Discussion Paper Series
RIEB
Kobe University

DP2026-20

On Extrapolation of Treatment Effects in Multiple-Cutoff Regression Discontinuity Designs

RIEB Junior Research Fellow
Yuta OKAMOTO
Yuuki OZAKI

June 5, 2026

* This Discussion Paper won the Kanematsu Prize (FY 2025).

Research Institute for Economics and Business Administration
Kobe University
2-1 Rokkodai, Nada, Kobe 657-8501 JAPAN

On Extrapolation of Treatment Effects in Multiple-Cutoff Regression Discontinuity Designs

Yuta Okamoto[♯] and Yuuki Ozaki[♭]

Graduate School of Economics, Kyoto University

Abstract

This paper investigates how to learn treatment effects away from the cutoff point in multiple-cutoff regression discontinuity designs. Using a microeconomic model, we demonstrate that the parallel-trend type assumption proposed in the literature is justified when cutoff positions are assigned as if randomly and the running variable is non-manipulable (e.g., parental income). However, when the running variable is partially manipulable (e.g., test scores), extrapolations based on that assumption can be biased. As a complementary strategy, we propose a novel partial identification approach based on empirically motivated assumptions. We also develop a uniform inference procedure and provide two empirical illustrations.

Keywords: Decision model, external validity, partial identification, regression discontinuity designs

JEL Classification: C14, C21, D00, D84

[♯]yuta.okamoto@r.hit-u.ac.jp (corresponding author). This version: 5 June 2026. This paper won the 2025 Kanematsu Prize from the Research Institution for Economics and Business Administration, Kobe University. Okamoto is grateful for financial support from JST SPRING, Grant Number JPMJSP2110.

1 Introduction

Regression discontinuity (RD) designs are among the most credible quasi-experimental methods for identifying causal effects. In the RD framework, treatment status changes discontinuously at a known threshold, allowing for identification of the treatment effect at the cutoff, provided a mild smoothness assumption holds (Hahn et al., 2001). This feature provides RD designs with strong internal validity.

However, this internal validity stems from the local nature of the RD designs, and hence, their external validity often remains uncertain. This is a primary limitation of RD designs, as the treatment effect estimated at the single cutoff is not necessarily the only parameter of interest. Researchers are often interested in the treatment effects away from the cutoff to guarantee a certain external validity of the local estimate (Cerulli et al., 2017). Furthermore, in some applications, policymakers may wish to understand the treatment effect at specific points or regions other than the current cutoff point—e.g., when considering moving the current cutoff to a different level (Dong and Lewbel, 2015). Standard RD designs, however, provide limited information for addressing such questions, and the extrapolation of RD treatment effects remains a "crucial open question" (Abadie and Cattaneo, 2018).

To address the issue of weaker external validity in RD designs, an additional source of information is required. A promising avenue is to leverage the presence of multiple cutoffs, a scenario frequently encountered in empirical work (Bertanha, 2020; Cattaneo et al., 2016). For example, eligibility thresholds for scholarships often vary by student background, such as gender, race, or cohort.

Although such multiple cutoffs have often been normalized in empirical studies, Cattaneo et al. (2021) recently proposed a novel strategy to identify the treatment effects away from the cutoff point by effectively utilizing these multiple cutoffs. Their identification strategy is as follows: Let $\mu_{0,l}(x)$ and $\mu_{1,l}(x)$ represent the conditional expectation functions under controlled and treated status, respectively, for a group with cut-

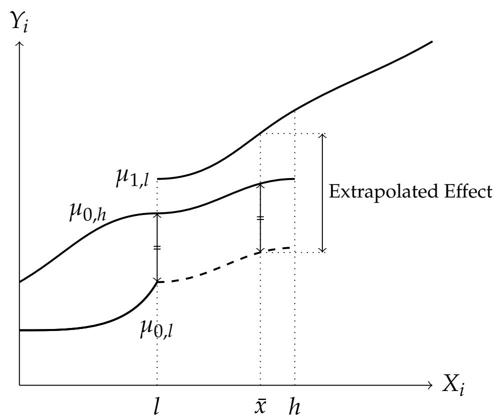

Figure 1: Extrapolation under Constant Bias

off l (see Figure 1). In standard RD, we typically identify and estimate $\mu_{1,l}(l) - \mu_{0,l}(l)$. The fundamental challenge in extrapolation is that, while $\mu_{1,l}(x)$ can be observed for $x > l$, $\mu_{0,l}(x)$ cannot. However, when another group with a higher cutoff $h(> l)$ is present, their regression function under control status, $\mu_{0,h}(x)$, can be observed for $x \in (l, h)$. Cattaneo et al. (2021) identifies the treatment effects for group l at $\bar{x} \in (l, h)$ as $\mu_{1,l}(\bar{x}) - \mu_{0,h}(\bar{x}) + \{\mu_{0,h}(l) - \mu_{0,l}(l)\}$—referred to as "Extrapolated Effect" in Figure 1—by introducing a "parallel trend"-type condition: that $\mu_{0,h}(x) - \mu_{0,l}(x)$ is constant over (l, h).

This parallel trend type assumption, termed the *constant bias assumption*, is undoubtedly useful for drawing additional policy insights from RD studies. Considering the popularity of difference-in-differences (DID) analysis in empirical economics, this extrapolation strategy has the potential to be widely used. At the same time, it appears somewhat ad hoc, and its empirical motivation and validity have not been fully explored. Consequently, researchers lack clear guidance on how to assess its plausibility, which may have restricted the potential to conduct extrapolation analysis relying on this assumption.

Our first goal is to investigate the plausibility of the constant bias assumption—when it holds and when it fails—thereby clarifying its scope of applicability. To do this, building on Fudenberg and Levine (2022), we develop an economic model that links multiple-cutoff RD designs to the decision-making behavior of rational agents.

The key implications of our analysis are twofold. First, the constant bias assumption is likely to hold when the distribution of the unobserved characteristics is similar across groups and the running variable is entirely *non-manipulable* by agents (e.g., parental income level). In other words, when the running variable is non-manipulable, the assumption is plausible if the cutoff position is assigned as if at random. This insight will be helpful for empirical researchers, as it links the validity of the constant bias assumption—that is, a functional form assumption—to a familiar random assignment-like condition often invoked in the causal inference literature. For example, if the income threshold for financial aid is revised from the previous year but unobserved characteristics are assumed to remain stable across cohorts, then the constant bias assumption among these two cohorts will plausibly hold.

Unfortunately, however, our model suggests that when the running variable is *partially manipulable* in the sense of McCrary (2008) (e.g., test scores), this assumption may fail even when the only difference between groups is the cutoff positions. The intuition behind this result is simple: if agents optimally choose their level of effort—which affects both the running variable and the outcome—then a shift in the cutoff position can induce a change in effort. This, in turn, alters the distribution of both the running variable (e.g., test scores) and the outcome variable (e.g., initial wage). These distributional shifts may undermine the plausibility of the constant bias assumption. Consequently, extrapolated estimates may be biased, even when the groups are otherwise identical.

To provide a practical solution for extrapolation when the constant bias assumption may fail, we introduce an alternative framework that is particularly well-suited to multi-cutoff RD designs and broadly applicable, especially in educational settings, one of the most common applications of RD designs (Villamizar-Villegas et al., 2022). Our strategy relies on a set of qualitative assumptions that are often easier to assess in practice. Specifically, our identification approach leverages a commonly employed *monotonicity assumption* for $\mu_{0,c}(x), c \in \{l, h\}$, along with a *dominance assumption*, $\mu_{0,h}(x) \geq \mu_{0,l}(x)$. As noted in Babii and Kumar (2023), the monotonicity assumption is often reasonable in RD applications. The dominance assumption is also plausible in several multi-cutoff RD settings, particularly when cutoffs are designed to reduce inequities or reflect pre-existing differences between groups. For example, consider an affirmative action scenario in which scholarship thresholds are relaxed for students from disadvantaged backgrounds—e.g., those from high-poverty regions (Melguizo et al., 2016). In such cases, it is reasonable to assume that the average outcome function for disadvantaged students is lower than that of more advantaged students.

Under these assumptions, we derive partial identification results that remain valid even when the constant bias assumption fails. These bounds offer a practical and robust alternative, complementing the point identification result of Cattaneo et al. (2021). In addition, we develop estimation and uniform inference procedures tailored to these bounds. The empirical relevance of our approach is illustrated through applications to two empirical examples.

Plan of the Article

The remainder of this article is organized as follows. The rest of this section reviews the related literature. In Section 2, we investigate the applicability and limitations of the constant bias assumption of Cattaneo et al. (2021) using an economic model. Section 3 proposes an alternative identification strategy under the sharp RD design. Empirical illustrations are provided in Section 4. Section 5 concludes this article. Proofs are collected in Appendix A.

Related Literature

The methodological RD literature has substantially expanded over the past few decades. In the standard single-cutoff RD setting, the theoretical foundations have been developed for nonparametric identification (Hahn et al., 2001), point estimation (Imbens and Kalyanaraman, 2011), robust bias-corrected inference procedure (Calonico et al., 2014, 2018, 2020), falsification tests (Arai et al., 2022; Bugni and Canay, 2021; Canay and Kamat, 2017; Cattaneo et al., 2020; Fusejima et al., 2025; McCrary, 2008), among other important extensions. For a comprehensive recent review of these and related developments, see Cattaneo and Titiunik (2022).

While these works focus on the internal validity of RD designs, recent studies have begun to address concerns regarding external validity. Angrist and Rokkanen (2015) proposes an extrapolation strategy applicable when the potential outcomes and the running variable are mean-independent conditional on some covariates. Bertanha and Imbens (2020) considers the case where the potential outcomes and compliance types are independent conditional on the running variable. Mehta (2019) derives bounds on the average treatment effect by assuming that the policymaker knows the treatment effect and sets the cutoff optimally. Deaner and Kwon (2025) proposes an extrapolation strategy that is applicable when an additional covariate satisfying the so-called comonotonicity assumption is available. Although useful, these strategies tend to require empirically strong assumptions or highly informative covariates, which may limit their practical applicability.

