Estimation and Inference in Boundary Discontinuity Designs: Location-Based Methods*

Matias D. Cattaneo† Rocio Titiunik‡ Ruiqi (Rae) Yu§

October 29, 2025

Abstract

Boundary discontinuity designs are used to learn about causal treatment effects along a continuous assignment boundary that splits units into control and treatment groups according to a bivariate location score. We analyze the statistical properties of local polynomial treatment effect estimators employing location information for each unit. We develop pointwise and uniform estimation and inference methods for both the conditional treatment effect function at the assignment boundary as well as for transformations thereof, which aggregate information along the boundary. We illustrate our methods with an empirical application. Companion general-purpose software is provided.

Keywords: regression discontinuity, treatment effects estimation, causal inference.

*We thank Alberto Abadie, Xiaohong Chen, Boris Hanin, Kosuke Imai, Oliver Linton, Xinwei Ma, Victor Panaretos, Jörg Stoye, and Jeff Wooldridge for comments and discussions. Cattaneo and Titiunik gratefully acknowledge financial support from the National Science Foundation (SES-2019432, DMS-2210561, and SES-2241575). Cattaneo gratefully acknowledge financial support from the Data-Driven Social Science initiative at Princeton University.

†Department of Operations Research and Financial Engineering, Princeton University.

‡Department of Politics, Princeton University.

§Department of Operations Research and Financial Engineering, Princeton University.

Contents

1 Introduction

We study estimation and inference in boundary discontinuity (BD) designs, where the goal is to learn about causal treatment effects along a continuous assignment boundary that splits units into control and treatment groups according to a bivariate location score. This setup is also known as a Multi-Score Regression Discontinuity (RD) design [Papay et al., 2011, Reardon and Robinson, 2012, Wong et al., 2013], a leading particular case being the Geographic RD design [Keele and Titiunik, 2015, Keele et al., 2015, Keele and Titiunik, 2016, Keele et al., 2017, Galiani et al., 2017, Rischard et al., 2021, Diaz and Zubizarreta, 2023]. As in the univariate RD design, under appropriate assumptions, the abrupt jump in treatment assignment along the boundary can be used to identify causal treatment effects, even if the treatment is specifically targeted to those units who need it the most. This makes BD designs a central tool in non-experimental policy evaluation. See Cattaneo and Titiunik [2022] for an overview of the RD design literature, Cattaneo et al. [2024c, Section 5] for a practical introduction to Multi-Dimensional RD designs, and Cattaneo et al. [2026] for a review of empirical practice employing BD designs.

BD designs are commonly used in the quantitative social, behavioral, and biomedical sciences when the interventions of interest are assigned based on bivariate scores. For example, Londoño-Vélez, Rodríguez, and Sánchez [2020] study the effects of *Ser Pilo Paga* (SPP), a government subsidy in Colombia that provided tuition support for post-secondary students to attend a government-certified higher education institution. Eligibility was based on both merit and economic need: in order to qualify for the program, students had to obtain a high grade in Colombia's national standardized high school exit exam, SABER 11, and they also had to come from economically disadvantaged families, as measured by the survey-based wealth index SISBEN. The eligibility rule was deterministic with a fixed bivariate cutoff: to qualify for the program, students had to obtain a SABER 11 score in the top 9 percent of scores, and their household's SISBEN index had to be below a region-specific threshold.

We study estimation and inference of causal treatment effect parameters in BD designs

where the location of each unit given by the two dimensions of the bivariate score is explicitly incorporated in (local) regression analyses—we refer to it as the *location-based* approach, to distinguish it from analyses where bivariate scores are not directly used but rather transformed according to some distance measure [Cattaneo et al., 2025b,a]. Although practitioners across the quantitative sciences employ the BD design to study treatment effects by adapting local polynomial methods developed for the univariate RD design, there is no foundational understanding of the statistical properties of these methods in the bivariate case, particularly when interest lies on uniformity or information aggregation over all boundary points. Our paper establishes the properties of local polynomial methods for BD estimation and inference using bivariate scores both pointwise and uniformly, and offers practical recommendations for the analysis and interpretation of BD designs in applications. Although our pointwise results are straightforward generalizations of prior results, our uniform methods are novel and offer researchers the ability to conduct valid inferences for the entire collection of average treatment effects at each boundary point, as well as for transformations thereof.

Our causal inference setup follows standard potential outcomes notation [see, e.g., Hernán and Robins, 2020]. The triplet $(Y_i(0), Y_i(1), \mathbf{X}_i^\top)^\top$, $i = 1, 2, \ldots, n$, is a random sample, where $Y_i(0)$ and $Y_i(1)$ denote the potential outcomes for unit i under control and treatment assignment, respectively, and the score $\mathbf{X}_i = (X_{1i}, X_{2i})^\top$ is a continuous bivariate vector with support $\mathscr{X} \subseteq \mathbb{R}^2$. Units are assigned to either the control group or the treatment group according to their location $\mathbf{X}_i$ relative to a known one-dimensional boundary curve $\mathscr{B}$ splitting the support $\mathscr{X}$ in two disjoint regions: $\mathscr{X} = \mathscr{A}_0 \cup \mathscr{A}_1$, with $\mathscr{A}_0$ and $\mathscr{A}_1$ the control and treatment disjoint (connected) regions, respectively, and $\mathscr{B} = \mathtt{bd}(\mathscr{A}_0) \cap \mathtt{bd}(\mathscr{A}_1)$, where $\mathtt{bd}(\mathscr{A}_t)$ denotes the boundary of the set $\mathscr{A}_t$. The observed outcome is $Y_i = \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_0) \cdot Y_i(0) + \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_1) \cdot Y_i(1)$. Without loss of generality, we assume that the boundary belongs to the treatment group, that is, $\mathtt{bd}(\mathscr{A}_1) \subset \mathscr{A}_1$ and $\mathscr{B} \cap \mathscr{A}_0 = \emptyset$.

In the SPP application, each student was assigned a bivariate score $\mathbf{X}_i = (\mathtt{SABER11}_i, \mathtt{SISBEN}_i)^\top$, where $X_{1i} = \mathtt{SABER11}_i$ recorded the SABER11 score and $X_{2i} = \mathtt{SISBEN}_i$ recorded the SISBEN wealth score. After recentering each variable at its corresponding threshold, the treatment

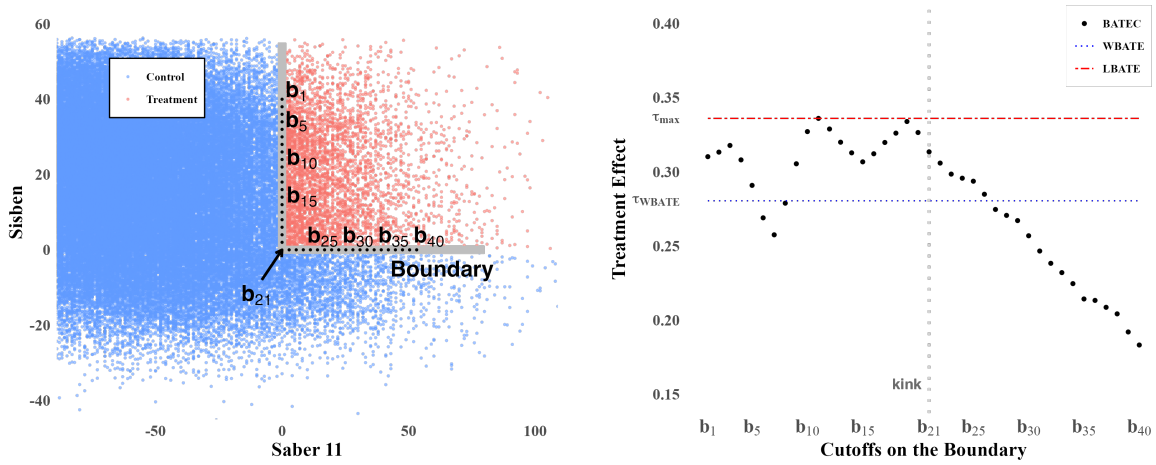

(a) Scatterplot and Boundary. (b) Treatment Effects Along Boundary.

Figure 1: Scatterplot, Assignment Boundary, and Treatment Effects Using SPP data.
Note: Panel (a) presents a scatterplot of the bivariate score $\mathbf{X}_i$ using the SPP data, and also plots the treatment boundary $\mathscr{B}$ with 40 marked grid points. Panel (b) presents causal treatment effect estimates over the 40 boundary grid points depicted in Panel (a). Specifically, the black solid dots correspond to $\widehat{\tau}(\mathbf{b}_j)$ (BATEC), the blue dotted line corresponds to $\widehat{\tau}_{\texttt{WBATE}}$ (WBATE), and the red dot-dash line corresponds to $\widehat{\tau}_{\texttt{LBATE}}$ (LBATE). The companion `R` software package `rd2d` is used for implementation; further details are available in the replication files and in Cattaneo et al. [2025c].

assignment boundary becomes $\mathscr{B} = \{(\texttt{SABER11}, \texttt{SISBEN}) : (\texttt{SABER11}, \texttt{SISBEN}) \in \{\texttt{SABER11} \geq 0 \text{ and } \texttt{SISBEN} = 0\} \cup \{\texttt{SABER11} = 0 \text{ and } \texttt{SISBEN} \geq 0\}\}$. Figure 1a presents a scatterplot of the bivariate score of the data of students in the 2014 cohort ($n = 363,096$ observations), and also plots the bivariate assignment boundary $\mathscr{B}$ together with 40 evenly-spaced cutoff points along the boundary, denoted by $\mathbf{b}_1, \ldots, \mathbf{b}_{40}$.

We begin by considering the functional causal parameter

$$\tau(\mathbf{x}) = \mathbb{E}[Y_i(1) - Y_i(0) | \mathbf{X}_i = \mathbf{x}], \qquad \mathbf{x} \in \mathscr{B},$$

which we call the *boundary average treatment effect curve* (BATEC) because it captures the average treatment effects at each point on the boundary. Our first goal is to conduct estimation and inference for $\tau(\mathbf{x})$, both pointwise for each $\mathbf{x} \in \mathscr{B}$ and uniformly over the entire boundary $\mathscr{B}$. In the SPP application, the outcome is $Y_i = 1$ if student i attended college or $Y_i = 0$ otherwise, and thus the causal parameter $\tau(\mathbf{x})$ captures the effect of SPP

on the probability of college education for students at the margin of program eligibility, as determined by their bivariate score $\mathbf{X}_i = (\texttt{SABER11}_i, \texttt{SISBEN}_i)^\top \in \mathscr{B}$. The parameter $\tau(\mathbf{x})$ captures policy-relevant heterogeneous treatment effects along the boundary $\mathscr{B}$: for example, in Figure 1a, $\tau(\mathbf{b}_1)$ is the average treatment effect at the boundary point $\mathbf{x} = \mathbf{b}_1$ where students have high SISBEN score (wealth) and low SABER11 score (academic), while $\tau(\mathbf{b}_{40})$ is the average treatment effect at the boundary point $\mathbf{x} = \mathbf{b}_{40}$, where students have low SISBEN score and high SABER11 score. Identification of these boundary treatment effects is analogous to identification in standard continuity-based univariate RD design [Hahn et al., 2001]: treatment assignment changes abruptly along the boundary $\mathscr{B}$, which implies that conditional expectations on each side of this boundary can be used to identify $\tau(\mathbf{x})$ whenever there is no systematic "sorting" of units into treatment and control, that is, whenever $\mathbb{E}[Y_i(0)|\mathbf{X}_i = \mathbf{x}]$ and $\mathbb{E}[Y_i(1)|\mathbf{X}_i = \mathbf{x}]$ are continuous for all $\mathbf{x} \in \mathscr{B}$ (Assumption 1 below). For more discussion, see Papay et al. [2011], Reardon and Robinson [2012], Wong et al. [2013], Keele and Titiunik [2015], and references therein.

