

Heterogeneous Treatment Effects in Regression Discontinuity Designs*

Sebastian Calonico[†] Matias D. Cattaneo[‡] Max H. Farrell[§] Filippo Palomba[¶]
Rocío Titiunik^{||}

March 19, 2025

Abstract

Empirical studies using Regression Discontinuity (RD) designs often explore heterogeneous treatment effects based on pretreatment covariates. However, the lack of formal statistical methods has led to the widespread use of ad hoc approaches in applications. Motivated by common empirical practice, we develop a unified, theoretically grounded framework for RD heterogeneity analysis. We show that a fully interacted local linear (in functional parameters) model effectively captures heterogeneity while still being tractable and interpretable in applications. The model structure holds without loss of generality for discrete covariates, while for continuous covariates our proposed (local functional linear-in-parameters) model can be potentially restrictive, but it nonetheless naturally matches standard empirical practice and offers a causal interpretation for RD applications. We establish principled bandwidth selection and robust bias-corrected inference methods to analyze heterogeneous treatment effects and test group differences. We provide companion software to facilitate implementation of our results. An empirical application illustrates the practical relevance of our methods.

Keywords: causal inference, regression discontinuity designs, covariate adjustment, heterogeneous treatment effects, local polynomial regression, robust bias correction.

*We thank participants at various seminars and conferences for comments. Cattaneo and Titiunik gratefully acknowledge financial support from the National Science Foundation (SES-2019432 and SES-2241575).

[†]Graduate School of Management, UC Davis.

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

[§]Department of Economics, UC Santa Barbara.

[¶]Department of Economics, Princeton University.

^{||}Department of Politics, Princeton University.

1 Introduction

Analyzing heterogeneous treatment effects based on pretreatment covariates is crucial in modern causal inference. While the average causal effect provides a broad measure of a treatment’s effectiveness and informs general policy decisions, uncovering covariate heterogeneity patterns is essential for evaluating substantive hypotheses and specific policy interventions. Covariate-based heterogeneous treatment effects can help understand fairness and differential impacts across subpopulations, develop targeted interventions, and optimize social policies.

Over the past decade, the regression discontinuity (RD) design has become a prominent tool for causal inference, leading to extensive methodological advancements for identification, estimation, and inference for average treatment effects in various contexts. However, despite its widespread adoption, there is a notable lack of rigorous methods for studying heterogeneity based on pretreatment covariates, resulting in many ad-hoc empirical approaches. We address this gap by developing practical, theoretically-grounded tools for heterogeneity analysis in RD designs, putting forward a unified methodology for empirical researchers studying heterogeneous treatment effects based on pretreatment covariates.

Extending RD estimation and inference to encompass heterogeneous treatment effects presents a high-dimensional, nonparametric challenge: while theoretically feasible, it is often impractical without model restrictions. Consequently, empirical researchers typically use semiparametric models, and thus explore heterogeneity in limited dimensions. Following empirical practice, we study the most common approach for RD treatment effect heterogeneity analysis: local least squares regression with linear-in-parameters interactions with pretreatment covariates. This model balances practical applicability with the flexible (nonparametric) nature of the RD design.

Our first contribution is to clarify the conditions under which a local regression with linear interactions accurately recovers meaningful causal heterogeneous RD treatment effects. We show that heterogeneous effects are identifiable when potential outcomes follow a local linear-in-parameters functional coefficient model in the pretreatment variables. This structure is both flexible and interpretable, and is without loss of generality for binary (and orthogonal) covariates. Without this structure, the linear interactions model leads to a less interpretable and useful probability limit, as the nonparametric nature of the estimation method hampers its usual causal interpretation

emerging from the best linear prediction viewpoint. The supplemental appendix computes the parameters emerging as the probability limit of a local linear regression with linear interactions under general misspecification, and demonstrates by example that a large smoothing bias emerges without our proposed identifying assumption. Therefore, from a practical perspective, we show that (i) heterogeneity analysis based on local polynomial regression methods with interactions is fully warranted without additional assumptions when the pretreatment covariates are binary orthogonal variables (e.g., indicator variables for each mutually exclusive level of a categorical variable), but (ii) heterogeneity analysis requires a semiparametric identifying assumption in the general case of discrete or continuous covariates for causal inference.

We then present formal methods for estimation and inference of heterogeneous RD treatment effects, developing optimal bandwidth selection, point estimation, and robust bias-corrected inference. Our key large-sample econometric results include consistency, a mean squared error expansion, a central limit theorem, and robust standard error estimators to both heteroskedasticity and clustering. We show that the optimal bandwidth depends on the specific inference target, implying different optimal bandwidths for average effects, subgroups, and group difference tests. Our goal is to standardize empirical practice, where bandwidth selection is often ad-hoc, particularly for heterogeneity analysis. From that perspective, we identify conditions under which a single bandwidth can be validly used for the entire analysis, thereby enhancing the practical appeal of our proposed methods.

Despite its major role in empirical work, covariate adjustment for heterogeneity analysis in RD designs has been surprisingly less explored in the literature than needed. The only antecedents we are aware of are [Hsu and Shen \(2019\)](#), who use moment condition methods to test for the presence of covariate-based heterogeneity, and [Reguly \(2021\)](#) and [Alcantara et al. \(2025\)](#), who leverage machine learning methods to recover nonparametric covariate-based heterogeneity. None of these contributions study the properties of the most common approach in empirical work: semi-linear local polynomial regression with pretreatment covariate interactions. Our paper addresses this gap in the literature by setting up a standard, most tractable local least squares framework with covariate-interactions, and then providing novel identification, estimation, and robust bias-corrected inference methods to support past, current, and future heterogeneity analysis in RD applications.

By introducing a comprehensive framework for identification, estimation, and inference in RD designs, we aim to unify the study of heterogeneous treatment effects in empirical research. The absence of such a framework has led to a variety of empirical approaches—some valid, others not—creating inconsistencies and confusion in the literature. To illustrate this issue, Section SA3 in the supplemental appendix compiles a sample of papers from AEA journals over the past decade that analyze covariate heterogeneity in RD settings, provided they offer sufficient methodological details. While not exhaustive or necessarily representative, this review reveals three common patterns. First, most studies examine heterogeneity for binary covariates, but often by discretizing categorical or continuous variables. Second, a single bandwidth, and often selected in an ad hoc manner, is typically used for both average and heterogeneous treatment effects analysis. Third, and most importantly, inference is rarely conducted properly, regardless of bandwidth choice or covariate granularity. Our proposed methods maintain the simplicity and interpretability of standard empirical practice while introducing objective procedures that bring validity and rigor.

