

rd2d: Causal Inference in Boundary Discontinuity Designs

Matias D. Cattaneo*

Rocio Titiunik[†]

Ruiqi (Rae) Yu[‡]

May 28, 2026

Abstract

Boundary Discontinuity (BD) designs are used in empirical research to learn about causal treatment effects along a continuous assignment boundary defined by a bivariate score. These designs are also known as multi-score regression discontinuity (RD) designs, and include geographic RD designs as a prominent example. This article introduces `rd2d`, a statistical software package for `R`, `Python`, and `Stata` that implements local polynomial estimation and inference for BD designs using either the bivariate score or a univariate signed distance-to-boundary score. The software covers sharp and fuzzy BD designs, providing automatic bandwidth selection, robust bias-corrected pointwise inference, uniform confidence bands, cluster-robust inference with joint or separate fitting conventions, covariate-adjusted efficiency improvements, mass-point checks, and covariance regularization, among other features. We illustrate the package with an empirical application to Opportunity Zones, where eligibility has a strong first-stage effect on designation but no significant effects on early workplace-job growth.

Keywords: boundary discontinuity designs, regression discontinuity designs, treatment effects, nonparametric inference, `R`, `Python`, `Stata`.

1 Introduction

Regression Discontinuity (RD) designs are commonly used for treatment effect estimation and causal inference in quantitative sciences [see Cattaneo and Titiunik, 2022, and references therein]. In their canonical form, each unit $i \in \{1, 2, \dots, n\}$ is assigned to control ($T_i = 0$) or treatment ($T_i = 1$) according to the discontinuous rule $T_i = \mathbf{1}(X_i \geq c)$, where X_i denotes a scalar score variable, c denotes a scalar cutoff, and $\mathbf{1}(\cdot)$ is the indicator function. The key idea underlying all RD designs is that units with a score near the cutoff determining treatment assignment are comparable in terms of all pretreatment observables and unobservable characteristics, the only difference being that some units are assigned to control ($X_i < c$) while others are assigned to treatment ($X_i \geq c$). Therefore,

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

[†]Department of Politics, Princeton University.

[‡]Department of Operations Research and Financial Engineering, Princeton University.

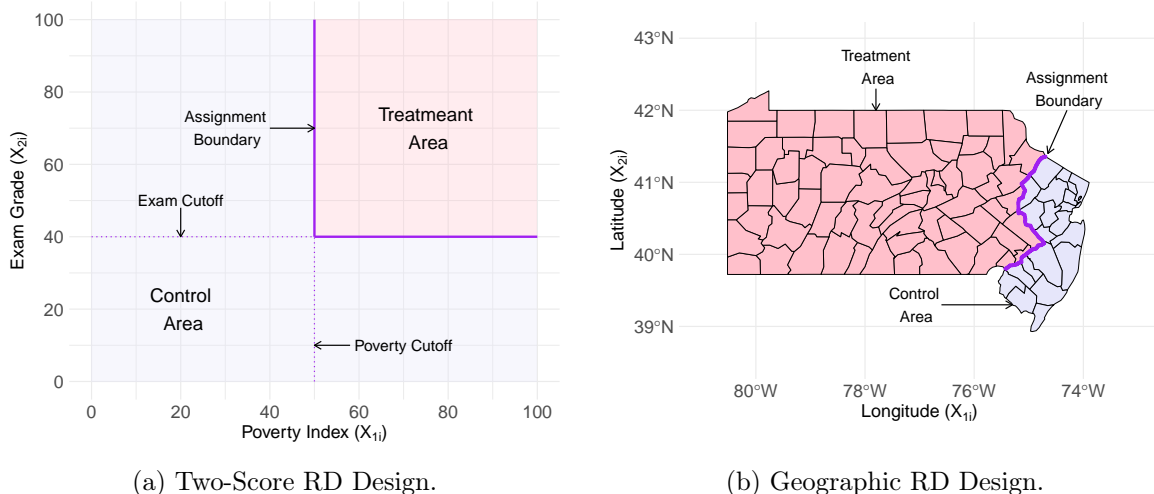


Figure 1: Two Examples of Boundary Discontinuity Designs

in the absence of score manipulation, units with scores near but on different sides of the cutoff can be used as counterfactuals for each other to learn about causal treatment effects.

Boundary discontinuity (BD) designs generalize the canonical RD design by allowing treatment assignment to be determined by a multidimensional score. The most common case is a bivariate score $\mathbf{X}_i = (X_{1i}, X_{2i})^\top$ and a one-dimensional assignment boundary \mathcal{B} that separates treated and control regions. These designs are also called multi-score RD designs [Papay et al., 2011, Reardon and Robinson, 2012, Wong et al., 2013], and include geographic RD designs as a prominent special case [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]. Figure 1a illustrates a two-score eligibility rule with a kinked boundary, while Figure 1b illustrates a generic geographic boundary. Empirical examples include education, taxation, political communication, labor markets, and place-based policy applications; see, among others, Londoño-Vélez et al. [2020] and Jardim et al. [2024]. See Cattaneo et al. [2020, 2024b] for a two-part practical introductory monograph.

While classical RD designs based on a scalar score are well understood, BD designs have only recently received systematic methodological treatment. Empirical work has used different reductions of the boundary problem, often without formal guidance about which causal parameter is being estimated, how treatment-effect heterogeneity along the boundary should be summarized, or when inference is valid uniformly over many boundary locations. Cattaneo et al. [2026a] and Cattaneo et al. [2026b] address this gap by studying two leading approaches:

- the *location-based approach*, which uses the bivariate score directly and fits local polynomial regressions on the two sides of \mathcal{B} ; and
- the *distance-based approach*, which replaces the bivariate score with a scalar signed distance to each point on \mathcal{B} and then applies univariate local polynomial methods.

Cattaneo et al. [2026a] develops location-based identification, estimation, bandwidth selection,

and inference results, both pointwise for each point on \mathcal{B} and uniformly over \mathcal{B} . Cattaneo et al. [2026b] develops the corresponding distance-based theory and clarifies when distance-to-boundary analyses recover the same causal targets. A central lesson is that distance-based methods can be useful and interpretable, but their validity depends more directly on the geometry of the boundary: kinks and other irregularities can generate additional bias even when the underlying bivariate conditional expectations are smooth. The location-based approach therefore provides the main benchmark whenever the bivariate score is available, while distance-based methods are valuable as a simpler implementation, a diagnostic, or a necessary alternative when only signed distances can be used.

This article introduces `rd2d`, software for `R`, `Python`, and `Stata` that implements these recent methodological developments for applied work. The package estimates the *Boundary Average Treatment Effect Curve* (BATEC), the *Weighted Boundary Average Treatment Effect* (WBATE), and the *Largest Boundary Average Treatment Effect* (LBATE), together with fuzzy BATEC, fuzzy WBATE, and fuzzy LBATE for settings with imperfect treatment take-up. It provides automatic bandwidth selection, robust bias-corrected pointwise confidence intervals, uniform confidence bands, heteroskedasticity-robust and cluster-robust variance estimators, joint and separate fitting conventions, covariate-adjusted efficiency improvements, mass-point checks, covariance regularization for uniform inference, and derivative estimation for location-based methods. The robust bias-correction and coverage-error ideas build on Calonico et al. [2014], Calonico et al. [2018], Calonico et al. [2020], and Calonico et al. [2022], adapted here to boundary-indexed local polynomial estimators. The `R` implementation centers on `rd2d()` and `rdbw2d()` for location-based analysis, and `rd2d.distance()` and `rdbw2d.distance()` for distance-based analysis; the `Python` and `Stata` implementations follow the same conceptual interface with language-specific syntax.

Two options deserve emphasis. First, the default option `fitmethod = "joint"` uses the treatment-interacted local regression representation when forming degrees-of-freedom adjustments, cluster-robust covariance matrices, and bandwidth-selection variance constants. This convention matters when a cluster, such as a school district, municipality, county, or tract group, can contain observations on both sides of the boundary: the cluster score should aggregate all local influence contributions from that cluster rather than treating the control and treated sides as independent samples. The option `fitmethod = "separate"` is retained for analyses that require the earlier side-specific convention. Second, the option `covs.eff` allows users to supply predetermined covariates for efficiency adjustment. These covariates enter with common coefficients across treatment sides and are propagated through both fixed-bandwidth estimation and automatic bandwidth selection.

The article is organized around the empirical steps in a BD analysis. We first clarify the causal target: a curve of local effects along the boundary, together with weighted and extremal summaries of that curve. We then show how the same targets are estimated using location-based and distance-based methods, how fuzzy designs decompose into reduced-form, first-stage, and Wald-ratio objects, and how software choices such as bandwidth selection, fitting convention, covariate adjustment, covariance storage, and boundary-grid construction affect empirical reporting. Throughout, an

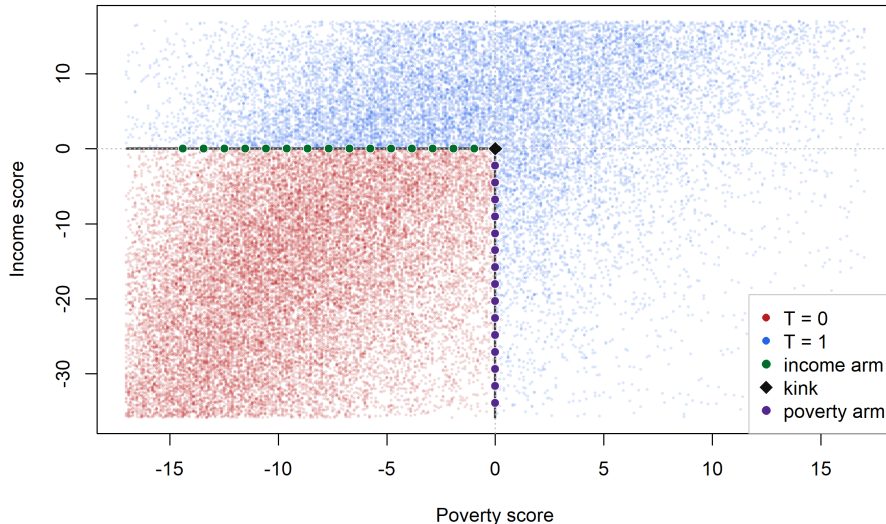


Figure 2: Opportunity Zones Running Example. The plot shows the final analysis sample, the low-income-community eligibility frontier, and the 31 evaluation points used in the empirical analysis.

Opportunity Zones application provides the notation, code snippets, figures, and numerical output.

1.1 Running Example: Opportunity Zones

We use the Opportunity Zones program as a running empirical example. The Tax Cuts and Jobs Act of 2017 allowed governors to nominate eligible low-income census tracts for Qualified Opportunity Zone designation. Following [Corinth and Feldman \[2024\]](#), we construct a bivariate boundary discontinuity design in which eligibility is determined by two tract-level scores: a poverty-rate threshold and a median-family-income threshold. Actual designation is not mechanically assigned to all eligible tracts, which makes this a fuzzy BD design.

The public analysis dataset contains one observation per census tract. In the notation used below, Y_i is the 2017–2019 change in $\log(1 + \text{workplace jobs})$, X_{1i} is the poverty rate minus 20, X_{2i} is 80 minus the tract’s median-family-income percentage, T_i is low-income-community eligibility, and W_i is Qualified Opportunity Zone designation. The assignment boundary is an L-shaped frontier with a kink at $(0, 0)$. We evaluate treatment-effect curves at 31 boundary points: the kink and 15 points on each arm of the boundary, avoiding the far extremes where local effective sample sizes are thin.

1.2 Organization and Background References

The rest of the paper proceeds as follows. Section 2 introduces the BD setup, the main causal parameters, and their interpretation. Section 3 discusses location-based estimation and inference, following the framework of [Cattaneo et al. \[2026a\]](#). Section 4 discusses distance-based methods and the additional geometric considerations studied by [Cattaneo et al. \[2026b\]](#). Section 5 returns to the Opportunity Zones example and compares location-based and distance-based estimates. Section

6 concludes. Appendix A summarizes the main R, Python, and Stata syntax. Replication codes, software documentation, and related information for rd2d are available at <https://rdpackages.github.io/>.