The boundary average treatment effect curve, $\tau(\mathbf{x})$ for $\mathbf{x} \in \mathscr{B}$, captures the heterogeneity of causal treatment effects along the boundary. In applications, researchers may also wish to summarize this heterogeneity in a single causal parameter. This can be achieved by considering functional transformations of $\tau(\mathbf{x})$. We consider two leading examples of aggregated causal effects along the boundary: the *weighted boundary average treatment effect* (WBATE), and the *largest boundary average treatment effect* (LBATE). The WBATE parameter is

$$\tau_{\texttt{WBATE}} = \frac{\int_{\mathscr{B}} \tau(\mathbf{x}) w(\mathbf{x}) d\mathfrak{H}(\mathbf{x})}{\int_{\mathscr{B}} w(\mathbf{x}) d\mathfrak{H}(\mathbf{x})},$$

where $\mathfrak{H}$ denotes the one-dimensional Hausdorff measure, and the integrals are defined via Carathéodory's Theorem following classical measure theory results [Simon et al., 1984, Federer, 2014]. If the boundary is "nice" enough, the integrals in the definition of $\tau_{\texttt{WBATE}}$ may sometimes be represented as simple line integrals. Intuitively, $\tau_{\texttt{WBATE}}$ "averages" all the (po-

tentially heterogeneous) treatment effects $\tau(\mathbf{x})$ at each boundary point. See Reardon and Robinson [2012] and Wong et al. [2013] for more discussion. The LBATE parameter is

$$\tau_{\texttt{LBATE}} = \sup_{\mathbf{x}\in\mathscr{B}} \tau(\mathbf{x}),$$

another aggregate causal parameter that captures the "best" causal treatment effect along the assignment boundary. See, e.g., Andrews et al. [2024] for a discussion on this type of extreme treatment effects for policy decisions. In the SPP application, $\tau_{\texttt{WBATE}}$ and $\tau_{\texttt{LBATE}}$ capture the average and largest causal effect of receiving the college subsidy across all students near the boundary of eligibility, respectively.

Motivated by the local-to-boundary identifiability of $\tau(\mathbf{x})$, it is natural to employ flexible regression methods using only observations whose scores are near each boundary point. Local polynomial methods are the preferred choice for estimation and inference because they rely on simple (weighted) linear regression estimates that intuitively incorporate localization to $\mathbf{x} \in \mathscr{B}$, while also retaining most of the familiar features of least squares regression. We study the statistical properties of these location-based treatment effect estimators, and present pointwise and uniform estimation and inference methods for $\tau(\mathbf{x})$ as well for $\tau_{\texttt{WBATE}}$ and $\tau_{\texttt{LBATE}}$.

As a first illustration, Figure 1b presents the estimated average treatment effects at 40 different boundary points depicted in Figure 1a, denoted by $\widehat{\tau}(\mathbf{b}_j)$, $j = 1, \ldots, 40$, as well point estimates of WBATE and LBATE, denoted by $\widehat{\tau}_{\texttt{WBATE}}$ and $\widehat{\tau}_{\texttt{LBATE}}$, respectively. This initial results motivate the importance of conducting estimation and inference for the boundary average treatment effect curve to capture heterogeneity: for students who are marginal in terms of merit, the program appears to have a roughly constant effect along wealth levels, but for students who are marginal on wealth, the program appears to have a decreasing effect as academic performance improves. This finding suggests that students from under-privileged backgrounds may benefit more from SPP than wealthier and high-achieving students. We offer formal estimation and inference methods to support this initial finding, and for aggre-

gation of heterogeneous treatment effects along the boundary.

1.1 Organization and Related Literature

Section 2 introduces the formal causal inference framework and the location-based estimation approach (Assumption 1 below), which generalizes the classical RD design setup to bivariate scores and treatment assignment rules. For a review of the RD literature see Cattaneo and Titiunik [2022], and for a practical introduction see Cattaneo et al. [2020, 2024c]. We provide foundational theory and methods for BD designs discussed in Cattaneo et al. [2024c, Chapter 5]; see also Jardim et al. [2024]. Seminal empirical examples in the social sciences include Card and Krueger [1994], Black [1999] and Dell [2010]; see Cattaneo et al. [2026] for many more empirical examples.

Section 3 presents pointwise (for each $\mathbf{x} \in \mathscr{B}$), integrated mean square (over $\mathscr{B}$), and uniform (over $\mathscr{B}$) estimation and inference methods for the location-based local polynomial treatment effect estimator. Our pointwise results are generalizations of well-known results in nonparametric literature [e.g., Härdle et al., 2004, and references therein]: the main innovation underlying those results is to present a new regularity conditions on the bivariate kernel function and boundary $\mathscr{B}$ (Assumption 2 below). We discuss the pointwise results as a building block for our integrated mean square and uniform results, which require new theoretical development and thus provide a novel contribution to the literature. The main technical issue is that our mean square and uniform results are established over the one-dimensional manifold $\mathscr{B}$, and thus its shape can affect the statistical validity of those estimation and inference methods. We also discuss new bandwidth selection methods based on mean square error (MSE) expansions, robust bias-corrected inference, and related implementation details. Our results provide natural generalizations of well-established results for univariate RD designs; see Calonico et al. [2020] for bandwidth selection, and Calonico et al. [2014, 2018, 2022] for robust bias correction. See also Papay et al. [2011] for an early methodological discussion.

Section 4 studies the aggregated causal parameter $\tau_{\texttt{WBATE}}$, which corresponds to an integral

over a one-dimensional manifold of a two-dimensional (difference of) conditional expectation functions. This type of object has only been studied in the literature very recently: in concurrent work, Chen and Gao [2025] developed estimation and inference methods for integral functionals on submanifolds using series/sieve estimation. Our theoretical results also concern integral functionals on submanifolds, but when using local polynomial regression as the underlying nonparametric ingredient. Therefore, our estimation and inference results for $\tau_{\texttt{WBATE}}$ are new to the literature. See Reardon and Robinson [2012] and Wong et al. [2013] for early methodological discussion. Our large-sample nonparametric results complement the design-based methods developed in Keele et al. [2015] and Diaz and Zubizarreta [2023], and the Bayesian methods developed in Rischard et al. [2021] for BD designs.

Section 5 studies the aggregated causal parameter $\tau_{\texttt{LBATE}}$, and presents estimation and inference methods based on the location-based local polynomial treatment effect estimator. These results offer new treatment effect estimation and causal inference in the context of BD designs, previously unavailable in the literature.

Section 6 deploys our theoretical and methodological results to the SPP data, revising the main results reported in Londoño-Vélez, Rodríguez, and Sánchez [2020]. In addition to providing further empirical evidence in favor of their empirical findings, we also find some evidence of treatment effect heterogeneity along the assignment boundary $\mathscr{B}$. All empirical results are obtained using the companion `R` package `rd2d`, and we provide complete replication files.

Section 7 discusses two extensions of our work. First, we consider imperfect compliance (fuzzy) BD designs. Second, we consider the role of pre-intervention covariates for efficiency improvements and heterogeneity analysis. We outline how our theoretical work can be directly deployed, or easily extended, to develop new point estimation and inference procedures for each generalization.

Section 8 concludes. The supplemental appendix presents more general theoretical results, reports all technical proofs, and gives other results that may be of independent interest. In particular, Section SA-6 presents a new strong approximation theorem for empirical pro-

cess with polynomial bounded moments, which generalizes recent work in Cattaneo and Yu [2025]. Finally, our companion software article [Cattaneo et al., 2025c] discusses the general-purpose `R` software package `rd2d` (`https://rdpackages.github.io/rd2d`) implementing the methods developed in this paper.

1.2 Notation

We employ standard concepts and notations from empirical process theory [van der Vaart and Wellner, 1996, Giné and Nickl, 2016] and geometric measure theory [Simon et al., 1984, Federer, 2014]. For a random variable $\mathbf{V}_i$, we write $\mathbb{E}_n[g(\mathbf{V}_i)] = n^{-1}\sum_{i=1}^n g(\mathbf{V}_i)$. For a Borel set $\mathcal{S} \subseteq \mathcal{X}$, the De Giorgi perimeter of $\mathcal{S}$ is $\texttt{perim}(\mathcal{S}) = \sup_{g\in\mathscr{D}_2(\mathcal{X})} \int_{\mathbb{R}^2} \mathbb{1}(\mathbf{x} \in \mathcal{S}) \operatorname{div} g(\mathbf{x}) d\mathbf{x} / \|g\|_\infty$, where div is the divergence operator, and $\mathscr{D}_2(\mathcal{X})$ denotes the space of C^∞ functions with values in $\mathbb{R}^2$ and with compact support included in $\mathcal{X}$. When $\mathcal{S}$ is connected, and the boundary $\texttt{bd}(\mathcal{S})$ is a smooth simple closed curve, $\texttt{perim}(\mathcal{S})$ simplifies to the curve length of $\texttt{bd}(\mathcal{S})$. A curve $\mathscr{B} \subseteq \mathbb{R}^2$ is a *rectifiable curve* if there exists a Lipschitz continuous function $\gamma : [0,1] \mapsto \mathbb{R}^2$ such that $\mathscr{B} = \gamma([0,1])$. $\mathfrak{H}$ denotes the one-dimensional Hausdorff (outer) measure, and integration against $\mathfrak{H}$ is defined using classical (geometric) measure theory. For a function $f : \mathbb{R}^2 \mapsto \mathbb{R}$, $\operatorname{Supp}(f)$ denotes closure of the set $\{\mathbf{x} \in \mathbb{R}^2 : f(\mathbf{x}) \neq 0\}$. For reals sequences $a_n = o(b_n)$ if $\limsup_{n\to\infty} \frac{|a_n|}{|b_n|} = 0$, $a_n \lesssim b_n$ if there exists some constant C and $N > 0$ such that $n > N$ implies $|a_n| \leq C|b_n|$. For sequences of random variables $a_n = o_{\mathbb{P}}(b_n)$ if $\operatorname{plim}_{n\to\infty} \frac{|a_n|}{|b_n|} = 0, |a_n| \lesssim_{\mathbb{P}} |b_n|$ if $\limsup_{M\to\infty} \limsup_{n\to\infty} \mathbb{P}[|\frac{a_n}{b_n}| \geq M] = 0$. Finally, $\Phi(x)$ denotes the standard Gaussian cumulative distribution function.

2 Setup

We impose the following basic conditions on the underlying data generating process.

Assumption 1 (Data Generating Process)**.** *Let $t \in \{0,1\}$, $p \geq 1$, and $v \geq 2$.*

(i) $(Y_1(t), \mathbf{X}_1^\top)^\top, \ldots, (Y_n(t), \mathbf{X}_n^\top)^\top$ *are independent and identically distributed random vectors.*

(ii) *The distribution of* $\mathbf{X}_i$ *has a Lebesgue density* $f_X(\mathbf{x})$ *that is continuous and bounded away from zero on its support* $\mathcal{X} = [\mathtt{L}, \mathtt{U}]^2$*, for* $-\infty < \mathtt{L} < \mathtt{U} < \infty$.

(iii) $\mu_t(\mathbf{x}) = \mathbb{E}[Y_i(t)|\mathbf{X}_i = \mathbf{x}]$ *is* $(p+1)$*-times continuously differentiable on* $\mathcal{X}$.

(iv) $\sigma_t^2(\mathbf{x}) = \mathbb{V}[Y_i(t)|\mathbf{X}_i = \mathbf{x}]$ *is bounded away from zero and continuous on* $\mathcal{X}$.

(v) $\sup_{\mathbf{x}\in\mathcal{X}} \mathbb{E}[|Y_i(t)|^{2+v}|\mathbf{X}_i = \mathbf{x}] < \infty$.

These conditions are analogous to the usual conditions imposed in the classical RD design with univariate score. Part (ii) goes beyond the usual compact support restriction and further assumes a tensor product structure on $\mathcal{X}$ to avoid technical difficulties in our strong approximations for uniform distribution theory; this condition is not practically restrictive because all methods considered in this paper localize to the boundary, which is in the interior of $\mathcal{X}$. Part (iii) imposes standard smoothness conditions on the bivariate conditional expectation functions of interest, which will play an important role in misspecification (smoothing) bias reduction in our upcoming results. Nonparametric identification of $\tau(\mathbf{x})$ follows directly from Assumption 1: see Papay et al. [2011], Reardon and Robinson [2012], Keele and Titiunik [2015], and references therein.