By contrast, Dong and Lewbel (2015) proposes a marginal extrapolation strategy that can be applied under only a mild smoothness condition. However, the extrapolation away from the cutoff point is not possible under their framework. Okamoto (2026) extends the Dong and Lewbel approach by introducing the locally linear treatment effects assumption. Although his identification result permits extrapolation farther from the cutoff, the confidence intervals widen rapidly as the extrapolation point moves away from the cutoff, so in practice the analysis remains most informative relatively close to the cutoff.

Cattaneo et al. (2021) explicitly leverages multiple cutoffs to extrapolate treatment effects, relying on a constant bias assumption. This approach serves as the foundation for the discussion in the present paper. Specifically, we connect their constant bias assumption to an economic model to clarify its applicability, and we further propose a complementary strategy by introducing a set of different assumptions based on an empirical motivation. Relatedly, Sun (2023) relaxes the constant bias assumption by introducing a bias bounding constant, following Manski and Pepper (2018). Her approach is also useful when the constant bias assumption seems implausible in multi-cutoff RD settings. One crucial difference between her strategy and ours is that the former requires researchers to specify a theoretical smoothness bounding constant, which can be a non-trivial task in practice, as emphasized in Cattaneo and Titiunik (2022, Section 3.4.3). In contrast, our identification result does not require researchers to specify such constants. Instead, it relies only on a set of qualitative assumptions, whose plausibility is often easier to assess in practice.

This paper is related to the literature on microeconomic analysis of econometric methods. Marx et al. (2024) analyzes the validity of the parallel trend assumption in the DID framework by embedding it within a dynamic choice model of rational agents. Motivated by this work, our paper interprets the constant bias assumption of Cattaneo et al. (2021) through the lens of individual decision-making in an RD environment, aiming to better understand when this assumption is empirically plausible. To this aim, we build on the work by Fudenberg and Levine (2022), which originally examines how agents' learning behavior influences causal estimands, showing that such effects are neutral in an RD design.

Finally, our work also contributes to the partial identification literature (e.g., Manski, 1997). In the context of program evaluation, numerous bounding strategies have been developed. Manski and Pepper (2018) and Rambachan and Roth (2023), for example, consider bounds on DID effects that are still valid when the parallel trend assumption is not satisfied. Within the RD context, Gerard et al. (2020) provides bounds that are robust to manipulation of the running variable.

2 The Constant Bias Assumption

2.1 Econometric Setup

We begin by introducing the notion of constant bias assumption using econometric terminology. For simplicity, we focus on a two-cutoff sharp RD design.

Let Y_i denote the outcome variable, X_i the running variable, $C_i \in \{l, h\}$ $(l < h)$ the cutoff, and $D_i \in \{0, 1\}$ be a treatment indicator, which takes one if i is treated. Treatment assignment follows the rule: $D_i = \mathbf{1}\{X_i \geq C_i\}$. Let $Y_i(d)$ denote the potential outcome under treatment status $D_i = d \in \{0, 1\}$. Define the conditional expectations and treatment effect functions as

$$\mu_{d,c}(x) := \mathbb{E}\left[Y_i(d)|X_i = x, C_i = c\right], \text{ and } \tau_c(x) := \mathbb{E}\left[Y_i(1) - Y_i(0)|X_i = x, C_i = c\right].$$

In typical RD settings, we identify $\tau_c(c)$ under a mild continuity assumption at c (Hahn et al., 2001):

Assumption 2.1 (Continuity). *$\mu_{d,c}(x)$ is continuous at $x = c$.*

Yet, researchers are often interested in treatment effects at other points, either to assess the external validity of the local RD estimate or to learn about subpopulations far from the cutoff (as discussed in Section 1). These parameters, however, are generally not identified under the standard continuity assumption alone.

To overcome this limitation, Cattaneo et al. (2021) propose the following assumption:

Assumption 2.2 (Constant Bias). *Let $B(x) = \mu_{0,h}(x) - \mu_{0,l}(x)$. It holds that $B(l) = B(x)$ for all $x \in (l, h)$.*

Under Assumptions 2.1 and 2.2, Cattaneo et al. (2021, Theorem 1) provides the following identification result:

$$\tau_{l,\mathrm{CB}}(\bar{x}) = \mu_{1,l}(\bar{x}) - \mu_{0,h}(\bar{x}) + B(l), \ \bar{x} \in (l, h), \tag{2.1}$$

where the CB stands for "Constant Bias."

The continuity assumption is a weak requirement. Hence, the key identifying condition is Assumption 2.2, namely the *constant bias assumption*. When, then, will this constant bias assumption be justified? Clarifying the conditions under which this assumption holds is crucial for understanding the scope and limitations of Cattaneo et al.'s (2021) result. When valid, this assumption provides a powerful extrapolation strategy, potentially yielding richer policy implications. However, their strategy can lead to a biased estimate if the regression functions of the two groups exhibit a different pattern.

To investigate the plausibility of this assumption, we turn in the next subsection to an economic model of a rational agent's behavior in a multi-cutoff RD environment, and examine its implications for Assumption 2.2.

2.2 A Decision Model of the Multi-Cutoff RD Environment

2.2.1 Economic Setup

To facilitate intuitive understanding, we describe the model using an educational setting—one of the most common applications—although the framework will apply broadly to other settings.[1]

Building on Fudenberg and Levine (2022), we assume that agents decide how much costly effort e_i to invest, influencing their short-term outcome S_i and future outcome $Y_i(0)$ under the control state. For example, in a typical educational setting, student i inputs their study effort e_i to achieve a higher test score S_i and higher future earnings $Y_i(0)$. In reality, the realized outcomes of S_i and $Y_i(0)$ are not solely determined by the effort and are also

[1]Depending on the empirical context, the model developed below may not be directly applicable. However, we believe that our analysis nonetheless provides potentially valuable insight for such cases. In particular, in settings that can be formulated as some agent's decision problem, a researcher could develop and analyze a variant of our model to examine the validity of the constant bias assumption in their context as well.

affected by stochastic shocks. We model this as follows:

$$Y_i(0) = y(e_i) + \gamma \mathbf{1}\{C_i = h\} + \eta_i^y, \text{ and } S_i = s(e_i) + \eta_i^s,$$

where $y(\cdot)$ and $s(\cdot)$ are structural functions, η_i^y and η_i^s are zero-mean stochastic errors that are independent of other factors, and γ represents a group-level difference in $Y_i(0)$. While our formulation is inspired by Fudenberg and Levine (2022), we depart from their framework in two ways that are particularly relevant for the questions we study: we do not assume linearity of $y(\cdot)$ and $s(\cdot)$ in e_i, and we allow the running variable to differ from S_i.

Effort incurs a cost $K(e_i, \epsilon_i)$, which also depends on the innate ability ϵ_i, known to the agent i. Thus, the cost of a given effort level varies across individuals.

Agents anticipate that they can receive an additional benefit $\tilde{\tau}_i$ in the future if their running variable X_i—which may or may not coincide with S_i—exceeds a predetermined threshold C_i, known to the individual. This $\tilde{\tau}_i$ represents the agent's subjective belief about the benefit from the treatment, which may differ from the actual one.

In this setup, we assume an agent decides the amount of effort to maximize expected utility:

$$\max_{e_i} \mathbb{E}\left[v\left(s(e_i) + \eta_i^s\right) - K(e_i, \epsilon_i) + \beta\left\{y(e_i) + \gamma \mathbf{1}\{C_i = h\} + \eta_i^y + \tilde{\tau}_i \mathbf{1}\{i \text{ is treated}\}\right\}\right],$$

where $v(\cdot)$ represents the utility derived from the short-run outcome $S_i = s(e_i) + \eta_i^s$, β denotes the discount factor, and the expectation is taken with respect to η_i^s and η_i^y. This is simplified as

$$\max_{e_i} \left[u\left(s(e_i)\right) - K(e_i, \epsilon_i) + \beta\left\{y(e_i) + \tilde{\tau}_i \mathbb{P}\left[X_i \geq C_i\right]\right\}\right], \tag{2.2}$$

where $u\left(s(e_i)\right) = \mathbb{E}\left[v\left(s(e_i) + \eta_i^s\right)\right]$.

2.2.2 Decision Problem and the Running Variable Manipulability

The optimization problem in Equation (2.2) reveals that the agent's decision depends on $\mathbb{P}[X_i \geq C_i]$—that is, the probability of crossing the threshold. The agent interprets this probability in a different way depending on whether the running variable X_i is influenced by their effort. Consider two examples:

- If financial aid eligibility is based on family wealth level, which is fixed and unaffected by student effort. In this case, $\mathbb{P}[X_i \geq C_i]$ is either 0 or 1—fully deterministic.
- If aid eligibility is determined by a test score, $X_i = S_i$, then the probability of crossing the threshold depends on effort via $s(e_i)$.

Following McCrary (2008), we say that the running variable is *partially manipulable* when it is under the agent's control but also affected by an idiosyncratic element, i.e., $X_i = S_i = s(e_i) + \eta_i^s$ in our formulation.[2] In contrast, we say that the running variable is *non-manipulable* when it is not a function of the agent's effort.

Thus, agents face one of two decision problems depending on whether the running variable is manipulable:

$$\max_{e_i} \left[u\left(s(e_i)\right) - K(e_i, \epsilon_i) + \beta y(e_i) \right] \text{ if } X_i \text{ is non-manipulable,} \tag{2.3a}$$

$$\max_{e_i} \left[u\left(s(e_i)\right) - K(e_i, \epsilon_i) + \beta \{ y(e_i) + \tilde{\tau}_i \mathbb{P}\left[s(e_i) + \eta_i^s \geq C_i\right] \} \right] \text{ if } X_i = S_i. \tag{2.3b}$$

We will write the optimal level of effort by e_i^*.

[2]Note that the so-called score manipulation problem, which undermines the validity of standard RD designs, does not arise when the running variable is only partially manipulable (Lee, 2008; McCrary, 2008). We also note that this article does not consider the case in which the running variable is *completely* manipulable, as such manipulability invalidates the identification of RD designs.

2.3 Main Results

We are now in a position to analyze the constant bias assumption within our model environment.