2 Setup

We employ standard potential outcomes notation [Hernán and Robins, 2020, and references therein]. Suppose that $(Y_i(0), Y_i(1), \mathbf{X}_i^\top)^\top$, $i = 1, 2, \dots, n$, is a random sample, where $Y_i(0)$ and $Y_i(1)$ denote the scalar potential outcomes for unit i under control and treatment assignment, respectively. Units are assigned to control group or treatment group according to their bivariate location score $\mathbf{X}_i = (X_{1i}, X_{2i})^\top$ relative to a known one-dimensional boundary \mathcal{B} splitting the support of \mathbf{X}_i in two disjoint regions: \mathcal{A}_0 denotes the control region, and \mathcal{A}_1 denotes the treatment region. Thus, $\mathcal{B} = \text{bd}(\mathcal{A}_0) \cap \text{bd}(\mathcal{A}_1)$, where $\text{bd}(\mathcal{A}_t)$ denotes the boundary of the set \mathcal{A}_t . We use the convention that boundary points belong to the treatment region, so $\mathcal{B} \subseteq \mathcal{A}_1$ and $\mathcal{B} \cap \mathcal{A}_0 = \emptyset$. The observed response variable is $Y_i = \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_0)Y_i(0) + \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_1)Y_i(1)$, and $T_i = \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_1)$ denotes treatment assignment. Figure 1 gives two graphical examples.

For implementation, the continuous assignment boundary \mathcal{B} is first discretized into J cutoff points $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_J)^\top$ with $\mathbf{b}_j \in \mathcal{B}$ for all $j = 1, \dots, J$. Then, the empirical analysis is conducted pointwise for each cutoff or uniformly over all cutoffs, employing either the bivariate location score \mathbf{X}_i directly, or an induced univariate distance to each cutoff point. See Cattaneo et al. [2024b, Section 5] for an introductory discussion.

In the Opportunity Zones example, $\mathbf{X}_i = (X_{1i}, X_{2i})^\top$ contains the poverty and income scores shown in Figure 2; \mathbf{b} contains the 31 evaluation points on the L-shaped frontier; T_i records eligibility; W_i records actual designation; and Y_i is early workplace-job growth. We use this notation in the code snippets below.

2.1 Sharp BD Designs

In a sharp BD design, assignment to the treatment region determines treatment exposure. The main causal parameter is the *Boundary Average Treatment Effect Curve* (BATEC):

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

The BATEC is a curve of local treatment effects indexed by points on the assignment boundary. It is useful when the policy rule creates many local comparisons rather than a single cutoff, because it allows researchers to inspect treatment-effect heterogeneity along the boundary before reporting lower-dimensional summaries.

In rd2d, the BATEC is obtained from the main output of a sharp BD design fit. The option `params.cov = "main"` stores the cross-boundary covariance matrix needed for uniform bands and aggregate summaries.

```
fit <- rd2d(Y = Y, X = X, assignment = T, b = b,
           params.cov = "main")
summary(fit, output = "main")
summary(fit, output = "main", cbands = "main")
```

The BATEC captures treatment effect heterogeneity along the assignment boundary, but researchers often want scalar summaries of that heterogeneity. For a weight function $w : \mathcal{B} \rightarrow [0, \infty)$, the *Weighted Boundary Average Treatment Effect* (WBATE) is

$$\tau_{\text{WBATE}} = \frac{\int_{\mathcal{B}} \tau(\mathbf{x})w(\mathbf{x}) \, d\mathbf{x}}{\int_{\mathcal{B}} w(\mathbf{x}) \, d\mathbf{x}}.$$

Different weights encode different policy or scientific targets. Equal weights over a pre-specified boundary grid produce a simple average of the BATEC values, while application-specific weights can emphasize boundary regions with more policy relevance or more population mass. A related scalar summary is the *Largest Boundary Average Treatment Effect* (LBATE),

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

which summarizes the largest treatment effect along the boundary.

The WBATE aggregates the BATEC into a weighted average effect, while the LBATE reports the largest value of the curve. Thus, WBATE is an aggregation target and LBATE is an extremal heterogeneity target. In `rd2d`, WBATE estimation and inference are implemented via the optional argument `WBATE = w` in the `summary()` method for `rd2d` and `rd2d.distance` objects, while LBATE summaries are requested with `LBATE = TRUE`. Uniform bands, WBATE, and LBATE require the relevant cross-boundary covariance matrix to be stored.

```
w <- rep(1, nrow(b))
summary(fit, output = "main", cbands = "main",
       WBATE = w, LBATE = TRUE)
```

2.2 Fuzzy BD Designs

In a fuzzy BD design, the assignment rule changes the probability of treatment receipt but does not perfectly determine it. This is the structure of the Opportunity Zones example: eligibility T_i changes the probability of designation W_i , but many eligible tracts are not designated and some policy choices are made after eligibility is determined. Let $T_i = \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_1)$ denote assignment. Following Cattaneo et al. [2026a], let $W_i(t)$ denote potential treatment status under assignment $t \in \{0, 1\}$, and let $Y_i(t, w)$ denote the potential outcome under assignment t and treatment status w . The observed treatment status and outcome are

$$\begin{aligned} W_i &= \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_0)W_i(0) + \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_1)W_i(1), \\ Y_i &= \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_0)Y_i(0, W_i(0)) + \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_1)Y_i(1, W_i(1)). \end{aligned}$$

The reduced-form or intention-to-treat BATEC, abbreviated ITT, and the first-stage BATEC, abbreviated FS, are

$$\begin{aligned}\tau_Y(\mathbf{x}) &= \mathbb{E}[Y_i(1, W_i(1)) - Y_i(0, W_i(0)) \mid \mathbf{X}_i = \mathbf{x}], \\ \tau_W(\mathbf{x}) &= \mathbb{E}[W_i(1) - W_i(0) \mid \mathbf{X}_i = \mathbf{x}], \quad \mathbf{x} \in \mathcal{B}.\end{aligned}$$

Under the standard continuity conditions for fuzzy BD designs, these curves are identified by the observed boundary discontinuities

$$\tau_A(\mathbf{x}) = \lim_{\mathbf{u} \rightarrow \mathbf{x}, \mathbf{u} \in \mathcal{A}_1} \mathbb{E}[A_i \mid \mathbf{X}_i = \mathbf{u}] - \lim_{\mathbf{u} \rightarrow \mathbf{x}, \mathbf{u} \in \mathcal{A}_0} \mathbb{E}[A_i \mid \mathbf{X}_i = \mathbf{u}], \quad A \in \{Y, W\}.$$

When the first stage is nonzero and the standard fuzzy BD design assumptions hold locally, the fuzzy BATEC is the Wald ratio

$$\zeta(\mathbf{x}) = \frac{\tau_Y(\mathbf{x})}{\tau_W(\mathbf{x})}, \quad \mathbf{x} \in \mathcal{B}.$$

We refer to $\tau_Y(\mathbf{x})$ as the ITT BATEC, $\tau_W(\mathbf{x})$ as the FS BATEC, and $\zeta(\mathbf{x})$ as the fuzzy BATEC. The ITT curve captures heterogeneity in the effect of assignment on the outcome, the FS curve captures heterogeneity in compliance or take-up induced by assignment, and the fuzzy BATEC rescales the outcome discontinuity by the local first stage. The Wald ratio is a fuzzy BD estimand; interpreting it as a complier-average treatment effect additionally requires local exclusion, monotonicity, and a nonzero first stage near the boundary point.

Fuzzy analogues of WBATE and LBATE replace $\tau(\mathbf{x})$ with $\zeta(\mathbf{x})$. For a weight function $w : \mathcal{B} \rightarrow [0, \infty)$, the fuzzy WBATE is

$$\zeta_{\text{WBATE}} = \frac{\int_{\mathcal{B}} \zeta(\mathbf{x}) w(\mathbf{x}) d\mathbf{x}}{\int_{\mathcal{B}} w(\mathbf{x}) d\mathbf{x}},$$

and the fuzzy LBATE is

$$\zeta_{\text{LBATE}} = \sup_{\mathbf{x} \in \mathcal{B}} \zeta(\mathbf{x}).$$

These summaries play the same role as in the sharp BD design: fuzzy WBATE aggregates the fuzzy BATEC over the boundary, while fuzzy LBATE reports the largest value of the fuzzy treatment-effect curve.

The package implements fuzzy BD designs by adding the treatment-receipt variable through `fuzzy`. The `output = "itt"` table reports the ITT BATEC, `output = "fs"` reports the FS BATEC, and `output = "main"` reports the fuzzy BATEC. The option `bwparam = "main"` targets automatic bandwidth selection to the fuzzy BATEC, while `bwparam = "itt"` targets the ITT BATEC.

```
fit_fuzzy <- rd2d(Y = Y, X = X, assignment = T, b = b,
                fuzzy = W,
```

```

      params.cov = c("main", "itt", "fs"))
summary(fit_fuzzy, output = "itt")
summary(fit_fuzzy, output = "fs")
summary(fit_fuzzy, output = "main")

```

Uniform bands and aggregate summaries for fuzzy BD designs are requested output by output. For example, the following calls report the ITT BATEC, FS BATEC, and fuzzy BATEC with uniform bands; the final call also reports fuzzy WBATE and fuzzy LBATE.

```

w <- rep(1, nrow(b))
summary(fit_fuzzy, output = "itt", cbands = "itt")
summary(fit_fuzzy, output = "fs", cbands = "fs")
summary(fit_fuzzy, output = "main", cbands = "main",
      WBATE = w, LBATE = TRUE)

```

3 Location-Based Methods

Location-based methods use the bivariate score \mathbf{X}_i directly and estimate side-specific local polynomial approximations at the boundary. In implementation these side-specific fits are represented through a joint treatment-interacted regression, which is equivalent for point estimation and is used for covariance calculations. Under the usual continuity assumptions invoked in canonical RD designs [Hahn et al., 2001], the BATEC in a sharp BD design is identified by

$$\tau(\mathbf{x}) = \mu_1(\mathbf{x}) - \mu_0(\mathbf{x}), \quad \mathbf{x} \in \mathcal{B},$$

where $\mu_t(\mathbf{x}) = \mathbb{E}[Y_i(t)|\mathbf{X}_i = \mathbf{x}] = \mathbb{E}[Y_i|\mathbf{X}_i = \mathbf{x}, T_i = t]$, $t \in \{0, 1\}$, are smooth functions on their corresponding sides of the boundary. This is the standard continuity argument for RD designs, adapted to a bivariate score and a one-dimensional assignment boundary [Papay et al., 2011, Reardon and Robinson, 2012, Wong et al., 2013, Keele and Titiunik, 2015]. Identification of WBATE and LBATE follows from identification of $\tau(\mathbf{x})$ over \mathcal{B} . In fuzzy BD designs, the same continuity argument identifies the reduced-form BATEC $\tau_Y(\mathbf{x})$ and the first-stage BATEC $\tau_W(\mathbf{x})$; when $\tau_W(\mathbf{x})$ is bounded away from zero on the target boundary, the fuzzy BATEC $\zeta(\mathbf{x}) = \tau_Y(\mathbf{x})/\tau_W(\mathbf{x})$ and its WBATE and LBATE analogues are identified, with the complier-average interpretation requiring the additional conditions discussed in Section 2. Following Cattaneo et al. [2026a], this section focuses on estimation, bandwidth selection, and uncertainty quantification for location-based methods.