The location-based estimator of the boundary average treatment effect curve $\tau(\mathbf{x})$ is

$$\widehat{\tau}(\mathbf{x}) = \mathbf{e}_1^\top \widehat{\boldsymbol{\beta}}_1(\mathbf{x}) - \mathbf{e}_1^\top \widehat{\boldsymbol{\beta}}_0(\mathbf{x}), \qquad \mathbf{x} \in \mathscr{B},$$

where, for $t \in \{0, 1\}$,

$$\widehat{\boldsymbol{\beta}}_t(\mathbf{x}) = \underset{\boldsymbol{\beta}\in\mathbb{R}^{\mathfrak{p}_p+1}}{\arg\min}\, \mathbb{E}_n\Big[\big(Y_i - \mathbf{r}_p(\mathbf{X}_i - \mathbf{x})^\top \boldsymbol{\beta}\big)^2 K_h(\mathbf{X}_i - \mathbf{x}) \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\Big], \tag{1}$$

with $\mathfrak{p}_p = (2+p)(1+p)/2 - 1$, $\mathbf{r}_p(\mathbf{u}) = (1, u_1, u_2, u_1^2, u_1u_2, u_2^2, \cdots, u_1^p, u_1^{p-1}u_2, \cdots, u_1u_2^{p-1}, u_2^p)^\top$ the pth order polynomial expansion of the bivariate vector $\mathbf{u} = (u_1, u_2)^\top$, and $K_h(\mathbf{u}) = K(u_1/h, u_2/h)/h^2$ for a bivariate kernel function $K(\cdot)$ and bandwidth parameter h. We employ the same bandwidth for each dimension of $\mathbf{X}_i$ only for simplicity, and because it is common practice to first standardize each dimension of the bivariate score before imple-

menting the estimator $\widehat{\tau}(\mathbf{x})$.

We impose the following assumption on the bivariate kernel function and assignment boundary.

Assumption 2 (Kernel and Boundary)**.** *Let* $t \in \{0,1\}$.

(i) $\mathscr{B}$ *is a rectifiable curve.*

(ii) $K : \mathbb{R}^2 \to [0,\infty)$ *is compact-supported and Lipschitz continuous, or* $K(\mathbf{u}) = \mathbb{1}(\mathbf{u} \in [-1,1]^2)$.

(iii) *There exists* $U \subseteq \mathbb{R}^2$*, such that* $K(\mathbf{u}) \geq \kappa > 0$ *for all* $\mathbf{u} \in U$*,* $\lambda_{\min}(\int_U \mathbf{r}_p(\mathbf{z})\mathbf{r}_p(\mathbf{z})^\top d\mathbf{z}) > 0$*, and the integration satisfies* $\liminf_{h\downarrow 0} \inf_{\mathbf{x}\in\mathscr{B}} \int_U K(\mathbf{u})\mathbb{1}(\mathbf{x}+h\mathbf{u} \in \mathscr{A}_t)d\mathbf{u} \gtrsim 1$.

Assumption 2(i) imposes minimal regularity on the one-dimensional assignment boundary $\mathscr{B}$, which is useful to compute integrals over that one-dimensional manifold, and to establish uniform and aggregated results. Assumption 2(ii) is standard in the literature. Assumption 2(iii) is new to the literature, and crucial for the asymptotic analysis. It restricts the geometry of $\mathscr{B}$ as well as it interacts with the kernel K; the density of the data $f_X(\mathbf{x})$ is not explicitly present because, by Assumption 1(ii), it is bounded away from zero. Heuristically, this assumption ensures that the kernel "splits" the (roughly locally uniform) mass between treatment and control regions "evenly" given the local shape of $\mathscr{B}$. In the supplemental appendix (Lemma SA-1), we carefully show that Assumptions 1(ii) and 2 imply that the population Gram matrix $\mathbf{\Gamma}_{t,\mathbf{x}} = \mathbb{E}\big[\mathbf{r}_p\big(\frac{\mathbf{X}_i-\mathbf{x}}{h}\big)\mathbf{r}_p\big(\frac{\mathbf{X}_i-\mathbf{x}}{h}\big)^\top K_h(\mathbf{X}_i-\mathbf{x})\mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\big] \gtrsim \int_{\mathscr{A}_t} \mathbf{r}_p\big(\frac{\mathbf{X}_i-\mathbf{x}}{h}\big)\mathbf{r}_p\big(\frac{\mathbf{X}_i-\mathbf{x}}{h}\big)^\top K_h(\mathbf{X}_i-\mathbf{x})$ is full rank for h small enough, uniformly in $\mathbf{x} \in \mathscr{B}$. This result is novel to the literature, and guarantees that the treatment effect estimator $\widehat{\tau}(\mathbf{x})$ will be well-defined in large samples, under the conditions imposed in this paper. Assumption 2(iii) is stated uniformly over $\mathscr{B}$ in anticipation of our uniform estimation and inference results, but it can be relaxed to hold only pointwise in $\mathbf{x} \in \mathscr{B}$ for our pointwise results; see the supplemental appendix for details.

Finally, to ensure that the aggregated weighted boundary average treatment effect $\tau_{\texttt{WBATE}}$ is well-defined, we impose the following conditions on the weight function.

Assumption 3 (Weight Function and Boundary). *Let* $w : \mathscr{B} \mapsto \mathbb{R}$ *with* $\sup_{\mathbf{x}\in\mathscr{B}} |w(\mathbf{x})| < \infty$, $\inf_{\mathbf{x}\in\mathscr{B}} |w(\mathbf{x})| > 0$, *and* $\int_{\mathscr{B}} |w(\mathbf{x})| d\mathfrak{H}(\mathbf{x}) < \infty$.

In particular, $w(\mathbf{x}) = f_X(\mathbf{x})$ satisfies Assumptions 3. In this case, the WBATE reduces to the *boundary average treatment effect* (BATE):

$$\tau_{\texttt{BATE}} = \int_{\mathscr{B}} \tau(\mathbf{x}) f(\mathbf{x}|\mathbf{X}_i \in \mathscr{B}) d\mathfrak{H}(\mathbf{x}), \qquad f(\mathbf{x}|\mathbf{X}_i \in \mathscr{B}) = \frac{f_X(\mathbf{x})}{\int_{\mathscr{B}} f_X(\mathbf{x}) d\mathfrak{H}(\mathbf{x})}.$$

This causal parameter corresponds to the density-weighted average causal effect along the assignment boundary, and is discussed by Wong et al. [2013]. See Cattaneo et al. [2025b] for an alternative regression-based approach commonly used in practice to estimate the BATE.

3 Boundary Average Treatment Effect Curve

Given Assumption 1, pointwise and uniform point estimation results for $(\tau(\mathbf{x}) : \mathbf{x} \in \mathscr{B})$ follow from standard local polynomial calculations and empirical process theory; the only technical issue is related to the geometry of the boundary $\mathscr{B}$, which is handled via Assumption 2. On the other hand, our integrated mean square expansion and uniform distribution results require new theoretical developments. All details are in the supplemental appendix.

3.1 Treatment Effect Estimation

Using standard concentration techniques from empirical process theory, we obtain the pointwise and uniform convergence rate of $\widehat{\tau}(\mathbf{x})$.

Theorem 1 (Convergence Rates). *Suppose Assumptions 1 and 2 hold. If* $nh^2/\log(1/h) \to \infty$ *and* $h \to 0$, *then*

(i) $|\widehat{\tau}(\mathbf{x}) - \tau(\mathbf{x})| \lesssim_{\mathbb{P}} \frac{1}{\sqrt{nh^2}} + \frac{1}{n^{\frac{1+v}{2+v}} h^2} + h^{p+1}$ *for* $\mathbf{x} \in \mathscr{B}$, *and*

(ii) $\sup_{\mathbf{x}\in\mathscr{B}} |\widehat{\tau}(\mathbf{x}) - \tau(\mathbf{x})| \lesssim_{\mathbb{P}} \sqrt{\frac{\log(1/h)}{nh^2}} + \frac{\log(1/h)}{n^{\frac{1+v}{2+v}} h^2} + h^{p+1}$.

This theorem immediately establishes consistency of the treatment effect estimator based on the bivariate location score. Notably, the theorem shows that the bias is of order h^{p+1}

regardless of whether there are kinks or other irregularities in $\mathscr{B}$. See Cattaneo et al. [2025a] for more discussion.

Given its standard structure, it is easy to establish a pointwise (conditional on $\mathbf{X}$) MSE expansion for the estimator $\widehat{\tau}(\mathbf{x})$ for each $\mathbf{x} \in \mathscr{B}$. Furthermore, employing tools from geometric measure theory, it is also possible to establish an integrated (conditional on $\mathbf{X}$) MSE expansion along the boundary. Using standard multi-index notation, the leading (pointwise) conditional bias is $B_{\mathbf{x}} = B_{1,\mathbf{x}} - B_{0,\mathbf{x}}$ with

$$B_{t,\mathbf{x}} = \mathbf{e}_1^\top \widehat{\mathbf{\Gamma}}_{t,\mathbf{x}}^{-1} \sum_{|\mathbf{k}|=p+1} \frac{\mu_t^{(\mathbf{k})}(\mathbf{x})}{\mathbf{k}!} \mathbb{E}\Big[\mathbf{r}_p\Big(\frac{\mathbf{X}_i - \mathbf{x}}{h}\Big)\Big(\frac{\mathbf{X}_i - \mathbf{x}}{h}\Big)^{\mathbf{k}} K_h(\mathbf{X}_i - \mathbf{x}) \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\Big]$$

and $\widehat{\mathbf{\Gamma}}_{t,\mathbf{x}} = \mathbb{E}_n\big[\mathbf{r}_p\big(\frac{\mathbf{X}_i-\mathbf{x}}{h}\big)\mathbf{r}_p\big(\frac{\mathbf{X}_i-\mathbf{x}}{h}\big)^\top K_h(\mathbf{X}_i - \mathbf{x})\mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\big]$, for $t \in \{0,1\}$. Similarly, the leading (pointwise) conditional variance is $V_{\mathbf{x}} = V_{1,\mathbf{x}} + V_{0,\mathbf{x}}$ with $V_{t,\mathbf{x}} = \mathbf{e}_1^\top \widehat{\mathbf{\Gamma}}_{t,\mathbf{x}}^{-1} \mathbf{\Sigma}_{t,\mathbf{x},\mathbf{x}} \widehat{\mathbf{\Gamma}}_{t,\mathbf{x}}^{-1} \mathbf{e}_1$ and

$$\mathbf{\Sigma}_{t,\mathbf{x},\mathbf{x}} = h^2 \mathbb{E}\Big[\mathbf{r}_p\Big(\frac{\mathbf{X}_i - \mathbf{x}}{h}\Big)\mathbf{r}_p\Big(\frac{\mathbf{X}_i - \mathbf{x}}{h}\Big)^\top K_h(\mathbf{X}_i - \mathbf{x})^2 \sigma_t^2(\mathbf{X}_i) \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\Big],$$

for $t \in \{0,1\}$. The following theorem gives the MSE expansions. Let $\mathbf{X} = (\mathbf{X}_1^\top, \ldots, \mathbf{X}_n^\top)$.

Theorem 2 (MSE Expansions)**.** *Suppose Assumptions 1, 2, and 3 hold. If* $nh^2/\log(1/h) \to \infty$ *and* $h \to 0$*, then*

(i) $\mathbb{E}[(\widehat{\tau}(\mathbf{x}) - \tau(\mathbf{x}))^2 | \mathbf{X}] = h^{2(p+1)} B_{\mathbf{x}}^2 + \frac{1}{nh^2} V_{\mathbf{x}} + o_{\mathbb{P}}(\mathfrak{R}_n)$*, and*

(ii) $\int_{\mathscr{B}} \mathbb{E}[(\widehat{\tau}(\mathbf{x}) - \tau(\mathbf{x}))^2 | \mathbf{X}] w(\mathbf{x}) d\mathfrak{H}(\mathbf{x}) = h^{2(p+1)} \int_{\mathscr{B}} B_{\mathbf{x}}^2 w(\mathbf{x}) d\mathfrak{H}(\mathbf{x}) + \frac{1}{nh^2} \int_{\mathscr{B}} V_{\mathbf{x}} w(\mathbf{x}) d\mathfrak{H}(\mathbf{x}) + o_{\mathbb{P}}(\mathfrak{R}_n)$,

with $\mathfrak{R}_n = h^{2p+2} + n^{-1}h^{-2} + n^{-\frac{2(1+v)}{2+v}} h^{-4}$.