We illustrate our methodological results by revisiting the analysis in [Akhtari et al. \(2022\)](#), who study how political turnover in mayoral elections in Brazil affects public service provision by local governments. The analysis exploits close elections, a canonical RD setting, using a rich data set of Brazilian mayoral elections. [Klašnja and Titiunik \(2017\)](#) originally analyzed this data and RD setting, documenting a negative incumbency advantage; see further references therein. We begin by replicating the results in [Akhtari et al. \(2022\)](#), who considered both standard RD treatment effect and a limited measure of covariate-heterogeneity based on municipality-level monthly household income. We recover their point estimates of heterogeneity for income dichotomized as above or below median, and also report valid uncertainty quantification; our results are qualitatively in line with theirs. Then, we further leverage our framework and methods to study income discretized by quartile and by decile, revealing further nuance in the heterogeneity of treatment effects by income level. To conclude, we also show that using income directly (as a continuous variable) provides an accurate summary of the heterogeneity in this application, which is (i) highly compatible with the discretized estimates, (ii) directly interpretable, and (iii) more efficient for estimation and inference. Specifically, we consider two local linear-in-parameters semiparametric heterogeneity models: one taking income linearly, while the other taking income quadratically. This empirical exercise demonstrates how our methods can be used to enhance RD heterogeneity analysis, while also high-

lighting the importance of having a clear understanding of how best to capture heterogeneity in the econometric model.

Our paper contributes to the methodological literature on RD designs (see [Cattaneo and Titiunik, 2022](#), for a review). In particular, it sheds new light on the role of pretreatment covariates for identification, estimation, and inference of heterogeneous RD treatment effects. [Cattaneo et al. \(2023\)](#) overviews covariate adjustment in RD designs and discusses its three most useful roles: (i) increasing estimation efficiency, (ii) changing the parameter of interest, and (iii) heterogeneity analysis. For efficiency gains, [Calonico et al. \(2019\)](#), [Arai et al. \(2025\)](#), [Ma and Yu \(2024\)](#), and references therein, give results leveraging pretreatment covariates to improve precision in the estimation of the average RD treatment effect. There is also a rich literature on the inclusion of covariates for identification of other RD treatment effects: for example, see [Angrist and Rokkanen \(2015\)](#) and [Cattaneo et al. \(2021\)](#) for extrapolation, [Peng and Ning \(2021\)](#) and [Caetano et al. \(2024\)](#) for weighted average treatment effects, [Grembi et al. \(2016\)](#) for difference-in-discontinuities, and [Frölich and Huber \(2019\)](#) for nonparametric methods, among many others. As mentioned above, the role of covariates for heterogeneity analysis has received the least attention in the literature, and our paper is the first to study it in the context of local semi-linear regression with interactions, the most common approach used in practice. Finally, it is worth mentioning that covariates are sometimes used in an attempt to “fix” a “broken” RD design. This is not valid without strong assumptions: it is analogous to employ covariate adjustment to “fix” a “broken” randomized assignment in experimental settings. Our paper does not speak to this practice, and we recommend against it. See [Cattaneo et al. \(2023\)](#) for further discussion.

The rest of the paper proceeds as follows. Section 2 gives details on the RD setup and estimation. Our main technical results are summarized in Section 3. We then present an empirical illustration of our proposed methods in Section 4. Section 5 concludes. The supplemental appendix contains more general, technical results, all proofs, implementation details, and additional numerical results. The companion software package `rdhte` offers general-purpose implementation of our methodological results, and is available at <https://rdpackages.github.io/rdhte/>. See [Calonico et al. \(2025\)](#) for details.

2 Setup

We consider the standard sharp regression discontinuity design setting with a continuous running variable (see [Cattaneo et al., 2020, 2024](#), for a practical introduction). The defining characteristic is that a binary treatment $T_i \in \{0, 1\}$ is assigned to each unit based on whether the continuous running variable X_i is above or below a known cutoff, denoted by c . Thus, $T_i = \mathbb{1}\{X_i \geq c\}$. The observed outcome is $Y_i = T_i Y_i(1) + (1 - T_i) Y_i(0)$, where $Y_i(t)$ is the potential outcome for unit i under treatment $T_i = t$. In addition to Y_i , T_i , and X_i , we observe a vector $\mathbf{W}_i \in \mathbb{R}^d$ of pretreatment covariates to be used for heterogeneity analysis. Distinct from \mathbf{W}_i , we may have an additional set of covariates \mathbf{Z}_i used for efficiency purposes only; see below. We assume a random sample $(Y_i, T_i, X_i, \mathbf{W}_i, \mathbf{Z}_i)$, for $i = 1, \dots, n$, is observed.

The usual parameter of interest in sharp RD is the average effect of T_i on Y_i at the cutoff. To simplify notation, we set the cutoff to be $c = 0$. Then, the average RD treatment effect is

$$\tau = \mathbb{E}[Y_i(1) - Y_i(0) | X_i = 0]. \quad (2.1)$$

Identification of τ follows from standard continuity assumptions ([Hahn et al., 2001](#)), while estimation and inference is typically done with local linear regression and robust bias correction ([Calonico et al., 2014](#)). More precisely, the point estimate $\hat{\tau}$ of the causal parameter τ is the coefficient on T_i in a least squares regression of Y_i on a constant, T_i , X_i , and $T_i \cdot X_i$, using only the observations such that $|X_i - 0| \leq h$, for a bandwidth h , and weighted by $K((X_i - 0)/h)$, for a kernel function $K(\cdot)$. This fit, and those below, are defined formally in the Supplemental Appendix (Section SA1). We denote the resulting estimated equation by

$$\hat{Y}_i = \hat{\mu} + \hat{\tau} T_i + \hat{\omega}_1 X_i + \hat{\omega}_2 T_i X_i. \quad (2.2)$$

The coefficient $\hat{\tau}$ is identical to the difference in intercepts from fitting separate (local, weighted) least squares regressions on either side of the cutoff. Conceptually, it is important to remember that (2.2) is formally a nonparametric local polynomial regression with the polynomial degree set to one. Finally, although point estimation is done with (local, weighted) linear models, we stress that inference is then conducted using robust bias correction, which we adapt for heterogeneity

below.

Calonico et al. (2019) studied the addition of the pretreatment covariates to this model, but aiming exclusively for more efficient estimation of the RD average treatment effect τ , not for heterogeneity analysis. They showed that efficiency gains could be achieved by including covariates linearly, but *not* interacting them with the treatment or running variable. Specifically, for pretreatment covariates \mathbf{Z}_i (used for efficiency only and hence necessarily distinct from \mathbf{W}_i), Calonico et al. (2019) studied $\tilde{\tau}$ from the (local, weighted) least squares fit given by

$$\tilde{Y}_i = \tilde{\mu} + \tilde{\tau}T_i + \tilde{\omega}_1X_i + \tilde{\omega}_2T_iX_i + \tilde{\gamma}'\mathbf{Z}_i. \quad (2.3)$$

Unlike $\hat{\tau}$, the estimated coefficient $\tilde{\tau}$ cannot be obtained from two separate fits because of the common coefficient on \mathbf{Z}_i . Despite being linear in both the running variable and covariates, it is important to note that (2.3) is nonparametric in X_i (local polynomial of degree one, localized using the bandwidth h) but it is a linear projection in \mathbf{Z}_i , and thus the interpretation of the estimated coefficient $\tilde{\gamma}$ is only as a best linear projection (unless the population mean function happens to be linear). That is, the nonparametric smoothing is only on the running variable. This respects the local nature of the RD design while being simple and tractable in applications. Calonico et al. (2019) proved that, nonetheless, $\tilde{\tau}$ recovers the average effect τ , characterized conditions for efficiency gains, and provided optimal bandwidth selection and valid inference using bias robust bias correction. Ma and Yu (2024) strengthened the efficiency argument for this regression fit. The exclusion of interactions with the covariates in (2.3) is crucial for correct estimation of the average effect τ of (2.1), but by construction does not allow for learning heterogeneity. Because our present goal is explicitly to learn about heterogeneous treatment effects, we need to allow some form of interaction between the pretreatment covariates and treatment indicator.