3.1 Treatment Effect Estimation: Methods

For each $\mathbf{x} \in \mathcal{B}$, the location-based estimator can be written as a single local polynomial regression with treatment-side interactions. Let $\mathbf{r}_p(\mathbf{u})$ collect all monomials $u_1^{a_1}u_2^{a_2}$ with $a_1 + a_2 \leq p$, starting with the constant term, and let $\mathbf{p}_p = (2+p)(1+p)/2$. With $T_i = \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_1)$, define the augmented basis

$$\mathbf{r}_{p,h}(\mathbf{u}; T) = \left(\mathbf{r}_p(\mathbf{u}/h) \quad T\mathbf{r}_p(\mathbf{u}/h) \right)^\top.$$

For any multi-index $\boldsymbol{\nu}$ with $|\boldsymbol{\nu}| \leq p$, let $\mathbf{e}_{\boldsymbol{\nu}}$ denote the length- $2p_p$ vector that selects the coefficient on the treatment-interaction monomial indexed by $\boldsymbol{\nu}$ in $\mathbf{r}_{p,h}(\mathbf{u}; T)$; in particular, \mathbf{e}_0 selects the treatment-interaction intercept. The local polynomial estimator is

$$\widehat{\boldsymbol{\delta}}(\mathbf{x}) = \arg \min_{\boldsymbol{\delta} \in \mathbb{R}^{2p_p}} \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}; T_i)^\top \boldsymbol{\delta})^2 K_h(\mathbf{X}_i - \mathbf{x}),$$

where $K_h(\mathbf{u}) = K(u_1/h, u_2/h)/h^2$. The common scalar bandwidth h is used for exposition; in implementation the score coordinates can be standardized, and separate bandwidths can be used by coordinate and by treatment side. For interpretation, decompose the coefficient vector as $\widehat{\boldsymbol{\delta}}(\mathbf{x}) = (\widehat{\boldsymbol{\delta}}_0(\mathbf{x})^\top, \widehat{\boldsymbol{\delta}}_1(\mathbf{x})^\top)^\top$ with $\widehat{\boldsymbol{\delta}}_0(\mathbf{x}), \widehat{\boldsymbol{\delta}}_1(\mathbf{x}) \in \mathbb{R}^{p_p}$. Then, $\widehat{\boldsymbol{\delta}}_0(\mathbf{x})$ describes the control-side local polynomial, while $\widehat{\boldsymbol{\delta}}_1(\mathbf{x})$ describes the effects due to treatment assignment. Thus the BATEC estimator is

$$\widehat{\tau}(\mathbf{x}) = \mathbf{e}_0^\top \widehat{\boldsymbol{\delta}}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{B}.$$

This interaction representation is useful for software implementation because additional terms, such as predetermined covariates, can be added to the same weighted regression while leaving the treatment-side intercept interaction as the treatment-effect parameter.

The leading MSE terms depend on the local design, conditional variance, and curvature of the regression functions on each side of the boundary. Let $\sigma_t^2(\mathbf{x}) = \mathbb{V}[Y_i(t) | \mathbf{X}_i = \mathbf{x}]$. Using the augmented basis, define

$$\begin{aligned} \boldsymbol{\Gamma}_{\mathbf{x}} &= \mathbb{E} \left[\mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}; T_i) \mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}; T_i)^\top K_h(\mathbf{X}_i - \mathbf{x}) \right], \\ \boldsymbol{\Sigma}_{\mathbf{x}_1, \mathbf{x}_2} &= h^2 \mathbb{E} \left[\mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}_1; T_i) \mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}_2; T_i)^\top K_h(\mathbf{X}_i - \mathbf{x}_1) K_h(\mathbf{X}_i - \mathbf{x}_2) \sigma_{T_i}^2(\mathbf{X}_i) \right]. \end{aligned}$$

Then the variance constant for the BATEC is

$$V_{\mathbf{x}} = \mathbf{e}_0^\top \boldsymbol{\Gamma}_{\mathbf{x}}^{-1} \boldsymbol{\Sigma}_{\mathbf{x}, \mathbf{x}} \boldsymbol{\Gamma}_{\mathbf{x}}^{-1} \mathbf{e}_0.$$

Using multi-index notation, the leading bias is

$$B_{\mathbf{x}} = \mathbf{e}_0^\top \boldsymbol{\Gamma}_{\mathbf{x}}^{-1} \sum_{t=0}^1 \sum_{|\mathbf{k}|=p+1} \frac{\mu_t^{(\mathbf{k})}(\mathbf{x})}{\mathbf{k}!} \mathbb{E} \left[\mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}; T_i) ((\mathbf{X}_i - \mathbf{x})/h)^{\mathbf{k}} K_h(\mathbf{X}_i - \mathbf{x}) \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_t) \right].$$

Under the regularity conditions in Cattaneo et al. [2026a], $\widehat{\tau}(\mathbf{x})$ is pointwise and uniformly consistent for $\tau(\mathbf{x})$ as $h \rightarrow 0$ and nh^2 grows. The pointwise conditional MSE expansion has the usual bias-variance form,

$$\mathbb{E}[(\widehat{\tau}(\mathbf{x}) - \tau(\mathbf{x}))^2 | \mathbf{X}] \approx_{\mathbb{P}} h^{2p+2} B_{\mathbf{x}}^2 + \frac{1}{nh^2} V_{\mathbf{x}}, \quad \mathbf{x} \in \mathcal{B},$$

with the leading bias and variance constants defined above. Integrating the same expansion over \mathcal{B}

gives

$$\int_{\mathcal{B}} \mathbb{E}[(\hat{\tau}(\mathbf{x}) - \tau(\mathbf{x}))^2 | \mathbf{X}] w(\mathbf{x}) d\mathbf{x} \approx_{\mathbb{P}} h^{2p+2} \int_{\mathcal{B}} B_{\mathbf{x}}^2 w(\mathbf{x}) d\mathbf{x} + \frac{1}{nh^2} \int_{\mathcal{B}} V_{\mathbf{x}} w(\mathbf{x}) d\mathbf{x}.$$

Ignoring higher-order terms gives the infeasible pointwise MSE and integrated MSE bandwidths

$$h_{\text{MSE},\mathbf{x}} = \left(\frac{2V_{\mathbf{x}}}{(2p+2)B_{\mathbf{x}}^2} \frac{1}{n} \right)^{1/(2p+4)}, \quad h_{\text{IMSE}} = \left(\frac{2 \int_{\mathcal{B}} V_{\mathbf{x}} w(\mathbf{x}) d\mathbf{x}}{(2p+2) \int_{\mathcal{B}} B_{\mathbf{x}}^2 w(\mathbf{x}) d\mathbf{x}} \frac{1}{n} \right)^{1/(2p+4)}.$$

The pointwise selector targets a particular boundary location, while the integrated selector targets the full boundary curve or a chosen subset of it.

3.2 Treatment Effect Estimation: Implementation

Location-based estimation is implemented by `rd2d()`, and bandwidth selection can be called separately through `rdbw2d()`. In the running example, `X` contains the poverty and income scores and `b` contains the boundary grid. A basic sharp BD design fit is

```
fit <- rd2d(Y = Y, X = X, assignment = T, b = b,
           params.cov = "main")
summary(fit, output = "main")
```

The resulting `summary()` table has one row per boundary evaluation point. It reports the order- p point estimate and standard error, the order- q bias-corrected estimate and standard error used for RBC inference, the associated test statistic, p -value, pointwise confidence interval, selected bandwidths, and side-specific effective sample sizes. When `cbands`, `WBATE`, or `LBATE` are requested, the same summary call appends uniform confidence-band columns or aggregate rows. The option `params.cov = "main"` stores the covariance matrix for the main sharp BD design estimand, which is needed for uniform bands and aggregate boundary summaries.

The implementation can report up to four bandwidths: h_{01} and h_{02} for the first and second score coordinates on the control side, and h_{11} and h_{12} for the corresponding treatment-side coordinates. The option `stdvars=TRUE`, the default, computes bandwidths after standardizing each score coordinate and then rescales them to the original units. The main `bwselect` choices are:

- `mserd` and `cerrd`: pointwise MSE and CER selectors with common bandwidths across treatment sides.
- `imserd` and `icerrd`: integrated MSE and CER selectors with common bandwidths across treatment sides.
- `msetwo` and `certwo`: pointwise MSE and CER selectors with separate control and treatment bandwidths.
- `imsetwo` and `icertwo`: integrated MSE and CER selectors with separate control and treatment bandwidths.

Unless overridden by the user, `rd2d()` uses a triangular product kernel, $p = 1$ for point estimation, $q = p + 1$ for RBC inference, `vce = "hc1"`, `fitmethod = "joint"`, `masspoints = "check"`, `level = 95`, and automatic bandwidth selection with `bwselect = "mserd"` and `method = "dpi"`. The option `method = "rot"` requests a rule-of-thumb implementation. If the user supplies `h`, the reported selector is "user provided".

The bandwidth selector can also be inspected before estimation:

```
bw <- rdbw2d(Y = Y, X = X, assignment = T, b = b,
            bwselect = "imserd", stdvars = TRUE)
summary(bw)

fit_imse <- rd2d(Y = Y, X = X, assignment = T, b = b,
               bwselect = "imserd", params.cov = "main")
```

The variance estimator is controlled by `vce` and `cluster`. When `cluster = NULL`, the package reports heteroskedasticity-consistent standard errors; when `cluster` is supplied, it reports cluster-robust standard errors. The `vce` choices are:

- `hc0`: plug-in residual variance with no degrees-of-freedom or leverage adjustment.
- `hc1`: degrees-of-freedom adjustment; this is the default.
- `hc2`: leverage-adjusted residual weighting.
- `hc3`: squared leverage-adjusted residual weighting.

The fitting convention used for variance estimation and bandwidth selection is controlled by `fitmethod`. The default, `fitmethod = "joint"`, implements the treatment-interacted regression described above. In designs without covariates, the joint and separate conventions give the same local polynomial point estimates, but they can give different standard errors and automatic bandwidths because the finite-sample degrees-of-freedom factors and cluster scores are formed differently. The joint convention is particularly relevant for clustered data because a cluster can contribute observations to both sides of the boundary and to several nearby boundary points. In that case, cluster-robust inference should add the control-side and treatment-side influence contributions before forming the cluster covariance. The option `fitmethod = "separate"` uses the earlier side-specific convention and is useful for replication or sensitivity checks.

```
# cluster_id can be, for example, a school, county, or tract group.
fit_cl <- rd2d(Y = Y, X = X, assignment = T, b = b,
             cluster = cluster_id,
             fitmethod = "joint",
             params.cov = "main")

fit_cl_separate <- rd2d(Y = Y, X = X, assignment = T, b = b,
                      cluster = cluster_id,
                      fitmethod = "separate")
```

The plug-in bandwidth rules also estimate score variances for standardization, preliminary residuals for variance constants, and higher-order derivatives of the conditional mean for bias constants. The preliminary bandwidths used for these quantities follow the direct plug-in procedure described in Cattaneo et al. [2026a] and its supplemental appendix.