2.3.1 Non-Manipulable Running Variable Case

We begin with the case where the running variable is non-manipulable, corresponding to equation (2.3a). An immediate implication from (2.3a) is that the optimal effort level e_i^* does not depend on the cutoff C_i. This means that differences across groups can only arise from differences in the distribution of ability ϵ_i (through the cost function) and from the group-level shift γ in outcomes. This leads to the following result:

Proposition 1. *Suppose that the running variable is non-manipulable. If the distribution of ϵ_i is identical between those with $C_i = l$ and those with $C_i = h$, Assumption 2.2 holds.* ♣

Statistically, Proposition 1 says that the constant bias assumption holds if $\epsilon_i \perp\!\!\!\perp C_i$, that is, when cutoff assignment is as good as random. This interpretation is appealing for empirical researchers, as it links the validity of the constant bias assumption to a familiar random assignment-like condition often invoked in the causal inference literature. It also implies that a conditional version of the assumption aligns with the logic of unconfoundedness, making such an extension a natural one (see Remark 1 below).

In this view, the proposition offers a useful reference point for assessing whether the constant bias assumption is reasonable in practical applications. If an economist believes that the cutoff positions are set "exogenously" and that groups are comparable in terms of unobserved characteristics, then the constant bias assumption is likely to hold in settings with a non-manipulable running variable. Consider, for example, a case in which the income threshold for financial aid is revised from the previous year—perhaps due to a change in budget constraints. Alternatively, imagine a newly introduced aid program that uses a wealth-based threshold. If unobserved characteristics are assumed to remain stable across cohorts, then the constant bias assumption may plausibly hold, and the extrapolation strategy of Cattaneo et al. (2021) can be applied. An empirical setting of this kind is found in Londoño-Vélez et al. (2020). In particular, Figure 5 of their paper seems to illustrate a context where the constant bias assumption appears credible.

That said, the assumption can fail when the cutoff is correlated with unobserved group characteristics. For instance, if more generous cutoffs are systematically applied to disadvantaged or minority groups, then the equality in ϵ_i distributions may not hold, even when the running variable itself is non-manipulable. Hence, when applying Cattaneo et al.'s (2021) method, we recommend that researchers justify the assumption of similarity in unobservables across groups. Balance tests on observed pre-treatment covariates may offer indirect evidence. For example, comparing covariate distributions or formally testing for differences can help assess plausibility.

Remark 1 (Unconfoundedness)**.** As discussed above, the constant bias assumption is likely to hold when cutoff assignment is as good as random, i.e., $C_i \perp\!\!\!\perp \epsilon_i$. While this assumption may not necessarily hold in all settings, researchers can still appeal to the constant bias assumption *conditional on covariates*, in a manner analogous to the unconfoundedness assumption in observational studies (e.g., Imbens and Rubin, 2015). Specifically, even if $\epsilon_i \perp\!\!\!\perp C_i$ is questionable, one might instead assume $\epsilon_i \perp\!\!\!\perp C_i \mid \boldsymbol{Z}_i$, where $\boldsymbol{Z}_i$ is a vector of predetermined covariates not influenced by effort e_i. In such cases, the conditional constant bias assumption, briefly discussed in Cattaneo et al. (2021, Supplemental Appendix, Section SA-4.1), provides a powerful alternative identification strategy. ⌟

2.3.2 Partially Manipulable Running Variable Case

We now turn to the case where the running variable is partially manipulable, as characterized by equation (2.3b). To proceed, we impose some conditions.

Assumption 2.3. *(i)* $u\left(s(e_i)\right) + \beta\left\{y(e_i) + \tilde{\tau}_i \mathbb{P}\left[s(e_i) + \eta_i^s \geq C_i\right]\right\}$ *is strictly concave in* e_i*,* *(ii)* $u(s(e_i))$, $K(e_i, \epsilon_i)$, $y(e_i)$, $s(e_i)$, *and* $\mathbb{P}\left[s(e_i) + \eta_i^s \geq C_i\right]$ *are continuously differentiable in* e_i. *Furthermore,* $s'(e_i) > 0$ *and* $K(e_i, \epsilon_i)$ *is continuous in* (e_i, ϵ_i).

(iii) The distribution of the shock η_i^s *has a density function* f_{η^s}.

(iv) $K(e_i, \epsilon_i)$ *is convex in* e_i,

(v) The support of ϵ_i *is an interval and identical across both groups. Furthermore,* $e_i \in [0, \bar{e}]$.

Assumption 2.4. $\tilde{\tau}_i = \tilde{\tau}$ *for every* i.

Assumption 2.3 comprises standard regularity conditions. Assumptions 2.3(i) and (ii) are smoothness and (high-level) concavity assumptions that ensure analytical traceability. Assumption 2.3(iii) is also a weak smoothness assumption. Assumption 2.3(iv) requires convexity of the cost function, covering commonly used specifications such as the linear and quadratic cost functions. Assumption 2.3(v) imposes a mild support condition on ability.[3]

Assumption 2.4 rules out heterogeneity in agents' beliefs about treatment benefits. While strong, the next result suggests that *even under this homogeneity assumption*, the constant bias assumption may not hold.

Proposition 2. *Suppose that the running variable is partially manipulable and that Assumptions 2.3-2.4 are satisfied. Suppose also that the solution* e_i^* *is an interior point of* $[0, \bar{e}]$ *for any* ϵ_i. *Then the optimal effort* e_i^* *does not depend on* C_i *for any* ϵ_i *if and only if the density function* f_{η^s} *is periodic with period* $h - l$, *i.e.,* $f_{\eta^s}(z) = f_{\eta^s}(z + h - l)$, *on the interval* $[l - \sup_\epsilon s(e^*(\epsilon)), h - \inf_\epsilon s(e^*(\epsilon))]$. ♣

In most empirical settings, the condition made on the density function is questionable. We typically have that $l < \sup_\epsilon s(e^*(\epsilon))$ and $\inf_\epsilon s(e^*(\epsilon)) < h$, so the interval

[3]To derive the sharpest theoretical prediction, we would need to rely on some parametric functional form assumptions on the structural functions. However, we hesitate to do so, as such implications could be driven more by auxiliary restrictions than by fundamental economic mechanisms. Instead, we maintain generality, while illustrating the intuition with numerical examples later in this section.

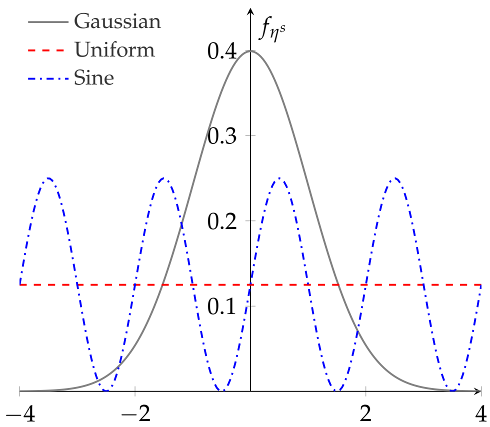

Figure 2: Examples of the Error Distribution f_{η^s}

Note: We set $\inf_\epsilon s(e^*(\epsilon)) = 2, \sup_\epsilon s(e^*(\epsilon)) = 8, l = 4$, and $h = 6$, which are defined in Proposition 2. The solid line indicates the density of the standard Gaussian distribution. In this case, e_i^* depends on C_i. The dashed line represents the density of the uniform distribution on $[-4, 4]$, while the dot-dashed line is a periodic sine function. In these cases, e_i^* does not depend on C_i.

$[l - \sup_\epsilon s(e^*(\epsilon)), h - \inf_\epsilon s(e^*(\epsilon))]$ includes zero. Now, recalling that f_{η^s} represents the distribution of idiosyncratic noise in the score, it is far more natural to assume that this density is unimodal and centered at zero, such as the Gaussian errors, rather than being periodic (see Figure 2). Therefore, the proposition states that the optimal effort generally *does* depend on the cutoff position C_i.

The dependence of the effort input e_i^* on C_i implies that the distribution of S_i and $Y_i(0)$, which are functions of e_i^*, can differ between the two groups. Hence, in contrast to the case with a non-manipulable running variable, the validity of the constant bias assumption is not guaranteed, even under random assignment, identical distributions of ϵ_i, and the strong homogeneity assumption $\tilde{\tau}_i = \tilde{\tau}$. Of course, the dependence of e_i^* on C_i does not rule out the possibility that the constant bias assumption holds; however, motivating this assumption in practice may be challenging, since its validity depends on the functional forms of the structural functions, which are not observable by economists. Furthermore, the example below illustrates that a deviation from constancy can be substantial:

Example 1. *Suppose that* $u(s(e_i)) = s(e_i) = 5\sqrt{e_i}$, $y(e_i) = 10\sqrt{e_i}$, *and* $K(e_i, \epsilon_i) = 15(2 -$

$\epsilon_i)e_i$. *The distribution of* η_i^s *is triangle, i.e.,* $f_{\eta^s}(z) = (1 - |z|)_+$, *which is made to obtain an explicit solution. The ability* ϵ *follows the uniform distribution,* $Uniform(0, 1)$. *The cutoffs are* $C_i \in \{2, 3\}$. *We set* $\tilde{\tau} = 1$, $\gamma = 0$, *and* $\beta = 1$. *In this setup, the optimal effort can be computed as*

$$e_i^* = \begin{cases} \dfrac{1}{4(2-\epsilon_i)^2} & \text{if } \epsilon_i \le \dfrac{4C_i - 9}{2(C_i - 1)} \text{ or } \dfrac{4C_i - 1}{2(C_i + 1)} < \epsilon_i \\ \dfrac{(4 + C_i)^2}{(17 - 6\epsilon_i)^2} & \text{if } \dfrac{6C_i - 10}{3C_i} < \epsilon_i \le \dfrac{4C_i - 1}{2(C_i + 1)} \\ \dfrac{(4 - C_i)^2}{(7 - 6\epsilon_i)^2} & \text{if } \dfrac{4C_i - 9}{2(C_i - 1)} < \epsilon_i \le \dfrac{6C_i - 10}{3C_i} \end{cases}, \tag{2.4}$$

which confirms the dependence of e_i^* *on* C_i. *Plugging in this optimal effort to* $s(e_i)$ *and* $y(e_i)$, *we can compute the conditional expectations* $\mathbb{E}[Y_i(0)|X_i]$ *for each group, which are shown in Figure 3a. The constant bias assumption is not satisfied, even though the two groups are similar in their ability and face the same decision problem except for the cutoff points.*

As illustrated in the example above, even when all determinants except the cutoff are identical across groups, the resulting regression functions can differ when the running variable is partially manipulable.

