

Limitless Regression Discontinuity

Adam C. Sales

University of Texas

Ben B. Hansen

University of Michigan

Conventionally, regression discontinuity analysis contrasts a univariate regression's limits as its independent variable, R , approaches a cut point, c , from either side. Alternative methods target the average treatment effect in a small region around c , at the cost of an assumption that treatment assignment, $I[R < c]$, is ignorable vis-à-vis potential outcomes. Instead, the method presented in this article assumes "residual ignorability," ignorability of treatment assignment vis-à-vis detrended potential outcomes. Detrending is effected not with ordinary least squares but with MM estimation, following a distinct phase of sample decontamination. The method's inferences acknowledge uncertainty in both of these adjustments, despite its applicability whether R is discrete or continuous; it is uniquely robust to leading validity threats facing regression discontinuity designs.

Keywords: causal inference; randomization inference; robust statistics; MM estimation; contamination sensitivity; higher education

1. Introduction

In a regression discontinuity design (RDD; Thistlethwaite & Campbell, 1960), treatment is allocated to subjects for whom a "running variable" R exceeds (or falls below) a predetermined cut point. Lee (2008) has argued that the RDD features "local randomization" of treatment assignment and is therefore "a highly credible and transparent way of estimating program effects" (Lee & Lemieux, 2010, p. 282).

Take the RDD found in Lindo, Sanders, and Oreopoulos (2010; hereafter LSO). LSO attempt to estimate the effect of "academic probation (AP)," an intervention for struggling college students, on students' subsequent grade point averages (GPAs). At one large Canadian university, students with first-year GPAs below a cutoff were put on probation. Comparing subsequent GPA (Y) for students with first-year GPA (R) just below and above the cutoff should reveal the effectiveness of the policy at promoting satisfactory grades.

LSO's data analysis, like that of most RDD studies, used ordinary regression analyses to target an extraordinary parameter. In Imbens and Lemieux's (2008) telling, for example, the target of estimation is not the average treatment effect (ATE) in any one region around the cutoff but rather the "local" average treatment effect (LATE): the limit of ATEs over concentric ever-shrinking regions, essentially an ATE over an infinitesimal interval. Following this "limit understanding," it is common to analyze RDDs using regression to estimate the functional relationships of r to $\mathbb{E}(Y|R = r)$ on either side of the cutoff. The difference between the two regression functions, as evaluated at the cut point, is interpreted as the treatment effect (e.g., Angrist & Lavy, 1999; Berk & Rauma, 1983).

However, the GPAs in the AP study are discrete, measured in 1/100s of a grade point; hence, limits of functions of GPA do not exist.¹ Further, reanalysis of LSO's RDD uncovers evidence of "social corruption" (Wong & Wing, 2016)—some students appear to have finely manipulated their GPAs to avoid probation. This necessitates excluding subjects immediately on or around the cut point—precisely those students to whom the LATE might most plausibly pertain. Either circumstance calls into question the appropriateness of limit-based methods.

Cattaneo, Frandsen, and Titiunik (2015) base RDD inference on the model that, in effect, the RDD is a randomized controlled trial (RCT), at least in sufficiently small neighborhoods of the cut point. Under this assumption, once attention is confined to such a region, the difference of Y means between subjects falling above and below the cut point estimates the ATE within that region. Despite being natural as a specification of Lee's local randomization concept, the RCT model involves an independence condition that is rarely plausible in RDDs. In the LSO example, the data refute this model—unless one rejects all but the small share of the sample contained in a narrow band of the cut point, sacrificing power and external validity.

To circumvent limitations of the simple RCT model, and of the limit understanding, this article weds parametric and local randomization ideas into a novel identifying assumption termed "residual ignorability." The residual ignorability assumption and corresponding ATE estimates pertain to all subjects in the data analysis sample; discrete GPAs do not pose a threat. Manipulation of the running variable remains a threat, but one that Section 3's combination of sample pruning and robust M estimation is uniquely equipped to address.

The remainder of Section 1 uses a public health example to introduce residual ignorability and to review distribution-free analysis of RCTs. Limitless RDD analysis combines these ideas with classical, wholly parametric methods for RDDs (Section 2.1) and RDD specification tests (Section 2.2). Section 3 adapts residual ignorability to data configurations typical of education studies and sets out an analysis plan anticipating common challenges of RDD analysis. Section 4 executes the plan with the LSO study, Section 5 explores the method's performance in simulations, and Section 6 takes stock. Replication materials, including R code, are available as a GitHub repository, <https://github.com/adamSales/lrd>.

1.1. The Death Toll From Hurricane Maria

Hurricane Maria struck the island of Puerto Rico on September 20, 2017. In spite of widespread devastation, for nearly a year, official statistics pegged the number of hurricane-induced deaths at just 64. Estimates from investigative journalists and academic researchers were higher. Santos-Lozada and Howard's (2018) analysis considered recorded mortality in months before and after the hurricane, estimating Maria to have caused 1,139 deaths in excess of those that would have occurred otherwise. (Rivera and Rolke's [2018] and Acosta and Irizarry's [2018] estimates also well exceeded the official count.) This section demonstrates the concept of residual ignorability, if not the scope and particulars of the method detailed in Section 3, in a reanalysis of these monthly death counts.