3.3 Statistical Inference: Methods

For a chosen bandwidth and polynomial order, define the feasible covariance kernel $\widehat{\Omega}_{\mathbf{x}_1, \mathbf{x}_2}$. For example, in the sharp BD design,

$$\widehat{\Omega}_{\mathbf{x}_1, \mathbf{x}_2} = \frac{1}{nh^2} \mathbf{e}_0^\top \widehat{\Gamma}_{\mathbf{x}_1}^{-1} \widehat{\Sigma}_{\mathbf{x}_1, \mathbf{x}_2} \widehat{\Gamma}_{\mathbf{x}_2}^{-1} \mathbf{e}_0.$$

Here

$$\widehat{\Gamma}_{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}; T_i) \mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}; T_i)^\top K_h(\mathbf{X}_i - \mathbf{x}),$$

and

$$\widehat{\Sigma}_{\mathbf{x}_1, \mathbf{x}_2} = \frac{h^2}{n} \sum_{i=1}^n [\mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}_1; T_i) \mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}_2; T_i)^\top K_h(\mathbf{X}_i - \mathbf{x}_1) K_h(\mathbf{X}_i - \mathbf{x}_2) \widehat{\varepsilon}_i(\mathbf{x}_1) \widehat{\varepsilon}_i(\mathbf{x}_2)],$$

where $\widehat{\varepsilon}_i(\mathbf{x})$ is the residual from the augmented local polynomial fit at \mathbf{x} . These objects are also what make uniform bands and aggregate standard errors depend on covariance across boundary points, not only on pointwise standard errors. The pointwise Wald interval is

$$\widehat{\text{CI}}_\alpha(\mathbf{x}) = \left[\widehat{\tau}(\mathbf{x}) - q_\alpha \widehat{\Omega}_{\mathbf{x}, \mathbf{x}}^{1/2}, \widehat{\tau}(\mathbf{x}) + q_\alpha \widehat{\Omega}_{\mathbf{x}, \mathbf{x}}^{1/2} \right], \quad \mathbf{x} \in \mathcal{B}.$$

For pointwise inference, $q_\alpha = \Phi^{-1}(1 - \alpha/2)$. For uniform inference over \mathcal{B} , the critical value is chosen to approximate the distribution of $\sup_{\mathbf{x} \in \mathcal{B}} |\widehat{Z}(\mathbf{x})|$, where $\widehat{Z}(\mathbf{x})$ is a conditionally mean-zero Gaussian process with covariance induced by $\widehat{\Omega}_{\mathbf{x}_1, \mathbf{x}_2}$. In practice, this process is simulated on the boundary grid $\mathbf{b}_1, \dots, \mathbf{b}_J$.

Because MSE-optimal bandwidths leave non-negligible smoothing bias, `rd2d` implements robust bias-corrected (RBC) inference, following the local-polynomial RD literature [Calonico et al., 2014, 2018, 2020, 2022]. With a p th order polynomial used for point estimation and a q th order polynomial, typically $q = p+1$, used for inference, let $\widehat{\tau}_{\text{bc}}(\mathbf{x})$ denote the bias-corrected estimator and let $\widehat{\Omega}_{\text{bc}, \mathbf{x}_1, \mathbf{x}_2}$ denote the corresponding covariance kernel, including the extra variability from estimating the leading bias. The RBC interval is

$$\widehat{\text{CI}}_{\text{bc}, \alpha}(\mathbf{x}) = \left[\widehat{\tau}_{\text{bc}}(\mathbf{x}) - q_\alpha \widehat{\Omega}_{\text{bc}, \mathbf{x}, \mathbf{x}}^{1/2}, \widehat{\tau}_{\text{bc}}(\mathbf{x}) + q_\alpha \widehat{\Omega}_{\text{bc}, \mathbf{x}, \mathbf{x}}^{1/2} \right].$$

The resulting intervals and bands are generally not centered at the original BATEC point estimator $\widehat{\tau}(\mathbf{x})$.

3.4 Statistical Inference: Implementation

Pointwise intervals are reported by default through the `summary()` method. Uniform bands are requested with `cbands`; the corresponding covariance matrix must be stored when the model is fit.

```
fit <- rd2d(Y = Y, X = X, assignment = T, b = b,
           params.cov = "main")
summary(fit, output = "main")
summary(fit, output = "main", cbands = "main")
```

The default polynomial orders are $p = 1$ for point estimation and $q = 2$ for RBC inference. The option `level` controls the nominal coverage level, and `repp` controls the number of Gaussian simulations used for uniform critical values. Finite-sample covariance regularization may be used when the simulated Gaussian covariance matrix is nearly singular, a situation that can arise with closely spaced boundary grid points or weak local support.

3.5 WBATE and LBATE: Estimation and Inference

The WBATE is fit by plugging the estimated BATEC into the boundary integral. With normalized weights $\int_{\mathcal{B}} w(\mathbf{x}) \, d\mathbf{x} = 1$,

$$\hat{\tau}_{\text{WBATE}} = \int_{\mathcal{B}} \hat{\tau}(\mathbf{x}) w(\mathbf{x}) \, d\mathbf{x}.$$

On a grid, this is approximated by $\sum_{j=1}^J \omega_j \hat{\tau}(\mathbf{b}_j) w(\mathbf{b}_j)$, with quadrature weights ω_j normalized so the weighted sum integrates to one. The WBATE covariance is the corresponding double integral of the BATEC covariance kernel,

$$\hat{\Omega}_{\text{WBATE}} = \int_{\mathcal{B}} \int_{\mathcal{B}} \hat{\Omega}_{\mathbf{x}_1, \mathbf{x}_2} w(\mathbf{x}_1) w(\mathbf{x}_2) \, d\mathbf{x}_1 \, d\mathbf{x}_2,$$

or its grid analogue. Aggregation changes the effective variance order: the pointwise BATEC estimator has variance of order $(nh^2)^{-1}$, while the WBATE has variance of order $(nh)^{-1}$ because only boundary locations within an $O(h)$ neighborhood have non-negligible covariance. Thus, the WBATE behaves like a one-dimensional nonparametric functional even though each local fit uses a two-dimensional score. Inference uses a normal critical value, with RBC implemented by applying the same plug-in formula to the bias-corrected BATEC and its covariance.

The LBATE is the supremum of the BATEC, with plug-in estimator

$$\hat{\tau}_{\text{LBATE}} = \sup_{\mathbf{x} \in \mathcal{B}} \hat{\tau}(\mathbf{x}) \approx \max_{1 \leq j \leq J} \hat{\tau}(\mathbf{b}_j).$$

Because LBATE is an extreme functional, Cattaneo et al. [2026a] recommends reporting it together with the full BATEC and WBATE. Its confidence interval is obtained by projecting the uniform

confidence band for $\tau(\mathbf{x})$:

$$\widehat{\text{CI}}_{\alpha, \text{LBATE}} = \left[\sup_{\mathbf{x} \in \mathcal{B}} \left(\widehat{\tau}(\mathbf{x}) - q_{\alpha} \widehat{\Omega}_{\mathbf{x}, \mathbf{x}}^{1/2} \right), \sup_{\mathbf{x} \in \mathcal{B}} \left(\widehat{\tau}(\mathbf{x}) + q_{\alpha} \widehat{\Omega}_{\mathbf{x}, \mathbf{x}}^{1/2} \right) \right],$$

using the uniform critical value q_{α} . The interval is generally conservative because it inherits simultaneous coverage for the full BATEC.

In the software, both summaries are produced from the stored covariance matrix:

```
w <- rep(1, nrow(b))
summary(fit, output = "main", cbands = "main",
        WBATE = w, LBATE = TRUE)
```

3.6 Imperfect Compliance

For fuzzy BD designs, the same local polynomial machinery is applied to the outcome Y_i and the treatment status W_i . Let $\widehat{\tau}_Y(\mathbf{x})$ and $\widehat{\tau}_W(\mathbf{x})$ denote the reduced-form and first-stage BATEC estimators. The fuzzy BATEC estimator is

$$\widehat{\zeta}(\mathbf{x}) = \frac{\widehat{\tau}_Y(\mathbf{x})}{\widehat{\tau}_W(\mathbf{x})}.$$

Here $\tau_Y(\mathbf{x})$ is the ITT BATEC and $\tau_W(\mathbf{x})$ is the FS BATEC. The ratio $\zeta(\mathbf{x})$ is the fuzzy BATEC parameter; it has a local complier-average causal interpretation only under additional fuzzy BD design assumptions such as a local exclusion restriction, monotonicity, and a nonzero first stage. For related discussion of causal identification in fuzzy multidimensional RD designs, see [Choi and Lee \[2023\]](#), [Schwarz et al. \[2025\]](#), [Jiang and Zhu \[2026\]](#), and [Cattaneo et al. \[2016, 2021\]](#); for univariate fuzzy RD designs, see [Arai et al. \[2022\]](#) and references therein.

[Cattaneo et al. \[2026a\]](#) uses the linearization

$$\widehat{\zeta}(\mathbf{x}) - \zeta(\mathbf{x}) = \mathbf{w}(\mathbf{x})^{\top} \mathfrak{F}(\mathbf{x}) + \mathfrak{R}_{\mathbf{F}}(\mathbf{x}), \quad \mathbf{w}(\mathbf{x}) = \left[\frac{1}{\tau_W(\mathbf{x})}, -\frac{\tau_Y(\mathbf{x})}{\tau_W(\mathbf{x})^2} \right]^{\top},$$

where $\mathfrak{F}(\mathbf{x}) = (\widehat{\tau}_Y(\mathbf{x}) - \tau_Y(\mathbf{x}), \widehat{\tau}_W(\mathbf{x}) - \tau_W(\mathbf{x}))^{\top}$. The remainder $\mathfrak{R}_{\mathbf{F}}(\mathbf{x})$ is negligible when the first stage is uniformly bounded away from zero and the reduced-form and first-stage estimators are uniformly consistent. Thus, estimation and inference for fuzzy BATEC, fuzzy WBATE, and fuzzy LBATE follow by applying the sharp BD design results to the linearized combination $\mathbf{w}(\mathbf{x})^{\top} \mathfrak{F}(\mathbf{x})$.

The fuzzy covariance kernel is

$$\widehat{\Omega}_{\mathbf{F}, \mathbf{x}_1, \mathbf{x}_2} = \widehat{\mathbf{w}}(\mathbf{x}_1)^{\top} \begin{pmatrix} \widehat{\Omega}_{Y, Y, \mathbf{x}_1, \mathbf{x}_2} & \widehat{\Omega}_{Y, W, \mathbf{x}_1, \mathbf{x}_2} \\ \widehat{\Omega}_{W, Y, \mathbf{x}_1, \mathbf{x}_2} & \widehat{\Omega}_{W, W, \mathbf{x}_1, \mathbf{x}_2} \end{pmatrix} \widehat{\mathbf{w}}(\mathbf{x}_2),$$

where $\widehat{\mathbf{w}}(\mathbf{x})$ replaces $\tau_Y(\mathbf{x})$ and $\tau_W(\mathbf{x})$ by their local polynomial estimates. Fuzzy BATEC intervals

and bands use $\widehat{\Omega}_{F, \mathbf{x}_1, \mathbf{x}_2}$ in place of the sharp BD design covariance kernel. Fuzzy WBATE uses

$$\widehat{\zeta}_{WBATE} = \int_{\mathcal{B}} \widehat{\zeta}(\mathbf{x}) w(\mathbf{x}) d\mathbf{x} \quad \text{and} \quad \widehat{\Omega}_{F, WBATE} = \int_{\mathcal{B}} \int_{\mathcal{B}} \widehat{\Omega}_{F, \mathbf{x}_1, \mathbf{x}_2} w(\mathbf{x}_1) w(\mathbf{x}_2) d\mathbf{x}_1 d\mathbf{x}_2,$$

with a normal critical value. Fuzzy LBATE uses $\widehat{\zeta}_{LBATE} = \sup_{\mathbf{x} \in \mathcal{B}} \widehat{\zeta}(\mathbf{x})$ and projects the fuzzy BATEC uniform band in the same way as for LBATE in a sharp BD design.

In `rd2d`, fuzzy BD designs are fit by supplying `fuzzy = W`. For Opportunity Zones, this means that the first-stage output estimates the effect of eligibility on designation, while the ITT output estimates the effect of eligibility on early workplace-job growth. The output `"itt"` reports $\widehat{\tau}_Y(\mathbf{x})$, `"fs"` reports $\widehat{\tau}_W(\mathbf{x})$, and `"main"` reports $\widehat{\zeta}(\mathbf{x})$. The option `bwparam = "main"` targets the fuzzy BATEC, while `bwparam = "itt"` targets the reduced-form BATEC.

```
fit_fuzzy <- rd2d(Y = Y, X = X, assignment = T, b = b,
                fuzzy = W,
                params.cov = c("main", "itt", "fs"))
w <- rep(1, nrow(b))
summary(fit_fuzzy, output = "itt", cbands = "itt")
summary(fit_fuzzy, output = "fs", cbands = "fs")
summary(fit_fuzzy, output = "main", cbands = "main",
        WBATE = w, LBATE = TRUE)
```

3.7 Regularization Strategies and Implementation Choices

The package `rd2d` implements several numerical safeguards and reporting choices that are important in empirical applications. These choices do not change the target parameters; they stabilize the finite-sample implementation of BATEC, WBATE, LBATE, and their fuzzy analogues, and make the resulting diagnostics easier to interpret.