Besides, the validity of the assumption becomes more unclear when the distribution of ϵ_i is supposed to differ, as illustrated in the next example:

Example 2. *Suppose the same structural functions as Example 1. We here assume that the group with* $C_i = 3$ *is more advantaged in that the ability* ϵ *for the group with* $C_i = 3$ *follows* $Uniform(2/3, 5/3)$. *The optimal effort is determined by* (2.4). *The regression functions are shown in Figure 3b. The constant bias assumption is, again, not satisfied.*

These findings highlight a concern for extrapolating treatment effects under the constant bias assumption in settings where the running variable is partially manipulable. When the structural functions are unknown—as is typically the case—justifying constancy becomes empirically difficult.

Remark 2. Even when the running variable partly reflects agent effort, such as a test score,

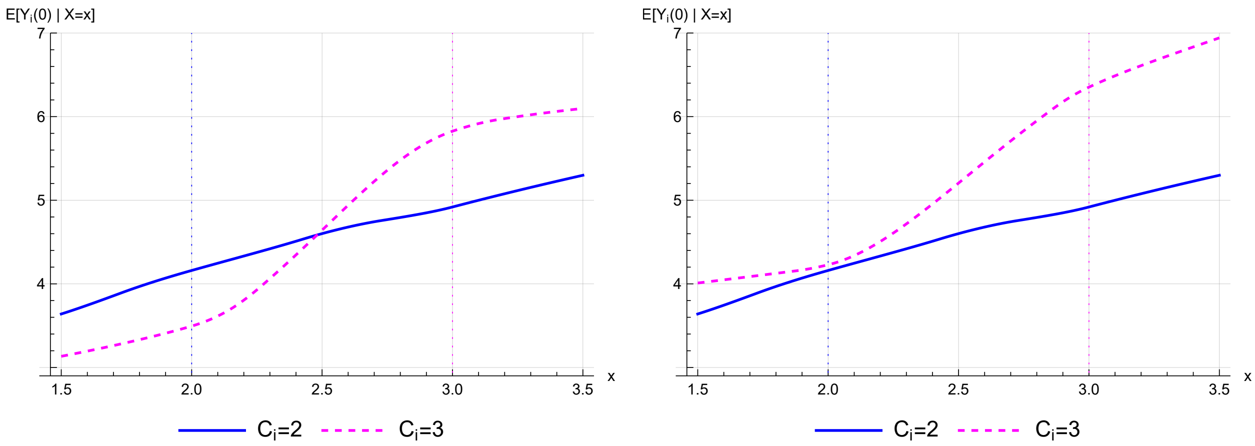

(a) Example 1

(b) Example 2

Figure 3: $\mathbb{E}\left[Y_i(0)|X_i = x\right]$ in Numerical Examples: Partially Manipulable Running Variable

the concerns discussed in this section may be mitigated depending on the design. For instance, if the introduction of a scholarship program or the implementation of multiple cutoffs is determined after the entrance exam has been administered, then all groups can be seen as facing the same decision problem at the time of their effort choice. In such cases, the distributional equivalence of ϵ_i implies the constant bias assumption. ⌟

3 Alternative Identification Results

In the previous section, we demonstrated that the constant bias assumption may be violated in some empirical settings. As a result, the extrapolation formula in equation (2.1), which relies on Assumption 2.2, can yield biased estimates. This motivates the need for an alternative approach when the plausibility of the constant bias assumption is in doubt.

One potential strategy is to fully specify the structural functions and distributional assumptions and then estimate the model structurally to recover $\mathbb{E}\left[Y(0)|X\right]$.[4] However, implementing such a strategy may require detailed individual-level data sufficient to iden-

[4]See Todd and Wolpin (2023) for a review of approaches that integrate structural modeling with causal inference, particularly in the context of randomized controlled trials.

tify underlying preference parameters and to serve as proxies for individual effort and belief—perhaps unavailable since many RD studies are observational and not designed to collect such granular information.

In light of these challenges, this section develops an alternative identification strategy that remains within a reduced-form framework but does not rely on the constant bias assumption. We begin by introducing a new set of assumptions and then derive identification results under these conditions. Our focus is on the two-cutoff case for expositional clarity, though the extension to multiple cutoffs is straightforward (see Remark 5). The extension to the one-sided fuzzy RD case is deferred to the Online Appendix.

3.1 Main Results

3.1.1 Assumptions and Identification Results

Our identification strategy is based on two empirically motivated assumptions. We begin with a commonly employed shape restriction:

Assumption 3.1 (Monotonicity). *$\mu_{0,c}(x)$ is weakly increasing in $x \in (l, h)$ for $c \in \{l, h\}$.*

This *monotonicity assumption* posits that the untreated potential outcome is a non-decreasing function of the running variable. Such a monotonicity assumption is standard in the partial identification literature (e.g., Manski, 1997).

It is plausible in many empirical RD settings. As Babii and Kumar (2023) wrote, "[r]egression discontinuity designs encountered in empirical practice are frequently monotone." For example, when X_i represents a test score and $Y_i(0)$ denotes future earnings in the absence of any treatment, it is natural to assume that $\mu_{0,c}(x)$ is increasing. A similar logic applies when X_i is family income level, and higher-income families are associated with greater expected earnings due to inherited ability or increased investment in human capital (e.g., Björklund et al., 2006).

Second, we introduce an alternative restriction to the constant bias assumption, one that relates the untreated outcome functions across groups:

Assumption 3.2 (Dominance). *$\mu_{0,l}(x) \le \mu_{0,h}(x)$ holds on $x \in (l, h)$.*

This *dominance assumption* assumes that the untreated conditional mean function of the lower cutoff group lies below that of the higher cutoff group. This is plausible in many multi-cutoff RD settings, especially when cutoffs are designed to reduce inequities or reflect pre-existing differences between groups.

For instance, scholarship thresholds are often relaxed for students from disadvantaged backgrounds—such as those from high-poverty regions—allowing them to qualify with lower test scores (Melguizo et al., 2016). Conversely, more academically prepared students may apply to competitive schools with higher cutoffs. In both examples, the cutoff reflects differences in group characteristics; that is, it is determined "endogenously." In such cases, the dominance assumption is more likely to hold.

Our main identification result is as follows:

Theorem 1 (Bounds on Extrapolated RD Effects). *Under Assumptions 2.1, 3.1, and 3.2, the treatment effect for group $C_i = l$ on $\bar{x} \in (l, h)$, $\tau_l(\bar{x})$, is bounded from below and above by*

$$\Delta_l(\bar{x}) = \mu_{1,l}(\bar{x}) - \mu_{0,h}(\bar{x}), \text{ and } \nabla_l(\bar{x}) = \mu_{1,l}(\bar{x}) - \mu_{0,l}(l).$$

These bounds $[\Delta_l(\bar{x}), \nabla_l(\bar{x})]$ are pointwise sharp. ♣

Corollary 1. *Suppose the "reverse" of Assumptions 3.1 and 3.2 hold instead, that is, $\mu_{0,c}(x)$ is weakly decreasing and $\mu_{0,l}(x) \ge \mu_{0,h}(x)$ on $x \in (l, h)$. Then, the sharp bounds are given by $[\nabla_l(\bar{x}), \Delta_l(\bar{x})]$.* ♣

The idea behind Theorem 1 is illustrated in Figure 4. By the monotonicity of $\mu_{0,l}(x)$, $\mu_{0,l}(\bar{x})$ can be bounded from below by $\mu_{0,l}(l)$. The dominance assumption ensures that it

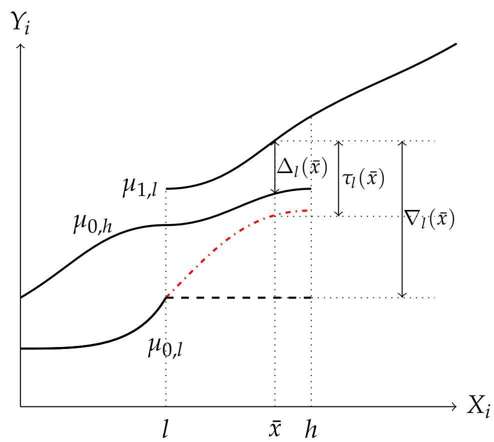

Figure 4: Bounds of Theorem 1

is bounded above by $\mu_{0,h}(\bar{x})$. Hence, we obtain that $\tau_l(\bar{x}) \in [\Delta_l(\bar{x}), \nabla_l(\bar{x})]$. Monotonicity of $\mu_{0,h}(x)$ guarantees the sharpness of the bounds (see also Remark 8 below).

The same reasoning applies across any points in (l, h), leading to the following corollary:

Corollary 2. *Take an arbitrary closed interval $\mathcal{X} \subset (l, h)$. Then, under the same assumptions in Theorem 1, the bounds $\Delta_l(x)$ and $\nabla_l(x)$ are both attainable as a function over the interval $\mathcal{X}$, in the sense that $\Delta_l(x)$ and $\nabla_l(x)$ are consistent with the observed data and maintained assumptions over $\mathcal{X}$.* ♣

This corollary states that $\Delta_l(x)$ and $\nabla_l(x)$ provide the tightest lower and upper bounds on $\tau_l(x)$ as a function over a closed interval $\mathcal{X}$ within (l, h). Unfortunately, the bounds $[\Delta_l(x), \nabla_l(x)]$ are *not* uniformly sharp in general, that is, there exists a function $\delta(x)$ that is inconsistent with our maintained assumptions although $\delta(x) \in [\Delta_l(x), \nabla_l(x)]$ over the interval $\mathcal{X}$. Such a counterexample can be easily constructed and is provided in the Online Appendix. Nevertheless, we emphasize that both the lower and upper bounds are *attainable* and cannot be rejected as representing a true treatment effect function over $\mathcal{X}$. Thus, the practical implication that the true treatment effect function lies inbetween $\Delta_l(x)$ and $\nabla_l(x)$ remains valid. In this view, the uniform (non)sharpness may be of limited practical consequence.