Ideally, one might try to recover the (local to $X_i = 0$) conditional average treatment effect (CATE) function fully flexibly in \mathbf{W}_i , which is defined by

$$\kappa(\mathbf{w}) = \mathbb{E}[Y_i(1) - Y_i(0)|X_i = 0, \mathbf{W}_i = \mathbf{w}]. \quad (2.4)$$

Absent further assumptions, $\kappa(\mathbf{w})$ is the difference between two $(1 + d)$ dimensional nonparametric

regressions evaluated at the point $(0, \mathbf{w}) \in \mathbb{R}^{1+d}$. While such a setting is possible to handle theoretically, it is a challenging problem in real-world settings due to the curse of dimensionality and the need to choose many tuning parameters. Following empirical practice, we aim for something more tractable.

Indeed, one possible explanation for the popularity of RD in empirical work is that it combines a credible causal identification design with perhaps the simplest nonparametric problem, namely *one-dimensional* nonparametric regression at a single point, $X_i = 0$. This is reflected in the simplicity of (2.2) and (2.3). Thus, rather than designing a fully general procedure for obtaining $\kappa(\mathbf{w})$, we focus on specifications that share the same practical appeal.

We consider the local weighted least squares regression with full interactions between the treatment indicator and the pretreatment covariates, which may be discrete or continuous. That is, our main task is analyzing the interpretation and properties of the regression fit:

$$\widehat{Y}_i = \widehat{\alpha} + \widehat{\theta}T_i + \widehat{\boldsymbol{\lambda}}'\mathbf{W}_i + \widehat{\boldsymbol{\xi}}'T_i\mathbf{W}_i + \widehat{\omega}_1X_i + \widehat{\omega}_2T_iX_i + \widehat{\boldsymbol{\omega}}_3'X_i\mathbf{W}_i + \widehat{\boldsymbol{\omega}}_4'T_iX_i\mathbf{W}_i, \quad (2.5)$$

and of the corresponding estimate of the CATE function $\kappa(\mathbf{w})$ at the cutoff given by

$$\widehat{\kappa}(\mathbf{w}) = \widehat{\theta} + \widehat{\boldsymbol{\xi}}'\mathbf{w}. \quad (2.6)$$

This regression is the natural next step in the progression of (2.2) and (2.3). However, the interpretation of the coefficients changes: the coefficient on T_i no longer recovers the average effect τ but rather $\kappa(\mathbf{0})$, the CATE for the baseline group $\mathbf{w} = \mathbf{0}$ (where $\mathbf{0}$ is a vector of zeroes), under appropriate identifying assumptions (Assumption 2 below). Moreover, while the coefficients on \mathbf{Z}_i in (2.3) are largely uninterpretable since they are for efficiency purposes only, here the coefficients $\widehat{\boldsymbol{\xi}}$ give the difference from the baseline group, either as the discrete change for binary covariates or, under our functional coefficient assumption below, as the slope of $\kappa(\mathbf{w})$ with respect to continuous \mathbf{w} . Although we omit it from our main discussion, the supplemental appendix provides results for estimation and inference on derivatives of $\kappa(\mathbf{w})$, covering kink RD designs and related settings (Card et al., 2015).

As with Equations (2.2) and (2.3), the fit (2.5) is a local polynomial (of degree 1) approximation

in X_i , with bandwidth h and kernel weights $K(X_i/h)$ around the cutoff $c = 0$. In this case, this is true for the main effects as well as the interaction terms: while $\tilde{\gamma}$ in (2.3) is a linear projection coefficient with no intended approximation power (or causal interpretation), the semiparametric fit (2.5) delivers local polynomial approximations to unknown functions of X_i only, on either side of the cutoff, that form linear interactions with \mathbf{W}_i . This is formalized in Assumption 2 below and in the supplemental appendix in full generality (Section SA1.3). Importantly, such an approximation cannot consistently estimate an arbitrary $\kappa(\mathbf{w})$; its causal interpretation crucially relies on Assumption 2 below. To reiterate, this is true because the regression (2.5) is nonparametric in X_i but not \mathbf{W}_i ; there is no approximation to unknown functions of \mathbf{W}_i .

In applications, the vector \mathbf{W}_i must therefore be specified by the researcher to capture the relevant heterogeneity. Changing the covariates, or their specification, included in the regression changes the target estimand and its interpretation. With discrete variables, the vector \mathbf{W}_i should contain indicators for every relevant category, unless the researcher decides ex-ante that certain categories are not to be studied. For example, if two binary variables $(B_{i,1}, B_{i,2})$ are present, then $\mathbf{W}_i = ((1 - B_{i,1})B_{i,2}, B_{i,1}(1 - B_{i,2}), B_{i,1}B_{i,2})'$, with $B_{i,1} = B_{i,2} = 0$ being the baseline category. A multi-level discrete variable would require that \mathbf{W}_i contain an indicator for each level beyond the baseline, which forms a collection of orthogonal binary variables. For example, if \mathcal{W} is the finite set of values the covariate of interest W_i takes (e.g., years of education), then $\mathbf{W}_i = (\mathbb{1}(W_i = w), w \in \mathcal{W})'$. In all of these cases, the structure of Assumption 2 holds without loss of generality. Continuous variables can be included directly into \mathbf{W}_i , in which case $\hat{\xi}$ have the standard interpretation of a slope coefficient in a linear model. Further, transformations of continuous variables (such as polynomials) can be added for additional flexibility. Alternatively, the researcher might estimate the heterogeneous treatment effects one variable at a time, which may not capture the full heterogeneity pattern but may provide a simple and intuitive summary of different dimensions of heterogeneity. Section 4 illustrates different approaches and interpretations empirically. (See Calonico et al. (2025) for further examples and discussion.)

As with (2.2), and in contrast with (2.3), the coefficients in (2.5) can be obtained by first fitting separate regressions on either side of the cutoff, and then taking appropriate differences. It is also possible to run separate models for different subgroups of discrete variables via (2.2) to recover the estimates in (2.5), but the joint estimation procedure is simple and intuitive and, thus, ap-

peeling for practice. Joint estimation and inference also simplify concerns of bandwidth selection, which otherwise may change depending on the group of interest or the hypothesis being tested. Furthermore, joint estimation is also natural if additional controls are added: following [Calonico et al. \(2019\)](#), it is possible to include some covariates for efficiency gains, while other covariates for heterogeneity analysis, thereby obtaining more efficient estimation and inference for $\kappa(\mathbf{w})$. This mixed covariate adjustment approach augments (2.5) with additional covariates \mathbf{Z}_i that are not interacted with T_i nor X_i , but interacted with \mathbf{W}_i . That is, one would fit

$$\check{Y}_i = \check{\alpha} + \check{\theta}T_i + \check{\lambda}'\mathbf{W}_i + \check{\xi}'T_i\mathbf{W}_i + \check{\omega}_1X_i + \check{\omega}_2T_iX_i + \check{\omega}'_3X_i\mathbf{W}_i + \check{\omega}'_4T_iX_i\mathbf{W}_i + \check{\gamma}'_1\mathbf{Z}_i + \check{\gamma}'_2(\mathbf{Z}_i \otimes \mathbf{W}_i). \quad (2.7)$$

When \mathbf{W}_i are binary, this approach is theoretically justified by the results presented in the upcoming sections coupled with the theory of [Calonico et al. \(2019\)](#). For the general case, the results in the supplemental appendix can be extended to accommodate the fit in (2.7), though the restriction on the coefficients of \mathbf{Z}_i would be slightly different than for category-wise estimation, and additional regularity conditions would be needed.