- *Boundary grid and target set.* Uniform bands, WBATE, and LBATE are computed on the user-supplied boundary grid `b`. For a piecewise linear boundary, a natural default is to place grid points approximately equally spaced in arc length, while forcing kinks, corners, or substantively important boundary points to be included. If inference is desired only on a trimmed subset $\mathcal{B}_0 \subseteq \mathcal{B}$, for example to avoid endpoints or poorly supported corners, the grid should represent \mathcal{B}_0 and the reported estimands should be interpreted as applying to that trimmed target set.
- *Grid sensitivity.* Since the theoretical uniform band is indexed by the continuum $\mathbf{x} \in \mathcal{B}$ but the software simulates the process on a finite grid, the grid should be fine enough that the discretized supremum is stable. A useful diagnostic is to refine the grid and check whether the simulated critical value, the confidence band, and aggregate summaries are essentially unchanged. In applications with heterogeneous effects or irregular boundaries, this check is more informative than relying on a fixed number of grid points.

- *Bandwidth choices along the boundary.* Pointwise selectors target a particular boundary location, whereas integrated selectors target the full boundary curve or a chosen subset. For uniform inference and graphical displays, a common boundary-wide bandwidth often gives smoother covariance estimates and simpler interpretation. Point-specific bandwidths can also be used, but the covariance matrix used for uniform bands and aggregate summaries must be computed using the same bandwidth choices. The default standardization option, `stdvars = TRUE`, helps make scalar bandwidths comparable when the two score coordinates have different units.
- *Joint versus separate fitting.* The option `fitmethod = "joint"` is the default because it matches the treatment-interacted regression used to define the estimator and because it aggregates cluster scores across both sides of the boundary. This distinction matters when the same cluster can contain control and treated units near a boundary point. The option `fitmethod = "separate"` is available when users want side-specific finite-sample adjustments for comparison with earlier implementations.
- *Covariate efficiency adjustment.* The option `covs.eff` supplies predetermined covariates used to improve precision without changing the treatment-effect estimand. The same adjustment is used in automatic bandwidth selection, pointwise inference, and covariance matrices for uniform bands and aggregate summaries. The safeguards `covs.drop` and `covs.tol` control how redundant covariate columns are handled after residualization.
- *Small bias regularization.* Approximate MSE and IMSE bandwidth choices require nonzero leading bias terms. A very small estimated bias can therefore produce an excessively large bandwidth. To avoid this problem, the software adds a regularization term to the estimated bias denominator:

$$h_{\text{MSE},\mathbf{x}} = \left(\frac{2\widehat{V}_{\mathbf{x}}}{(2p+2)(\widehat{B}_{\mathbf{x}}^2 + s \cdot \mathbb{V}[\widehat{B}_{\mathbf{x}}])} \frac{1}{n} \right)^{1/(2p+4)}.$$

The analogous IMSE selector regularizes $\int_{\mathcal{B}} \widehat{B}_{\mathbf{x}}^2 w(\mathbf{x}) \, d\mathbf{x}$ by the integrated variance of the bias estimator. The factor s is controlled by `scaleregul`; its default is 3 when bandwidth selection is called through `rd2d()` and 1 in standalone calls to `rdbw2d()`.

- *Minimum sample size.* The option `bwcheck` enlarges the selected or user-provided bandwidth until a minimum number of observations is included in the estimation region on each side of the cutoff. For `rd2d()`, the default is `bwcheck = 50 + p + 1`; standalone calls to `rdbw2d()` use `bwcheck = 20`.
- *Mass points in \mathbf{X}_i .* The option `masspoints = "check"` reports the number of unique score values and warns when duplication is substantial. With `masspoints = "adjust"`, bandwidths are regularized so that kernels contain a minimum number of unique observations. With `masspoints = "off"`, the potential presence of mass points is ignored.

- *Covariance regularization for uniform bands.* The estimated covariance matrix is converted to a correlation matrix before Gaussian simulation. In finite samples this matrix can be nearly singular or slightly fail to be positive semidefinite because of numerical precision, closely spaced grid points, or weak local support. The implementation symmetrizes the matrix, truncates eigenvalues below a small tolerance, rescales the diagonal entries to one, and warns when the adjustment is nonnegligible. This eigenvalue regularization follows the same logic as other covariance regularization methods for uniform inference [Cattaneo et al., 2024a].
- *Gaussian simulation error.* The option `repp` controls the number of Gaussian simulations used for uniform critical values and LBATE inference. The default is `repp = 1000` in the summary methods. When the critical value is central to the empirical conclusion, increasing `repp` and checking that the reported band is stable is a simple numerical diagnostic.
- *Local support and design stability.* Boundary points with little one-sided support, very unbalanced kernel mass across the two sides, or poorly conditioned local design matrices can lead to unstable BATEC estimates and covariance estimates. The `bwcheck` and `masspoints` options address two common symptoms, but researchers should also inspect the score scatterplot, side-specific local sample sizes, and sensitivity to trimming weakly supported regions of the boundary. These checks matter most near kinks, corners, and endpoints.
- *Reporting practice.* Empirical analyses should report the target boundary set, the evaluation grid, the bandwidth selector, the fitting convention, any clustering or covariate adjustment, whether covariance regularization produced a warning, and the simulation count used for uniform bands. Basic sensitivity checks with alternative grid sizes, bandwidth choices, and, when relevant, boundary trims help separate genuine treatment-effect heterogeneity from discretization, weak support, or numerical instability.

The empirical implementation can be organized as follows. First, construct and plot the boundary grid, making sure that kinks, endpoints, and substantively important locations are represented. Second, estimate the location-based BATEC with covariance storage for the outputs that will be used for uniform bands, WBATE, or LBATE, using `fitmethod = "joint"` unless the goal is a comparison with the earlier convention. Third, add `cluster` and `covs.eff` when the design calls for clustered inference or pre-treatment efficiency adjustment. Fourth, in fuzzy BD designs, inspect the FS BATEC before interpreting the fuzzy Wald curve, because weak or sign-changing first stages make the ratio unstable. Fifth, report pointwise intervals, uniform bands, WBATE, and LBATE for the main target, and use the side-specific effective sample sizes and bandwidth table to flag poorly supported boundary regions. Finally, compare the location-based analysis with distance-based estimates and with modest changes to the grid, bandwidth selector, and trimmed target set.

3.8 Derivative Estimation

Location-based methods can also be used to estimate derivatives of the side-specific regression functions and their discontinuities at the boundary. For a multi-index $\boldsymbol{\nu} = (\nu_1, \nu_2)$ with $|\boldsymbol{\nu}| = \nu_1 + \nu_2 \leq p$, define the derivative BATEC

$$\tau^{(\boldsymbol{\nu})}(\mathbf{x}) = \mu_1^{(\boldsymbol{\nu})}(\mathbf{x}) - \mu_0^{(\boldsymbol{\nu})}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{B},$$

where $\mu_t^{(\boldsymbol{\nu})}(\mathbf{x})$ denotes the partial derivative of $\mu_t(\mathbf{x})$ of order $\boldsymbol{\nu}$. These derivative targets describe how the discontinuity changes locally with the score coordinates; for example, $\boldsymbol{\nu} = (1, 0)$ and $\boldsymbol{\nu} = (0, 1)$ correspond to the two coordinate-specific slopes of the BATEC at the boundary.

The plug-in estimator uses the same local polynomial coefficients as the level estimator. Using the selector $\mathbf{e}_{\boldsymbol{\nu}}$ introduced above,

$$\hat{\tau}^{(\boldsymbol{\nu})}(\mathbf{x}) = \boldsymbol{\nu}! h^{-|\boldsymbol{\nu}|} \mathbf{e}_{\boldsymbol{\nu}}^\top \hat{\boldsymbol{\delta}}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{B},$$

where $\boldsymbol{\nu}! = \nu_1! \nu_2!$. Directional derivatives, including derivatives along the boundary, are obtained by applying the corresponding linear combination to the first partial derivatives.

In `rd2d`, derivative targets are requested with `deriv`; for example, `deriv = c(1, 0)` estimates the discontinuity in the first partial derivative with respect to the first score coordinate. Directional derivatives along user-supplied vectors, such as boundary tangent vectors, are requested with `tangvec`, a $J \times 2$ matrix with one vector for each boundary evaluation point; when supplied, `tangvec` overrides `deriv`. In fuzzy BD designs, the derivative options also apply to the reduced-form and first-stage output tables. Ratios involving derivative discontinuities should be interpreted separately from the level fuzzy BATEC $\zeta(\mathbf{x}) = \tau_Y(\mathbf{x})/\tau_W(\mathbf{x})$.

3.9 Covariate Adjustment

Predetermined covariates can improve precision in boundary discontinuity designs when they explain residual variation in the outcome. The implementation in `rd2d` follows the covariate adjustment logic developed for canonical RD designs by Calonico et al. [2019]. Let \mathbf{Z}_i be a d_Z -vector of predetermined covariates. We use a tilde to denote covariate-adjusted quantities. The covariate-adjusted fit uses the same augmented local-polynomial basis as above and adds $\mathbf{Z}_i^\top \boldsymbol{\lambda}$ to the weighted least-squares regression:

$$\left(\tilde{\boldsymbol{\delta}}(\mathbf{x}), \tilde{\boldsymbol{\lambda}}(\mathbf{x}) \right) = \arg \min_{\boldsymbol{\delta} \in \mathbb{R}^{2p}, \boldsymbol{\lambda} \in \mathbb{R}^{d_Z}} \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{r}_{p,h}(\mathbf{X}_i - \mathbf{x}; T_i)^\top \boldsymbol{\delta} - \mathbf{Z}_i^\top \boldsymbol{\lambda})^2 K_h(\mathbf{X}_i - \mathbf{x}).$$

The covariate-adjusted BATEC estimator is $\tilde{\tau}(\mathbf{x}) = \mathbf{e}_0^\top \tilde{\boldsymbol{\delta}}(\mathbf{x})$. Under the usual smoothness and no-manipulation conditions, this adjustment does not change the estimand; it can reduce variance when the covariates are predictive of the outcome.

In the software, covariates are supplied through `covs.eff`. The covariate coefficient is common

across treatment sides, while the local polynomial basis remains side-specific. The implementation residualizes the covariates on the local polynomial basis and carries the resulting adjusted influence functions through pointwise inference, uniform bands, WBATE, LBATE, fuzzy ratio inference, and automatic bandwidth selection. Thus, when `h` is omitted, the same covariate adjustment used in estimation is also used by `rdbw2d()`. The options `covs.drop` and `covs.tol` control safeguards for rank-deficient covariate matrices after residualization.

```
# Z is an n by d matrix or data frame of pre-treatment covariates.
fit_cov <- rd2d(Y = Y, X = X, assignment = T, b = b,
              fuzzy = W, covs.eff = Z,
              fitmethod = "joint",
              params.cov = c("main", "itt", "fs"))
summary(fit_cov, output = "main", cbands = "main")

bw_cov <- rdbw2d(Y = Y, X = X, assignment = T, b = b,
                fuzzy = W, covs.eff = Z,
                bwparam = "main")
```

The boundary setting differs from the scalar RD setting because the local fit is repeated over points on \mathcal{B} and because covariance across boundary points matters for uniform bands, WBATE, and LBATE. Thus, covariate adjustment must be carried through the same covariance calculations used for the unadjusted BATEC. The same command structure can also be used for falsification exercises by treating predetermined covariates as outcomes, but that diagnostic task is distinct from covariate-adjusted treatment-effect estimation. The empirical application below keeps the fuzzy BD analysis unadjusted to focus on the default bivariate and distance-based estimators.