In this example, let $i = 1, \dots, 12$ denote the months of 2017 and let Puerto Rico's monthly death counts constitute the outcome, Y_i . The running variable $R_i = i$ is month order, and months are "treated," $Z_i = 1$, if and only if $R \geq 9$, Hurricane Maria having occurred in September. Following Neyman (1923/1990) and Rubin (1974), we may then take each i to have two potential outcomes: Y_{Ti} , a potential response under the treatment condition (the number of deaths that would occur were i to fall after Maria), and Y_{Ci} , a potential response to control (the death count were i to fall before Maria). For each i , at most one of Y_{Ci} and Y_{Ti} is observed, depending on Z_i ; observed responses (Y) coincide with $ZY_T + (1 - Z)Y_C$. Differences, $\tau_i = Y_{Ti} - Y_{Ci}$, $i = 9, \dots, 12$, represent mortality caused by Maria. We will discuss RDD estimation of total excess mortality for 2017, $\sum_{i \geq 9} \tau_i$, in due course. The remainder of this section demonstrates how to test the hypothesis $\tau_i \equiv 0$, all i , using a Fisher randomization test—but without assuming the following independence property:

Assumption (Strong Ignorability; Rosenbaum & Rubin, 1983).

$$Y_C \perp Z. \tag{1}$$

Although RCTs validate Equation 1 as a matter of course, the assumption is implausibly strong for the mortality series surrounding Maria. For Equation 1 to hold, Z must be independent of monthly death counts that would have been observed in the absence of exposure to Maria: The distribution underlying the September through December counts must be no different than that of the year's first 8 months. Monthly mortality in Puerto Rico for the period 2010 through 2017, shown in the left panel of Figure 1, shows there is no precedent for such an equivalence. Rather, a marked seasonal trend is apparent, with death counts being higher in the winter months than during the rest of the year; Equation 1 cannot be sustained. (For RDD methodology nonetheless founded on Equation 1, see Cattaneo et al. 2015 or Mattei and Mealli 2016.)

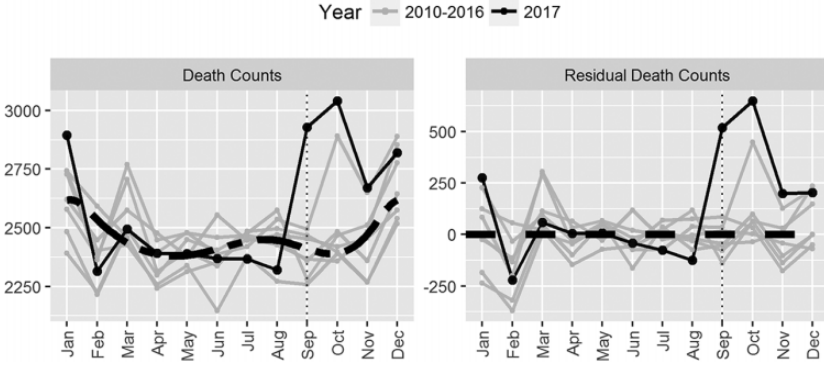


FIGURE 1. Monthly death counts in Puerto Rico from years 2010–2017, before and after residualization. The plot on the left shows monthly death counts, adjusted for month length. A vertical dotted line denotes September, the month Maria hit. The fit of the robust model described in the text is shown as a dashed black line. The plot on the right shows the monthly residuals of the model fit, with a dashed line denoting zero.

Dependence between R and Y_C , violating Equation 1, is common in RDDs; when present, it must be addressed. Inspection of the 2010 through 2016 mortality series (Figure 1) reveals a periodic, nonlinear relationship between calendar month and death count, with some years appearing to be more hazardous than others. To accommodate these factors, we regressed 2010 through 2016 monthly death counts on dummy variables for year and a periodic b -spline for month order, with knots at February, May, August, and November. There are several outlying observations, one of which Santos-Lozada and Howard (2018) remove from the sample prior to analysis. Rather than identifying and removing outliers informally, we fit the regression model using a robust redescending M estimator, which systematically down-weights and sometimes rejects outliers that would otherwise be influential (Maronna, Martin, & Yohai, 2006). This model fit is displayed as a dashed black line in Figure 1.

Now let $\hat{Y}_C(R_i)$ be that model’s prediction for month i in 2017, and let $e_\theta(Y_i|R_i) = Y_i - \hat{Y}_C(R_i)$ be the prediction residual, with potential values $e_\theta(Y_{Ci}|R_i) = Y_{Ci} - \hat{Y}_C(R_i)$ and $e_\theta(Y_{Ti}|R_i) = Y_{Ti} - \hat{Y}_C(R_i)$. Instead of Equation 1, we assume only that the model we have fit to pre-2017 monthly death counts captured and removed seasonal mortality trends, such that the potential residuals $\{E_{Ci} : i\} \equiv \{e_\theta(Y_{Ci}|R_i) : i\}$ can be regarded as random, at least as far as Z is concerned:

$$E_C \perp Z. \tag{2}$$

The right panel of Figure 1 shows residualized death counts $E = e_\theta(Y|R)$ as a function of month order R . Seasonal mean trends are no longer in evidence; Equation 2 is thus more plausible than the standard ignorability assumption

(Equation 1). Aside from a technical elaboration that will be necessary to apply our method in the general case (Section 3.1), Equation 2 is residual ignorability, this article’s alternative to strong ignorability as a basis for analysis of RDDs.