Unlike Equation (2.3), our proposed fit (2.5) does not yield an estimator of the average treatment effect (2.1). This is implicit in Theorem 1 below, and was already shown by [Calonico et al. \(2019\)](#). There are two potential remedies to this, neither of which can be recommended in practice. First, the covariates can be demeaned on each side of the cutoff. This would follow the logic for linear regression adjustments in randomized experiments, but in this case the demeaning would amount to Nadaraya–Watson estimation at a boundary point with \mathbf{W}_i as the outcome, thereby hampering the entire RD estimation due to the presence of smoothing bias. Second, one could average the fitted CATEs, but this again essentially amounts to Nadaraya–Watson estimation at a boundary point, now using $\hat{\kappa}(\mathbf{w})$ as the outcome. In both cases, the convergence rate will be slow because the misspecification bias is of order h , and is thus impractical for inference (e.g., severe undersmoothing would be needed). Therefore, in applications, researchers should only use our proposed methods for heterogeneity analysis and rely on the established methods underlying (2.2) or (2.3) to estimate the average RD treatment effect.

Finally, a related motivation for the specification (2.7) is to control for panel or group structures. It is common in some RD applications to include fixed effects for individuals, groups, or time

periods. Identification, estimation, and inference then follow from the same arguments presented in the upcoming sections. In this way, our proposed RD methodology, although developed for the cross-sectional setting, can be easily applied to panel data, provided weighted linear least squares methods with additive or interactive fixed effect are used.

3 Main Results

We turn to the theoretical properties of the regression fit in (2.5), and particularly the CATE estimator (2.6). We provide a complete set of practical tools for identification, estimation, and inference. The following standard regularity conditions are imposed; recall that $c = 0$ without loss of generality, and let $x_L < 0 < x_U$. All omitted details are given in the supplemental appendix (see Section SA1.3).

Assumption 1 (Sharp RD Design). *The data $(Y_i, X_i, \mathbf{W}'_i), i = 1, \dots, n$, where $T_i = \mathbb{1}\{X_i \geq 0\}$, $Y_i = T_i Y_i(1) + (1 - T_i) Y_i(0)$, and the \mathbf{W}_i are pretreatment, is a random sample obeying the following regularity conditions:*

- (a) X_i admits a continuous, bounded, and bounded away from zero Lebesgue density;
- (b) For $t \in \{0, 1\}$, $\mathbf{w} \in \mathbb{R}^d$, and all $x \in [x_L, x_U]$: $\mathbb{V}[Y_i(t)|X_i = x]$, $\mathbb{E}[|Y_i(t)|^4|X_i = x, \mathbf{W}_i = \mathbf{w}]$, $\mathbb{E}[\mathbf{W}_i|X_i = x]$, $\mathbb{E}[\mathbf{W}_i \mathbf{W}'_i|X_i = x]$, $\mathbb{E}[|\mathbf{W}_i|^4|X_i = x]$, $\mathbb{E}[\mathbf{W}_i \mathbb{V}[Y_i(t)|X_i, \mathbf{W}_i]|X_i = x]$, and $\mathbb{E}[\mathbf{W}_i \mathbf{W}'_i \mathbb{V}[Y_i(t)|X_i, \mathbf{W}_i]|X_i = x]$ are continuous in x ;
- (c) $\mathbb{E}[K(X_i/h) \mathbf{A}_i \mathbf{A}'_i]$ is invertible, where $\mathbf{A}_i = (1, X_i, X_i^2)' \otimes (1, T_i, \mathbf{W}'_i, T_i \mathbf{W}'_i)'$; and
- (d) $K(\cdot)$ is nonnegative and continuous on its support $[-1, 1]$.

These assumptions are on par with prior RD literature, only adding minimal extra regularity to accommodate the covariates \mathbf{W}_i for heterogeneity analysis via interactions. The most important substantive condition we require is that \mathbf{W}_i are pretreatment, which is formalized in the supplemental appendix (Section SA1.3). Most commonly this holds when \mathbf{W}_i is determined before treatment is assigned. Parts (b) and (c) are regularity conditions needed for valid estimation and (robust bias-corrected) inference. Part (c) is a familiar rank condition for least squares regression, adapted to (2.5) and its bias-corrected counterpart. The regularity conditions are slightly stronger than those

imposed by [Calonico et al. \(2019\)](#) because of the more flexible regression required for recovery of heterogeneous treatment effects, $\kappa(\mathbf{w})$, compared to the average treatment effect τ considered in that paper (compare (2.3) to (2.5)).

To characterize exactly when the CATE function can be recovered based on a local least squares fit with interactions, we introduce the following identifying assumption.

Assumption 2 (Local CATE Structure). *For $\mathbf{w} \in \mathbb{R}^d$, and all $x \in [x_L, x_U]$:*

$$\mathbb{E}[Y_i(1) - Y_i(0)|X_i = x, \mathbf{W}_i = \mathbf{w}] = \theta(x) + \boldsymbol{\xi}(x)' \mathbf{w}$$

and

$$\mathbb{E}[Y_i(0)|X_i = x, \mathbf{W}_i = \mathbf{w}] = \alpha(x) + \boldsymbol{\lambda}(x)' \mathbf{w},$$

where the functions $\theta(x)$, $\boldsymbol{\xi}(x)$, $\alpha(x)$, and $\boldsymbol{\lambda}(x)$ are thrice continuously differentiable.

This assumption requires the potential outcomes to obey a local semilinear-in- \mathbf{W} -parameters model, that is, a local partially linear functional coefficient model, which naturally aligns with the estimation procedure in (2.5). It follows from this assumption that the CATE function (at the cutoff $x = c = 0$) is linear in \mathbf{w} :

$$\kappa(\mathbf{w}) = \mathbb{E}[Y_i(1) - Y_i(0)|X_i = 0, \mathbf{W}_i = \mathbf{w}] = \theta(0) + \boldsymbol{\xi}(0)' \mathbf{w}.$$

From an econometric point of view, Assumption 2 means that estimation of the CATE function remains a one-dimensional nonparametric problem and, further, that the bias can be removed up to the same order as learning the average treatment effect in standard sharp RD settings. This is important for maintaining the empirical tractability of RD analyses. When \mathbf{W}_i contains continuous variables, the linearity (in parameters) places a restriction on the “long” regression functions $\mathbb{E}[Y(t)|X_i = x, \mathbf{W}_i = \mathbf{w}]$, and thus caution must be taken when interpreting the results. Note that the assumption requires linearity in parameters, and so transformations, such as polynomials, can be included in \mathbf{W}_i to weaken the implied identifying restrictions. Alternatively, when the covariates are binary, $\mathbf{W}_i \in \{0, 1\}^d$, Assumption 2 is automatically satisfied. Thus, our proposed joint

estimation and inference methods are automatically valid for causal heterogeneous RD treatment effects based on discrete covariates (via the construction of orthogonal binary variables for each category), thereby offering a formal causal interpretation for the covariate-interacted local least squares estimates in RD designs.

