{x}$ are close enough), in which case we account for such correlation in our inference results. More precisely, $\mathbb{V}[\hat{\tau}_\ell(\bar{x})|\mathbf{X}] = \mathbb{V}[\hat{\mu}_{1,\ell}(\bar{x})|\mathbf{X}] + \mathbb{V}[\hat{\mu}_{0,\hbar}(\bar{x})|\mathbf{X}] + \mathbb{V}[\hat{\mu}_{0,\ell}(\ell)|\mathbf{X}] + \mathbb{V}[\hat{\mu}_{0,\hbar}(\ell)|\mathbf{X}] - 2\mathbb{C}\text{ov}(\hat{\mu}_{0,\hbar}(\ell), \hat{\mu}_{0,\hbar}(\bar{x})|\mathbf{X})$, where $\mathbf{X} = (X_1, X_2, \ldots, X_n)'$.

Precise regularity conditions for large sample validity of our estimation and inference methods can be found in the references given above. The replication files contain details on practical implementation.

3.3. Assessing the Validity of the Identifying Assumption

Assessing the validity of our extrapolation strategy should be a key component of empirical work using these methods. In general, while the assumption of constant bias is not testable, this assumption can be tested indirectly via falsification. While a falsification test cannot demonstrate that an assumption holds, it can provide persuasive evidence that an assumption is implausible. We now discuss two strategies for falsification tests to probe the credibility of the constant bias assumption that is at the center of our extrapolation approach.

The first falsification approach relies on a global polynomial regression. We test globally whether the conditional expectation functions of the two control groups are parallel below the lowest cutoff. One way to implement this idea, given the two cutoff points $\ell' < \hbar$, is to test $\boldsymbol{\delta} = \mathbf{0}$ based on the regression model

$$\begin{aligned} Y_i =& \alpha + \beta\mathbb{1}(C_i = \hbar) + \mathbf{r}_p(X_i)'\boldsymbol{\gamma} + \mathbb{1}(C_i = \hbar)\mathbf{r}_p(X_i)'\boldsymbol{\delta} + u_i, \\ & \mathbb{E}[u_i|X_i, C_i] = 0, \end{aligned}$$

only for units with $X_i < \ell'$. In words, we employ a pth order global polynomial model to estimate the two regression functions $\mathbb{E}[Y_i|X_i = x, X_i < \ell', C_i = \ell']$ and $\mathbb{E}[Y_i|X_i = x, X_i < \ell', C_i = \hbar]$, separately, and construct a hypothesis test for whether they are equal up to a vertical shift (i.e., the null hypothesis is $\mathsf{H}_0 : \boldsymbol{\delta} = \mathbf{0}$). This approach is valid under standard regularity conditions for parametric least squares regression. This approach could also be justified from a nonparametric series approximation perspective, under additional regularity conditions.

The second falsification approach employs nonparametric local polynomial methods. We test for equality of the derivatives

of the conditional expectation functions for values $x < \ell$. Specifically, we test for $\mu_\ell^{(1)}(x) = \mu_\hbar^{(1)}(x)$ for all $x < \ell$, where $\mu_\ell^{(1)}(x)$ and $\mu_\hbar^{(1)}(x)$ denote the derivatives of $\mathbb{E}[Y_i|X_i = x, X_i < \ell, C_i = \ell]$ and $\mathbb{E}[Y_i|X_i = x, X_i < \ell, C_i = \hbar]$, respectively. This test can be implemented using several evaluation points, or using a summary statistic such as the supremum. Validity of this approach is also justified using nonparametric estimation and inference results in the literature, under regularity conditions.

4. Extrapolating the Effect of Loan Access on College Enrollment

We use our proposed methods to investigate the external validity of the ACCES program RD effects. As mentioned above, our sample has observations exposed to two cutoffs, $\ell = -850$ and $\hbar = -571$. We begin by extrapolating the effect to the point $\bar{x} = -650$; our focus is thus the effect of eligibility for ACCES on whether the student enrolls in a higher education program for the subpopulation exposed to cutoff 850 when their SABER 11 score is 650.