1.2. Using Equation 2 to Test the Hypothesis of Strictly No Effect

In parallel with Fisherian analysis of RCTs (Fisher, 1935), which can be regarded as conditioning on the potential outcome random vector $\mathbf{Y}_C \equiv \{Y_{Ci} : i\}$, our Maria analysis conditions on the potential residual vector $\mathbf{E}_C \equiv \{E_{Ci}\}_{i=1}^{12}$. (Here and throughout, boldface indicates the concatenation of n variables or constants.) The approach applies to the testing of “strict” null hypotheses, hypotheses that designate a value for each $E_{Ti} - E_{Ci}$ with $Z_i = 1$, not just $\mathbb{E}(E_T - E_C | Z = 1)$. This includes the hypothesis of strictly no effect, $H_0 : Y_T \equiv Y_C$, under which $\mathbf{E}_T \equiv \mathbf{E}_C = \mathbf{E}$.

In this analysis Y , R , and Z data for years 2010 through 2016 are treated as fixed—formally, inference will be made after conditioning not only on \mathbf{E}_C but also on the values of $\{(Y_i, R_i, Z_i) : i \leq 0\}$. To make use of Fisher’s (1935) permutation technique, we likewise condition on the realized sizes N_1 and N_0 of the treatment and control group samples,² where $N_j \equiv \sum_{i=1}^{12} \mathcal{I}[Z_i = j]$. Such conditioning is appropriate because the conditioning statistic $\mathbf{A}^* = \{\mathbf{E}_C, N_1, N_0; (Y_{-71}, R_{-71}), \dots, (Y_{-1}, R_{-1})\}$ is ancillary to, that is, carries no information about,³ the target of estimation $\mathbb{E}(E_T - E_C)$.

Under H_0 , we may exactly enumerate the sampling distribution of any test statistic $t(\mathbf{E}_C, \mathbf{Z}) = t(\mathbf{E}, \mathbf{Z})$ conditional on \mathbf{A}^* ; the permutational p value is found by comparing a test statistic to its conditional distribution thus enumerated. In this analysis, $\hat{Y}_C(R_i)$ cannot itself be influenced by Maria since it is based on a model fit to pre-Maria death counts (cf. Sales, Hansen, & Rowan, 2018). Hence, the effect of Maria on E (i.e., $E_T - E_C$) is exactly equal to its effect on Y . The null hypothesis H_0 states that Hurricane Maria caused precisely no change to each month’s death count nor to its residual. Under H_0 , we condition on \mathbf{A}^* and calculate the sampling distribution of the treatment group residual mean $t(\mathbf{e}, \mathbf{Z}) = \bar{e}_{Z=1}$ by calculating its value for all $\binom{N_0 + N_1}{N_1} = \binom{12}{4} = 495$ possible permutations of \mathbf{Z} . The null distribution of $t(\mathbf{e}, \mathbf{Z})$ is simply that of the mean of a without-replacement sample of size 4 from $\{e_1, \dots, e_{12}\}$. It turns out that only two permutations of Z result in test statistic values higher than the realized value $\bar{e}_{z=1} = 1,569$ (which is unique in the distribution). This implies a two-sided “mid” p value (Agresti & Gottard, 2005) of $2(2 + 0.5)/495 \approx 0.01$ for H_0 .

This combination of regression and permutation testing applies just as readily to test the hypothesis $E_T = E_C + c$, for any constant c . In the Maria example, no such hypothesis is sustainable at level .05 unless $c \geq 170$, corresponding to 680

excess deaths due to the hurricane. Upper confidence limits and Hodges–Lehmann-type estimates of the effect can also be obtained in this way. Rather than pursuing this approach further, we now turn to developing a residual ignorability-based procedure using M estimation (Huber, 1964; Stefanski & Boos, 2002; also called “generalized estimating equations” or “generalized method of moments”), which is better adapted to data scenarios without the luxury of a separate sample for estimation of trends in the absence of treatment.

2. Review of Selected RDD Methods

The method presented in this article builds on existing methods for RDDs. This section selectively reviews relevant literature.

Let $Z \in \{0, 1\}$ indicate assignment to treatment ($Z = 1$) as opposed to control ($Z = 0$). For the remainder of the article, let R be the *centered* running variable—the difference between the running variable and the RDD threshold c —so that $Z \equiv \mathcal{I}[R < 0]$, $\mathcal{I}[R \leq 0]$, $\mathcal{I}[R \geq 0]$, or $\mathcal{I}[R > 0]$, depending on how intervention eligibility relates to the threshold, where $\mathcal{I}[x] = 1$ if x is true and 0 otherwise. Let Y represent the outcome of interest. For simplicity assume noninterference, the model that a subject’s response may depend on his but not also on other subjects’ treatment assignments (Cox, 1958; Rubin, 1978). Thus, we may take each i to have two potential outcomes, y_{Ti} and y_{Ci} , at most one of which is observed; observed responses Y coincide with $ZY_T + (1 - Z)Y_C$.