Ignoring the higher-order terms, the approximate MSE-optimal and IMSE-optimal bandwidth choices are

$$h_{\mathtt{MSE},\mathbf{x}} = \Big(\frac{2V_{\mathbf{x}}}{(2p+2)B_{\mathbf{x}}^2}\frac{1}{n}\Big)^{1/(2p+4)} \quad \text{and} \quad h_{\mathtt{IMSE}} = \Big(\frac{2\int_{\mathscr{B}} V_{\mathbf{x}} w(\mathbf{x}) d\mathfrak{H}(\mathbf{x})}{(2p+2)\int_{\mathscr{B}} B_{\mathbf{x}}^2 w(\mathbf{x}) d\mathfrak{H}(\mathbf{x})}\frac{1}{n}\Big)^{1/(2p+4)},$$

provided that $B_{\mathbf{x}} \neq 0$ and $\int_{\mathscr{B}} B_{\mathbf{x}}^2 \omega(\mathbf{x}) d\mathfrak{H}(\mathbf{x}) \neq 0$, respectively. These choices are infeasible because a preliminary bandwidth, as well as estimates of the the conditional variances and higher-order derivatives of the conditional mean, are needed. We discuss implementation in Section 3.3, and in our companion software article [Cattaneo et al., 2025c]. It follows from Theorems 1 and 2 that, for an appropriate choice of tuning parameters, the estimator $\widehat{\tau}(\mathbf{x})$ can achieve the usual nonparametric optimal convergence rates [see, e.g., Tsybakov, 2008, for a textbook review]. Therefore, MSE-optimal and IMSE-optimal location-based treatment effect estimators of $\tau(\mathbf{x})$ can be constructed using $h_{\texttt{MSE},\mathbf{x}}$ and $h_{\texttt{IMSE}}$, respectively.

Most of the results in this section are standard in the literature. The one exception is the integrated MSE expansion in Theorem 2, which requires additional care to handle the integral over the one-dimensional manifold $\mathscr{B}$. We present these point estimation results because they will play an important role in our upcoming uniform (over $\mathscr{B}$) and aggregation (along $\mathscr{B}$) analyses of the boundary average treatment effect curve.

3.2 Uncertainty Quantification

Given a bandwidth choice, the feasible t-statistic is

$$\widehat{\mathrm{T}}(\mathbf{x}) = \frac{\widehat{\tau}(\mathbf{x}) - \tau(\mathbf{x})}{\sqrt{\widehat{\Omega}_{\mathbf{x},\mathbf{x}}}}, \qquad \mathbf{x} \in \mathscr{B},$$

where, using standard least squares algebra, the variance estimator is

$$\widehat{\Omega}_{\mathbf{x}_1,\mathbf{x}_2} = \widehat{\Omega}_{0,\mathbf{x}_1,\mathbf{x}_2} + \widehat{\Omega}_{1,\mathbf{x}_1,\mathbf{x}_2}, \qquad \widehat{\Omega}_{t,\mathbf{x}_1,\mathbf{x}_2} = \frac{1}{nh^2} \mathbf{e}_1^\top \widehat{\mathbf{\Gamma}}_{t,\mathbf{x}_1}^{-1} \widehat{\mathbf{\Sigma}}_{t,\mathbf{x}_1,\mathbf{x}_2} \widehat{\mathbf{\Gamma}}_{t,\mathbf{x}_2}^{-1} \mathbf{e}_1$$

with

$$\widehat{\mathbf{\Sigma}}_{t,\mathbf{x}_1,\mathbf{x}_2} = h^2 \mathbb{E}_n\Big[\mathbf{r}_p\Big(\frac{\mathbf{X}_i - \mathbf{x}_1}{h}\Big)\mathbf{r}_p\Big(\frac{\mathbf{X}_i - \mathbf{x}_2}{h}\Big)^\top K_h(\mathbf{X}_i - \mathbf{x}_1) K_h(\mathbf{X}_i - \mathbf{x}_2) \widehat{\varepsilon}_i(\mathbf{x}_1) \widehat{\varepsilon}_i(\mathbf{x}_2) \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\Big]$$

and $\widehat{\varepsilon}_i(\mathbf{x}) = Y_i - \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_0)\mathbf{R}_p(\mathbf{X}_i - \mathbf{x})^\top \widehat{\boldsymbol{\beta}}_0(\mathbf{x}) - \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_1)\mathbf{R}_p(\mathbf{X}_i - \mathbf{x})^\top \widehat{\boldsymbol{\beta}}_1(\mathbf{x})$, for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathscr{B}$ and $t \in \{0, 1\}$.

Wald-type feasible confidence intervals and confidence bands over $\mathscr{B}$ take the form:

$$\widehat{\mathrm{I}}_\alpha(\mathbf{x}) = \left[\, \widehat{\tau}(\mathbf{x}) - \mathscr{q}_\alpha \sqrt{\widehat{\Omega}_{\mathbf{x},\mathbf{x}}} \;,\; \widehat{\tau}(\mathbf{x}) + \mathscr{q}_\alpha \sqrt{\widehat{\Omega}_{\mathbf{x},\mathbf{x}}} \,\right], \qquad \mathbf{x} \in \mathscr{B},$$

for any $\alpha \in (0,1)$, and where $\mathscr{q}_\alpha$ denotes the appropriate quantile depending on the desired inference procedure. For pointwise inference, it is a textbook exercise to show that $\sup_{t\in\mathbb{R}} |\mathbb{P}[\widehat{\mathrm{T}}(\mathbf{x}) \leq t] - \Phi(t)| \to 0$ for each $\mathbf{x} \in \mathscr{B}$, under standard regularity conditions, and provided that the "small bias" condition $nh^{2p+4} \to 0$ holds. This result can then be used to construct the usual confidence intervals for $\tau(\mathbf{x})$.

For uniform inference (over $\mathscr{B}$), two challenges arise: (i) the stochastic process $(\widehat{\mathrm{T}}(\mathbf{x}) : \mathbf{x} \in \mathscr{B})$ is not asymptotically tight, and thus it does not converge weakly in the space of uniformly bounded real functions supported on $\mathscr{B}$ and equipped with the uniform norm [van der Vaart and Wellner, 1996, Giné and Nickl, 2016]; and (ii) the geometry of the manifold $\mathscr{B}$ can affect the validity of the distributional approximation (this is a new problem specific to this paper). To circumvent both problems, we first approximate the distribution of the entire non-Donsker stochastic process $(\widehat{\mathrm{T}}(\mathbf{x}) : \mathbf{x} \in \mathscr{B})$, and we then deduce a distributional approximation for $\sup_{\mathbf{x}\in\mathscr{B}} |\widehat{\mathrm{T}}(\mathbf{x})|$. This approach enable us to construct asymptotically valid confidence bands because

$$\mathbb{P}\big[\tau(\mathbf{x}) \in \widehat{\mathrm{I}}_\alpha(\mathbf{x}) \;,\; \text{for all } \mathbf{x} \in \mathscr{B}\big] = \mathbb{P}\Big[\sup_{\mathbf{x}\in\mathscr{B}} \big|\widehat{\mathrm{T}}(\mathbf{x})\big| \leq \mathscr{q}_\alpha\Big].$$

Our technical results are established via a new strong approximation theorem (Section SA-7 in the supplemental appendix), combined with technical results from Cattaneo and Yu [2025], Cattaneo et al. [2024a], Chernozhukov et al. [2014a,b], Chernozhuokov et al. [2022], and Dudley [2014]. Let $\mathbf{W} = (Y_1, \ldots, Y_n, \mathbf{X})$.

Theorem 3 (Confidence Intervals and Bands)**.** *Suppose Assumptions 1 and 2 hold.*

(i) *For all* $\mathbf{x} \in \mathscr{B}$*, if* $n^{\frac{v}{2+v}} h^2 \to \infty$ *and* $nh^{2p+4} \to 0$*, then*

$$\mathbb{P}\big[\tau(\mathbf{x}) \in \widehat{\mathrm{I}}_\alpha(\mathbf{x})\big] \to 1 - \alpha,$$

for $\mathcal{q}_\alpha = \Phi^{-1}(1-\alpha/2)$.

(ii) *If* $n^{\frac{v}{2+v}}h^2/\log n \to \infty$, $\liminf_{n\to\infty}\frac{\log h}{\log n} > -\infty$, $nh^{2p+4} \to 0$, *and* $\texttt{perim}(\{\mathbf{y} \in \mathscr{A}_t : (\mathbf{y}-\mathbf{x})/h \in \mathrm{Supp}(K)\}) \lesssim h$ *for all* $\mathbf{x} \in \mathscr{B}$ *and* $t \in \{0,1\}$, *then*

$$\mathbb{P}\big[\tau(\mathbf{x}) \in \widehat{\mathrm{I}}_\alpha(\mathbf{x}),\ \textit{for all}\ \mathbf{x} \in \mathscr{B}\big] \to 1-\alpha,$$

for $\mathcal{q}_\alpha = \inf\{c > 0 : \mathbb{P}[\sup_{\mathbf{x}\in\mathscr{B}}|\widehat{Z}_n(\mathbf{x})| \geq c|\mathbf{W}] \leq \alpha\}$, *where* $(\widehat{Z}_n : \mathbf{x} \in \mathscr{B})$ *is a Gaussian process conditional on* $\mathbf{W}$ *with* $\mathbb{E}[\widehat{Z}_n(\mathbf{x}_1)|\mathbf{W}] = 0$ *and* $\mathbb{E}[\widehat{Z}_n(\mathbf{x}_1)\widehat{Z}_n(\mathbf{x}_2)|\mathbf{W}] = \widehat{\Omega}_{\mathbf{x}_1,\mathbf{x}_2}/\sqrt{\widehat{\Omega}_{\mathbf{x}_1,\mathbf{x}_1}\widehat{\Omega}_{\mathbf{x}_2,\mathbf{x}_2}}$, *for all* $\mathbf{x}_1, \mathbf{x}_2 \in \mathscr{B}$.

This theorem gives asymptotically valid pointwise and uniform uncertainty quantification for the conditional treatment effect $\tau(\mathbf{x})$ using the location-based estimator $\widehat{\tau}(\mathbf{x})$. As expected in a nonparametric smoothing setting, the undersmoothing condition $nh^{2p+4} \to 0$ rules out the (I)MSE-optimal point estimator of $\tau(\mathbf{x})$. Thus, for implementation of both pointwise and uniform inference, we address this issue via robust bias correction [Calonico et al., 2014, 2018, 2022]; see Section 3.3 for details. For uniform inference, our method requires an additional technical restriction on $\mathscr{B}$ in order to avoid an overly "wiggly" assignment boundary that would lead to invalid statistical inference. Intuitively, this condition rules out one-dimensional manifolds that may be "smooth and nice" but nonetheless "too long".

3.3 Implementation

The bivariate local polynomial estimator $\widehat{\tau}(\mathbf{x})$ and associated t-statistic $\widehat{\mathrm{T}}(\mathbf{x})$ are fully adaptive to kinks and other irregularities of the boundary $\mathscr{B}$, provided Assumption 2(ii) holds and the boundary is not too "wiggly". Therefore, it is straightforward to implement local and global bandwidth selectors based on Theorem 2. In particular, replacing the asymptotic bias and variance constants, $B_\mathbf{x}$ and $V_\mathbf{x}$, with preliminary estimators, we obtain the feasible

plug-in bandwidth selectors

$$\widehat{h}_{\mathtt{MSE},\mathbf{x}} = \Big(\frac{2\widehat{V}_{\mathbf{x}}}{(2p+2)\widehat{B}_{\mathbf{x}}^2}\frac{1}{n}\Big)^{1/(2p+4)} \quad \text{and} \quad \widehat{h}_{\mathtt{IMSE}} = \Big(\frac{2\int_{\mathscr{B}}\widehat{V}_{\mathbf{x}}w(\mathbf{x})d\mathfrak{H}(\mathbf{x})}{(2p+2)\int_{\mathscr{B}}\widehat{B}_{\mathbf{x}}^2 w(\mathbf{x})d\mathfrak{H}(\mathbf{x})}\frac{1}{n}\Big)^{1/(2p+4)},$$

where, for a preliminary bandwidth choice $a \to 0$, $\widehat{B}_{\mathbf{x}} = \widehat{B}_{1,\mathbf{x}} - \widehat{B}_{0,\mathbf{x}}$ is constructed using

$$\widehat{B}_{t,\mathbf{x}} = \mathbf{e}_1^\top \widehat{\mathbf{\Gamma}}_{t,\mathbf{x}}^{-1} \sum_{|\mathbf{k}|=p+1} \frac{\widehat{\mu}_t^{(\mathbf{k})}(\mathbf{x})}{\mathbf{k}!}\mathbb{E}_n\Big[\mathbf{r}_p\Big(\frac{\mathbf{X}_i-\mathbf{x}}{a}\Big)\Big(\frac{\mathbf{X}_i-\mathbf{x}}{a}\Big)^{\mathbf{k}} K_a(\mathbf{X}_i-\mathbf{x})\mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\Big],$$

with $\widehat{\mathbf{\Gamma}}_{t,\mathbf{x}}$ computed using the preliminary bandwidth a, and where $\widehat{\mu}_t^{(\mathbf{k})}(\mathbf{x})$ a preliminary estimator of $\mu_t^{(\mathbf{k})}(\mathbf{x})$, and $\widehat{V}_{\mathbf{x}} = na^2\widehat{\mathbf{\Omega}}_{\mathbf{x},\mathbf{x}}$ is constructed using the variance estimator with the preliminary bandwidth a. Omitted implementation details are discussed in Cattaneo et al. [2025c]; see also Calonico et al. [2020] for a review on modern bandwidth selection methods in RD designs with univariate score.