More generally, Assumption 2 gives a causal interpretation to the commonly used estimation approach in (2.5), but Assumption 1 would be sufficient for all our estimation and inference results if researchers were interested in studying a best mean square approximation of the mean functions $\mathbb{E}[Y(t)|X = c, \mathbf{W} = \mathbf{w}]$ —that is, if the parameter of interest were simply the probability limit of the estimated coefficients in (2.5). However, those are often neither useful nor interpretable as heterogeneous causal effects. For more details, see the supplemental appendix.

We summarize our first identification and estimation result in the following theorem. Let $\rightarrow_{\mathbb{P}}$ denote convergence in probability as $n \rightarrow \infty$.

Theorem 1. *Suppose Assumptions 1 and 2 hold, $h \rightarrow 0$, and $nh \rightarrow \infty$. Then, $\widehat{\kappa}(\mathbf{w}) \rightarrow_{\mathbb{P}} \kappa(\mathbf{w})$.*

This result establishes that the regression fit in (2.5) correctly captures the heterogeneous causal effects in RD designs. Without Assumption 2, the probability limit of the coefficients in (2.5) do not appear to have a causal interpretation or useful closed-form expressions, and $\widehat{\kappa}(\mathbf{w})$ does not recover a useful object (see the supplemental appendix). This is in contrast to recovering the average RD effect, τ , using $\tilde{\tau}$ from (2.3), where further restrictions on the underlying data generating process are not needed despite the inclusion of pretreatment covariates (for efficiency purposes).

3.1 Mean Squared Error and Bandwidth Selection

Estimation and inference rely crucially on selecting a bandwidth h that appropriately localizes to the cutoff $c = 0$. This problem is well studied in the methodological literature on RD designs (Arai and Ichimura, 2018; Calonico et al., 2020). The choice of kernel weights is typically less consequential, and typically a linear fit is used for point estimation (see Pei et al., 2021, for more discussion). We develop a bandwidth selection that is optimal in a mean squared error (MSE) sense for estimation of heterogeneous RD treatment effects. That bandwidth yields optimal point estimates of heterogeneous RD treatment effects, making it an appropriate choice for reporting point estimates in empirical applications. However, using such a bandwidth yields invalid inference

in general, and thus, in the next section, we develop inference methods using the standard approach of robust bias correction (Calonico et al., 2014).

To characterize the optimal bandwidth for treatment effect estimation, we establish an MSE expansion of the coefficients $\widehat{\boldsymbol{\varsigma}} = (\widehat{\theta}, \widehat{\boldsymbol{\xi}})'$ in the local regression (2.5) around their population counterparts. The probability limit of this vector is, under Assumption 2, $\boldsymbol{\varsigma} = (\theta(0), \boldsymbol{\xi}(0)')$. Linear combinations of these yield the CATE function for different groups or contrasts thereof. Let \mathbf{s} be the vector that selects or combines the appropriate elements. For example, in the case of a binary W_i (so that $d = 1$), the base category subgroup can be studied as $\widehat{\theta} = \mathbf{s}'\widehat{\boldsymbol{\varsigma}}$ with $\mathbf{s} = (1, 0)'$, whereas testing group differences would correspond to $\widehat{\boldsymbol{\xi}} = \mathbf{s}'\widehat{\boldsymbol{\varsigma}}$ with $\mathbf{s} = (0, 1)'$. In general, setting $\mathbf{s} = (1, \mathbf{w})'$ for a given value \mathbf{w} yields $\mathbf{s}'\widehat{\boldsymbol{\varsigma}} = \widehat{\theta} + \widehat{\boldsymbol{\xi}}'\mathbf{w} = \widehat{\kappa}(\mathbf{w})$, recovering the CATE estimator of (2.6).

Here we give the bandwidth selection result. The general MSE expansion is given in the supplemental appendix (Section SA2.6) and is more cumbersome notationally. As it is customary, after some approximations, it depends on the leading variance and squared bias: $\text{Var}[\mathbf{s}'\widehat{\boldsymbol{\varsigma}}] = \frac{1}{nh}\mathbf{V}_{\mathbf{s}}$ and $\text{Bias}^2[\mathbf{s}'\widehat{\boldsymbol{\varsigma}}] = h^4\mathbf{B}_{\mathbf{s}}^2$, respectively. The rates reflect the specific structure of the estimator in (2.5), that is, a local linear approximation based only on the scalar X_i . The constants $\mathbf{V}_{\mathbf{s}}$ and $\mathbf{B}_{\mathbf{s}}$ also capture the specific features of the estimator and target estimand (hence a function \mathbf{s}). See the supplemental appendix for omitted details. The following theorem summarizes our bandwidth selection result.

Theorem 2 (MSE-optimal Bandwidth). *Suppose Assumptions 1 and 2 hold, and $h \rightarrow 0$ and $nh^3 \rightarrow \infty$. Then, provided that $\mathbf{B}_{\mathbf{s}} \neq 0$, the MSE-optimal bandwidth for estimating $\mathbf{s}'\widehat{\boldsymbol{\varsigma}}$ is*

$$h_{\mathbf{s}}^* = \left(\frac{\mathbf{V}_{\mathbf{s}}}{4\mathbf{B}_{\mathbf{s}}^2} \frac{1}{n} \right)^{-1/5}, \quad (3.1)$$

where $\mathbf{V}_{\mathbf{s}}$ and $\mathbf{B}_{\mathbf{s}}$ are given in the supplemental appendix (Section SA2.6).

While the optimal bandwidth depends on the specific target estimand through the constant terms, its convergence rate is not affected because Equation (2.5) is still a one-dimensional nonparametric (in X_i) estimation problem. Again, consider the leading case of a binary covariate. If one group is substantially smaller than the other, then all else being equal, the variance will be higher for that group. It could also be that the two groups have conditional expectations with different local curvature, resulting in different bias. Thus, ideally, the bandwidth should be different for each

estimand. However, this is not necessarily practical, and furthermore, the specific constant used in bandwidth selection is often less important empirically than ensuring that h is roughly the right magnitude. In the context of heterogeneous effects in RD designs, this perspective supports the common practice of employing the same bandwidth for the estimation of all heterogeneous RD treatment effects of interest. Section 4 provides further discussion in the context of our empirical application.