2.1. The Analysis of Covariance (ANCOVA) Model for RDDs

The classical ANCOVA model for groups $i = 1, \dots, k$, each including subjects $j = 1, \dots, n_i$, says that $Y_{ij} = \alpha_i + \beta X_{ij} + \varepsilon_{ij}$, where $\varepsilon_{ij} \sim \text{Normal}(0, \sigma^2)$ is independent of the continuous covariate X_{ij} . In the classical development of RDDs, ANCOVA with $k = 2$ groups—treated and untreated—is a leading option among statistical models (Thistlethwaite & Campbell, 1960). A potential outcomes version of the model is $Y_{Ci} = \alpha_0 + \beta R_i + \varepsilon_{Ci}$ and $Y_{Ti} = \alpha_1 + \beta R_i + \varepsilon_{Ti}$, with $\varepsilon_{Ci} \sim \text{Normal}(0, \sigma^2)$ and $\varepsilon_{Ti} \sim \text{Normal}(0, \sigma^2)$. In marked contrast to RCTs, it is not required that $(Y_T, Y_C) \perp Z$: To the contrary, both Y_C and Y_T are presumed to associate with R , which in turn determines Z . Nonetheless, under this model, the estimated Z coefficient from the model

$$Y_i = \alpha + \beta R_i + \tau Z_i + \varepsilon_i, \tag{3}$$

fit using ordinary least squares (OLS), is unbiased for $\alpha_1 - \alpha_0$. Under the ANCOVA model, this estimation target coincides with $\lim_{r \downarrow 0} \mathbb{E}(Y|R = r) - \lim_{r \uparrow 0} \mathbb{E}(Y|R = r)$ and, simultaneously, limit-free estimation targets such as $\mathbb{E}Y_T - \mathbb{E}Y_C$.

The OLS approach estimates τ as a parameter in regression model (Equation 3). In contrast, the analysis of Section 1.2 took place in two separate steps: first,

adjust outcomes for R ; then, test hypotheses by contrasting adjusted outcomes of treated and untreated subjects. OLS and the ANCOVA model can also be used for hypothesis testing, with steps paralleling those of Section 1.2; this brings an important advantage to be described in Section 2.1.1. Consider the hypothesis $H : Y_T = Y_C + \tau$. Define $Y_H = Y - \tau Z$ (so that under H , $Y_H = Y_C$) and $e_{(a,b)}(\mathbf{y}_H | \mathbf{r}) = \mathbf{y}_H - a - b\mathbf{r}$. Finally, test H with statistic

$$d(\mathbf{Y}_H, \mathbf{Z}) = \overline{e_{(\hat{\alpha}, \hat{\beta})}(Y_H | R)}_{Z=1} - \overline{e_{(\hat{\alpha}, \hat{\beta})}(Y_H | R)}_{Z=0}, \quad (4)$$

where $\hat{\alpha}$ and $\hat{\beta}$ are estimated from an OLS fit of the variant of Equation 3 with dependent variable \mathbf{y}_H . A more essential difference between the current section's procedure and the permutational method of Section 1.2 is that the null distribution of Equation 4 is not tractable. (In Section 1.2, test statistics' permutation distributions were straightforwardly enumerable because slope and intercept parameters had been estimated from a separate sample; in Equation 4, one cannot consider alternate realizations of Z without also considering alternate realizations of $(\hat{\alpha}, \hat{\beta})$.) However, under the parametric ANCOVA model, with conditioning on \mathbf{R} rather than on (N_0, N_1, \mathbf{Y}_C) as in Section 1.2, $d(\mathbf{Y}_H, \mathbf{Z})$ is straightforwardly Normal, with variance equal to the classical OLS variance of the coefficient on Z .

In general, the set $\{c : H_c : Y_T = Y_C + c \text{ is not rejected at level } \alpha\}$, which can be seen to be an interval, is a $100(1 - \alpha)\%$ confidence interval for τ of the Rao score type (Agresti, 2011); the c solving $d(\mathbf{y}_H, \mathbf{z}) = 0$, which can be seen to be unique, is an M estimate of τ under both the classical ANCOVA model and various of its generalizations. In fact, the estimate for τ corresponding to these statistical tests is algebraically equal to the Z coefficient from an OLS estimate of Equation 3, and the two-sided 95% confidence interval induced in this manner is the familiar $\hat{\tau} \pm 1.96SE(\hat{\tau})$. However, these equivalences do not necessarily extend to estimation strategies outside of OLS, such as the robust estimators of Section 3.3.