As described in Section 2.1, the observations facing cutoff ℓ correspond to years 2000–2008, whereas observations facing cutoff $\hbar$ correspond to years 2009 and 2010. Our identification assumption allows these two groups to differ in observable and unobservable characteristics so long as the difference between the conditional expectations of their control potential outcomes is constant as a function of the running variable. In addition, our approach relies on the assumption that the underlying population does not change over time (which is implicit in our notation). We offer empirical support for these assumptions in two ways. First, we implement the tests discussed in Section 3.3 to assess the plausibility of Assumption 2. In addition, Section SA-2 in the supplemental appendix shows that our results remain qualitatively unchanged when restricting the empirical analysis to the period 2007–2010, which reduces the (potential) heterogeneity of the underlying populations over time.

We begin by assessing the validity of our constant bias assumption with the methods described in Section 3.3. The results can be seen in Tables 3 and 4. Specifically, Table 3 reports results employing global polynomial regression, which does not reject the null hypothesis of parallel trends. Figure 4(a) offers a graphical illustration. Table 4 shows the results for the local polynomial approach, which again does not reject the null hypothesis. Additionally, Figure 4(b) plots the difference in derivatives (solid line) between groups estimated nonparametrically at ten evaluation points below ℓ, along with pointwise robust bias-corrected confidence intervals (dashed lines). The

Table 3. Parallel trends test: global polynomial approach.

	Estimate	p-value
Constant	5.534	0.159
Score	0.010	0.220
Score^2	0.000	0.245
$\mathbb{1}(C = \hbar)$	5.732	0.779
$\mathbb{1}(C = \hbar) \times$ Score	0.012	0.790
$\mathbb{1}(C = \hbar) \times$ Score^2	0.000	0.795
N	257	
F-test		0.919

NOTES: Global (quadratic) polynomial regression with interactions to test for parallel trends between control regression functions for low ($C = \ell$) and high ($C = \hbar$) cutoffs. Estimation and inference is conducted using standard parametric linear least squares methods. *F*-test refers to a joint significance test that the coefficients associated with $\mathbb{1}(C = \hbar)$, $\mathbb{1}(C = \hbar) \times$ Score and $\mathbb{1}(C = \hbar) \times$ Score^2 are simultaneously equal to zero.

Table 4. Parallel trends test: local polynomial approach.

			Robust BC inference	
	Estimate	Bw	*p*-value	95%CI
$\mu_\ell^{(1)}(\ell)$	−0.00025	58.9	0.986	[−0.0179, 0.0176]
$\mu_\hbar^{(1)}(\ell)$	0.00042	154.6	0.977	[−0.0015, 0.0014]
Difference	−0.00066		0.988	[−0.0180, 0.0177]

NOTES: Local polynomial methods for testing equality of first derivatives of control regression functions for low ($C = \ell$) and high ($C = \hbar$) cutoffs, over a grid of points below the low ($C = \ell$) cutoff. Estimation and robust bias corrected inference is conducted using methods in Calonico, Cattaneo, and Farrell (2018, 2020a), implemented via the general purpose software described in Calonico, Cattaneo, and Farrell (2019).