The bandwidth choices $\widehat{h}_{\mathtt{MSE},\mathbf{x}}$ and $\widehat{h}_{\mathtt{IMSE}}$ can be used to implement (I)MSE-optimal $\widehat{\tau}(\mathbf{x})$ treatment effect estimators, both pointwise and uniformly over $\mathscr{B}$. Furthermore, leveraging the results in Theorem 3, a simple application of robust bias-corrected inference proceeds by employing the same (I)MSE-optimal bandwidth (for pth order point estimation), but then constructing the t-statistic $\widehat{\mathrm{T}}(\mathbf{x})$ with a $(p+1)$th polynomial order instead of pth polynomial order. The core idea is to simultaneously (i) debias the (I)MSE-optimal point estimator $\widehat{\tau}(\mathbf{x})$, and (ii) adjust the variance estimator to incorporate the additional uncertainty introduced by the bias correction. This inference approach has several theoretical advantages [Calonico et al., 2014, 2018, 2022], and has been validated empirically [Hyytinen et al., 2018, De Magalhães et al., 2025].

Finally, regarding the computation of the Gaussian process conditional on $\mathbf{W}$, $(\widehat{Z}_n(\mathbf{x}) : \mathbf{x} \in \mathscr{B})$, there are two methodological issues to consider. First, simulation is implemented over a grid of points forming a discretization of the index set of the continuous stochastic process $\widehat{Z}_n$; it is not difficult to show that as the number of points in the mesh increases, the approximation becomes more accurate. Second, the estimated (discretized) covariate

function may not be positive definite in finite samples, but this finite-sample issue can be easily fixed via regularization; see Cattaneo et al. [2024b] for a discussion and related technical results.

Our companion software package `rd2d` implements the procedures described in this section, see Cattaneo et al. [2025c]. In Section 6, we use our proposed estimation and inference methods to re-analyze the effects of the SPP program on college attendance.

4 Weighted Boundary Average Treatment Effect

Without loss of generality, we set $\int_{\mathscr{B}} w(\mathbf{x}) d\mathfrak{H}(\mathbf{x}) = 1$, and thus consider the (normalized) causal parameter

$$\tau_{\texttt{WBATE}} = \int_{\mathscr{B}} \tau(\mathbf{x}) w(\mathbf{x}) d\mathfrak{H}(\mathbf{x}),$$

where the weight function $w : \mathscr{B} \mapsto \mathbb{R}$ satisfies Assumption 3. The WBATE aggregates the heterogeneous treatment effects $(\tau(\mathbf{x}) : \mathbf{x} \in \mathscr{B})$ at each boundary point according to the chosen weighting scheme. Our results allow for similar causal parameters defined over a region of the boundary as in Reardon and Robinson [2012].

The (plug-in, location-based) WBATE estimator is

$$\widehat{\tau}_{\texttt{WBATE}} = \int_{\mathscr{B}} \widehat{\tau}(\mathbf{x}) w(\mathbf{x}) d\mathfrak{H}(\mathbf{x}).$$

In practice, this estimator can be computed by either analytic integration (when $\mathscr{B}$ is "simple" enough, e.g., via a line integral), or by forming a discretization of the assignment boundary (e.g., as in Figure 1a). For the latter approach, letting $(\mathbf{b}_j : j = 1, \ldots, J)$ be points on $\mathscr{B}$, the estimator can be computed as $\widehat{\tau}_{\texttt{WBATE}} \approx \sum_{j=1}^{J} \widehat{\tau}(\mathbf{b}_j) w(\mathbf{b}_j)$, where in some applications $w(\mathbf{b}_j)$ could be replaced by some data-driven quantity of interest such as $w(\mathbf{b}_j) = \frac{N_j}{\sum_{j=1}^{J} N_j}$ with $N_j = \sum_{i=1}^{n} \mathbb{1}(\|\mathbf{X}_i - \mathbf{b}_j\| \leq h)$. Furthermore, the parameter and estimator may be defined only for a segment of $\mathscr{B}$. See Cattaneo et al. [2025c] for more discussion on imple-

mentation.

Our first result establishes a MSE expansion for $\widehat{\tau}_{\texttt{WBATE}}$. We employ the following notation for the leading (integrated) bias and variance:

$$B_{\texttt{WBATE}} = B_{1,\texttt{WBATE}} - B_{0,\texttt{WBATE}}, \qquad B_{t,\texttt{WBATE}} = \int_{\mathscr{B}} B_{t,\mathbf{x}} w(\mathbf{x}) d\mathfrak{H}(\mathbf{x}),$$

and

$$\Omega_{\texttt{WBATE}} = \Omega_{1,\texttt{WBATE}} + \Omega_{0,\texttt{WBATE}}, \qquad \Omega_{t,\texttt{WBATE}} = \int_{\mathscr{B}}\int_{\mathscr{B}} \Omega_{t,\mathbf{x}_1,\mathbf{x}_2} w(\mathbf{x}_1) w(\mathbf{x}_2) d\mathfrak{H}(\mathbf{x}_1) d\mathfrak{H}(\mathbf{x}_2),$$

for $t \in \{0,1\}$.

Theorem 4 (MSE Expansion: WBATE)**.** *Suppose Assumptions 1, 2 and 3 hold. If* $nh^2/\log(1/h) \to \infty$ *and* $h \to 0$*, then*

$$\mathbb{E}\big[(\widehat{\tau}_{\texttt{WBATE}} - \tau_{\texttt{WBATE}})^2 \big| \mathbf{X}\big] = \Omega_{\texttt{WBATE}} + h^{2p+2} B^2_{\texttt{WBATE}} + o_{\mathbb{P}}(\mathfrak{R}_n),$$

where $(nh)^{-1} \lesssim \Omega_{\texttt{WBATE}} \lesssim (nh)^{-1}$*, and* $\mathfrak{R}_n = (nh)^{-1} + h^{2p+2}$.

This theorem immediately establishes consistency of the estimator, that is, $\widehat{\tau}_{\texttt{WBATE}} = \tau_{\texttt{WBATE}} + o_{\mathbb{P}}(1)$. In addition, Theorem 4 shows that the convergence rate of the estimator is improved due to the aggregation along the boundary $\mathscr{B}$: while the pointwise estimator $\widehat{\tau}(\mathbf{x})$ had "variance terms" of order $(nh^2)^{-1}$ (Theorems 1 and 2), the estimator $\widehat{\tau}_{\texttt{WBATE}}$ has a "variance term" of order $(nh)^{-1}$. Intuitively, the estimator $\widehat{\tau}_{\texttt{WBATE}}$ corresponds to a "one-dimensional" nonparametric estimate, thereby having a faster (and optimal) convergence rate.

The infeasible MSE-optimal bandwidth selector is

$$h_{\texttt{WBATE}} = \Big(\frac{2V_{\texttt{WBATE}}}{(2p+2)B^2_{\texttt{WBATE}}}\frac{1}{n}\Big)^{1/(2p+4)},$$

and a feasible counterpart can be constructed using plug-in estimators of $B_{\texttt{WBATE}}$ and $V_{\texttt{WBATE}}$, as we discussed in Section 3 for the estimator of the boundary average treatment effect

curve, $\widehat{\tau}(\mathbf{x})$. Our companion software package `rd2d` implements this bandwidth selector to construct an MSE-optimal point estimator $\widehat{\tau}_{\texttt{WBATE}}$ of $\tau_{\texttt{WBATE}}$, see Cattaneo et al. [2025c].

For inference, we consider the usual feasible t-statistics

$$\widehat{\mathrm{T}}_{\texttt{WBATE}} = \frac{\widehat{\tau}_{\texttt{WBATE}} - \tau_{\texttt{WBATE}}}{\sqrt{\widehat{\Omega}_{\texttt{WBATE}}}},$$

where the variance estimator is

$$\widehat{\Omega}_{\texttt{WBATE}} = \widehat{\Omega}_{1,\texttt{WBATE}} + \widehat{\Omega}_{0,\texttt{WBATE}}, \qquad \widehat{\Omega}_{t,\texttt{WBATE}} = \int_{\mathscr{B}}\int_{\mathscr{B}} \widehat{\Omega}_{t,\mathbf{x}_1,\mathbf{x}_2} w(\mathbf{x}_1)w(\mathbf{x}_2) d\mathfrak{H}(\mathbf{b}_1) d\mathfrak{H}(\mathbf{x}_2),$$

for $t \in \{0,1\}$.

Theorem 5 (Distributional Approximation: WBATE)**.** *Suppose Assumptions 1, 2, and 3 hold. If* $\frac{\log(1/h)}{n^{\frac{v}{2+v}} h^d} = o(1)$ *and* $nh^{2p+3} = o(1)$*, then*

$$\sup_{u\in\mathbb{R}} \left|\mathbb{P}(\widehat{\mathrm{T}}_{\texttt{WBATE}} \leq u) - \Phi(u)\right| = o(1).$$

Asymptotically valid hypothesis testing procedures and confidence interval estimators can be constructed directly from this result. For example, under the conditions of the theorem, a valid confidence interval estimator is

$$\widehat{\mathrm{I}}_{\alpha,\texttt{WBATE}} = \left[\, \widehat{\tau}_{\texttt{WBATE}} - \Phi^{-1}(1-\alpha/2)\sqrt{\widehat{\Omega}_{\texttt{WBATE}}} \,,\, \widehat{\tau}_{\texttt{WBATE}} + \Phi^{-1}(1-\alpha/2)\sqrt{\widehat{\Omega}_{\texttt{WBATE}}} \,\right],$$

for any $\alpha \in (0,1)$, and because $\mathbb{P}\big[\tau_{\texttt{WBATE}} \in \widehat{\mathrm{I}}_{\alpha,\texttt{WBATE}}\big] \to 1-\alpha$. For implementation, the MSE-optimal point estimator can be used along with robust bias correction [Calonico et al., 2014, 2018, 2022], as discussed after Theorem 3. Our companion software package `rd2d` also implements this inference procedure.

5 Largest Boundary Average Treatment Effect

As an alternative to the WBATE, $\tau_{\texttt{WBATE}}$, researchers may be interested in learning about the best (or worst) treatment effect along the assignment boundary: $\tau_{\texttt{LBATE}} = \sup_{\mathbf{x}\in\mathscr{B}} \tau(\mathbf{x})$. A (plug-in) location-based estimator of the LBATE is

$$\widehat{\tau}_{\texttt{LBATE}} = \sup_{\mathbf{x}\in\mathscr{B}} \widehat{\tau}(\mathbf{x}).$$

Implementation of this estimator can be done over a grid of points discretizing $\mathscr{B}$, that is, $\widehat{\tau}_{\texttt{LBATE}} \approx \max_{1\leq j\leq J} \widehat{\tau}(\mathbf{b}_j)$, where $(\mathbf{b}_j : j = 1, \ldots, J)$ are points on $\mathscr{B}$. If desired, the parameter and estimator may be defined only for a region of $\mathscr{B}$.