2.1.1. Addressing the Wald interval's shortcomings for fuzzy RDDs. RDDs susceptible to noncompliance—where subjects' actual treatments may differ from Z —are called “fuzzy.” In these cases, let D indicate whether treatment was actually received. This D is an intermediate outcome, so there are corresponding potential outcomes D_C and D_T , with $D \equiv ZD_T + (1 - Z)D_C$. Subject i is a non-complier if $D_{Ci} = 1$ or $D_{Ti} = 0$, though we will assume the monotonicity condition $D_C \equiv 0$; there may be subjects assigned to the treatment who avoid it, but no one gets treatment without being assigned to it. We shall also posit the exclusion restriction, that Z influences Y only by way of its effect on D (Angrist, Imbens, & Rubin, 1996; Bloom, 1984; Imbens & Rosenbaum, 2005). Our focus of estimation is the “treatment-on-treated” effect (TOTE), $\mathbb{E}(Y_T - Y_C | D_T = 1)$.

Statistical hypotheses about the TOTE take the form $H_\tau : Y_T = Y_C + D\tau$. To test H_τ under noncompliance, let $Y_H = Y - \tau D$, designate $t(\mathbf{y}_H, \mathbf{r})$ as test statistic

and compare its value to a standard Normal distribution. (The only difference between hypothesis testing for a “strict” RDD, one with full compliance, versus a fuzzy RDD, is in the formulation of hypothesis H, and the construction of Y_H —the rest of the process remains unchanged [Rosenbaum, 1996].) When compliance is imperfect, this iterative method yields confidence intervals with better coverage than Wald-type confidence intervals—that is, intervals of form $\hat{\tau} \pm q_* SE(\hat{\tau})$ with $SE(\hat{\tau})$, a single, hypothesis-independent quantity (Baiocchi, Cheng, & Small, 2014; Imbens & Rosenbaum, 2005, section 7).

2.1.2. Robust standard error estimation. The ANCOVA model for (Y_T, Y_C) is not readily dispensed with, but it may be relaxed. OLS estimates of $\alpha_1 - \alpha_0$ and β remain unbiased under non-Normality, provided the ε s have expectation 0 and bounded variances. The ordinary ANCOVA standard error does not require Normality of the $(\varepsilon_i : i)$, either, for use in large samples, although it does require that they have a common variance. To test $\mathbb{E}\{e_{\hat{\theta}}(Y_H|R)|Z = 1\} = \mathbb{E}\{e_{\hat{\theta}}(Y_H|R)|Z = 0\}$ under potential heteroscedasticity, one estimates $\text{Var}\{d(\mathbf{Y}_H, \mathbf{Z})\}$ using a sandwich or Huber–White estimator, $SE_s^2\{d(\mathbf{Y}_H, \mathbf{Z})\}$ (Bell & McCaffrey, 2002; Huber, 1967; Long & Ervin, 2000; MacKinnon & White, 1985; Pustejovsky & Tipton, 2018) and refers $d(\mathbf{Y}_H, \mathbf{Z})/SE_s$ to a t or standard Normal reference distribution. Sandwich standard errors confer robustness to misspecification of $\text{Var}(e_{\hat{\theta}}(Y_H|R)|R)$, not of $\mathbb{E}(Y_H|R)$ (Freedman, 2006), the latter being the topic of the following section.

2.2. Threats to RDD Validity and Some Remedies

The ANCOVA model for RDDs encodes additional assumptions, beyond normality and homoskedasticity of regression errors and full compliance with treatment assignment, which are not so easily dispensed with. Methodological RDD literature has responded with specification tests to detect these threats or with flexible estimators that seek to avoid them.

2.2.1. Covariate balance tests. Analysis of RCTs and quasi-experiments often hinges on assumptions of independence of \mathbf{Z} from $(\mathbf{X}, \mathbf{Y}_C, \mathbf{Y}_T)$. Although neither $\mathbf{Z} \perp \mathbf{Y}_C$ nor $\mathbf{Z} \perp \mathbf{Y}_T$ can be directly tested, since potential outcomes are only partly observed, assumptions of form $\mathbf{Z} \perp \mathbf{X}$ are falsifiable: Researchers can conduct placebo tests for effects of \mathbf{Z} on \mathbf{X} . Of course, treatment cannot affect pretreatment variables; this is model-checking (Cox, 2006, section 5.13).

Writing in the RDD context, Cattaneo et al. (2015) test for marginal associations of \mathbf{Z} with covariates \mathbf{X}_i , $i = 1, \dots, k$, using the permutational methods that are applied in Fisherian analysis of RCTs (also see Li, Mattei, & Mealli, 2015). Relatedly, Lee and Lemieux (2010) recommend a test for conditional association, given R , of \mathbf{Z} and \mathbf{X} , by fitting models like those discussed in Section 2.1 for impact estimation but with covariates rather than outcomes as independent

variables. Viewing the R slopes and intercepts as simultaneously estimated nuisance parameters, these are balance tests applied to the covariates' residuals rather than to the covariates themselves.

If there are multiple covariates, there will be several such tests. To summarize their findings with a single p value, the regressions themselves may be fused within a “seemingly unrelated regressions” model (Lee & Lemieux, 2010); however, to our knowledge, current software implementations do not support the combination of linear and generalized linear models, such as when covariates are of mixed type. Alternatives include hierarchical Bayesian modeling (Li et al., 2015) or combining separate tests' p values using the Bonferroni principle or other multiple comparison corrections.