3.1.2 Remarks

We conclude this subsection with several important remarks:

Remark 3 (Flexibility of the Treatment Effect Function)**.** Both of our main assumptions impose restrictions only on the control state, and no assumptions are made about the functional form under the treated status. Thus, the treatment effect function itself is left unrestricted, in line with Cattaneo et al. (2021). ⌟

Remark 4 ((Non-)Sensitivity to Transformation)**.** In RD studies, it is common practice to apply a monotonic transformation to the outcome variable—for example, log transformation of annual earnings as in Oreopoulos (2006). The constant bias assumption can be sensitive to such transformations. This type of sensitivity to transformations has been pointed out by Roth and Sant'Anna (2023) in the DID setting. Our identification conditions are invariant to monotonic transformations. This robustness makes the bounds especially valuable when the transformed outcome is of interest, but the plausibility of constant bias in the transformed space is unclear. ⌟

Remark 5 (Multiple Cutoff Points)**.** Suppose there are $J + 1 (\geq 3)$ cutoff points, denoted by $c_0, c_1, \ldots, c_J$. Without loss of generality, we focus on $\tau_{c_0}(x)$. A natural extension of the dominance assumption (Assumption 3.2) is a sequential dominance: $\mu_{0,c_0}(x) \leq \mu_{0,c_1}(x) \leq \cdots \leq \mu_{0,c_J}(x)$. Under this assumption, we can derive a sharp lower bound for $\tau_{c_0}(\bar{x})$ as $\mu_{1,c_0}(\bar{x}) - \mu_{0,c_K}(\bar{x})$ when $\bar{x} \in (c_{K-1}, c_K)$. The upper bound remains the same as in Theorem 1. Hence, using the notation from the main text, the bounds over the interval (l, h) are (weakly) tightened when there exists an "intermediate" group m satisfying $l < m < h$. Estimation and inference proceed analogously to those described in the following section. ⌟

Remark 6 (Effect of Changing Threshold)**.** $\tau_l(\bar{x})$ should be understood as the treatment effect in an environment where agents make decisions under cutoff l. Consequently, when an economist is interested in the effect of shifting the threshold from l to $\bar{x}$, an additional assumption is required. To interpret $\tau_l(\bar{x})$ in this context, a policy invariance assumption, akin to the local policy invariance assumption in Dong and Lewbel (2015), becomes necessary. Under such a condition, the derived bounds characterize the effect of adjusting the threshold to $\bar{x}$. ⌟

Remark 7 (Limitation of the Multi-Cutoff RD Designs)**.** Our bounds do not provide any information about $\tau_l(x)$ on $x < l$, which may also be of interest. This limitation mirrors the challenge discussed by Cattaneo et al. (2021, p. 1947). Exploring identification strategies in this region is an important area for future work. We note, however, that the marginal extrapolation strategy proposed by Dong and Lewbel (2015) remains applicable in this setting. ⌟

Remark 8 (Testability and Falsification of Assumptions)**.** Under Assumption 2.1, the dominance assumption is refuted if $\lim_{x \uparrow l} \mu_{0,l}(x) > \lim_{x \uparrow l} \mu_{0,h}(x)$, which is directly testable. In general, the monotonicity of $\mu_{0,l}$ is not testable, while the monotonicity of $\mu_{0,h}$ is directly testable by, for example, Chetverikov's (2019) procedure. Technically, the falsification of the monotonicity of $\mu_{0,h}$ affects only the assertion of pointwise sharpness or the attainability of the lower bound, and does not immediately invalidate the bounds themselves. However, in many applications, the rejection of the monotonicity of $\mu_{0,h}$ may serve as indirect evidence for the rejection of the monotonicity of $\mu_{0,l}$, suggesting the potential invalidity of the upper bound. In such cases, researchers may focus on the lower bound $\Delta(x)$, which requires only Assumption 3.2. While this bound is only partially informative, it still offers valuable insight into the external validity of RD estimates. ⌟

Remark 9 (Sensitivity Analysis)**.** In empirical studies, it is common practice to report layered estimates—such as point or partially identified intervals nested within one another—to assess the identifying power of different sets of assumptions and examine the sensitivity of the results (e.g., Kreider et al., 2012). The bounds obtained in Theorem 1 well align with this purpose. Under Assumptions 3.1 and 3.2, it follows that $\tau_{l,\mathrm{CB}}(x) \in$

$[\Delta_l(x), \nabla_l(x)]$. Thus, researchers can use our bounds to assess the identification power of the constant bias assumption, given that the shape restrictions hold. However, this nesting implies that the bounds cannot be used to test the constant bias assumption itself. ⌟

3.2 Estimation and Inference

3.2.1 Local Linear Estimation

To construct the bounds estimates, we need to estimate $\mu_{1,l}(x)$, $\mu_{0,h}(x)$, and $\mu_{0,l}(l)$. These quantities are all consistently estimable by standard nonparametric regression techniques. Following the recent RD literature, we employ the local linear smoothing with mean-squared error (MSE) optimal bandwidth selector (see Fan and Gijbels, 1996 for a comprehensive review). Specifically, we estimate each function $\mu_{d,c}(x)$ by $\widehat{\mu}_{d,c}(x) := (1,0)\widehat{\beta}_{d,c}(x)$, where

$$\widehat{\beta}_{d,c}(x) = \underset{(b_0,b_1)^\top \in \mathbb{R}^2}{\arg\min} \sum_{i:D_i=d,C_i=c} \{Y_i - (b_0 + b_1(X_i - x))\}^2 K\left(\frac{X_i - x}{b}\right), \tag{3.1}$$

where K is the kernel function and b is the bandwidth. Note that we only use the observations with $D_i = d$ and $C_i = c$ to estimate $\mu_{d,c}(x)$. Note also that K and b can differ in each estimation. One can estimate the bounds by $\widehat{\Delta}_l(x) = \widehat{\mu}_{1,l}(x) - \widehat{\mu}_{0,h}(x)$ and $\widehat{\nabla}_l(x) = \widehat{\mu}_{1,l}(x) - \widehat{\mu}_{0,l}(l)$.

3.2.2 Pointwise Inference

In some applications, researchers may have a specific point of interest, $\bar{x}$. In this case, pointwise uncertainty quantification will be useful.

To address the bias introduced by kernel smoothing, we employ the robust bias-corrected inference procedure developed by Calonico et al. (2014, 2018, 2022). This method corrects for the leading-order smoothing bias and adjusts the variance induced by this bias correction, thereby enabling valid inference under MSE-optimal bandwidth selection.

Let $B_{d,c}(x)$ denote the asymptotic smoothing bias of $\widehat{\mu}_{d,c}(x)$ due to the local linear regression (3.1) and $\widehat{B}_{d,c}(x)$ be its estimator, typically computed via local quadratic regression using a bandwidth b_{bias}. A common and practical choice for b_{bias} is to set $b_{\text{bias}} = b$ (Calonico et al., 2018, p. 773; Calonico et al., 2019, p. 11), which we adopt hereafter.

Define the bias corrected estimator $\widehat{\mu}^{\text{BC}}_{d,c}(x) = \widehat{\mu}_{d,c}(x) - \widehat{B}_{d,c}(x)$, and let $\mathcal{V}_{d,c}(x) = \mathbb{V}\left[\widehat{\mu}_{d,c}(x) - \widehat{B}_{d,c}(x)|X_1, \ldots, X_n\right]$, with $\widehat{\mathcal{V}}_{d,c}(x)$ as its estimator. Under standard smoothness and regularity conditions (see Calonico et al., 2014, 2018), we have that

$$\widehat{\mathcal{V}}^{-1/2}_{d,c}(x)\left\{\widehat{\mu}^{\text{BC}}_{d,c}(x) - \mu_{d,c}(x)\right\} \to_d \mathcal{N}(0,1),$$

where $\to_d$ denotes the convergence in distribution. Put $\widehat{\Delta}^{\text{BC}}_l(\bar{x}) = \widehat{\mu}^{\text{BC}}_{1,l}(\bar{x}) - \widehat{\mu}^{\text{BC}}_{0,h}(\bar{x})$ and $\widehat{\nabla}^{\text{BC}}_l(\bar{x}) = \widehat{\mu}^{\text{BC}}_{1,l}(\bar{x}) - \widehat{\mu}^{\text{BC}}_{0,l}(l)$. Then, since $\widehat{\mu}^{\text{BC}}_{1,l}(\bar{x})$, $\widehat{\mu}^{\text{BC}}_{0,h}(\bar{x})$, and $\widehat{\mu}^{\text{BC}}_{0,l}(l)$ are independent by the assumption of random sampling, $\widehat{\Delta}^{\text{BC}}_l(\bar{x})$ and $\widehat{\nabla}^{\text{BC}}_l(\bar{x})$ are asymptotically normal with the asymptotic variances $\widehat{\mathcal{V}}_L = \widehat{\mathcal{V}}_{1,l}(\bar{x}) + \widehat{\mathcal{V}}_{0,h}(\bar{x})$, and $\widehat{\mathcal{V}}_U = \widehat{\mathcal{V}}_{1,l}(\bar{x}) + \widehat{\mathcal{V}}_{0,l}(l)$. Using them, we can draw confidence intervals (CIs) for each lower and upper bound in the usual manner. One can also construct the CI that covers $[\Delta_l(\bar{x}), \nabla_l(\bar{x})]$, or the CI for $\tau_l(\bar{x})$ using the method of Imbens and Manski (2004) and Stoye (2009).

3.2.3 Uniform Confidence Band

In other applications, researchers are interested in a range of values for the running variable rather than a specific point. In such cases, constructing a uniform confidence band over the region of interest provides a useful tool for inference. This subsection proposes a procedure for uniform inference based on the multiplier bootstrap, following the ideas of Fan et al. (2022) and Imai et al. (2025). In the Online Appendix, we establish the asymptotic validity of the proposed procedure by combining the results of Chernozhukov et al.

(2013, 2014a,b).

Let $\mathcal{I} \subset (l, h)$ be a closed interval of interest. We proceed with the following steps:

1. Obtain $\widehat{\mu}_{1,l}^{\mathrm{BC}}(x)$ and $\widehat{\mu}_{0,h}^{\mathrm{BC}}(x)$ using local linear regressions with integrated MSE (IMSE) optimal bandwidths $b_{1,l}$ and $b_{0,h}$ (e.g., Calonico et al., 2019). Construct $\widehat{\mu}_{0,l}^{\mathrm{BC}}(l)$ similarly but with an MSE-optimal bandwidth $b_{0,l}$. For all components, we apply bias correction using the same bandwidths for both the main and bias estimation: $b_{\mathrm{bias},d,c} = b_{d,c}$, and we employ the same kernel function throughout.