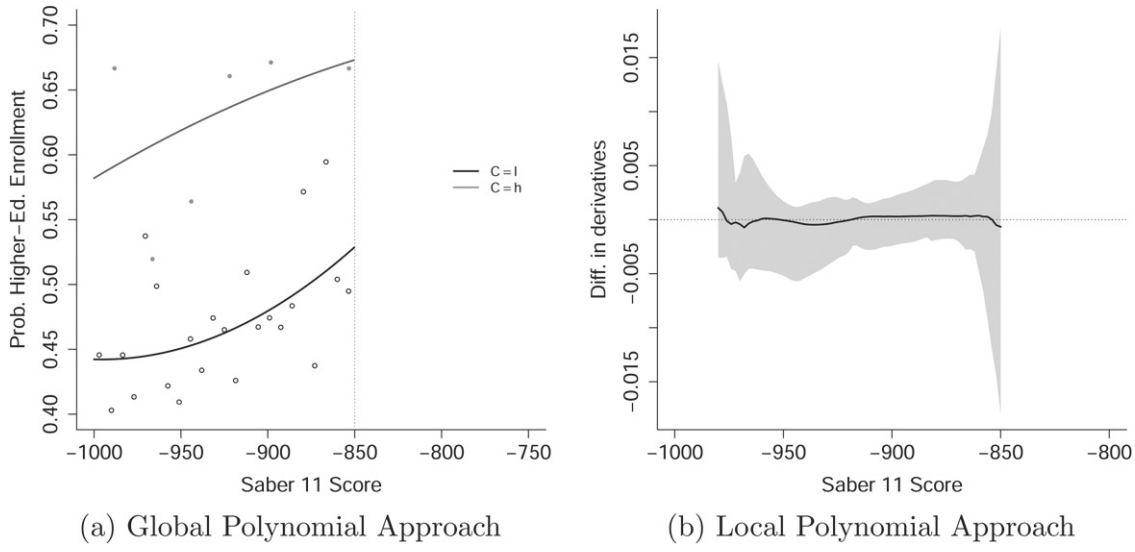

(a) Global Polynomial Approach (b) Local Polynomial Approach

Figure 4. Parallel trends test. NOTES: Panel (a) plots regression functions estimated using a quadratic global polynomial regression. Panel (b) plots the difference in derivatives at several points, estimated using a local quadratic polynomial regression (solid line). The gray area represents the RBC 95% (pointwise) confidence intervals.

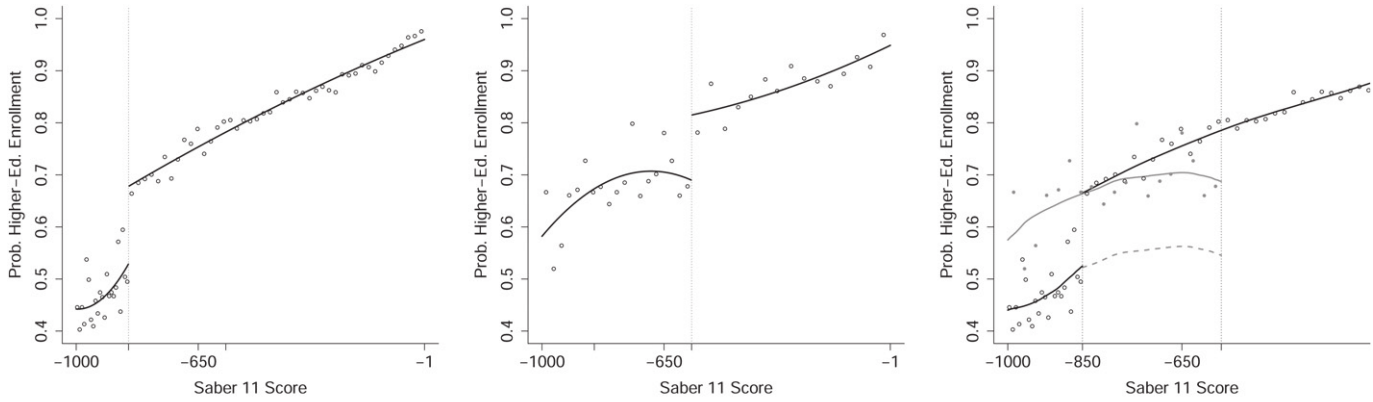

Figure 5. RD and extrapolation effects of ACCES loan eligibility on higher education enrollment. NOTES: Panels (a) and (b) show the RD plots for the cutoff-specific effects at the low and high cutoff, respectively. Panel (c) shows the nonparametric local-polynomial estimates of the regression functions for the low-cutoff (solid black line) and high-cutoff (gray line) groups. The dashed line represents the nonparametric local-polynomial estimate of the imputed regression function for control units facing the low cutoff.

figure reveals that the difference in derivatives is not significantly different from zero.