Theorem 1 establishes consistency of the largest average treatment effect estimator along the boundary, $\widehat{\tau}_{\texttt{LBATE}} = \tau_{\texttt{LBATE}} + o_{\mathbb{P}}(1)$; this theorem can also be used to deduce the convergence rate for $\widehat{\tau}_{\texttt{LBATE}}$. Valid uncertainty quantification is established using our new strong approximation theorem (Section SA-7 in the supplemental appendix) for the non-Donsker t-statistic stochastic process $(\widehat{\mathrm{T}}(\mathbf{x}) : \mathbf{x} \in \mathscr{B})$. Specifically, recall from Theorem 3 that $(\widehat{Z}(\mathbf{x}) : \mathbf{x} \in \mathscr{B})$ is a (conditionally on $\mathbf{W}$) mean-zero Gaussian process with feasible (conditional) covariance function $\mathbb{C}\mathrm{ov}\big[\widehat{Z}(\mathbf{x}_1), \widehat{Z}(\mathbf{x}_2)\big|\mathbf{W}\big] = \widehat{\Omega}^{-1/2}_{\mathbf{x}_1,\mathbf{x}_1} \widehat{\Omega}_{\mathbf{x}_1,\mathbf{x}_2} \widehat{\Omega}^{-1/2}_{\mathbf{x}_2,\mathbf{x}_2}$ for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathscr{B}$, and thus define the confidence interval estimator

$$\widehat{\mathrm{I}}_{\alpha,\texttt{LBATE}} = \left[\sup_{\mathbf{x}\in\mathscr{B}} \left(\widehat{\tau}(\mathbf{x}) - \mathfrak{q}_\alpha \sqrt{\widehat{\Omega}_{\mathbf{x},\mathbf{x}}} \right) , \sup_{\mathbf{x}\in\mathscr{B}} \left(\widehat{\tau}(\mathbf{x}) + \mathfrak{q}_\alpha \sqrt{\widehat{\Omega}_{\mathbf{x},\mathbf{x}}} \right) \right],$$

where $\mathfrak{q}_\alpha = \inf\big\{c > 0 : \mathbb{P}\big(\sup_{\mathbf{x}\in\mathscr{B}} \big|\widehat{Z}(\mathbf{x})\big| \geq c\big|\mathbf{W}\big) \leq \alpha\big\}$.

Theorem 6 (Confidence Interval: Largest Treatment Effect)**.** *Suppose the assumptions and conditions in Theorem 3 hold. If* $\liminf_{n\to\infty} \frac{\log h}{\log n} > -\infty$, $\frac{(\log n)^3}{n^{\frac{v}{2+v}} h^2} = o(1)$ *and* $nh^{2p+4} = o(1)$, *then*

$$\mathbb{P}\Big[\tau_{\texttt{LBATE}} \in \widehat{\mathrm{I}}_{\alpha,\texttt{LBATE}}\Big] \geq 1 - \alpha + o(1).$$

Implementation of $\widehat{\mathrm{I}}_{\alpha,\mathtt{LBATE}}$ is straightforward following the same approaches outlined for $\widehat{\mathrm{I}}_{\alpha}(\mathbf{x})$ and $\widehat{\mathrm{I}}_{\alpha,\mathtt{WBATE}}$. In practice, finite-sample regularization may be useful to ensure that $\inf_{\mathbf{x}\in\mathscr{B}} \widehat{\Omega}_{\mathbf{x},\mathbf{x}}$ is bounded away from zero.

6 The Causal Effects of SPP on College Attendance

We illustrate our proposed causal inference methodology for BD designs with the SPP application introduced in Section 1. Recall that the dataset has $n = 363,096$ complete observations for the first cohort of the program (2014), where each observation corresponds to one student, and the bivariate score is $\mathbf{X}_i = (X_{1i}, X_{2i})^\top = (\mathtt{SABER11}_i, \mathtt{SISBEN}_i)^\top$, where the first dimension is the student's SABER11 test score (ranging from -310 to 172) and the second dimension is the student's SISBEN wealth index (ranging from -103.41 to 127.21). Without loss of generality, each dimension of the score is recentered at its corresponding cutoff for program eligibility, so that the treatment assignment boundary is as shown in Figure 1a. All the results in this section are implemented using our companion `R` software package `rd2d`, and omitted details are given in the replication files.

The outcome variable of interest is *college enrollment*, with $Y_i = 1$ if the student enrolled in college and $Y_i = 0$ otherwise. Figure 2 presents the results for the location-based estimators of the boundary average treatment effect curve $\widehat{\tau}(\mathbf{x})$ estimated at 40 evenly-spaced grid points $\mathbf{b}_j \in \mathscr{B}$, with $j \in \{1, \ldots, 40\}$ depicted in Figure 1a, using a data-driven implementation of $h_{\mathtt{MSE},\mathbf{x}}$ as discussed in Section 3.3. Without loss of generality, each dimension of $\mathbf{X}_i$ is standardized in order to accommodate a common bandwidth h. The point estimates coincide with those reported in Figure 1b, but Figure 2a also includes confidence intervals (CI) and confidence bands (CB) as developed in Section 3.2. The average treatment effects at the chosen boundary points are highly statistically significant (different from zero), indicating roughly homogeneous treatment effects along poverty ($\mathbf{b}_1$ through $\mathbf{b}_{21}$) and heterogeneous treatment effects along academic performance ($\mathbf{b}_{21}$ through $\mathbf{b}_{40}$). The average treatment effect on (the probability of) college attendance remains roughly constant as marginally

academically achieving students become wealthier ($\widehat{\tau}(\mathbf{x}) \approx 0.3$ for $\mathbf{x} \in \{\mathbf{b}_1, \ldots, \mathbf{b}_{21}\}$), but it decreases as the wealthiest eligible students increase their academic performance (from $\widehat{\tau}(\mathbf{b}_{21}) \approx 0.3$ to $\widehat{\tau}(\mathbf{b}_{40}) \approx 0.18$). Figures 2b and 2c offer "heat maps" for the point estimates and associated robust bias-corrected p-values along the assignment boundary $\mathscr{B}$.

[FIGURE 2 AROUND HERE]

To demonstrate the credibility of the BD design, we repeat the empirical analysis using a pre-treatament covariate, the education level (measured in years) of the student's mother, as the outcome variable. This corresponds to a standard "placebo" analysis on a variable that is known to be unaffected by the policy, where the treatment effect is therefore expected to be statistically indistinguishable from zero. See Cattaneo et al. [2020, Section 5] for more discussion of falsification tests in RD designs. As expected in a valid BD design, Figure 3 shows that the average treatment effects at the boundary points considered are all statistically indistinguishable from zero, both pointwise and uniformly.

[FIGURE 3 AROUND HERE]

Table 1 presents the numerical results underlying our figures. The table reports the boundary average treatment effect curve estimated for a subset of boundary points, and the estimators of the aggregate parameters WBATE and LBATE. This table reports results for $\tau(\mathbf{b}_j)$ with $j = 1, 5, 10, 15, 20, 25, 30, 35, 40$ to streamline the presentation, together with results for the WBATE $\tau_{\texttt{WBATE}}$, and the LBATE $\tau_{\texttt{LBATE}}$. The WATE estimator $\widehat{\tau}_{\mathscr{B}}$ is implemented with equal weighting: $w(\mathbf{b}_j) = 1/J$ for all $j = 1, \ldots, J$ with $J = 40$. Table 1 also reports the data-driven MSE-optimal bandwidths used (after undoing the standardization of the bivariate location score $\mathbf{X}_i$), and measures of uncertainty quantification (p-values and confidence intervals).

The point estimation results in Table 1 corresponds to those reported in Figures 1b and 2a; see also Figure 2b for a heatmap plot. The confidence intervals corresponds to those plotted in Figure 2a, while the p-values were graphically reported in Figure 2c. The numerical findings confirm the lack of heterogeneous treatment effects along the first portion of the

Method	$(h_{\texttt{MSE},1}, h_{\texttt{MSE},2})$	Estimate	T-stat	p-value	CI
$\tau(\mathbf{b}_1)$	(26.9, 11.9)	0.3103	8.2791	0.0000	(0.2021, 0.4283)
$\tau(\mathbf{b}_5)$	(21.7, 9.6)	0.2910	7.0627	0.0000	(0.1518, 0.3722)
$\tau(\mathbf{b}_{10})$	(18.1, 8.0)	0.3272	7.1519	0.0000	(0.1733, 0.4194)
$\tau(\mathbf{b}_{15})$	(20.3, 9.0)	0.3068	7.9734	0.0000	(0.1747, 0.3822)
$\tau(\mathbf{b}_{20})$	(18.1, 8.0)	0.3266	9.3204	0.0000	(0.2264, 0.4381)
$\tau(\mathbf{b}_{25})$	(20.2, 9.0)	0.2938	8.6263	0.0000	(0.2100, 0.4306)
$\tau(\mathbf{b}_{30})$	(25.3, 11.2)	0.2570	6.6527	0.0000	(0.1570, 0.4105)
$\tau(\mathbf{b}_{35})$	(24.1, 10.7)	0.2144	3.9771	0.0001	(0.0588, 0.4059)
$\tau(\mathbf{b}_{40})$	(31.5, 14.0)	0.1834	3.3760	0.0007	(0.0228, 0.3568)
$\tau_{\texttt{WBATE}}$		0.2806	15.3726	0.0000	(0.2467, 0.3188)
$\tau_{\texttt{LBATE}}$		0.3361			(0.2610, 0.4771)

Table 1: Treatment Effect Analysis Along the Boundary.

assignment boundary corresponding to variation in wealth levels as measured by SISBEN (vertical-axis in Figure 1a), for students who met the minimum criteria of academic performance as measured by SABER11 (horizontal-axis in Figure 1a). The numerical results also confirm the presence of heterogeneous treatment effects for the wealthiest eligible students as their academic performance increases, that is, $\widehat{\tau}(\mathbf{x})$ decreases as SABER11 increases. These empirical findings are consistent with the expected behavioral response of students: the program has a relatively smaller causal effect among the wealthiest eligible students with higher academic performance, presumably because this type of student is more likely to be able to enroll in college in the absence of the subsidy.

Table 1 also reports point estimates and uncertainty quantification for $\tau_{\texttt{WBATE}}$ and $\tau_{\texttt{LBATE}}$; the point estimates coincide with those reported in Figure 1b. The estimated WATE is $\widehat{\tau}_{\texttt{WBATE}} = 0.2806$, indicating that on average students who are eligible for SPP 28 percentage points more likely to enroll in higher education. As shown, this aggregate effect masks considerable heterogeneity along the boundary, which in this application ranges from $\widehat{\tau}(\mathbf{b}_1) = 0.3103$ to $\widehat{\tau}(\mathbf{b}_{40}) = 0.1834$. The estimated LBATE is $\widehat{\tau}_{\texttt{LBATE}} = 0.3361$, corresponding to students in the mid-range of poverty as measured by SISBEN (vertical-axis in Figure 1a) and with the lowest eligible academic performance (horizontal-axis in Figure 1a); see Figure 1b.

7 Extensions

The theoretical results in the supplemental appendix also consider derivative estimation for each point on the assignment boundary $\mathscr{B}$, which is needed to extend our findings to regression kink designs [Card et al., 2015]. This section briefly discusses two other extensions of our work that are useful for treatment effect estimation and causal inference in BD designs.

7.1 Imperfect Compliance

Our results can easily be extended to settings with imperfect treatment compliance, where treatment assignment and treatment status may not be equal for some units; [see, e.g., Hernán and Robins, 2020]. To formally describe this setup, we need to modify the potential outcomes notation. Let $D_i = \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_0) \cdot D_i(0) + \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_1) \cdot D_i(1)$ be the observed treatment status, where $D_i(t)$ denotes the potential treatment status under treatment assignment $t \in \{0,1\}$ for each unit. Accordingly, the observed outcome is now $Y_i = \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_0) \cdot Y_i(0, D_i(0)) + \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_1) \cdot Y_i(1, D_i(1))$, where the potential outcomes are now a function of two arguments: $Y_i(t,d)$ denotes the potential outcome for unit i when this unit is assigned to treatment $t \in \{0,1\}$ and takes treatment status $d \in \{0,1\}$.