Our result offers clear practical guidance and justification for several approaches. First, researchers may want to judge if different groups have substantially different bandwidths. If not, it may be fine to conduct all analyses using a common bandwidth. This is easily accommodated by the joint fitting in (2.5). This common bandwidth could be selected as the same bandwidth used for obtaining the average effect in (2.2). One could also use the median of all the relevant bandwidths or, for obtaining robust inference, the smallest of the set. Using different bandwidths for different tasks may yield more accurate inference in some applications.

3.2 Robust Bias-Corrected Inference

We now state our main results for inference on heterogeneous RD treatment effects. The supplemental appendix gives more general theoretical results (see Section SA2.6), but here we restrict to the case of estimating the joint regression (2.5) using an MSE-optimal bandwidth for point estimation and then conducting inference using robust bias correction. Our main inference procedures rely on a Gaussian distributional approximation for the robust bias corrected versions of $\hat{\boldsymbol{\zeta}} = (\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\xi}})'$ of the local regression (2.5), along with valid standard errors. This can then be specialized to different use cases depending on which coefficients are combined or contrasted via the continuous mapping theorem.

Using an MSE-optimal bandwidth means that the limiting distribution of $\hat{\boldsymbol{\zeta}}$ is not centered at $\boldsymbol{\varsigma} = (\boldsymbol{\theta}(0), \boldsymbol{\xi}(0))'$, under Assumption 2, a fact that renders inference invalid. More generally, it is shown in the supplemental appendix that

$$\sqrt{nh}(\hat{\boldsymbol{\zeta}} - \boldsymbol{\varsigma} - h^2\mathbf{B}) \rightsquigarrow \text{Normal}(0, \mathbf{V}),$$

where \rightsquigarrow denotes convergence in distribution as $h \rightarrow 0$ and $nh \rightarrow \infty$, and \mathbf{B} and \mathbf{V} the leading

asymptotic bias and asymptotic variance of the estimator $\hat{\zeta}$. In particular, $V_s = \mathbf{s}'V\mathbf{s}$ and $B_s = \mathbf{s}'B$. It thus follows that a choice of bandwidth proportional to $n^{-1/5}$ (e.g., h_s^* developed above) leads to a MSE-optimal point estimator $\hat{\zeta}$ but invalid inference due to a first-order bias in the distributional approximation.

For inference in RD applications, robust bias correction is a popular way to deal with this problem. In brief, there are two key ingredients to robust bias correction: (i) an estimate of h^2B is subtracted from the point estimate $\hat{\zeta}$, which amounts to bias correction, and (ii) the standard errors are adjusted to account for the additional estimation error of the bias correction, which amounts to variance correction. Robust bias-correction continues to rely on a Gaussian distributional approximation, but inference is based on the corrected statistics

$$\hat{\zeta}_{\text{rbc}} = \hat{\zeta} - h^2\hat{B} \quad \text{and} \quad \hat{V}_{\text{rbc}} = \widehat{\text{var}}(\hat{\zeta}_{\text{rbc}}), \quad (3.2)$$

where \hat{B} is an estimator of the bias of $\hat{\zeta}$, and \hat{V}_{rbc} is an estimator of the variance of $\hat{\zeta}_{\text{rbc}}$.

The robust bias correction method for RD design has been discussed in detail: see [Cattaneo et al. \(2020\)](#), and references therein. See also the supplemental appendix for omitted technical details, including results on heteroskedastic- and cluster-robust standard errors ([Cameron and Miller, 2015](#); [MacKinnon et al., 2023](#)). The method is known to have excellent theoretical and numerical properties ([Calonico et al., 2018, 2022](#)), and has been shown to yield correct inference in extensive validation empirical studies ([Hyttinen et al., 2018](#); [De Magalhães et al., 2025](#)).

We obtain the following asymptotic normality result, based on the estimates and standard errors in (3.2), for any specific estimand defined by \mathbf{s} . The exact formulas for \hat{B} and \hat{V}_{rbc} are given in the supplemental appendix (Section SA2.2).

Theorem 3 (Asymptotic Normality). *Suppose Assumptions 1 and 2 hold, and $h = h_s^*$. Then,*

$$\frac{\mathbf{s}'\hat{\zeta}_{\text{rbc}} - \mathbf{s}'\zeta}{\sqrt{\mathbf{s}'\hat{V}_{\text{rbc}}\mathbf{s}}} \rightsquigarrow \text{Normal}(0, 1).$$

The supplemental appendix gives a more general result, which includes a Gaussian distributional approximation without Assumption 2 (and hence with a different centering). Theorem 3 can be used to carry out valid inference for any of the estimands we have discussed, including the CATE function

itself and comparisons between different values of \mathbf{w} , for discrete or continuous \mathbf{W}_i . For example, an approximate robust bias-corrected 95% confidence interval for $\kappa(\mathbf{w})$ based on $\mathbf{s} = (1, \mathbf{w})'$ is

$$\left[\hat{\kappa}_{\text{rbc}}(\mathbf{w}) - 1.96\sqrt{\mathbf{s}'\hat{\mathbf{V}}_{\text{rbc}}\mathbf{s}}, \hat{\kappa}_{\text{rbc}}(\mathbf{w}) + 1.96\sqrt{\mathbf{s}'\hat{\mathbf{V}}_{\text{rbc}}\mathbf{s}} \right].$$

Robust hypothesis tests and p -values can also be constructed in the standard manner. For example, assuming $\mathbf{W}_i \in \{0, 1\}$, a hypothesis test of no heterogeneous RD treatment effect boils down to testing $H_0 : \xi(0) = 0$, which corresponds to employing $\mathbf{s} = (0, 1)'$. Such inference procedures should be reported along with the point estimates derived from $\hat{\zeta}$ (as opposed to reporting $\hat{\zeta}_{\text{rbc}}$).

4 Application

To illustrate our methodological results we revisit the analysis by [Akhtari et al. \(2022\)](#), which examines the impact of political turnover in Brazilian mayoral elections on local public service provision using an RD design based on close elections. [Klašnja and Titiunik \(2017\)](#) first studied this data and RD design, and documented large negative RD effects of party incumbency based on Brazilian mayoral elections, comparing municipalities where the incumbent party barely loses (resulting in political turnover) to those where the incumbent party barely wins (no turnover): a bare victory results in a large reduction in the probability of victory in the following election, which they attribute to deficient accountability resulting from weak parties and term limits.

Using a similar RD design based on close elections, [Akhtari et al. \(2022\)](#) find that political turnover leads to a sudden expansion of municipal bureaucracy, with new personnel appointed across multiple sectors at both managerial and non-managerial levels. They also investigate educational outcomes, reporting increased replacement rates among school personnel in municipally controlled schools, along with declines in student test scores. We focus on heterogeneous effects on headmaster replacement (i.e., whether a headmaster is new to the school) by municipality income (we rescale income into hundreds to improve the numerical presentation in [Table 1](#) below). See their [Table A.21](#).