As discussed in Section 2.2 and Table 1, the pooled RD estimated effect is 0.125 with a RBC confidence interval of [0.012, 0.219]. The single-cutoff effect at −850 is 0.137 with 95% RBC confidence interval of [0.036, 0.232], and the effect at −571 is somewhat higher at 0.169, with 95% RBC confidence interval of [−0.039, 0.428]. These estimates based on single-cutoffs are illustrated in Figures 5(a) and (b), respectively.

In finite samples, the pooled estimate may not be a weighted average of the cutoff-specific estimates as it contains an additional term that depends on the bandwidth used for estimation and small sample discrepancies between the estimated slopes for each group. This is evident in Table 1, where the pooled estimate does not lie between the cutoff specific estimates. This additional term vanishes as the sample size grows and the bandwidths converge to zero, yielding the result in Cattaneo et al. (2016). To provide further evidence on the overall effect of the program, we also estimated a weighted average of cutoff-specific effects using estimated weights. This average effect equals 0.156 with a RBC confidence interval of [0.025, 0.314]. Since this estimate is a proper weighted average of cutoff-specific effects, it may give a more accurate assessment of the overall effect of the program.

The extrapolation results are illustrated in Figure 5(c) and reported in the last two panels of Table 1. At the −650 cutoff, the treated population exposed to cutoff −850 has an enrollment rate of 0.756, while the control population exposed to cutoff −571 has a rate of 0.706. This naive comparison, however, is likely biased due to unobservable differences between both subpopulations. The bias, which is estimated at the low cutoff −850, is −0.142, showing that the control population exposed to the −850 cutoff has lower enrollment rates at that point than the population exposed to the high cutoff −571 (0.525 vs. 0.667). The extrapolated effect in the last row corrects the naive comparison according to Theorem 1. The resulting extrapolated effect is $0.756 - (0.706 - 0.142) = 0.191$ with RBC confidence interval of [0.080, 0.336].

The choice of the point −650 is simply for illustration purposes, and indeed considering a set of evaluation points for extrapolation can give a much more complete picture of the impact of the program away from the cutoff point. In Figures 6(a) and (b), we conduct this analysis by estimating the extrapolated effect at 14 equidistant points between −840 and −580. The effects are statistically significant, ranging from around 0.14 to 0.25.

5. Simulations

We report results from a simulation study aimed to assess the performance of the local polynomial methods described in Section 3.2. We construct $\mu_{0,\ell}(x)$ as a fourth-order polynomial where the coefficients are calibrated using the data from our empirical application, and $\mu_{0,\ell}(x) = \mu_{0,\ell}(x) + \Delta$. Based on our empirical findings, we set $\Delta = -0.14$ and an extrapolated treatment effect of $\tau_\ell(\bar{x}) = 0.19$. We consider three sample sizes: $N = 1000$ ("small N"), $N = 2000$ ("moderate N"), and $N = 5000$ ("large N"). To assess the effect of unbalanced sample sizes across evaluation points/cutoffs, our simulation model ensures that some evaluation points/cutoffs have fewer observations than others. In particular, the available sample size to estimate $\mu_\ell(\ell)$ is always less than a third of the sample size available to estimate $\mu_\ell(\bar{x})$. We provide all details in the supplemental appendix to conserve space.

The results are shown in Table 5. The robust bias-corrected 95% confidence interval for $\tau_\ell(\bar{x})$ has an empirical coverage rate of around 91% in the "small N" case. This is because one of the parameters, $\mu_\ell(\ell)$, is estimated using very few observations. The empirical coverage rate increases slightly to 92% in the "moderate N" case, and to about 94% in the "large N" case. In sum, in our Monte Carlo experiment, we find that local polynomial methods can yield estimators with little bias and RBC confidence intervals with accurate coverage rates for RD extrapolation.