The usual fuzzy estimand and estimators are, respectively,

$$\zeta(\mathbf{x}) = \frac{\tau_Y(\mathbf{x})}{\tau_D(\mathbf{x})} \qquad \text{and} \qquad \widehat{\zeta}(\mathbf{x}) = \frac{\widehat{\tau}_Y(\mathbf{x})}{\widehat{\tau}_D(\mathbf{x})},$$

where, for each $\mathbf{x} \in \mathscr{B}$, $\tau_Y(\mathbf{x}) = \mathbb{E}[Y_i(1, D_i(1)) - Y_i(0, D_i(0)) | \mathbf{X}_i = \mathbf{x}]$ and $\tau_D(\mathbf{x}) = \mathbb{E}[D_i(1) - D_i(0) | \mathbf{X}_i = \mathbf{x}]$, and $\widehat{\tau}_Y(\mathbf{x}) = \mathbf{e}_1^\top \widehat{\boldsymbol{\beta}}_{Y,1}(\mathbf{x}) - \mathbf{e}_1^\top \widehat{\boldsymbol{\beta}}_{Y,0}(\mathbf{x})$ and $\widehat{\tau}_D(\mathbf{x}) = \mathbf{e}_1^\top \widehat{\boldsymbol{\beta}}_{D,1}(\mathbf{x}) - \mathbf{e}_1^\top \widehat{\boldsymbol{\beta}}_{D,0}(\mathbf{x})$, with $\widehat{\boldsymbol{\beta}}_{A,1}(\mathbf{x})$ denoting the local polynomial fit (1) when using the outcome variable $A \in \{Y, D\}$.

Under regularity conditions, our results immediately imply that $\widehat{\tau}_Y(\mathbf{x}) = \tau_Y(\mathbf{x}) + o_{\mathbb{P}}(1)$ and $\widehat{\tau}_D(\mathbf{x}) = \tau_D(\mathbf{x}) + o_{\mathbb{P}}(1)$, pointwise and uniformly over $\mathbf{x} \in \mathscr{B}$. Therefore, using the exact

second-order "linearization",

$$\widehat{\zeta}(\mathbf{x}) - \zeta(\mathbf{x}) = \frac{1}{\tau_D(\mathbf{x})}\big(\widehat{\tau}_Y(\mathbf{x}) - \tau_Y(\mathbf{x})\big) - \frac{\tau_Y(\mathbf{x})}{\tau_D(\mathbf{x})^2}\big(\widehat{\tau}_D(\mathbf{x}) - \tau_D(\mathbf{x})\big) + \mathfrak{R}_n(\mathbf{x})$$

with

$$\mathfrak{R}_n(\mathbf{x}) = \frac{\tau_Y(\mathbf{x})}{\tau_D(\mathbf{x})^2\widehat{\tau}_D(\mathbf{x})}\big(\widehat{\tau}_D(\mathbf{x}) - \tau_D(\mathbf{x})\big)^2 - \frac{1}{\tau_D(\mathbf{x})\widehat{\tau}_D(\mathbf{x})}\big(\widehat{\tau}_Y(\mathbf{x}) - \tau_Y(\mathbf{x})\big)\big(\widehat{\tau}_D(\mathbf{x}) - \tau_D(\mathbf{x})\big),$$

it follows that $\widehat{\zeta}(\mathbf{x})-\zeta(\mathbf{x})$ is a linear combination of $(\widehat{\tau}_Y(\mathbf{x})-\tau_Y(\mathbf{x}))$ and $(\widehat{\tau}_D(\mathbf{x})-\tau_D(\mathbf{x}))$. The remainder is negligible because $\sup_{\mathbf{x}\in\mathscr{B}}|\mathfrak{R}_n(\mathbf{x})| \lesssim_{\mathbb{P}} \sup_{\mathbf{x}\in\mathscr{B}}|\widehat{\tau}_D(\mathbf{x})-\tau_D(\mathbf{x})|^2+\sup_{\mathbf{x}\in\mathscr{B}}|\widehat{\tau}_Y(\mathbf{x})-\tau_Y(\mathbf{x})|\sup_{\mathbf{x}\in\mathscr{B}}|\widehat{\tau}_D(\mathbf{x})-\tau_D(\mathbf{x})|$, and hence Theorem 1 can be applied for each of the two BD estimators (one with outcome variable Y_i, the other with outcome variable D_i), under additional regularity conditions (e.g., $\inf_{\mathbf{x}\in\mathscr{B}}\tau_D(\mathbf{x}) > 0$). Pointwise and uniform inference can be established using the results in the supplemental appendix. The causal interpretation of the fuzzy estimand $\zeta(\mathbf{x})$, for each $\mathbf{x}\in\mathscr{B}$, can be obtained under additional assumptions; see Arai et al. [2022] and references therein.

Aggregation of causal effects along the boundary under imperfect compliance is also possible. For example, in the case of the weighted boundary average treatment effect, the estimand is $\zeta_{\texttt{WBATE}} = \int_{\mathscr{B}}\zeta(\mathbf{x})w(\mathbf{x})d\mathfrak{H}(\mathbf{x}) = \int_{\mathscr{B}}\frac{\tau_Y(\mathbf{x})}{\tau_D(\mathbf{x})}w(\mathbf{x})d\mathfrak{H}(\mathbf{x})$, where again the "linearization" described above can be used to establish valid estimation and inference methods based on the estimator $\widehat{\zeta}_{\texttt{WBATE}} = \int_{\mathscr{B}}\widehat{\zeta}(\mathbf{x})w(\mathbf{x})d\mathfrak{H}(\mathbf{x})$. However, the causal interpretation of $\zeta_{\texttt{WBATE}}$ depends on specific assumptions, and may not be straightforward in general; this is an open question for future research.

7.2 Pre-treatment Covariates

In the context of standard (univariate) RD designs, it is common to incorporate pre-treatment covariates in the estimation either for efficiency improvements [Calonico et al., 2019] or for heterogeneity analysis [Calonico et al., 2025]. Following that literature, we can extend our results to incorporate pre-treatment covariates in the analysis of BD designs. To conserve

space, we only briefly illustrate the approach in the basic (sharp) setup of Section 2, but the methods could also be extended to accommodate imperfect compliance.

Suppose that $\mathbf{Z}_1, \ldots, \mathbf{Z}_n$ are pre-intervention covariates of dimension $d_Z \geq 1$. For efficiency improvements, the covariate-adjusted location-based conditional treatment effect estimator of $\tau(\mathbf{x})$ is

$$\widetilde{\tau}(\mathbf{x}) = \mathbf{e}_{\mathfrak{p}_p+2}^\top \widetilde{\boldsymbol{\gamma}}(\mathbf{x})$$

with

$$\widetilde{\boldsymbol{\gamma}}(\mathbf{x}) = \underset{\boldsymbol{\gamma}\in\mathbb{R}^{\mathfrak{q}}}{\arg\min}\, \mathbb{E}_n\Big[\big(Y_i - \widetilde{\mathbf{r}}_p(\mathbf{X}_i - \mathbf{x}, T_i, \mathbf{Z}_i)^\top \boldsymbol{\gamma}\big)^2 K_h(\mathbf{X}_i - \mathbf{x}) \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\Big],$$

where $T_i = \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_1)$, and $\widetilde{\mathbf{r}}_p(\mathbf{X}_i - \mathbf{x}, T_i, \mathbf{Z}_i) = [\mathbf{r}_p(\mathbf{X}_i - \mathbf{x})^\top, T_i \cdot \mathbf{r}_p(\mathbf{X}_i - \mathbf{x})^\top, \mathbf{Z}_i^\top]^\top$ contains the full polynomial basis function for each treatment group but the pre-intervention covariates are not interacted with the treatment indicator; hence its dimension is $\mathfrak{q} = 2(\mathfrak{p}_p + 1) + d_Z$. For heterogeneity analysis, the covariate-heterogeneous location-based boundary average treatment effect curve estimator is

$$\check{\kappa}(\mathbf{x}, \mathbf{z}) = \mathbf{e}_{\mathfrak{p}_p+2}^\top \check{\boldsymbol{\gamma}}(\mathbf{x}) + \check{\boldsymbol{\gamma}}(\mathbf{x})^\top \begin{bmatrix} \mathbf{0}_{2(\mathfrak{p}_p+1)+d_Z+(\mathfrak{p}_p+1)d_Z \times d_Z} \\ \mathbf{I}_{d_Z} \\ \mathbf{0}_{\mathfrak{p}_p d_Z \times d_Z} \end{bmatrix} \mathbf{z}$$

with

$$\check{\boldsymbol{\gamma}}(\mathbf{x}) = \underset{\boldsymbol{\gamma}\in\mathbb{R}^{2(\mathfrak{p}_p+1)+d_Z}}{\arg\min}\, \mathbb{E}_n\Big[\big(Y_i - \check{\mathbf{r}}_p(\mathbf{X}_i - \mathbf{x}, T_i, \mathbf{Z}_i)^\top \boldsymbol{\gamma}\big)^2 K_h(\mathbf{X}_i - \mathbf{x}) \mathbb{1}(\mathbf{X}_i \in \mathscr{A}_t)\Big],$$

where $\mathbf{I}_d$ denotes the $(d \times d)$ identity matrix, $\mathbf{0}_{d_1 \times d_2}$ denotes the $(d_1 \times d_2)$ matrix of zeros, $\mathbf{z}$ takes values on the support of $\mathbf{Z}_i$, and $\check{\mathbf{r}}_p(\mathbf{X}_i - \mathbf{x}, T_i, \mathbf{Z}_i) = [\widetilde{\mathbf{r}}_p(\mathbf{X}_i - \mathbf{x}, T_i, \mathbf{Z}_i)^\top, (\mathbf{r}_p(\mathbf{X}_i - \mathbf{x}) \otimes \mathbf{Z}_i)^\top, T_i \cdot (\mathbf{r}_p(\mathbf{X}_i - \mathbf{x}) \otimes \mathbf{Z}_i)^\top]^\top$ contains the full interaction between the polynomial basis function, the treatment assignment indicator, and the pre-intervention covariates.

8 Conclusion

We developed pointwise and uniform estimation and inference methods for estimation of the boundary average treatment effect curve, a causal functional parameter that captures the heterogeneity of treatment effects in BD designs where assignment to treatment is a deterministic function of a unit's bivariate location with respect to a boundary. We also studied point estimation and inference for the weighted boundary treatment effect (WBATE) and the largest boundary average treatment effect (LBATE), which offer natural aggregation measures of treatment effect heterogeneity. The main technical challenge in our analysis was accounting for the effect of the geometry of the assignment boundary $\mathscr{B}$, which is a one-dimensional manifold on the plane. Our uniform inference results also relied on a novel strong approximation theorem that may be of independent interest (Section SA-7 in the supplemental appendix). Finally, we implement all our methods in the software package `rd2d`, available at `https://rdpackages.github.io/rd2d/`; see Cattaneo et al. [2025c] for additional software details.