The analysis of [Akhtari et al. \(2022\)](#) highlights two key methodological difficulties researchers have faced when studying heterogeneous treatment effects in RD without methodological and the-

oretical guidance in the literature, as discussed above. First, Akhtari et al. (2022) use an optimal data-driven bandwidth selection method (Calonico et al., 2019) to localize at the cutoff. However, both estimation and inference rely solely on local linear fitting. Although estimation using linear regression (specifically, their implementation of (2.3)) is correct, valid inference requires robust bias correction, as standard errors from the linear fit are invalid when paired with the optimal bandwidth. In our replication, we revisit their heterogeneity analysis using robust bias-corrected inference. Our results, presented in Table 1, report confidence intervals instead of standard errors, following best practices in RD analysis.

The second key aspect is that the heterogeneity variable of interest, income, is continuous. Akhtari et al. (2022) binarize income and compare municipalities above and below the median income level. We first replicate this exercise using our methods. The results are shown in the top two rows of Panel (b), Table 1. Our point estimates are nearly identical to theirs (0.373 and 0.122 compared 0.389 and 0.126 in their Table A.21), where the differences are due to slight bandwidth variations. Next, we incorporate valid uncertainty measures using robust bias correction, revealing that the treatment effect is statistically significant only for low (below-median) income municipalities. We also formally test the null hypothesis that the treatment effect is the same across income groups, finding strong evidence against it: our robust inference yields a 95% confidence interval of $[-0.422, -0.078]$ for the difference, which excludes zero thus revealing statistically significant heterogeneity. We also extend this discretized analysis by studying treatment effects by income quartiles and deciles, rather than a binary median split. This refinement adds nuance to the original study, showing that variation in treatment effects is more pronounced across the median than within the upper or lower groups. These results further strengthen the original findings.

Finally, we analyze income directly, as a continuous variable. The specification remains (2.5), but its interpretation is now governed by Assumption 2: with an intercept and slope coefficient. In Table 1, Panel (c) reports estimated intercept and slope based on a linear-in-income model with $\mathbf{W}_i = \text{income}_i$, while Panel (d) considers a quadratic-in-income model with $\mathbf{W}_i = (\text{income}_i, \text{income}_i^2)$. Comparing the discretized heterogeneity with this linear and quadratic models, as shown graphically in Figure 1, reveals that the two continuous models effectively capture the overall heterogeneity in this application. Notably, the continuous specifications highlight differences in income distribution within each group, particularly emphasizing that the top quartile has a much wider range than the

others—an important finding for policy implications.

Ultimately, the decision of how to best capture heterogeneity is up to the researcher. The discretized and continuous versions have their strengths and limitations. The continuous approach, particularly when compared to a fine discretization, may yield substantial precision improvements, but requires the structure of Assumption 2. The choice ultimately depends on the specific application and the researcher’s objectives, but in any case, it is important to have rigorous estimation and inference methods for both.

5 Conclusion

We have shown how treatment effect heterogeneity can be effectively analyzed in RD settings using a tractable and transparent linear-interactions model. Our approach avoids the complexities of high-dimensional nonparametric regression while seamlessly integrating into standard RD practice. By incorporating principled bandwidth selection and robust bias-corrected inference, we aim to bring the same level of rigor and standardization to heterogeneous effects as currently exists for average treatment effects. Our analysis has focused exclusively on sharp RD designs, including kink RD designs in the supplemental appendix. Extending these methods to other settings—particularly fuzzy designs—remains an important direction for future research.


Table 1: Overall and Heterogeneous RD Treatment Effects by Income: Political Turnover and Headmaster Replacement in Municipalities.

	Point Estimate	RBC 95% CI	RBC p -value	Sample Size	h
Panel (a): Overall Average Effect					
τ	0.275	[0.177 ; 0.357]	0.000	14,622	0.151
Panel (b): Heterogeneity with Discretize Income					
<i>Binary</i>					
Below median	0.373	[0.284 ; 0.482]	0.000	9,454	0.188
Above median	0.122	[-0.043 ; 0.238]	0.174	6,167	0.154
<i>By Quartile</i>					
1	0.416	[0.253 ; 0.527]	0.000	4,354	0.151
2	0.334	[0.201 ; 0.537]	0.000	4,132	0.166
3	0.164	[-0.164 ; 0.367]	0.453	2,330	0.151
4	0.095	[-0.070 ; 0.230]	0.297	4,495	0.204
<i>By Decile</i>					
1	0.425	[0.121 ; 0.626]	0.004	1,757	0.150
2	0.392	[0.133 ; 0.662]	0.003	1,677	0.132
3	0.472	[0.358 ; 0.727]	0.000	1,928	0.221
4	0.257	[0.023 ; 0.521]	0.032	1,733	0.158
5	0.302	[0.080 ; 0.688]	0.013	1,458	0.161
6	0.087	[-0.661 ; 0.556]	0.866	956	0.190
7	0.288	[0.018 ; 0.535]	0.036	1,064	0.150
8	0.067	[-0.276 ; 0.149]	0.559	912	0.150
9	0.068	[-0.127 ; 0.526]	0.231	1,325	0.134
10	0.124	[-0.066 ; 0.335]	0.188	2,774	0.259
Panel (c): Heterogeneity with Linear Income					
$\theta(0)$	0.471	[0.306 ; 0.635]	0.000	14,410	0.152
$\xi(0)$	-0.217	[-0.385 ; -0.068]	0.005	14,410	0.152
Panel (d): Heterogeneity with Quadratic Income					
$\theta(0)$	0.625	[0.231 ; 0.923]	0.001	14,410	0.152
$\xi_1(0)$	-0.586	[-1.328 ; 0.359]	0.260	14,410	0.152
$\xi_2(0)$	0.161	[-0.287 ; 0.513]	0.580	14,410	0.152

Notes:

- (i) Panel (a) shows the estimated RD treatment effect $\hat{\tau}$ from (2.2); Panel (b) shows $\hat{\kappa}(\cdot)$ from (2.5)–(2.6), where income is discretized and entered as a vector of indicator variables for each category, excluding the first category; Panel (c) shows $\hat{\kappa}(\text{income})$ from (2.5)–(2.6) with income used linearly as a continuous variable, $W_i = \text{income}_i$; and Panel (d) shows $\hat{\kappa}(\text{income})$ from (2.5)–(2.6) with income used quadratically as a continuous variable, $W_i = (\text{income}_i, \text{income}_i^2)'$.
- (ii) Column “Point Estimate” reports local linear regression ($p = 1$) with a triangular kernel MSE-optimal treatment effect point estimators, implemented using the data-driven bandwidth given in Column “ h ”. For more details, see Theorems 1 and 2, and the supplemental appendix.
- (iii) Columns “RBC 95% CI” and “RBC p -value” report robust bias-corrected 95% confidence intervals and p -values, respectively, implemented using the data-driven bandwidth given in Column “ h ”. Variance Estimators are cluster-robust at the municipality level (as in the original study, Akhtari et al., 2022, column 1 of Table A.21). For more details, see Theorem 3, and the supplemental appendix.
- (iv) Column “Sample Size” shows the effective sample size used for each estimation and robust bias-corrected inference procedure, as determined by the data-driven bandwidth reported in column “ h ”. The total sample size is $n = 26,099$.
- (v) For more details, see the supplemental appendix, and the replication files at <https://rdpackages.github.io/replication/>.