6. Extension to Fuzzy RD Designs

The main idea underlying our extrapolation methods can be extended in several directions that may be useful in other applications. We briefly discuss an extension to fuzzy RD designs employing a continuity-based approach. In the supplemental

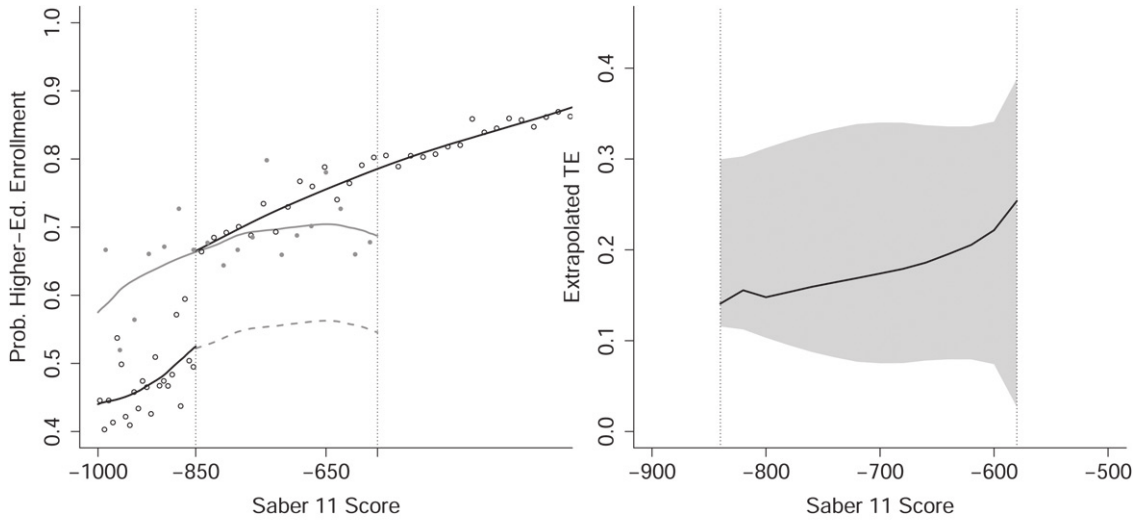

(a) Estimated regression functions. (b) Extrapolation at multiple points.

Figure 6. Extrapolation treatment effects. NOTES: Panel (a) shows local-linear estimates of the regression functions using an IMSE-optimal bandwidth for the control and treated groups facing cutoff ℓ (black solid lines) and for the control group facing cutoff $\hbar$ (solid gray line). The dashed line represents the extrapolated regression function for the control group facing cutoff ℓ. Panel (b) shows local-linear extrapolation treatment effects estimates at 14 equidistant evaluation points between −840 and −580. The gray area represents the RBC 95% (pointwise) confidence intervals.

Table 5. Simulation results.

		Point estimation			RBC inference
	Eff. *N*	Bias	Var	RMSE	Cov.(95%)
Small *N*					
$\mu_\ell(\ell)$	23.4	0.001	0.0213	0.146	0.890
$\mu_\ell(\bar{x})$	78.0	−0.007	0.0015	0.039	0.937
$\mu_\hbar(\ell)$	140.9	−0.001	0.0008	0.029	0.945
$\mu_\hbar(\bar{x})$	101.6	−0.010	0.0012	0.036	0.937
$\tau_\ell(\bar{x})$		0.002	0.0247	0.157	0.905
Moderate *N*					
$\mu_\ell(\ell)$	43.0	−0.001	0.0134	0.116	0.905
$\mu_\ell(\bar{x})$	136.9	−0.005	0.0008	0.029	0.945
$\mu_\hbar(\ell)$	279.4	−0.000	0.0004	0.020	0.950
$\mu_\hbar(\bar{x})$	213.7	−0.012	0.0005	0.026	0.949
$\tau_\ell(\bar{x})$		0.008	0.0150	0.123	0.917
Large *N*					
$\mu_\ell(\ell)$	108.6	0.001	0.0050	0.071	0.933
$\mu_\ell(\bar{x})$	288.0	−0.003	0.0004	0.020	0.945
$\mu_\hbar(\ell)$	681.9	−0.000	0.0002	0.013	0.953
$\mu_\hbar(\bar{x})$	517.5	−0.012	0.0002	0.019	0.951
$\tau_\ell(\bar{x})$		0.007	0.0058	0.076	0.939

NOTES: Local polynomial regression estimation with MSE-optimal bandwidth selectors and robust bias corrected inference. See Calonico, Cattaneo, and Titiunik (2014) and Calonico, Cattaneo, and Farrell (2018) for methodological details, and Calonico et al. (2017) and Cattaneo, Titiunik, and Vazquez-Bare (2020a) for implementation. "Eff. *N*" indicates the effective sample size, that is, the sample size within the MSE-optimal bandwidth. Results from 10,000 simulations.

appendix, we discuss other extensions: covariate adjustments (i.e., ignorable cutoff bias), score adjustments (i.e., polynomial-in-score cutoff bias), many multiple cutoffs, and multiple scores and geographic RD designs.