2. Choose a large number of bootstrap replications M (e.g., $M = 1000$). For each $m = 1, \ldots, M$, draw an i.i.d. random variable $\{\xi_i^m\}_{i=1}^n$ from Mammen's (1993) two-point distribution.[5] Compute the local *quadratic* regression estimators, $\widehat{\mu}_{d,c}^{\star m}(x) = (1, 0, 0)\widehat{\beta}_{d,c}^{\star m}(x)$, where $\widehat{\beta}_{d,c}^{\star m}(x)$'s are defined as

$$\underset{(b_0,b_1,b_2)^\top \in \mathbb{R}^3}{\arg\min} \sum_{i:D_i=d,C_i=c} (\xi_i^m + 1) \left\{ Y_i - b_0 - b_1(X_i - x) - b_2(X_i - x)^2 \right\}^2 K\left(\frac{X_i - x}{b_{d,c}}\right).$$

   Note that the bandwidths $b_{d,c}$ are the same as the ones used in step 1 in every iteration. Define $\widehat{\Delta}_l^{\star m}(x) = \widehat{\mu}_{1,l}^{\star m}(x) - \widehat{\mu}_{0,h}^{\star m}(x)$ and $\widehat{\nabla}_l^{\star m}(x) = \widehat{\mu}_{1,l}^{\star m}(x) - \widehat{\mu}_{0,l}^{\star m}(l)$.

3. For each replication m, calculate the studentized maximum deviations:

$$S_L^\star(m) = \sup_{x \in \mathcal{I}} \frac{\widehat{\Delta}_l^{\star m}(x) - \widehat{\Delta}_l^{\mathrm{BC}}(x)}{\widehat{\mathcal{V}}_L^{1/2}(x)} \quad \text{and} \quad S_U^\star(m) = \sup_{x \in \mathcal{I}} \frac{\widehat{\nabla}_l^{\star m}(x) - \widehat{\nabla}_l^{\mathrm{BC}}(x)}{\widehat{\mathcal{V}}_U^{1/2}(x)}.$$

   In practice, the supremum is approximated by the maximum over some fine grid points.

[5] Mammen's two-point distribution is defined as $\xi_i = (1 - \sqrt{5})/2$ with probability $(1 + \sqrt{5})/(2\sqrt{5})$ and $\xi_i = (\sqrt{5} + 1)/2$ with probability $(\sqrt{5} - 1)/(2\sqrt{5})$. Theoretically, the Gaussian multiplier can be used, but in small samples, the Gaussian multiplier may occasionally cause the matrix $\boldsymbol{R}^\top \mathrm{diag}((\xi_i^m + 1)K((X_i - x)/b_{d,c}))\boldsymbol{R}$ to become (nearly) singular due to the possible negativity of the weights $\xi_i^m + 1$. To avoid this, we recommend using Mammen's weights, guaranteeing $\xi_i^m + 1 > 0$, and the computation becomes stabler.

Given a confidence level $1-\alpha$, compute the critical values

$$c_V^\star(1-\alpha/2) := \text{ the } (1-\alpha/2)\text{-quantile of } \{S_V^\star(m) : m = 1, \ldots, M\},\ V \in \{L, U\}.$$

4. Construct a confidence band

$$\widehat{\mathcal{C}}(x) = \left[\widehat{\Delta}_l^{\mathrm{BC}}(x) - c_L^\star(1-\alpha/2)\widehat{\mathcal{V}}_L^{-1/2}(x), \widehat{\nabla}_l^{\mathrm{BC}}(x) + c_U^\star(1-\alpha/2)\widehat{\mathcal{V}}_U^{-1/2}(x)\right].$$

Then, under assumptions made in the Online Appendix Section S1.1.1, it holds that

$$\lim_{n\to\infty} \mathbb{P}\left[[\Delta_l(x), \nabla_l(x)] \subseteq \widehat{\mathcal{C}}(x) \text{ for all } x \in \mathcal{I}\right] \geq 1-\alpha.$$

Trivially, it also holds that $\lim_{n\to\infty} \mathbb{P}\left[\tau_l(x) \in \widehat{\mathcal{C}}(x) \text{ for all } x \in \mathcal{I}\right] \geq 1-\alpha$.

4 Empirical Illustrations

In this section, we present two empirical applications to demonstrate the potential usefulness of the bounds derived in the previous section.

4.1 SPP Program (Non-Manipulable Running Variable Case)

We begin with a financial aid program setting with a non-manipulable running variable, originally studied by Londoño-Vélez et al. (2020). Our primary focus is on the intent-to-treat (ITT) effects; we defer issues related to the incomplete compliance to the Online Appendix.

4.1.1 Empirical Context

We investigate the effect of *Ser Pilo Paga* (SPP), a financial aid program introduced in Colombia in 2014. Eligibility for full scholarship loans from SPP is determined by a combination of merit- and need-based criteria, using national test scores and a family wealth index as running variables.

Londoño-Vélez et al. (2020) exploited these eligibility rules to estimate standard RD effects on higher education enrollment. Following their setup, we focus on merit-eligible students—those who meet the minimum test score threshold—and estimate what they term "frontier-specific" RD effects on enrollment rates. Notably, the merit-based cutoff is constant across all students, while the need-based threshold varies across geographic areas. This creates a multi-cutoff RD setting.

In this setup, the running variable is the wealth index, which ranges from 0 (poorer) to 100 (wealthier). For consistency with the stylized RD settings, we multiply the wealth index by -1, so that larger values correspond to poorer households, with -100 representing the wealthiest. Under this transformation, students are eligible for aid if they come from households whose index exceeds a given threshold (i.e., sufficiently poor households). The cutoff for students from metropolitan areas is -57.21, while that for rural areas is -40.75, resulting in a multi-cutoff RD setting.

4.1.2 Discussion of Assumptions

In our analysis, the outcome variable is the enrollment rate, and the running variable is the family wealth level. Given this context, the "reverse" versions of Assumptions 3.1 and 3.2 appear plausible in our context. The monotonicity assumption posits that $\mu_{0,c}$ is decreasing, meaning that the probability of enrollment declines as family wealth decreases. This aligns with economic intuition and empirical patterns. The dominance assumption (in reverse) posits that the regression function for students from rural areas (who face a higher cutoff) lies below that for students from metropolitan areas. This is a natural assumption, as students in rural areas may be more disadvantaged in terms of access to educational resources and may also receive less parental support or guidance regarding the benefits of pursuing higher education.

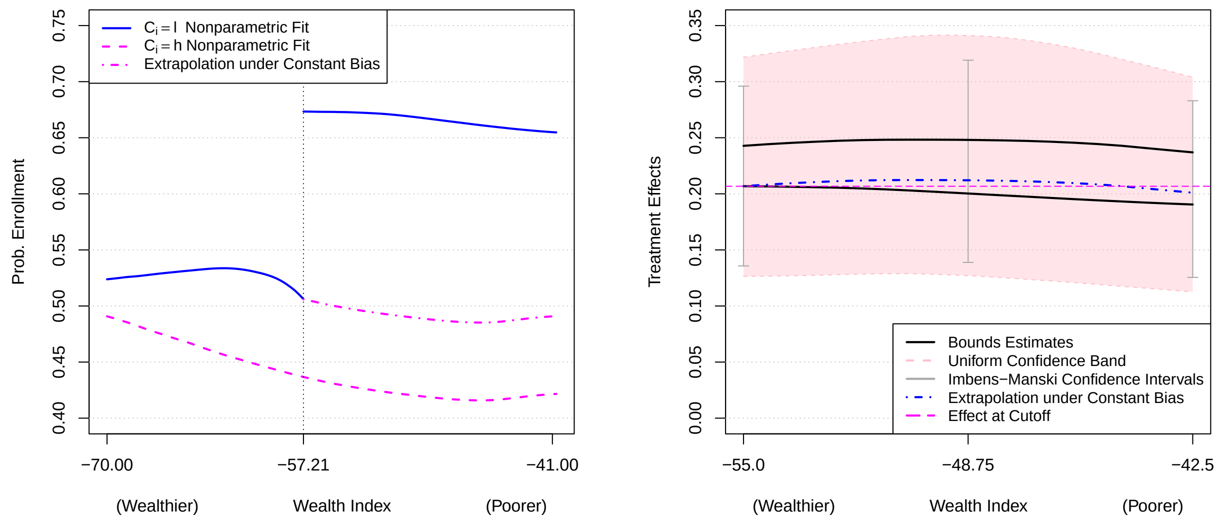

(a) Conditional Expectations

(b) Bounds on Treatment Effects

Figure 5: SPP ($N = 20551$)

In the present context, where the running variable is non-manipulable, concerns about the constant-bias assumption may be less severe than in the partially manipulable case, although certain issues remain—for instance, differences in access to infrastructure or social support between urban and rural students. Focusing exclusively on merit-eligible students, however, may help ensure more comparable unobserved characteristics across groups, akin to conditioning on covariates.

In the next subsection, we compute the bounds derived in Corollary 1 and evaluate potential deviations from constancy by comparing them to the extrapolated RD effects under the constant bias assumption.

4.1.3 Results and Implications

We consider the closed interval $\mathcal{I} = [-55.0, -42.5] \subset (-57.21, -40.75)$ in our analysis. Figure 5a presents the estimated regression functions, and Figure 5b displays the estimated bounds over $\mathcal{I}$, along with the extrapolated RD effect function under the constant bias assumption (dot-dashed line). The dashed line indicates the level of the treatment effect at the cutoff, -57.21.

Several empirically important findings emerge. First, the estimated bounds indicate positive treatment effects throughout the interval $\mathcal{I}$. These bounds are sufficiently narrow to draw meaningful policy implications, suggesting that the effect lies approximately between 0.20 and 0.25. Second, the horizontal line indicating the RD effect at the cutoff consistently falls within the (tight) bounds, implying that the hypothesis of constant average treatment effects is not rejected. These observations offer strong support for the *external validity* of the standard RD estimate. Furthermore, the finding that similarly sized effects persist across other points in $\mathcal{I}$ provides useful guidance for future policy adjustments, such as modifying eligibility thresholds.