4 Distance-Based Methods

Distance-based methods replace the bivariate score \mathbf{X}_i by a scalar signed distance score constructed separately for each boundary point. For $\mathbf{x} = (x_1, x_2)^\top \in \mathcal{B}$, let

$$D_i(\mathbf{x}) = (\mathbf{1}(\mathbf{X}_i \in \mathcal{A}_1) - \mathbf{1}(\mathbf{X}_i \in \mathcal{A}_0)) \cdot \mathcal{d}(\mathbf{X}_i, \mathbf{x}),$$

where $\mathcal{d}(\cdot, \cdot)$ is a distance function such as the Euclidean distance $\mathcal{d}(\mathbf{X}_i, \mathbf{x}) = \|\mathbf{X}_i - \mathbf{x}\|$. Thus, $D_i(\mathbf{x}) \geq 0$ identifies observations on the treated side of the boundary and $D_i(\mathbf{x}) < 0$ identifies observations on the control side. Conditional on \mathbf{x} , the problem resembles a univariate RD design with cutoff zero, while varying \mathbf{x} over \mathcal{B} recovers a curve of local effects. In the Opportunity Zones example, each column of the distance matrix is the signed distance from a tract's poverty-income score to one of the 31 boundary points shown in Figure 2.

This construction is attractive when the full bivariate score is unavailable, difficult to release, or inconvenient to smooth directly. It also gives an intuitive set of RD plots and diagnostics for each boundary point. The cost is that distance-based methods use less location information than the location-based procedures in Section 3; as a result, their interpretation depends more explicitly on

the geometry of \mathcal{B} and on the chosen distance function.

Distance-based identification requires an additional step because the scalar score averages potential outcomes over level sets of the distance function. For $\mathbf{x} \in \mathcal{B}$, define

$$\theta_{t,\mathbf{x}}(r) = \mathbb{E}[Y_i | D_i(\mathbf{x}) = r] = \mathbb{E}[Y_i(t) | \mathcal{d}(\mathbf{X}_i, \mathbf{x}) = |r|, \mathbf{X}_i \in \mathcal{A}_t], \quad r \in \mathcal{I}_t,$$

where $\mathcal{I}_0 = (-\infty, 0)$ and $\mathcal{I}_1 = [0, \infty)$. The induced signed-distance treatment-effect parameter is

$$\vartheta(\mathbf{x}) = \theta_{1,\mathbf{x}}(0) - \theta_{0,\mathbf{x}}(0), \quad \mathbf{x} \in \mathcal{B}.$$

[Cattaneo et al. \[2026b\]](#), Theorem 1] gives conditions on the data-generating process, the assignment boundary, the kernel, and the distance function under which

$$\vartheta(\mathbf{x}) = \tau(\mathbf{x}), \quad \mathbf{x} \in \mathcal{B}.$$

The result is obtained by representing the induced conditional expectations as integrals over shrinking submanifolds around \mathbf{x} ; without those restrictions on the distribution, the boundary geometry, and the distance function, the distance-based estimand need not coincide with the BATEC.

The same logic identifies aggregate and fuzzy targets. Since $\vartheta(\mathbf{x}) = \tau(\mathbf{x})$ pointwise along \mathcal{B} under the distance-based identification conditions, WBATE and LBATE can be represented by applying the same integral and supremum functionals to $\vartheta(\mathbf{x})$. In fuzzy BD designs, applying the signed-distance construction to $A \in \{Y, W\}$ gives

$$\vartheta_A(\mathbf{x}) = \theta_{A,1,\mathbf{x}}(0) - \theta_{A,0,\mathbf{x}}(0), \quad \xi(\mathbf{x}) = \frac{\vartheta_Y(\mathbf{x})}{\vartheta_W(\mathbf{x})}.$$

Under the fuzzy identification conditions in [Cattaneo et al. \[2026b\]](#), Section SA-6], $\vartheta_A(\mathbf{x}) = \tau_A(\mathbf{x})$ for $A \in \{Y, W\}$, and if the first stage is nonzero then $\xi(\mathbf{x}) = \zeta(\mathbf{x})$. Thus, aggregating $\xi(\mathbf{x})$ or taking its supremum gives distance-based representations of the causal fuzzy WBATE and fuzzy LBATE.

4.1 BATEC Estimation and Inference

The distance-based local polynomial estimator is computed separately for every evaluation point. In sharp BD designs, the analysis replaces \mathbf{X} with an n by J matrix of signed distance scores, one column per boundary point. For Opportunity Zones, the j th column of \mathbf{D} is the signed Euclidean distance from each tract's poverty-income score to \mathbf{b}_j , with positive values for eligible tracts:

```
# D is an n by J signed-distance matrix, one column per boundary point.
fit_dist <- rd2d.distance(Y = Y, distance = D, b = b,
                        cbands = TRUE)
summary(fit_dist, output = "main", cbands = "main")
```

The point estimator is

$$\hat{\vartheta}(\mathbf{x}) = \mathbf{e}_0^\top \hat{\boldsymbol{\gamma}}_1(\mathbf{x}) - \mathbf{e}_0^\top \hat{\boldsymbol{\gamma}}_0(\mathbf{x}), \quad \mathbf{x} \in \mathcal{B},$$

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

$$\hat{\boldsymbol{\gamma}}_t(\mathbf{x}) = \arg \min_{\boldsymbol{\gamma} \in \mathbb{R}^{p+1}} \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{r}_p(D_i(\mathbf{x})/h)^\top \boldsymbol{\gamma})^2 K_h(D_i(\mathbf{x})) \mathbf{1}(D_i(\mathbf{x}) \in \mathcal{J}_t),$$

with \mathbf{e}_0 the intercept selector, $\mathbf{r}_p(u) = (1, u, u^2, \dots, u^p)^\top$ the usual univariate polynomial basis, $K_h(u) = K(u/h)/h^2$, and $\mathcal{J}_0 = (-\infty, 0)$ and $\mathcal{J}_1 = [0, \infty)$. The normalization by h^2 reflects the dimension of the original bivariate score; it does not change the weighted least squares solution because it multiplies all weights by the same positive constant.

Cattaneo et al. [2026b] establishes pointwise and uniform consistency, distributional approximations, and confidence intervals and bands for $\hat{\vartheta}(\mathbf{x})$. The main practical distinction relative to location-based methods is the bias behavior near nonsmooth regions of the boundary. If \mathcal{B} is sufficiently smooth, the distance-induced regression functions have the usual local-polynomial approximation properties, and robust bias correction proceeds with $q = p + 1$ by default. If \mathcal{B} has a kink or another local irregularity, the induced univariate regression can be only Lipschitz even when the underlying bivariate regression functions are smoother. Then the leading bias is of order h uniformly near the irregularity, and increasing p does not remove it.

4.2 WBATE and LBATE

The causal WBATE and LBATE remain the functionals of the BATEC defined in Section 2. Under the distance-based identification conditions, $\tau(\mathbf{x}) = \vartheta(\mathbf{x})$ for $\mathbf{x} \in \mathcal{B}$, so the same targets can be written using the signed-distance estimand. With normalized weights $\int_{\mathcal{B}} w(\mathbf{x}) \, d\mathbf{x} = 1$,

$$\tau_{\text{WBATE}} = \int_{\mathcal{B}} \tau(\mathbf{x}) w(\mathbf{x}) \, d\mathbf{x} = \int_{\mathcal{B}} \vartheta(\mathbf{x}) w(\mathbf{x}) \, d\mathbf{x}, \quad \hat{\vartheta}_{\text{WBATE}} = \int_{\mathcal{B}} \hat{\vartheta}(\mathbf{x}) w(\mathbf{x}) \, d\mathbf{x}.$$

On the evaluation grid, $\hat{\vartheta}_{\text{WBATE}}$ is implemented as a weighted average of the distance-based BATEC estimates. The covariance is the corresponding double integral of the distance-based covariance kernel,

$$\hat{\Xi}_{\text{WBATE}} = \int_{\mathcal{B}} \int_{\mathcal{B}} \hat{\Xi}_{\mathbf{x}_1, \mathbf{x}_2} w(\mathbf{x}_1) w(\mathbf{x}_2) \, d\mathbf{x}_1 \, d\mathbf{x}_2,$$

or its grid analogue. As shown in Cattaneo et al. [2026b, Section SA-4], the WBATE has effective variance order $(nh)^{-1}$ under local boundary-measure regularity, even though each pointwise distance-based fit has the bivariate variance order $(nh^2)^{-1}$ when $d = 2$.

Similarly, under distance-based identification, the causal LBATE can be written as

$$\tau_{\text{LBATE}} = \sup_{\mathbf{x} \in \mathcal{B}} \tau(\mathbf{x}) = \sup_{\mathbf{x} \in \mathcal{B}} \vartheta(\mathbf{x}), \quad \widehat{\vartheta}_{\text{LBATE}} = \sup_{\mathbf{x} \in \mathcal{B}} \widehat{\vartheta}(\mathbf{x}).$$

Its confidence interval is obtained by projecting the uniform confidence band for $\vartheta(\mathbf{x})$, as in [Cattaneo et al. \[2026b, Section SA-5\]](#). The resulting interval is conservative in general because it inherits simultaneous coverage for the full distance-based BATEC.

In the software, both summaries are produced by the `summary()` method once the relevant covariance matrix has been stored:

```
w <- rep(1, nrow(b))
summary(fit_dist, output = "main", cbands = "main",
        WBATE = w, LBATE = TRUE)
```

Under the distance-based identification conditions, these are plug-in estimators of τ_{WBATE} and τ_{LBATE} .

4.3 Imperfect Compliance

Distance-based methods extend to fuzzy BD designs by applying the same signed-distance local polynomial procedure to the outcome and to treatment receipt. For $A \in \{Y, W\}$, let

$$\widehat{\vartheta}_A(\mathbf{x}) = \mathbf{e}_0^\top \widehat{\gamma}_{A,1}(\mathbf{x}) - \mathbf{e}_0^\top \widehat{\gamma}_{A,0}(\mathbf{x}),$$

where $\widehat{\gamma}_{A,t}(\mathbf{x})$ is the local polynomial fit above with A_i replacing Y_i . The distance-based fuzzy BATEC estimator is the Wald ratio

$$\widehat{\xi}(\mathbf{x}) = \frac{\widehat{\vartheta}_Y(\mathbf{x})}{\widehat{\vartheta}_W(\mathbf{x})}, \quad \mathbf{x} \in \mathcal{B}.$$

The reduced-form output estimates $\vartheta_Y(\mathbf{x})$, the first-stage output estimates $\vartheta_W(\mathbf{x})$, and the main fuzzy output estimates $\xi(\mathbf{x})$. Under the identification conditions summarized above, these correspond to $\tau_Y(\mathbf{x})$, $\tau_W(\mathbf{x})$, and the fuzzy BATEC $\zeta(\mathbf{x})$.

Inference is based on the ratio linearization in [Cattaneo et al. \[2026b, Section SA-6\]](#). With

$$\mathbf{w}(\mathbf{x}) = \left[\frac{1}{\vartheta_W(\mathbf{x})}, -\frac{\vartheta_Y(\mathbf{x})}{\vartheta_W(\mathbf{x})^2} \right]^\top,$$

the first-order component of $\widehat{\xi}(\mathbf{x}) - \xi(\mathbf{x})$ is the linear combination of the reduced-form and first-stage errors indexed by $\mathbf{w}(\mathbf{x})$. Thus, the fuzzy covariance kernel is obtained by sandwiching the joint covariance kernel of $(\widehat{\vartheta}_Y(\mathbf{x}), \widehat{\vartheta}_W(\mathbf{x}))^\top$ with this gradient vector. This delta-method covariance is used for pointwise intervals, uniform confidence bands, fuzzy WBATE inference, and fuzzy LBATE inference.