In the fuzzy RD design, treatment compliance is imperfect, which is common in empirical applications. For simplicity, we focus on the case of one-sided (treatment) noncompliance: units assigned to the control group comply with their assignment but units assigned to treatment status may not. This case is relevant for a wide array of empirical applications in which program administrators are able to successfully exclude units from the treatment, but cannot force units to actually comply with it.

We employ the fuzzy multi-cutoff RD framework of Cattaneo et al. (2016), which builds on the canonical framework of Angrist, Imbens, and Rubin (1996). Let $D_i(x,c)$ be the binary treatment indicator and $\underline{x} \leq \bar{x}$. We define compliers as units with $D_i(\underline{x},c) < D_i(\bar{x},c)$, always-takers as units with $D_i(\underline{x},c) = D_i(\bar{x},c) = 1$, never-takers as units with $D_i(\underline{x},c) = D_i(\bar{x},c) = 0$, and defiers as units with $D_i(\underline{x},c) > D_i(\bar{x},c)$. We assume the following conditions:

Assumption 3 (Fuzzy RD design).

1. Continuity: $\mathbb{E}[Y_i(0)|X_i = x, C_i = c]$ and $\mathbb{E}[(Y_i(1) - Y_i(0))D_i(x,c)|X_i = x, C_i = c]$ are continuous in x for all c.
2. Constant bias: $B(\ell') = B(x)$ for all $x \in (\ell', \hbar)$.
3. Monotonicity: $D_i(\underline{x},c) \leq D_i(\bar{x},c)$ for all i and for all $\underline{x} \leq \bar{x}$.
4. One-sided noncompliance: $D_i(x,c) = 0$ for all $x < c$.

The conditions are standard in the fuzzy RD literature and used to identify the local average treatment effect (LATE), which is the treatment effect for units that comply with the RD assignment. The following result shows how to recover a LATE-type extrapolation parameter in this fuzzy RD setting.

Theorem 2. Under Assumption 3,

$$\frac{\mu_\ell(\bar{x}) - [\mu_\hbar(\bar{x}) + B(\ell')]}{\mathbb{E}[D_i|X_i = x, C_i = \ell']} = \mathbb{E}[Y_i(1) - Y_i(0)|X_i = x, C_i = \ell', D_i(x,\ell') = 1].$$

The left-hand side can be interpreted as an "adjusted" Wald estimand, where the adjustment allows for extrapolation away from the cutoff point ℓ'. More precisely, this theorem shows that under one-sided (treatment) noncompliance we can recover the average extrapolated effect on compliers by dividing the

adjusted intention-to-treat parameter by the proportion of compliers.

7. Conclusion

We introduced a new framework for the extrapolation of RD treatment effects when the RD design has multiple cutoffs. Our approach relies on the assumption that the average outcome difference between control groups exposed to different cutoffs is constant over a chosen extrapolation region. Our method does not require any information external to the design, and can be used whenever two or more cutoffs are used to assign the treatment for different subpopulations, which is a very common feature in many RD applications. Our main extrapolation idea can also be used in settings with more than two cutoffs, multi-scores RD designs (Papay, Willett, and Murnane 2011; Reardon and Robinson 2012), and geographic RD designs (Keele and Titiunik 2015). In addition, our main idea can be extended to the RD local randomization framework introduced by Cattaneo, Frandsen, and Titiunik (2015) and Cattaneo, Titiunik, and Vazquez-Bare (2017). These additional results are reported in the supplemental appendix for brevity.