2.2.2. The McCrary density test. McCrary's (2008) test for manipulation of treatment assignments can be understood as a placebo test with the density of R as the independent variable. The test's purpose is to expose the circumstance of subjects finely manipulating their R values in order to secure or avoid assignment to treatment. Absent such a circumstance, if R has a density, then it should appear to be roughly the same just below and above the cut point. McCrary's test statistic is the difference in logs of two estimates of R 's density at 0, based on observations with $R < 0$ and $R > 0$, respectively. Manipulation is expected to generate a clump just beside the cut point, on one side of it but not the other, and this in turn engenders imbalance in terms of distance from the cut point.

2.2.3. Reducing the bandwidth. In practice, specification test failures inform sample exclusions. When balance tests fail, Lee and Lemieux (2010) would select a bandwidth $b > 0$, restrict analysis to observations with $R \in \mathcal{W} \subseteq [-b, b]$, and repeat the test on $\{i : r_i \in \mathcal{W}\}$. If that test fails, the process may be repeated with a new bandwidth $b' < b$ and perhaps repeated again until arriving at suitable bandwidth. This may seem to call for a further layer of multiplicity correction, since any number of bandwidths may have been tested before identifying a suitable b , but it so happens that this form of sequential testing implicitly corrects for multiplicity, according to the sequential intersection union principle (SIUP; Rosenbaum, 2008, proposition 1; Hansen & Sales, 2015). Li, Mattei, and Mealli (2015) and Cattaneo et al. (2015) also suggest the use of covariate balance to select a bandwidth.

Restricting analysis to data within a bandwidth may change the interpretation of the result. The ATE and the TOTE refer to a discrete population, and reducing the bandwidth likewise reduces those populations—the new target populations consist of subjects for whom $|R| \leq b$. (In contrast, the definition of the LATE is unaffected by bandwidth choice.)

Failures of the density test are addressed by restricting estimation to observations with $|R| > a$, some $a \geq 0$ (e.g., Barreca, Guldi, Lindo, & Waddell, 2011; Eggers, Fowler, Hainmueller, Hall, & Snyder, 2015), and repeating the test. If

this test rejects, we repeat the process with a new $a' > a$, terminating the process when the p value from the density test exceeds a preset threshold. By a second application of the SIUP, the size of this test sequence is equal to the size of each individual density test. Taken together, placebo and McCrary tests restrict the sample to $\mathcal{W} = (-b, b)$ or $(-b, -a) \cup (a, b)$.

2.2.4. Nonlinear models for Y as a function of R . The methods discussed in Section 2.1 continue to apply if $\mathbb{E}(Y_C|R) = \alpha + R\beta$ is relaxed to $\mathbb{E}(Y_C|R) = \alpha + \mathbf{f}(R)\boldsymbol{\beta}$, for $\mathbf{f}(\cdot)$ a $1 \times k$ vector valued function, and $\boldsymbol{\beta}$ a $k \times 1$ vector of coefficients. Unfortunately, if the model is fit by OLS, then such relaxation of assumptions can have the unwelcome side effect of undercutting the robustness of the analysis. The reasons have to do with mechanics of regression fitting.

Polynomial specifications $\mathbb{E}(Y|R = r) = \sum_{j=0}^J r^j \beta_j$ are common but often problematic; in combination with OLS fitting, they implicitly assign case weights that can vary widely and counterintuitively (Gelman & Imbens, 2019). This liability is already in evidence with $J = 1$, the linear specification, where leverage increases with the square of $r - \bar{r}$. If analysts select a bandwidth b that is slightly too large, then the analysis sample will include problematic observations near its outer boundaries, precisely where leverage is at its highest. If the analysis sample is contaminated near the cut point, the bad data may not threaten linear specifications, but with $J > 1$, they can still bear undue leverage. In order to identify leverage points that are also influential, OLS fitting is sometimes combined with specialized diagnostics such as plots of Cook's distances (Cook & Weisberg, 1982). Section 3.3 will discuss an alternate remedy.

3. Randomness and Regression in RDDs

The analysis of Section 1.1 mounted an analogy between the Hurricane Maria RDD and a hypothetical RCT but only after a preparatory step of modeling and removing the outcomes' nonrandom component. In Section 1.1, these two steps used two different data sets—we regressed Y on R using data from years prior to 2017, when Maria hit, and then used 2017 data to estimate effects, under the assumption of residual ignorability, Equation 2. This luxury is unavailable in the typical RDD, in which both steps must use the same data, as in Section 2.1. This section will describe a generalization of residual ignorability (Equation 2) to the typical case, along with robust analysis techniques incorporating the specification tests reviewed in Section 2.2.

3.1. An Analytic Model for RDDs

This section will formalize residual ignorability for the typical RDD, which relies on a single data set including variables Y , R , and Z . The assumption is that,

after a suitable residual transformation, potential outcomes Y_C are conditionally independent of Z . Hence, causal inference in an RDD may take the perspective that Z is random due to randomness in R .