With the same weight normalization, the causal fuzzy WBATE and fuzzy LBATE remain the

functionals of $\zeta(\mathbf{x})$ defined in Section 2,

$$\zeta_{\text{WBATE}} = \int_{\mathcal{B}} \zeta(\mathbf{x})w(\mathbf{x}) d\mathbf{x}, \quad \zeta_{\text{LBATE}} = \sup_{\mathbf{x} \in \mathcal{B}} \zeta(\mathbf{x}).$$

The corresponding distance-based plug-in estimators replace $\zeta(\mathbf{x})$ by $\hat{\xi}(\mathbf{x})$:

$$\hat{\zeta}_{\text{WBATE}} = \int_{\mathcal{B}} \hat{\xi}(\mathbf{x})w(\mathbf{x}) d\mathbf{x}, \quad \hat{\zeta}_{\text{LBATE}} = \sup_{\mathbf{x} \in \mathcal{B}} \hat{\xi}(\mathbf{x}).$$

Under fuzzy distance-based identification, $\xi(\mathbf{x}) = \zeta(\mathbf{x})$, so these are plug-in estimators of ζ_{WBATE} and ζ_{LBATE} . In the Opportunity Zones application, the FS curve measures how eligibility changes the probability of designation, the ITT curve measures how eligibility changes early workplace-job growth, and the fuzzy curve rescales the latter by the former.

In the software, supplying fuzzy changes `main` to the fuzzy Wald effect, while `itt` and `fs` report the reduced-form and first-stage BATECs. The option `bwparam = "main"` selects automatic bandwidths for the linearized fuzzy ratio, while `bwparam = "itt"` targets the reduced-form outcome.

```
fit_dist_fuzzy <- rd2d.distance(
  Y = Y, distance = D, b = b, fuzzy = W,
  params.cov = c("main", "itt", "fs"))

summary(fit_dist_fuzzy, output = "itt", cbands = "itt")
summary(fit_dist_fuzzy, output = "fs", cbands = "fs")
summary(fit_dist_fuzzy, output = "main", cbands = "main",
  WBATE = w, LBATE = TRUE)
```

4.4 Implementation Details

The distance-based commands share most options with the location-based commands. Bandwidth selection uses `rdbw2d.distance()`; if `h` is omitted, `rd2d.distance()` calls this selector internally. A scalar `h` applies the same bandwidth to all boundary points and both sides, while a J by 2 matrix supplies boundary-specific control and treated bandwidths. Unless overridden, `rd2d.distance()` uses a triangular kernel, $p = 1$, $q = p + 1$ under the smooth-boundary default, `vce = "hc1"`, `fitmethod = "joint"`, `masspoints = "check"`, `level = 95`, and automatic bandwidth selection with `bwselect = "mserd"`. The available `bwselect` choices are `mserd`, `cerrd`, `imserd`, `icerrd`, `msetwo`, `certwo`, `imsetwo`, and `icertwo`; the default regularization constants are `scaleregul = 1`, `qct = 0.5`, and `bwcheck = 50 + p + 1` for `rd2d.distance()`.

Cluster-robust inference and covariate efficiency adjustment use the same option names and interpretation as in the location-based commands. The default `fitmethod = "joint"` aggregates cluster scores across both signs of the distance score, which is the appropriate convention when clusters can appear on both sides of a local distance cutoff. The option `covs.eff` supplies predetermined covariates for the common-coefficient efficiency adjustment and is also passed to `rdbw2d.distance()` when automatic bandwidths are selected.

```
fit_dist_cov <- rd2d.distance(
  Y = Y, distance = D, b = b, fuzzy = W,
  cluster = cluster_id,
  covs.eff = Z,
  fitmethod = "joint",
  params.cov = c("main", "itt", "fs"))
```

The package exposes the boundary-geometry distinction in Cattaneo et al. [2026b] through two options. The default setting treats the boundary locally as smooth and uses the usual robust bias-corrected construction. When kinks may be present but their locations are not being used, `kink.unknown` can shrink the automatic bandwidths to the rates justified for unknown irregularities and set the default bias-correction order to $q = p$. When kink locations are known on the evaluation grid, `kink.position` accepts either a logical vector or integer indices of the kink boundary points, and automatic bandwidth selection adapts to proximity to those locations:

```
fit_kink <- rd2d.distance(Y = Y, distance = D, b = b,
  kink.position = 16)
summary(fit_kink, output = "main", cbands = "main")

fit_ukink <- rd2d.distance(Y = Y, distance = D, b = b,
  kink.unknown = c(TRUE, TRUE))
summary(fit_ukink, output = "main", cbands = "main")
```

Additional output can be requested when needed. In sharp BD designs, `params.other` can request `main.0` and `main.1`, the two side-specific fits. In fuzzy BD designs, `params.other` can request `itt.0`, `itt.1`, `fs.0`, and `fs.1`; `params.cov` controls which covariance matrices are stored for uniform bands and aggregate inference. The printed and returned `summary()` tables mirror the location-based output: point estimates, RBC standard errors, test statistics, p -values, pointwise intervals, optional uniform-band columns, bandwidths, and effective sample sizes are reported at each boundary point, with optional WBATE and LBATE rows appended. Standard error options include `vce = "hc0"`, `"hc1"`, `"hc2"`, or `"hc3"`. Finally, `masspoints` controls checking and adjustment for repeated signed-distance values, which can matter when locations or distances are discretized. These options make distance-based analysis useful both as a standalone strategy and as a robustness comparison to the richer location-based methods.

5 Empirical Results

We now return to the Opportunity Zones running example introduced in Section 1.1. The empirical analysis reports the three objects in a fuzzy BD design: the reduced-form BATEC for early workplace-job growth, the first-stage BATEC for Qualified Opportunity Zone designation, and the fuzzy BATEC obtained by dividing the reduced form by the first stage. We report the fuzzy curve as a Wald estimand; a complier-average causal interpretation would require the additional local exclusion, monotonicity, and nonzero-first-stage conditions discussed in Section 2. For each object, we compare the location-based fit, which uses the bivariate poverty-income score directly, with

the distance-based fit, which uses the signed Euclidean distance to each evaluation point on the boundary. Both fits use the package’s automatic default choices, including `fitmethod = "joint"`; the only additional request is to store covariance matrices for uniform bands and aggregate summaries. We do not use `covs.eff` in this illustration, so that the reported results correspond to the unadjusted fuzzy BD estimators.

Figure 3 gives the main visual comparison. The rows correspond to the reduced form, first stage, and fuzzy effect; the columns correspond to the location-based and distance-based methods. Both methods find a strong first stage along most of the boundary: eligibility sharply increases the probability of Qualified Opportunity Zone designation. In contrast, the reduced-form BATEC for early workplace-job growth is small relative to its uncertainty, and the fuzzy BATEC is correspondingly imprecise. The distance-based first stage is slightly less uniform near some boundary locations, which is consistent with the sensitivity to boundary geometry discussed in Section 4.

Tables 1–3 report selected numerical output underlying Figure 3. To conserve space, each table reports every fifth boundary point, the kink, the two endpoints, and the WBATE and LBATE rows generated by `summary()`. The full tables are generated by the replication script and written to the `inputs` folder.


The boundary-wide summaries reinforce the same interpretation. For the first stage, all 31 location-based pointwise estimates are statistically significant and all 31 uniform bands exclude zero; the distance-based first stage gives 30 pointwise significant estimates and 29 uniform bands excluding zero. Neither method finds statistically significant WBATE evidence for the reduced-form job-growth outcome or for the fuzzy effect. The location-based fuzzy WBATE is -0.066 with a p -value of 0.263, while the distance-based fuzzy WBATE is -0.099 with a p -value of 0.124.

The core empirical calls are short. The full script `CTY_2026_rd2d_empapp.R` reads the clean dataset `CTY_2026_rd2d_data.csv`, constructs the distance matrix, runs both analyses, and writes the figures and tables used in this section to the `inputs` folder.


```
fit_loc <- rd2d(Y = Y, X = X, assignment = T, b = b, fuzzy = W,
              fitmethod = "joint",
              params.cov = c("main", "itt", "fs"))
fit_dist <- rd2d.distance(Y = Y, distance = D, b = b, fuzzy = W,
                        fitmethod = "joint",
                        params.cov = c("main", "itt", "fs"))
summary(fit_loc, output = "main", cbands = "main")
summary(fit_dist, output = "main", cbands = "main")
```

6 Conclusion

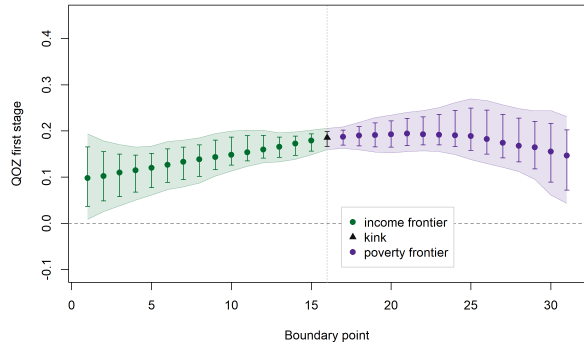
This article introduced `rd2d`, software for causal inference in boundary discontinuity designs. The package implements location-based and distance-based local-polynomial methods, robust bias-corrected pointwise inference, uniform confidence bands, BATEC, WBATE, LBATE, fuzzy BATEC, fuzzy WBATE, and fuzzy LBATE. It also includes joint cluster-robust fitting, the earlier separate fit-




(a) Location-Based ITT




(b) Distance-Based ITT



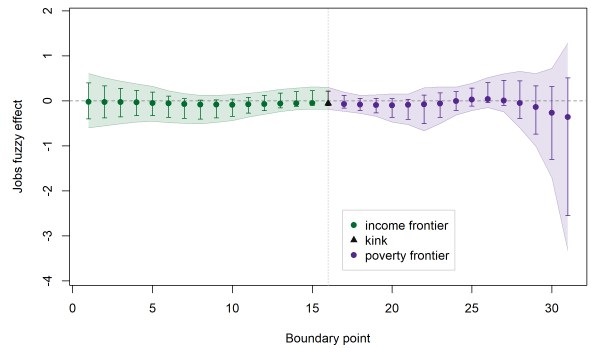
(c) Location-Based First Stage



(d) Distance-Based First Stage



(e) Location-Based Fuzzy BATEC



(f) Distance-Based Fuzzy BATEC

Figure 3: Opportunity Zones Empirical Results. Rows report the reduced-form BATEC, first-stage BATEC, and fuzzy BATEC; columns compare location-based and distance-based methods. Bands are uniform confidence bands over the 31-point boundary grid shown in Figure 2.

Table 1: Opportunity Zones Reduced-Form BATEC

	Location-Based			Distance-Based		
	Estimate	<i>p</i> -value	95% CI	Estimate	<i>p</i> -value	95% CI
b₁	-0.019	0.471	(-0.070, 0.032)	0.000	0.996	(-0.041, 0.041)
b₅	0.011	0.475	(-0.019, 0.042)	-0.008	0.611	(-0.037, 0.022)
b₁₀	-0.024	0.050	(-0.048, -0.000)	-0.022	0.115	(-0.050, 0.005)
b₁₅	0.006	0.696	(-0.026, 0.038)	0.013	0.442	(-0.019, 0.045)
b₁₆	0.010	0.637	(-0.031, 0.050)	0.011	0.510	(-0.022, 0.045)
b₂₀	-0.014	0.382	(-0.044, 0.017)	-0.029	0.125	(-0.066, 0.008)
b₂₅	-0.021	0.294	(-0.061, 0.018)	0.023	0.406	(-0.031, 0.077)
b₃₀	-0.021	0.492	(-0.081, 0.039)	-0.059	0.167	(-0.143, 0.025)
b₃₁	-0.020	0.524	(-0.083, 0.042)	-0.077	0.040	(-0.150, -0.004)
WBATE	-0.011	0.227	(-0.029, 0.007)	-0.011	0.185	(-0.027, 0.005)
LBATE	0.012		(-0.035, 0.074)	0.047		(-0.037, 0.135)

Notes. The outcome is the 2017–2019 change in log workplace jobs. The WBATE estimator uses equal weights over the boundary grid points.