We also confirm that the extrapolation under the constant bias assumption performs well, although the upper bound suggests that the true effect may be slightly larger. Overall, the tightness of the bounds indicates that any bias from deviations from the constant bias assumption is likely limited and not practically severe.

4.2 ACCES Program (Partially Manipulable Running Variable Case)

We now turn to a different context involving a financial aid program, in which the running variable is partially manipulable.

4.2.1 Empirical Context and Background

We revisit the empirical analysis of Cattaneo et al. (2021). They investigated the extrapolated effect of the *Acceso con Calidad a la Educación Superior* (ACCES) program—a national merit-based financial aid initiative—on higher education enrollment among Colombian students, originally studied by Melguizo et al. (2016).

Eligibility for ACCES requires students to score above a specific threshold on the national high school exit exam (SABER 11). The score ranges from 1 (best) to 1000 (worst), and the eligibility cutoff was fixed at 850 prior to 2008. Beginning in 2009, however, region-specific cutoffs were introduced, resulting in a multi-cutoff RD setting. For consistency, we multiply the scores by -1, so that higher achievement corresponds to larger values.

Cattaneo et al. (2021) leveraged this design to estimate extrapolated RD effects on college enrollment, focusing on two cohorts: students who applied between 2000 and 2008 (cutoff -850) and those who applied between 2009 and 2010 in a region with a cutoff of -571. They reported that the extrapolated RD effect at -650 was 0.191, larger than the local RD estimate at -850, which was 0.137.

4.2.2 Discussion on Assumptions

The plausibility of the constant bias assumption in this context is debatable. First, the running variable is partially manipulable, suggesting that the constant bias assumption is not guaranteed even when two groups are similar (Section 2.3.2). Second, the 2009 reform introduced cutoffs in a progressive manner: regions with greater disadvantage received lower thresholds, while more advantaged areas were subject to higher ones. Notably, the region with a cutoff of -571 is among the most advantaged, with a very low share of students from low socioeconomic backgrounds (Melguizo et al., 2016, Figure 1). As a result, the two groups under comparison may differ substantially in unobservable characteristics.

Another potential concern regarding the constant bias assumption arises from how it was assessed in Cattaneo et al. (2021). In their analysis, the authors fitted separate quadratic regression functions on the left side of the lower cutoff at -850, and compared them. Building on this approach, we can also extrapolate the fitted quadratic model to the right of the cutoff, as depicted by the dotted line in Figure 6a. For comparison, the extrapolation based on the constant bias assumption is shown as the dot-dashed line. The two curves diverge notably just above the cutoff, especially in their slopes, suggesting potential violations of the constant bias assumption. Furthermore, the parametric

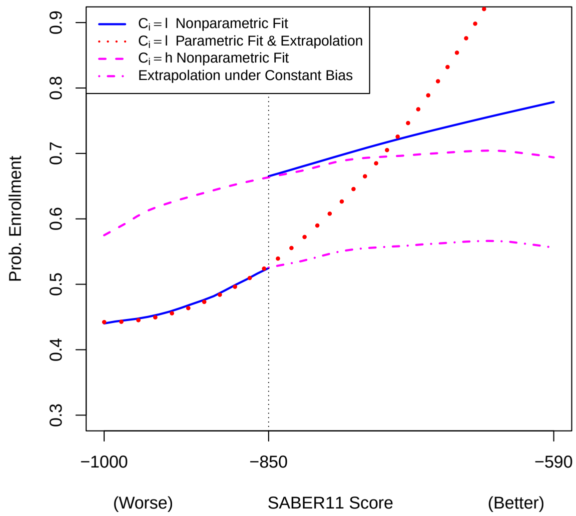

(a) Conditional Expectations

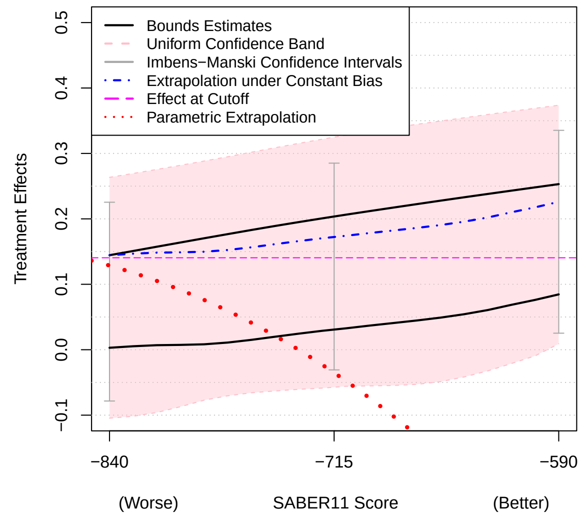

(b) Bounds on Treatment Effects

Figure 6: ACCES ($N = 1365$)

extrapolation implies a sizable negative treatment effect farther from the cutoff, yielding markedly different conclusions depending on the extrapolation method employed.

In this application, the monotonicity and dominance assumptions appear reasonable: students with higher SABER 11 scores are more likely to enroll in college, and those from more advantaged regions (i.e., $C_i = h$) are expected to have higher average enrollment rates than those from disadvantaged regions.

4.2.3 Results and Implications

Figure 6b displays the estimation results over the interval $\mathcal{I} = [-840, -590]$. The bounds suggest that the ACCES program likely has a nonnegative effect overall, and the effects are roughly inbetween $[0.00, 0.15]$ just above the cutoff and about $[0.10, 0.25]$ far away from the cutoff.

While the bounds are admittedly wide, they nonetheless deliver meaningful information about treatment effects away from the cutoff. Importantly, they rule out large negative effects implied by the parametric extrapolation, which lie well outside the identified set. This provides reassurance that strong pessimistic conclusions are inconsistent with the maintained assumptions. At the same time, it remains possible that the true effect is smaller than what is implied by extrapolation under the constant bias assumption—for instance, closer in magnitude to the effect estimated at the cutoff. In the absence of additional institutional knowledge, this possibility should also be borne in mind when evaluating the ACCES program.

5 Conclusion

This paper explored when and how the treatment effect in RD designs can be extrapolated away from cutoff points in multi-cutoff settings. We began by examining the plausibility of the constant bias assumption proposed by Cattaneo et al. (2021), interpreting it through the lens of rational agent behavior. We found that the assumption is indeed plausible when the two groups are composed of agents with comparable characteristics and when the running variable is non-manipulable. However, this justification may fail when the running variable is partially manipulable by the agent, potentially resulting in biased estimates.

To address this issue, we proposed an alternative identification strategy grounded in empirically motivated assumptions—monotonicity and dominance—which do not require the constant bias assumption. We derived sharp bounds on extrapolated treatment effects under these assumptions and established a uniform inference procedure. Our empirical applications illustrated the potential usefulness of these bounds.

A Proofs

A.1 Proofs of the Main Results

Proof of Proposition 1. The objective function (2.3a) does not depend on C_i, implying the desired result. $\square$

Proof of Proposition 2. Let $e_c^*(\epsilon_i)$ denote the optimal effort of those with ϵ_i in the group $c \in \{l, h\}$. Take an arbitrary ϵ_i and suppose $e_l^*(\epsilon_i) = e_h^*(\epsilon_i) =: e_i^*$. Then, by the first-order condition, we have

$$
\begin{aligned}
&u'\left(s(e_i^*)\right) s'(e_i^*) - \frac{\partial K}{\partial e}(e_i^*, \epsilon_i) + \beta(y'(e_i^*) + \tilde{\tau} s'(e_i^*) f_{\eta^s}(l - s(e_i^*))) \\
&\quad = u'(s(e_i^*) s'(e_i^*) - \frac{\partial K}{\partial e}(e_i^*, \epsilon_i) + \beta(y'(e_i^*) + \tilde{\tau} s'(e_i^*) f_{\eta^s}(h - s(e_i^*))) = 0.
\end{aligned}
$$

This implies $f_{\eta^s}(l - s(e_i^*)) = f_{\eta^s}(h - s(e_i^*))$. The objective function in (2.3b) is continuous in (e, ϵ) and the choice set $[0, \bar{e}]$ is compact. Hence, by Berge's maximum theorem, the argmax correspondence is upper hemicontinuous and nonempty. Moreover, by strict concavity in e, the maximizer is unique, so the optimal effort $e_c^*(\epsilon_i)$ is continuous in ϵ_i. Since $e_l^*(\epsilon_i) = e_h^*(\epsilon_i) =: e_i^*$ by assumption, the map $\epsilon_i \mapsto s(e^*(\epsilon_i))$ is continuous. Then, recalling that $f_{\eta^s}(l - s(e_i^*)) = f_{\eta^s}(h - s(e_i^*))$, f_{η^s} must be periodic with period $h - l$ on the interval $[l - b, h - a]$, where $a = \inf_\epsilon s(e^*(\epsilon))$ and $b = \sup_\epsilon s(e^*(\epsilon))$.