Suppose the statistician to have selected a *detrending procedure*: a trend fitter, that is, a function of $\{(y_i, d_i, r_i)\}_{i=1}^n$ returning fitted parameters $\hat{\theta}$ in a sufficiently regular fashion, along with a family $\{e_\theta(\cdot) : \theta\}$ of residual or partial residual transformations, each mapping data (\mathbf{y}, \mathbf{r}) to residuals $\{e_\theta(y_i|r_i)\}_{i=1}^n$. Appendix A states the needed regularity condition, which is ordinarily met by OLS and always met with our preferred fitters (Section 3.3). Then, causal inference in an RDD proceeds from the following assumption:

Assumption (Residual Ignorability). Given a detrending procedure $(\hat{\theta}, e_\theta(y|r))$,

$$e_{\hat{\theta}}(Y_C|R) \perp Z | \{R \in \mathcal{W}\}. \quad (5)$$

Here $Z = f(R)$, for some deterministic f (such as $f(r) = \mathcal{I}[r < 0]$); θ_0 is a constant such that $\sqrt{n}(\hat{\theta} - \theta_0)$ is bounded in probability (and thus $\hat{\theta} \xrightarrow{P} \theta_0$); \mathcal{W} satisfies $\Pr(R \in \mathcal{W}) > 0$ and $0 < \Pr(Z = 1 | R \in \mathcal{W}) < 1$.

Residual ignorability states that, though Y_C may not be independent of Z , it admits a residual transformation bringing about such independence. With $e_{\hat{\theta}}(Y_C|R)$ a suitable partial residual, residual ignorability is entailed by the ANCOVA model (Section 2.1) or by the combination of any parametric model for $\mathbb{E}(Y_C|R)$ with a strict null H relative to which the value of Y_C can be reconstructed from the values of Y, D , and Z (Section 1.2). (In either of these cases, $e_{\theta_0}(Y_C|R)$ is independent not only of Z but also R , a modest strengthening of Equation 5.)

Assuming residual ignorability, inference about treatment effects is made conditionally, on $\mathbf{A} = (e_{\theta_0}(\mathbf{Y}_C|\mathbf{R}), \mathbf{D}_T, \{(Y_{Ti}, Y_{Ci}, D_{Ti}, R_i)\mathcal{I}[R_i \in \mathcal{W}]\}_{i=1}^n)$. (D_{Cs} are omitted because of our assumption $D_C \equiv 0$; see Section 2.1.1.) Conditioning on the full data vector when $R \in \mathcal{W}$ excludes observations for which Equation 5 is not assumed. Conditioning on $e_{\theta_0}(\mathbf{Y}_C|\mathbf{R})$ removes little of the randomness in \mathbf{R} , leaving it available as a basis for inference. Uncoupled to Y_{Ts} , the detrended Y_{Cs} , $e_{\theta_0}(\mathbf{Y}_C|\mathbf{R})$, are in themselves uninformative about $\mathbb{E}(Y_T - Y_C)$, so the variables comprising \mathbf{A} are jointly ancillary, just as \mathbf{A}^* was seen to be in Section 1.2. As in Fisher-style randomization inference for RCTs, some conditioning variables are unobserved, but this is not an impediment, at least for large-sample inferences.

Causal inference based on residual ignorability takes place in four steps: (1) choosing and validating the analysis sample or bandwidth, (2) choosing an appropriate fitting procedure (we recommend robust fitters), (3) treatment effect estimation and inference, and (4) postfitting diagnostics. We will discuss each of these steps in sequence.

3.2. Pre-Fitting Diagnostics and Bandwidth Choice

If subject matter knowledge suggests that the ATE or TOTE would be most relevant for subjects with $|R| \leq b$, then b might form an initial bandwidth choice. But it is also sensible to subject this choice to specification testing (Section 2.2).

Covariate balance or placebo tests for RDDs (Section 2.2.1) assess residual ignorability with a multivariate “outcome” Y^* combining the actual outcome Y with covariates X —Equation 5 with $\mathbf{Y}_C^* = (\mathbf{X}, Y_C)$ in place of Y_C .

Of the placebo testing procedures discussed in Section 2.2.1, that of Lee and Lemieux (2010) is best suited to this conception. In effect, it begins with preliminary detrending procedures—mechanisms to decompose X into components that are systematic or unpredictable, vis-à-vis \mathbf{R} , just as Y_C will later be decomposed. Our analysis of the LSO data posits systematic components that are linear and logistic linear in R , depending on whether X is a measurement or binary variable. The placebo check adds Z to the specification and tests whether its coefficient is zero. We implement these checks as Wald tests with heteroscedasticity-robust standard errors, as in Section 2.1, using the Bonferroni method to combine placebo checks across covariates. To ensure adequate power to detect misspecification, we test at level .15, not .05.

We use sequential balance tests to adjust the bandwidth b , alongside McCrary density tests to further refine the analysis sample \mathcal{W} (Section 2.2.3). These specification tests rely on covariates, R and Z , but not on Y ; therefore, selection of \mathcal{W} is objectivity-preserving in the sense of Rubin (2007).