Table 2: Opportunity Zones First-Stage BATEC

	Location-Based			Distance-Based		
	Estimate	<i>p</i> -value	95% CI	Estimate	<i>p</i> -value	95% CI
b₁	0.101	0.002	(0.037, 0.165)	0.103	0.000	(0.057, 0.149)
b₅	0.114	0.000	(0.077, 0.151)	0.113	0.000	(0.080, 0.147)
b₁₀	0.156	0.000	(0.126, 0.187)	0.147	0.000	(0.114, 0.181)
b₁₅	0.175	0.000	(0.156, 0.193)	0.196	0.000	(0.162, 0.231)
b₁₆	0.183	0.000	(0.167, 0.199)	0.204	0.000	(0.172, 0.237)
b₂₀	0.194	0.000	(0.165, 0.222)	0.182	0.000	(0.133, 0.231)
b₂₅	0.204	0.000	(0.158, 0.250)	0.269	0.000	(0.184, 0.355)
b₃₀	0.153	0.000	(0.089, 0.216)	0.121	0.016	(0.022, 0.219)
b₃₁	0.137	0.000	(0.072, 0.203)	0.075	0.085	(-0.010, 0.161)
WBATE	0.163	0.000	(0.143, 0.182)	0.171	0.000	(0.152, 0.191)
LBATE	0.205		(0.162, 0.271)	0.269		(0.174, 0.397)

Notes. The outcome is Qualified Opportunity Zone designation. The WBATE estimator uses equal weights over the boundary grid points.

ting convention, automatic bandwidth selection, and covariate efficiency adjustment. The current software ecosystem includes R, Python, and Stata implementations with aligned syntax and output conventions. The Opportunity Zones application illustrates how the package can separate the reduced-form BATEC, the first-stage BATEC, and the fuzzy BATEC along a multi-dimensional assignment boundary. Replication codes and related information are available at <https://rdpackages.github.io/>.

7 Acknowledgments

We thank Alberto Abadie, Kosuke Imai, Xinwei Ma and Filippo Palomba for insightful discussions. Cattaneo and Titiunik gratefully acknowledge financial support from the National Science Foundation (SES-2019432, DMS-2210561, SES-2241575 and SES-2342226). Cattaneo gratefully acknowledges financial support from the National Institute for Food and Agriculture (NIFA) through grant 2024-67023-42704, the Data-Driven Social Science initiative at Princeton University, and a 2026 Guggenheim Fellowship.

A Multi-Platform Syntax

The three implementations share the same conceptual inputs. Users supply an outcome, a two-dimensional score or distance-to-boundary score, an assignment variable, evaluation points on the boundary, and optional fuzzy treatment receipt. In the code snippets below, Z denotes an optional matrix or data frame of predetermined covariates, and `cluster_id` denotes an optional cluster identifier. The most visible differences across platforms are naming conventions and return objects.

A.1 R

In R, location-based methods use `rd2d()` and `rdbw2d()`, while distance-based methods use `rd2d.distance()` and `rdbw2d.distance()`.

```
install.packages("rd2d")
library(rd2d)

fit <- rd2d(Y = Y, X = X, assignment = T, b = b,
           fuzzy = W, params.cov = c("main", "itt", "fs"))
summary(fit, output = "main", cbands = "main")
summary(fit, output = "itt", cbands = "itt")
summary(fit, output = "fs", cbands = "fs")

fit_cov <- rd2d(Y = Y, X = X, assignment = T, b = b,
              fuzzy = W, cluster = cluster_id,
              covs.eff = Z, fitmethod = "joint",
              params.cov = c("main", "itt", "fs"))

fit_dist <- rd2d.distance(Y = Y, distance = D, b = b, fuzzy = W,
```

```

cluster = cluster_id,
covs.eff = Z, fitmethod = "joint",
params.cov = c("main", "itt", "fs"))

```

The R interface uses dotted option names such as `params.cov` and `covs.eff`, and returns S3 objects with `print()`, `summary()`, `plot()`, `coef()`, `vcov()`, and `confint()` methods. Uniform bands and aggregate summaries require the relevant covariance matrices to be stored through `params.cov` or related options. The main automatic bandwidth selectors are `mserd`, `cerrd`, `inserd`, `icerrd`, `msetwo`, `certwo`, `imsetwo`, and `icertwo`; user-supplied bandwidths are recorded as "user provided". The default fitting convention is `fitmethod = "joint"`; `fitmethod = "separate"` requests the earlier side-specific convention. Distance-based kink handling uses `kink.unknown` and `kink.position`.

A.2 Python

The Python implementation mirrors the R API but follows Python naming conventions.

```

from rd2d import rd2d, rdbw2d, rd2d_dist, rdbw2d_dist

fit = rd2d(Y, X, T, b, fuzzy=W,
           params_cov=["main", "itt", "fs"])
fit.summary(output="main", cbands="main")
fit.summary(output="itt", cbands="itt")
fit.summary(output="fs", cbands="fs")

fit_cov = rd2d(Y, X, T, b, fuzzy=W,
               cluster=cluster_id,
               covs_eff=Z, fitmethod="joint",
               params_cov=["main", "itt", "fs"])

fit_dist = rd2d_dist(Y, D, b=b, fuzzy=W,
                     cluster=cluster_id,
                     covs_eff=Z, fitmethod="joint",
                     params_cov=["main", "itt", "fs"])

```

The main differences are snake-case option names, such as `kernel_type`, `params_cov`, `covs_eff`, `kink_unknown`, and `kink_position`, and Python return objects that expose tabular outputs through familiar array or data-frame structures. The option `fitmethod="joint"` is the default, and `fitmethod="separate"` is available for comparisons with the earlier convention. Distance-based aliases use `rd2d_dist` and `rdbw2d_dist`; longer aliases are also available in the package.

A.3 Stata

The Stata implementation follows Stata command syntax. Location-based methods use `rd2d` and `rdbw2d`. Distance-based methods use `rd2d_dist` and `rdbw2d_dist`; aliases ending in `_distance` are also available.

```

net install rd2d, from(https://raw.githubusercontent.com/rdpackages/rd2d/main/
  stata) replace

rd2d y x1 x2 assignment, b(0 0 0 1) fuzzy(w)
rdbw2d y x1 x2 assignment, b(0 0 0 1) fuzzy(w) bwparam(main)

rd2d y x1 x2 assignment, b(0 0 0 1) fuzzy(w) ///
  cluster(cluster_id) covseff(z1 z2) fitmethod(joint)

rd2d_dist y dist1 dist2, b(0 0 0 1) fuzzy(w) ///
  cluster(cluster_id) covseff(z1 z2) fitmethod(joint)

```

Stata stores results in `e()` matrices, including covariance matrices for the main, reduced-form, and first-stage tables when available. In fuzzy BD designs, the main table reports the fuzzy BATEC, while the reduced-form and first-stage BATECs are stored separately in `e(itt)` and `e(fs)`. The joint fitting convention is requested with `fitmethod(joint)`, the earlier side-specific convention with `fitmethod(separate)`, and covariate efficiency adjustment with `covseff()`. Distance-based kink options are written as `kinkunknown()` and `kinkposition()`.

To summarize the main syntax across platforms:

- In R, the main commands are `rd2d()`, `rdbw2d()`, and `rd2d.distance()`. The interface returns S3 objects and uses dotted option names such as `params.cov` and `covs.eff`.
- In Python, the main commands are `rd2d()`, `rdbw2d()`, and `rd2d_dist()`. The interface follows Python conventions with snake-case option names such as `params_cov` and `covs_eff`.
- In Stata, the main commands are `rd2d`, `rdbw2d`, and `rd2d_dist`. Results are stored in `e()` returns, covariate adjustment is requested with `covseff()`, and distance-based aliases with `distance` are available.

Table 3: Opportunity Zones Fuzzy BATEC

	Location-Based			Distance-Based		
	Estimate	p -value	95% CI	Estimate	p -value	95% CI
b₁	-0.186	0.475	(-0.696, 0.324)	0.001	0.996	(-0.399, 0.401)
b₅	0.098	0.480	(-0.173, 0.369)	-0.067	0.612	(-0.326, 0.192)
b₁₀	-0.153	0.055	(-0.309, 0.003)	-0.152	0.121	(-0.344, 0.040)
b₁₅	0.036	0.696	(-0.146, 0.219)	0.064	0.443	(-0.100, 0.228)
b₁₆	0.053	0.637	(-0.168, 0.275)	0.055	0.511	(-0.109, 0.218)
b₂₀	-0.070	0.383	(-0.227, 0.087)	-0.160	0.131	(-0.367, 0.048)
b₂₅	-0.104	0.299	(-0.299, 0.092)	0.085	0.407	(-0.116, 0.285)
b₃₀	-0.137	0.498	(-0.533, 0.259)	-0.492	0.233	(-1.300, 0.317)
b₃₁	-0.149	0.530	(-0.612, 0.315)	-1.020	0.192	(-2.552, 0.511)
WBATE	-0.066	0.263	(-0.180, 0.049)	-0.099	0.124	(-0.226, 0.027)
LBATE	0.110		(-0.217, 0.609)	0.183		(-0.150, 1.282)

Notes. The fuzzy BATEC divides the reduced-form BATEC by the first-stage BATEC. The WBATE estimator uses equal weights over the boundary grid points.

References

- 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.
- 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.
- Matias D. Cattaneo and Rocio Titiunik. Regression discontinuity designs. *Annual Review of Economics*, 14:821–851, 2022.
- Matias D. Cattaneo, Luke Keele, Rocio Titiunik, and Gonzalo Vazquez-Bare. Interpreting regression discontinuity designs with multiple cutoffs. *Journal of Politics*, 78(4):1229–1248, 2016.
- Matias D. Cattaneo, Nicolás Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press, 2020.
- Matias D. Cattaneo, Luke Keele, Rocio Titiunik, and Gonzalo Vazquez-Bare. Extrapolating treatment effects in multi-cutoff regression discontinuity designs. *Journal of the American Statistical Association*, 116(536):1941–1952, 2021.
- 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, 2024a.
- Matias D. Cattaneo, Nicolás Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Extensions*. Cambridge University Press, 2024b.
- Matias D. Cattaneo, Rocio Titiunik, and Ruiqi (Rae) Yu. Estimation and inference in boundary discontinuity designs: Location-based methods. *arXiv:2505.05670*, 2026a.

- Matias D. Cattaneo, Rocio Titiunik, and Ruiqi (Rae) Yu. Estimation and inference in boundary discontinuity designs: Distance-based methods. *Journal of Econometrics*, forthcoming, 2026b.
- Jin-young Choi and Myoung-jae Lee. Complier and monotonicity for fuzzy multi-score regression discontinuity with partial effects. *Economics Letters*, 228:111169, 2023.
- Kevin Corinth and Naomi Feldman. Are opportunity zones an effective place-based policy? *Journal of Economic Perspectives*, 38(3):113–136, 2024. doi: 10.1257/jep.38.3.113.
- 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.
- 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.
- 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.
- Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC, 2020.
- 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.
- Weiwei Jiang and Rong JB Zhu. Extrapolating treatment effects in multiple-score regression discontinuity designs. *Journal of Educational and Behavioral Statistics*, page 10769986261426399, 2026.
- 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.
- Philipp Alexander Schwarz, Oliver Schacht, Sven Klaassen, Johannes Oberpriller, and Martin Spindler. Effect identification and unit categorization in the multi-score regression discontinuity design with application to led manufacturing. *arXiv preprint arXiv:2508.15692*, 2025.
- 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.