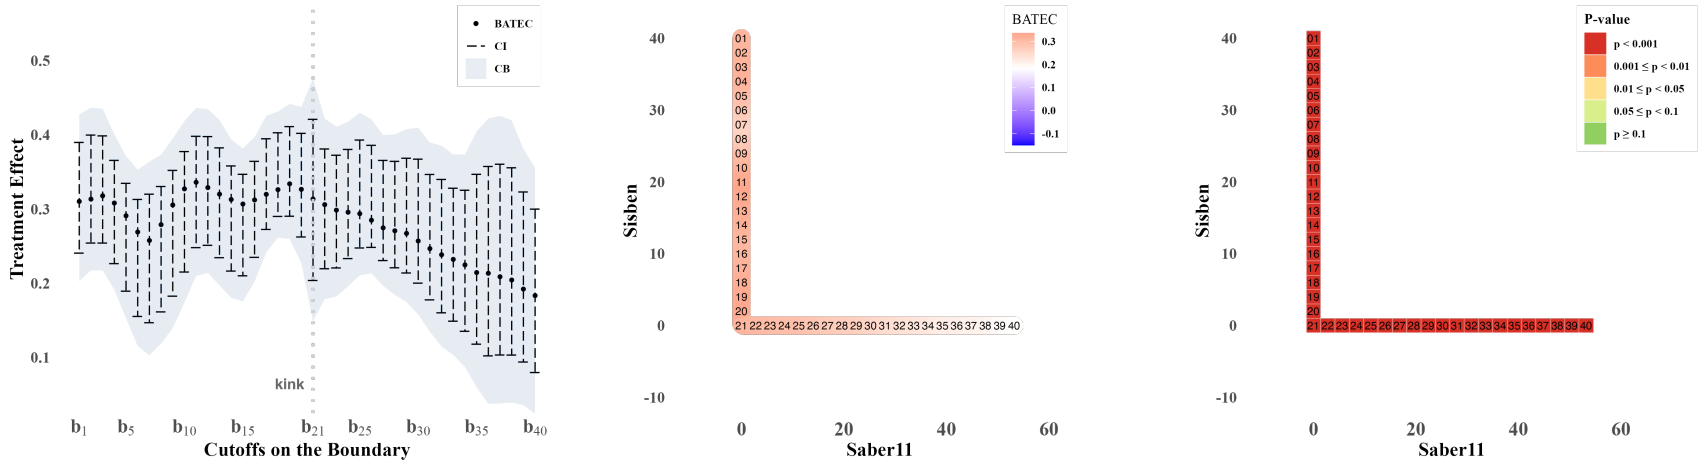

(a) Confidence Intervals and Bands (b) Treatment Effects Heatmap (c) p-values Heatmap

Figure 2: Estimation and Inference. Outcome: College Enrollment.

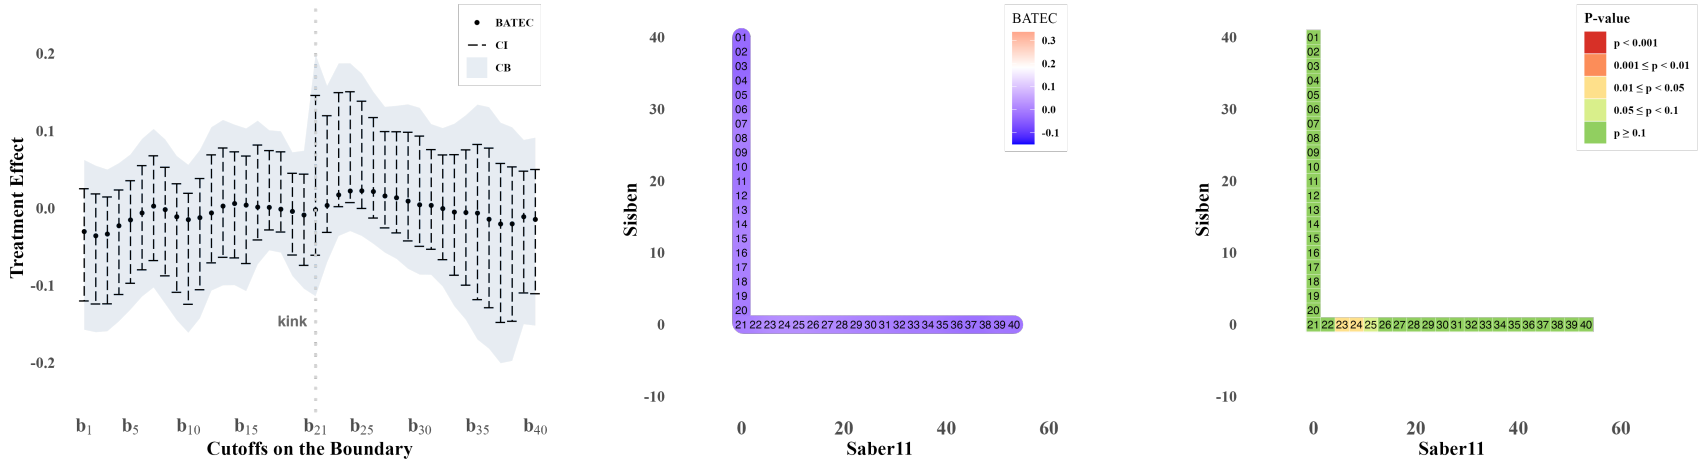

(a) Confidence Interval and Bands (b) Heatmap of Treatment Effect (c) Heatmap of P-value

Figure 3: Estimation and Inference. Outcome: Mother's education (placebo).

References

Isaiah Andrews, Toru Kitagawa, and Adam McCloskey. Inference on winners. *Quarterly Journal of Economics*, 139(1):305–358, 2024.

Yoichi Arai, Yu-Chin Hsu, Toru Kitagawa, Ismael Mourifié, and Yuanyuan Wan. Testing identifying assumptions in fuzzy regression discontinuity designs. *Quantitative Economics*, 13(1):1–28, 2022.

Sandra E Black. Do better schools matter? parental valuation of elementary education. *Quarterly Journal of Economics*, 114(2):577–599, 1999.

Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014.

Sebastian Calonico, Matias D. Cattaneo, and Max H. Farrell. On the effect of bias estimation on coverage accuracy in nonparametric inference. *Journal of the American Statistical Association*, 113(522):767–779, 2018.

Sebastian Calonico, Matias D. Cattaneo, Max H. Farrell, and Rocio Titiunik. Regression discontinuity designs using covariates. *Review of Economics and Statistics*, 101(3):442–451, 2019.

Sebastian Calonico, Matias D. Cattaneo, and Max H. Farrell. Optimal bandwidth choice for robust bias corrected inference in regression discontinuity designs. *Econometrics Journal*, 23(2):192–210, 2020.

Sebastian Calonico, Matias D. Cattaneo, and Max H. Farrell. Coverage error optimal confidence intervals for local polynomial regression. *Bernoulli*, 28(4):2998–3022, 2022.

Sebastian Calonico, Matias D. Cattaneo, Max H. Farrell, Filippo Palomba, and Rocio Titiunik. Treatment effect heterogeneity in regression discontinuity designs. *arXiv preprint arXiv:2503.13696*, 2025.

David Card and Alan B. Krueger. Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *American Economic Review*, 84(4): 772–793, 1994.

David Card, David S. Lee, Zhuan Pei, and Andrea Weber. Inference on causal effects in a generalized regression kink design. *Econometrica*, 83(6):2453–2483, 2015.

Matias D. Cattaneo and Rocio Titiunik. Regression discontinuity designs. *Annual Review of Economics*, 14:821–851, 2022.

Matias D. Cattaneo and Ruiqi (Rae) Yu. Strong approximations for empirical processes indexed by lipschitz functions. *Annals of Statistics*, 53(3):1203–1229, 2025.

Matias D. Cattaneo, Nicolás Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press, 2020.

Matias D. Cattaneo, Rajita Chandak, Michael Jansson, and Xinwei Ma. Boundary adaptive local polynomial conditional density estimators. *Bernoulli*, 30(4):3193–3223, 2024a.

Matias D. Cattaneo, Yingjie Feng, and William G. Underwood. Uniform inference for kernel density estimators with dyadic data. *Journal of the American Statistical Association*, 119 (524):2695–2708, 2024b.

Matias D. Cattaneo, Nicolás Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Extensions*. Cambridge University Press, 2024c.

Matias D. Cattaneo, Rocio Titiunik, and Ruiqi (Rae) Yu. Estimation and inference in boundary discontinuity designs: Distance-based methods. *Working paper*, 2025a.

Matias D. Cattaneo, Rocio Titiunik, and Ruiqi (Rae) Yu. Estimation and inference in boundary discontinuity designs: Pooling-based methods. *Working paper*, 2025b.

Matias D. Cattaneo, Rocio Titiunik, and Ruiqi (Rae) Yu. `rd2d`: Causal inference in boundary discontinuity designs. *Working paper*, 2025c.

Matias D. Cattaneo, Rocio Titiunik, and Ruiqi (Rae) Yu. Boundary discontinuity designs: Theory and practice. In *Invited book chapter for the 2025 Econometric Society World Congress*, volume 1, chapter 2. Cambridge University Press, 2026.

Xiaohong Chen and Wayne Yuan Gao. Semiparametric learning of integral functionals on submanifolds. *arXiv preprint arXiv:2507.12673*, 2025.

Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Anti-concentration and honest, adaptive confidence bands. *Annals of Statistics*, 42(5):1787–1818, 2014a.

Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Gaussian approximation of suprema of empirical processes. *Annals of Statistics*, 42(4):1564–1597, 2014b.

Victor Chernozhuokov, Denis Chetverikov, Kengo Kato, and Yuta Koike. Improved central limit theorem and bootstrap approximations in high dimensions. *Annals of Statistics*, 50 (5):2562–2586, 2022.

Leandro De Magalhães, Dominik Hangartner, Salomo Hirvonen, Jaakko Meriläinen, Nelson A Ruiz, and Janne Tukiainen. When can we trust regression discontinuity design estimates from close elections? evidence from experimental benchmarks. *Political Analysis*, 2025.

Melissa Dell. The persistent effects of peru's mining mita. *Econometrica*, 78(6):1863–1903, 2010.

Juan D Diaz and Jose R Zubizarreta. Complex discontinuity designs using covariates for policy impact evaluation. *Annals of Applied Statistics*, 17(1):67–88, 2023.

Richard M Dudley. *Uniform central limit theorems*, volume 142. Cambridge university press, 2014.

Herbert Federer. *Geometric measure theory*. Springer, 2014.

Sebastian Galiani, Patrick J. McEwan, and Brian Quistorff. External and internal validity of a geographic quasi-experiment embedded in a cluster-randomized experiment. In Matias D.

Cattaneo and Juan Carlos Escanciano, editors, *Regression Discontinuity Designs: Theory and Applications (Advances in Econometrics, volume 38)*, pages 195–236. Emerald Group Publishing, 2017.

Evarist Giné and Richard Nickl. *Mathematical Foundations of Infinite-dimensional Statistical Models*. Cambridge University Press, New York, 2016.

Jinyong Hahn, Petra Todd, and Wilbert van der Klaauw. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209, 2001.

Wolfgang Härdle, Marlene Müller, Stefan Sperlich, and Axel Werwatz. *Nonparametric and Semiparametric Models*. Springer, Heidelberg, 2004.

Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC, 2020.

Ari Hyytinen, Jaakko Meriläinen, Tuukka Saarimaa, Otto Toivanen, and Janne Tukiainen. When does regression discontinuity design work? evidence from random election outcomes. *Quantitative Economics*, 9(2):1019–1051, 2018.

Ekaterina Jardim, Mark C Long, Robert Plotnick, Jacob Vigdor, and Emma Wiles. Local minimum wage laws, boundary discontinuity methods, and policy spillovers. *Journal of Public Economics*, 234:105131, 2024.

Luke Keele and Rocio Titiunik. Natural experiments based on geography. *Political Science Research and Methods*, 4(1):65–95, 2016.

Luke J. Keele and Rocio Titiunik. Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1):127–155, 2015.

Luke J. Keele, Rocio Titiunik, and Jose Zubizarreta. 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(1):223–239, 2015.

Luke J. Keele, Scott Lorch, Molly Passarella, Dylan Small, and Rocio Titiunik. An overview of geographically discontinuous treatment assignments with an application to children's health insurance. In Matias D. Cattaneo and Juan Carlos Escanciano, editors, *Regression Discontinuity Designs: Theory and Applications (Advances in Econometrics, volume 38)*, pages 147–194. Emerald Group Publishing, 2017.

Juliana Londoño-Vélez, Catherine Rodríguez, and Fabio Sánchez. 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, 2020.

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

Sean F Reardon and Joseph P Robinson. Regression discontinuity designs with multiple rating-score variables. *Journal of Research on Educational Effectiveness*, 5(1):83–104, 2012.

Maxime Rischard, Zach Branson, Luke Miratrix, and Luke Bornn. Do school districts affect nyc house prices? identifying border differences using a bayesian nonparametric approach to geographic regression discontinuity designs. *Journal of the American Statistical Association*, 116(534):619–631, 2021.

Leon Simon et al. *Lectures on geometric measure theory*. Centre for Mathematical Analysis, Australian National University Canberra, 1984.

A.B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2008.

Aad W. van der Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes*. Springer, 1996.

Vivian C Wong, Peter M Steiner, and Thomas D Cook. Analyzing regression-discontinuity designs with multiple assignment variables: A comparative study of four estimation methods. *Journal of Educational and Behavioral Statistics*, 38(2):107–141, 2013.