Figure 1: Graphical Presentation of Overall and Heterogeneous RD Treatment Effects by Income



Notes:

- (i) The horizontal dashed line is the estimated overall RD treatment effect: $\hat{\tau}$ of (2.2). See Table 1, Panel (a), for the precise empirical result plotted.
- (ii) The dots show the heterogeneous treatment effects for discretized income, with vertical bars representing their associated robust bias-corrected 95% confidence intervals. Subfigure (a) reports discretized heterogeneity below and above median, subfigure (b) reports discretized heterogeneity by quartiles, and subfigure (c) reports discretized heterogeneity by deciles. See Table 1, Panel (b), for the precise empirical results plotted.
- (iii) The solid line shows the estimated semi-linear heterogeneous treatment effect using income as a continuous variable linearly: $\hat{\kappa}(\text{income}_i) = \hat{\theta} + \hat{\xi} \cdot \text{income}_i$. See Table 1, Panel (c), for the precise intercept and slope estimates.
- (iv) The nonlinear dashed line shows the estimated semi-linear heterogeneous treatment effect using income as a continuous variable quadratically: $\hat{\kappa}(\text{income}_i) = \hat{\theta} + \hat{\xi}_1 \cdot \text{income}_i + \hat{\xi}_2 \cdot \text{income}_i^2$. See Table 1, Panel (d), for the precise estimates.
- (v) For more details, see the footnote of Table 1, and the replication files at <https://rdpackages.github.io/replication/>.

References

- Akhtari, M., Moreira, D., and Trucco, L. (2022), “Political Turnover, Bureaucratic Turnover, and the Quality of Public Services,” *American Economic Review*, 112, 442–93.
- Alcantara, R., Hahn, P. R., Carvalho, C., and Lopes, H. (2025), “Learning Conditional Average Treatment Effects in Regression Discontinuity Designs Using Bayesian Additive Regression Trees,” *arXiv preprint arXiv:2503.00326*.
- Angrist, J. D., and Rokkanen, M. (2015), “Wanna get away? Regression discontinuity estimation of exam school effects away from the cutoff,” *Journal of the American Statistical Association*, 110, 1331–1344.
- Arai, Y., and Ichimura, H. (2018), “Simultaneous Selection of Optimal Bandwidths for the Sharp Regression Discontinuity Estimator,” *Quantitative Economics*, 9, 441–482.
- Arai, Y., Otsu, T., and Seo, M. H. (2025), “Regression Discontinuity Design with Potentially Many Covariates,” *Econometric Theory*.
- Caetano, C., Caetano, G., and Escanciano, J. C. (2024), “Robust Identification in Regression Discontinuity Designs with Covariates,” *working paper*.
- Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2018), “On the Effect of Bias Estimation on Coverage Accuracy in Nonparametric Inference,” *Journal of the American Statistical Association*, 113, 767–779.
- (2020), “Optimal Bandwidth Choice for Robust Bias Corrected Inference in Regression Discontinuity Designs,” *Econometrics Journal*, 23, 192–210.
- (2022), “Coverage Error Optimal Confidence Intervals for Local Polynomial Regression,” *Bernoulli*, 28, 2998–3022.
- Calonico, S., Cattaneo, M. D., Farrell, M. H., Palomba, F., and Titiunik, R. (2025), “rdhte: Learning Conditional Average Treatment Effects in RD Designs,” *working paper*.
- Calonico, S., Cattaneo, M. D., Farrell, M. H., and Titiunik, R. (2019), “Regression Discontinuity Designs using Covariates,” *Review of Economics and Statistics*, 101, 442–451.

- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014), “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 82, 2295–2326.
- Cameron, A. C., and Miller, D. L. (2015), “A practitioner’s guide to cluster-robust inference,” *Journal of human resources*, 50, 317–372.
- Card, D., Lee, D. S., Pei, Z., and Weber, A. (2015), “Inference on Causal Effects in a Generalized Regression Kink Design,” *Econometrica*, 83, 2453–2483.
- Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2020), *A Practical Introduction to Regression Discontinuity Designs: Foundations*, Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge University Press.
- (2024), *A Practical Introduction to Regression Discontinuity Designs: Extensions*, Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge University Press.
- Cattaneo, M. D., Keele, L., and Titiunik, R. (2023), “Covariate Adjustment in Regression Discontinuity Designs,” in *Handbook of Matching and Weighting in Causal Inference*, eds. D. S. S. J. R. Zubizarreta, E. A. Stuart and P. R. Rosenbaum, chapter 8, Boca Raton, FL: Chapman & Hall, pp. 153–168.
- Cattaneo, M. D., Keele, L., Titiunik, R., and Vazquez-Bare, G. (2021), “Extrapolating Treatment Effects in Multi-Cutoff Regression Discontinuity Designs,” *Journal of the American Statistical Association*, 116, 1941–1952.
- Cattaneo, M. D., and Titiunik, R. (2022), “Regression Discontinuity Designs,” *Annual Review of Economics*, 14, 821–851.
- De Magalhães, L., Hangartner, D., Hirvonen, S., Meriläinen, J., Ruiz, N. A., and Tukiainen, J. (2025), “When Can We Trust Regression Discontinuity Design Estimates from Close Elections? Evidence from Experimental Benchmarks,” *Political Analysis*.
- Frölich, M., and Huber, M. (2019), “Including covariates in the regression discontinuity design,” *Journal of Business & Economic Statistics*, 37, 736–748.

- Grembi, V., Nannicini, T., and Troiano, U. (2016), “Do fiscal rules matter?” *American Economic Journal: Applied Economics*, 8, 1–30.
- Hahn, J., Todd, P. E., and der Klaauw, W. V. (2001), “Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design,” *Econometrica*, 69, 201–209.
- Hsu, Y.-C., and Shen, S. (2019), “Testing treatment effect heterogeneity in regression discontinuity designs,” *Journal of Econometrics*, 208, 468–486.
- Hyytinen, A., Meriläinen, J., Saarimaa, T., Toivanen, O., and Tukiainen, J. (2018), “When does regression discontinuity design work? Evidence from random election outcomes,” *Quantitative Economics*, 9, 1019–1051.
- Klašnja, M., and Titiunik, R. (2017), “The incumbency curse: Weak parties, term limits, and unfulfilled accountability,” *American Political Science Review*, 111, 129–148.
- Ma, J., and Yu, Z. (2024), “Empirical Likelihood Covariate Adjustment for Regression Discontinuity Designs,” *arXiv preprint arXiv:2008.09263*.
- MacKinnon, J. G., Nielsen, M. Ø., and Webb, M. D. (2023), “Cluster-Robust Inference: A Guide to Empirical Practice,” *Journal of Econometrics*, 232, 272–299.
- Pei, Z., Lee, D. S., Card, D., and Weber, A. (2021), “Local Polynomial Order in Regression Discontinuity Designs,” *Journal of Business & Economic Statistics*, 31, 507–524.
- Peng, S., and Ning, Y. (2021), “Regression Discontinuity Design under Self-selection,” in *International Conference on Artificial Intelligence and Statistics*, PMLR, pp. 118–126.
- Reguly, Á. (2021), “Heterogeneous Treatment Effects in Regression Discontinuity Designs,” *arXiv preprint arXiv:2106.11640*.