Conversely, suppose that f_{η^s} is periodic with period $h - l$ on $[l - b, h - a]$. Again, the first-order condition for $C_i = l$ suggests

$$
u'\left(s(e_i^*)\right) s'(e_i^*) - \frac{\partial K}{\partial e}(e_i^*, \epsilon_i) + \beta(y'(e_i^*) + \tilde{\tau} s'(e_i^*) f_{\eta^s}(l - s(e_i^*))) = 0.
$$

By definition, $s(e^*(\epsilon_i)) \in [a, b]$, hence $l - s(e^*(\epsilon_i)) \in [l - b, l - a] \subset [l - b, h - a]$ and $h - s(e^*(\epsilon_i)) = (l - s(e^*(\epsilon_i))) + (h - l) \in [h - b, h - a] \subset [l - b, h - a]$. Therefore, periodicity on $[l - b, h - a]$ implies $f_{\eta^s}(l - s(e_i^*)) = f_{\eta^s}(h - s(e_i^*))$. Substituting this equality into the first-order condition shows that the same $e^*(\epsilon_i)$ also satisfies the first-order condition under $C_i = h$. Hence, the optimal effort under $C_i = h$ must coincide with that under $C_i = l$. Because ϵ_i was arbitrary, $e^*(\epsilon_i)$ does not depend on C_i for any ϵ_i. □

Proof of Theorem 1. The validity of the bounds is illustrated in the main text. Hence, it suffices to show the sharpness. Fix $\bar{x} \in (l, h)$ arbitrarily. Define $m_a(\bar{x}) = \mu_{1,l}(\bar{x}) - \{a\Delta_l(\bar{x}) + (1-a)\nabla_l(\bar{x})\}$ with $a \in [0,1]$, and

$$\mu_{0,l}(x) = \begin{cases} \min\left\{\frac{m_a(\bar{x})-\mu_{0,l}(l)}{\bar{x}-l}(x-l) + \mu_{0,l}(l), \mu_{0,h}(x)\right\} & \text{if } l < x < \bar{x} \\ m_a(\bar{x}) & \text{if } x \geq \bar{x} \end{cases}.$$

Then, Assumptions 2.1, 3.1, and 3.2 hold. Further, note that $\tau_l(\bar{x}) = a\Delta_l(\bar{x}) + (1-a)\nabla_l(\bar{x})$ in this case, so that every value in the bounds is attainable since $a \in [0,1]$ is arbitrary. This proves the statement. Corollary 1 can be shown similarly. □

Proof of Corollary 2. Suppose that $\mu_{0,l}(x) = \mu_{0,l}(l)$ over $x \in (l, h)$, then all the assumptions hold and $\nabla_l(x)$ is attained on $\mathcal{X}$. Write $\mathcal{X} = [l + \xi_1, h - \xi_2]$, where ξ_1, ξ_2 are some positive constant. Suppose instead that

$$\mu_{0,l}(x) = \begin{cases} \min\left\{\frac{\mu_{0,h}(l+\xi_1)-\mu_{0,l}(l)}{\xi_1}(x-l) + \mu_{0,l}(l), \mu_{0,h}(x)\right\} & \text{if } l < x < l + \xi_1 \\ \mu_{0,h}(x) & \text{if } x \geq l + \xi_1 \end{cases},$$

then, all the assumptions are satisfied and $\Delta_l(x)$ is attained over $\mathcal{X}$. □

References

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

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(512):1331–1344.

Arai, Y., Hsu, Y., Kitagawa, T., Mourifié, I., and Wan, Y. (2022). Testing identifying assumptions in fuzzy regression discontinuity designs. *Quantitative Economics*, 13(1).

Babii, A. and Kumar, R. (2023). Isotonic Regression Discontinuity Designs. *Journal of Econometrics*, 234(2):371–393.

Bertanha, M. (2020). Regression Discontinuity Design with Many Thresholds. *Journal of Econometrics*, 218(1):216–241.

Bertanha, M. and Imbens, G. W. (2020). External Validity in Fuzzy Regression Discontinuity Designs. *Journal of Business & Economic Statistics*, 38(3):593–612.

Björklund, A., Lindahl, M., and Plug, E. (2006). The Origins of Intergenerational Associations: Lessons from Swedish Adoption Data. *The Quarterly Journal of Economics*, 121(3):999–1028.

Bugni, F. A. and Canay, I. A. (2021). Testing Continuity of a Density via G-order Statistics in the Regression Discontinuity Design. *Journal of Econometrics*, 221(1):138–159.

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(522):767–779.

Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2019). nprobust: Nonparametric Kernel-Based Estimation and Robust Bias-Corrected Inference. *Journal of Statistical Software*, 91(8):1–33.

Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2020). Optimal Bandwidth Choice for Robust Bias-Corrected Inference in Regression Discontinuity Designs. *The Econometrics Journal*, 23(2):192–210.

Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2022). Coverage Error Optimal Confidence Intervals for Local Polynomial Regression. *Bernoulli*, 28(4):2998–3022.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326.

Canay, I. A. and Kamat, V. (2017). Approximate Permutation Tests and Induced Order Statistics in the Regression Discontinuity Design. *The Review of Economic Studies*, 85(3):1577–1608.

Cattaneo, M. D., Jansson, M., and Ma, X. (2020). Simple Local Polynomial Density Estimators. *Journal of the American Statistical Association*, 115(531):1449–1455.

Cattaneo, M. D., Keele, L., Titiunik, R., and Vazquez-Bare, G. (2016). Interpreting Regression Discontinuity Designs with Multiple Cutoffs. *The Journal of Politics*, 78(4):1229–1248.

Cattaneo, M. D., Keele, L., Titiunik, R., and Vazquez-Bare, G. (2021). Extrapolating Treatment Effects in Multi-Cutoff Regression Discontinuity Designs. *Journal of the American Statistical Association*, 116(536):1941–1952.

Cattaneo, M. D. and Titiunik, R. (2022). Regression Discontinuity Designs. *Annual Review of Economics*, 14:821–851.

Cerulli, G., Dong, Y., Lewbel, A., and Poulsen, A. (2017). Testing Stability of Regression Discontinuity Models. In *Regression Discontinuity Designs*, volume 38 of *Advances in Econometrics*, pages 317–339. Emerald Publishing Limited.

Chernozhukov, V., Chetverikov, D., and Kato, K. (2013). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6):2786 – 2819.

Chernozhukov, V., Chetverikov, D., and Kato, K. (2014a). Anti-concentration and honest, adaptive confidence bands. *The Annals of Statistics*, 42(5):1787 – 1818.

Chernozhukov, V., Chetverikov, D., and Kato, K. (2014b). Gaussian approximation of suprema of empirical processes. *The Annals of Statistics*, 42(4):1564 – 1597.

Chetverikov, D. (2019). TESTING REGRESSION MONOTONICITY IN ECONOMETRIC MODELS. *Econometric Theory*, 35(4):729–776.

Deaner, B. and Kwon, S. (2025). Extrapolation in Regression Discontinuity Design Using Comonotonicity. *arXiv:2507.00289*.

Dong, Y. and Lewbel, A. (2015). Identifying the Effect of Changing the Policy Threshold in Regression Discontinuity Models. *The Review of Economics and Statistics*, 97(5):1081–1092.

Fan, J. and Gijbels, I. (1996). *Local Polynomial Modelling and Its Applications*. Chapman & Hall/CRC.

Fan, Q., Hsu, Y.-C., Lieli, R. P., and Zhang, Y. (2022). Estimation of Conditional Average Treatment Effects With High-Dimensional Data. *Journal of Business & Economic Statistics*, 40(1):313–327.

Fudenberg, D. and Levine, D. K. (2022). Learning in Games and the Interpretation of Natural Experiments. *American Economic Journal: Microeconomics*, 14(3):353–77.

Fusejima, K., Ishihara, T., and Sawada, M. (2025). A Unified Diagnostic Test for Regression Discontinuity Designs. *Journal of Econometrics*, 251:106074.

Gerard, F., Rokkanen, M., and Rothe, C. (2020). Bounds on Treatment Effects in Regression Discontinuity Designs with a Manipulated Running Variable. *Quantitative Economics*, 11(3):839–870.

Hahn, J., Todd, P., and Van der Klaauw, W. (2001). Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica*, 69(1):201–209.

Imai, S., Qin, L., and Yanagi, T. (2025). Doubly Robust Uniform Confidence Bands for Group-Time Conditional Average Treatment Effects in Difference-in-Differences. *arXiv:2305.02185*.

Imbens, G. and Kalyanaraman, K. (2011). Optimal Bandwidth Choice for the Regression Discontinuity Estimator. *The Review of Economic Studies*, 79(3):933–959.

Imbens, G. W. and Manski, C. F. (2004). Confidence Intervals for Partially Identified Parameters. *Econometrica*, 72(6):1845–1857.

Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.

Kreider, B., Pepper, J. V., Gundersen, C., and Jolliffe, D. (2012). Identifying the Effects of SNAP (Food Stamps) on Child Health Outcomes When Participation Is Endogenous and Misreported. *Journal of the American Statistical Association*, 107(499):958–975.

Lee, D. S. (2008). Randomized experiments from non-random selection in U.S. House elections. *Journal of Econometrics*, 142(2):675–697. The regression discontinuity design: Theory and applications.

Londoño-Vélez, J., Rodríguez, C., and Sánchez, F. (2020). Upstream and Downstream Impacts of College Merit-Based Financial Aid for Low-Income Students: Ser Pilo Paga in Colombia. *American Economic Journal: Economic Policy*, 12(2):193–227.

Mammen, E. (1993). Bootstrap and Wild Bootstrap for High Dimensional Linear Models. *The Annals of Statistics*, 21(1):255–285.

Manski, C. F. (1997). Monotone Treatment Response. *Econometrica*, 65(6):1311–1334.

Manski, C. F. and Pepper, J. V. (2018). How Do Right-to-Carry Laws Affect Crime Rates? Coping with Ambiguity Using Bounded-Variation Assumptions. *The Review of Economics and Statistics*, 100(2):232–244.

Marx, P., Tamer, E., and Tang, X. (2024). Parallel Trends and Dynamic Choices. *Journal of Political Economy Microeconomics*, 2(1):129–171.

McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics*, 142(2):698–714.

Mehta, N. (2019). An Economic Approach to Generalizing Findings from Regression-Discontinuity Designs. *Journal of Human Resources*, 54(4):953–985.

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.

Okamoto, Y. (2026). Joint Inference for the Regression Discontinuity Effect and Its External Validity. *Journal of Applied Econometrics*.

Oreopoulos, P. (2006). Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter. *American Economic Review*, 96(1):152–175.

Rambachan, A. and Roth, J. (2023). A More Credible Approach to Parallel Trends. *The Review of Economic Studies*, 90(5):2555–2591.

Roth, J. and Sant'Anna, P. H. C. (2023). When Is Parallel Trends Sensitive to Functional Form? *Econometrica*, 91(2):737–747.

Stoye, J. (2009). More on Confidence Intervals for Partially Identified Parameters. *Econometrica*, 77(4):1299–1315.

Sun, Y. (2023). Extrapolating Away from the Cutoff in Regression Discontinuity Designs. *arXiv:2311.18136*.

Todd, P. E. and Wolpin, K. I. (2023). The Best of Both Worlds: Combining Randomized Controlled Trials with Structural Modeling. *Journal of Economic Literature*, 61(1):41–85.

Villamizar-Villegas, M., Pinzon-Puerto, F. A., and Ruiz-Sanchez, M. A. (2022). A comprehensive history of regression discontinuity designs: An empirical survey of the last 60 years. *Journal of Economic Surveys*, 36(4):1130–1178.