Supplementary Materials

The supplemental appendix includes a description of RD empirical papers with multiple cutoffs, further extensions and generalizations of our main methodological results, and omitted proofs and derivations. In addition, it outlines an extension of our methods to the RD local randomization framework.

Acknowledgments

We are very grateful to Fabio Sanchez and Tatiana Velasco for sharing the dataset used in the empirical application. We also thank Josh Angrist, Sebastian Calonico, Sebastian Galiani, Nicolas Idrobo, Xinwei Ma, Max Farrell, and seminar participants at various institutions for their comments. We also thank the co-editor, Regina Liu, an associate editor, and a reviewer for their comments.

Funding

Cattaneo and Titiunik gratefully acknowledge financial support from the National Science Foundation (SES 1357561).

References

Abadie, A. (2005), "Semiparametric Difference-in-Differences Estimators," *Review of Economic Studies*, 72, 1–19. [1944]

Abadie, A., and Cattaneo, M. D. (2018), "Econometric Methods for Program Evaluation," *Annual Review of Economics*, 10, 465–503. [1944]

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996), "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 91, 444–455. [1950]

Angrist, J. D., and Rokkanen, M. (2015), "Wanna Get Away? Regression Discontinuity Estimation of Exam School Effects Away From the Cutoff," *Journal of the American Statistical Association*, 110, 1331–1344. [1944]

Bertanha, M. (2020), "Regression Discontinuity Design With Many Thresholds," *Journal of Econometrics* (forthcoming). [1944]

Bertanha, M., and Imbens, G. W. (2020), "External Validity in Fuzzy Regression Discontinuity Designs," *Journal of Business & Economic Statistics* (forthcoming). [1944]

Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2018), "On the Effect of Bias Estimation on Coverage Accuracy in Nonparametric Inference," *Journal of the American Statistical Association*, 113, 767–779. [1943,1947,1948,1950]

——— (2019), "`nprobust`: Nonparametric Kernel-Based Estimation and Robust Bias-Corrected Inference," *Journal of Statistical Software*, 8, 1–33. [1947,1948]

——— (2020a), "Coverage Error Optimal Confidence Intervals for Local Polynomial Regression," arXiv no. 1808.01398. [1943,1947,1948]

——— (2020b), "Optimal Bandwidth Choice for Robust Bias Corrected Inference in Regression Discontinuity Designs," *Econometrics Journal*, 23, 192–210. . [1943,1947]

Calonico, S., Cattaneo, M. D., Farrell, M. H., and Titiunik, R. (2017), "`rdrobust`: Software for Regression Discontinuity Designs," *Stata Journal*, 17, 372–404. [1943,1947,1950]

——— (2019), "Regression Discontinuity Designs Using Covariates," *Review of Economics and Statistics*, 101, 442–451. [1943,1947]

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014), "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82, 2295–2326. [1943,1947,1950]

——— (2015), "Optimal Data-Driven Regression Discontinuity Plots," *Journal of the American Statistical Association*, 110, 1753–1769. [1943]

Cattaneo, M. D., Frandsen, B., and Titiunik, R. (2015), "Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate," *Journal of Causal Inference*, 3, 1–24. [1944,1951]

Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2019), *A Practical Introduction to Regression Discontinuity Designs: Foundations*, Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge: Cambridge University Press. [1944]

——— (2020), *A Practical Introduction to Regression Discontinuity Designs: Extensions*, Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge: Cambridge University Press (to appear). [1944]

Cattaneo, M. D., Keele, L., Titiunik, R., and Vazquez-Bare, G. (2016), "Interpreting Regression Discontinuity Designs With Multiple Cutoffs," *Journal of Politics*, 78, 1229–1248. [1942,1943,1944,1949,1950]

Cattaneo, M. D., Titiunik, R., and Vazquez-Bare, G. (2017), "Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality," *Journal of Policy Analysis and Management*, 36, 643–681. [1944,1951]

——— (2020a), "Analysis of Regression Discontinuity Designs With Multiple Cutoffs or Multiple Scores," *Stata Journal*, forthcoming. [1943,1950]

——— (2020b), "The Regression Discontinuity Design," in *Handbook of Research Methods in Political Science and International Relations*, eds. L. Curini and R. J. Franzese, Thousand Oaks, CA: SAGE Publications. [1944]

Dong, Y., Lee, Y.-Y., and Gou, M. (2020), "Regression Discontinuity Designs With a Continuous Treatment," SSRN Working Paper No. 3167541. [1947]

Dong, Y., and Lewbel, A. (2015), "Identifying the Effect of Changing the Policy Threshold in Regression Discontinuity Models," *Review of Economics and Statistics*, 97, 1081–1092. [1944]

Fan, J., and Gijbels, I. (1996), *Local Polynomial Modelling and Its Applications*, New York: Chapman & Hall/CRC. [1944]

Ganong, P., and Jäger, S. (2018), "A Permutation Test for the Regression Kink Design," *Journal of the American Statistical Association*, 113, 494–504. [1944,1947]

Hyytinen, A., Meriläinen, J., Saarimaa, T., Toivanen, O., and Tukiainen, J. (2018), "When Does Regression Discontinuity Design Work? Evidence From Random Election Outcomes," *Quantitative Economics*, 9, 1019–1051. [1947]

Imbens, G. W., and Lemieux, T. (2008), "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142, 615–635. [1944]

Imbens, G. W., and Rubin, D. B. (2015), *Causal Inference in Statistics, Social, and Biomedical Sciences*, Cambridge: Cambridge University Press. [1944]

Keele, L. J., and Titiunik, R. (2015), "Geographic Boundaries as Regression Discontinuities," *Political Analysis*, 23, 127–155. [1951]

Keele, L. J., Titiunik, R., and Zubizarreta, J. (2015), "Enhancing a Geographic Regression Discontinuity Design Through Matching to Estimate

the Effect of Ballot Initiatives on Voter Turnout," *Journal of the Royal Statistical Society*, Series A, 178, 223–239. [1944]

Li, F., Mattei, A., and Mealli, F. (2015), "Evaluating the Causal Effect of University Grants on Student Dropout: Evidence From a Regression Discontinuity Design Using Principal Stratification," *Annals of Applied Statistics*, 9, 1906–1931. [1944]

Mealli, F., and Rampichini, C. (2012), "Evaluating the Effects of University Grants by Using Regression Discontinuity Designs," *Journal of the Royal Statistical Society*, Series A, 175, 775–798. [1944]

Melguizo, T., Sanchez, F., and Velasco, T. (2016), "Credit for Low-Income Students and Access to and Academic Performance in Higher Education in Colombia: A Regression Discontinuity Approach," *World Development*, 80, 61–77. [1941,1942,1943]

Papay, J. P., Willett, J. B., and Murnane, R. J. (2011), "Extending the Regression-Discontinuity Approach to Multiple Assignment Variables," *Journal of Econometrics*, 161, 203–207. [1951]

Reardon, S. F., and Robinson, J. P. (2012), "Regression Discontinuity Designs With Multiple Rating-Score Variables," *Journal of Research on Educational Effectiveness*, 5, 83–104. [1951]

Rokkanen, M. (2015), "Exam Schools, Ability, and the Effects of Affirmative Action: Latent Factor Extrapolation in the Regression Discontinuity Design," Unpublished Manuscript. [1944]

Rosenbaum, P. R. (2010), *Design of Observational Studies*, New York: Springer. [1944]

Sekhon, J. S., and Titiunik, R. (2016), "Understanding Regression Discontinuity Designs as Observational Studies," *Observational Studies*, 2, 174–182. [1944]

——— (2017), "On Interpreting the Regression Discontinuity Design as a Local Experiment," in *Regression Discontinuity Designs: Theory and Applications*, Advances in Econometrics (Vol. 38), eds. M. D. Cattaneo and J. C. Escanciano, Bingley: Emerald Group Publishing, pp. 1–28. [1944]

Wing, C., and Cook, T. D. (2013), "Strengthening the Regression Discontinuity Design Using Additional Design Elements: A Within-Study Comparison," *Journal of Policy Analysis and Management*, 32, 853–877. [1944]