3.3. Robust Fitters

Observations in the analysis sample that do not satisfy residual ignorability can undercut the validity of an RDD analysis. Even moderate amounts of such contamination—specifically, contamination of a $O(n^{-1/2})$ -sized share of the sample that happens to contain influential observations—can defeat OLS-based estimation strategies, rendering them inconsistent. Indeed, even some robust regression methods—those engineered to meet objectives other than bounding the influence function—may be misled (Stefanski, 1991). The inclusion of problematic observations in the analysis sample can be due to misspecification of the model for $\mathbb{E}(Y_C|R)$ (Section 2.2.4), manipulation of treatment assignments (Section 2.2.2), or other violations of residual ignorability, coupled with the failure of specification tests to detect these problems. However, no specification test is powerful enough to reliably detect moderate contamination; if the probability of a false alarm is controlled, then power to detect anomalies affecting only $O(n^{-1/2})$ of the sample can only tend to a number strictly less than 1. However b and \mathcal{W} are selected, at least some contamination may remain in the sample.

Accordingly, consistent estimation of θ_0 requires robust M estimators, in Yohai and Zamar’s (1997) sense, a class excluding maximum likelihood estimation while including modern MM, SM, and other estimators with bounded influence function (see also He, 1991, thm. 3). In MM estimation as in OLS, coefficients $\boldsymbol{\beta}$ of a linear specification solve estimating equations $\sum_i \psi\{(y_i - \mathbf{x}_i^T \boldsymbol{\beta})/s\} \mathbf{x}_i^T = \mathbf{0}$, where $s > 0$ and $\psi(\cdot)$ is an odd function satisfying $\psi(0) = 0$, $\psi'(0) = 1$, and $t\psi(t) \geq 0$; bounded influence fitters replace OLS’s $s \equiv 1$ with resistant preliminary estimates of residual scale and OLS’s $\psi(t) = t$ with a continuous ψ that satisfies $\int_0^\infty \psi(t)dt < \infty$. This limits the loss incurred by the fitter for failing to adapt itself to a small portion of aberrant observations; it is permitted to systematically down-weight them instead.

The analyses and simulations presented below use MM estimators with bisquare ψ and “fast S” initialization (Salibian-Barrera & Yohai, 2006). We are not aware of prior work addressing potential contamination of an RDD sample with the assistance of bounded influence MM estimation. Surprisingly, given their common origins in Huber (1964), MM estimation is not routinely paired with sandwich estimates of variance, as in Section 2.1.2. Exceptions include Stata’s `mmregress` and R’s `lmrob`, which optionally provide Huber–White standard errors (Rousseeuw et al., 2015; Verardi & Croux, 2009); our analyses use the latter.

3.4. Treatment Effect Estimation and Inference

For inference about τ under the model $Y_T = Y_C + \tau D_T$, select a specification $\mu_\beta(\cdot)$ for $\mathbb{E}(Y_C|R)$ such as the linear model $\mu_\beta(R) = \alpha + R\beta$, a window of analysis \mathcal{W} , and a fitter. Then, separately for each hypothesis $H : \tau = \tau_0$ under consideration, calculate $\mathbf{y}_H = \mathbf{y} - \mathbf{d}\tau_0$ and apply the chosen specification and fitter to $(\mathbf{y}_H, \mathbf{r})$. The combination of the data, the model fit, and the residual transformation $e_\theta(\cdot|\cdot)$ give rise to residuals $e_\theta(\mathbf{y}_H|\mathbf{r})$, completing the detrending procedure. Whether H is rejected or sustained is determined by the value of the sandwich-based ANCOVA t statistic in Section 2.1.2.

In practice, it is expedient to use a near-equivalent test by modifying the detrending procedure. When regressing Y_H on R , include an additive contribution from Z , so that $\mu_\theta(R) = \alpha + R\beta$ is replaced with $\mu_{(\theta,\gamma)}(R) = \alpha + \beta R + \gamma Z$. With sandwich estimates of $\text{Cov}\{\hat{\theta}_H, \hat{\gamma}_H\}$, the t ratio comparing $\hat{\gamma}_H$ to $SE_s(\hat{\gamma}_H)$ induces a generalized score test (Boos, 1992). Implicitly it is a two-sample t statistic with covariance adjustment for R (with fitting via OLS, this correspondence would be exact, as noted in Section 2.1.2; with the robust MM estimation we favor, the correspondence is one of large-sample equivalence; see Appendix A.2).

As in Section 2.1, the corresponding M estimate of the TOTE is the value of τ_0 making $\hat{\gamma}_H/SE_s(\hat{\gamma}_H)$ equal 0; those τ_0 for which $H : \tau = \tau_0$ is not rejected at level α constitute a $100(1 - \alpha)\%$ confidence interval. Iteration is facilitated by regressing y on \mathbf{r} and \mathbf{z} with offset variable $\mathbf{d}\tau_0$; then only the offset needs to be modified to test $H : \tau = \tau_1, \tau_1 \neq \tau_0$.

Strictly speaking, this estimation procedure relies on the assumption of a constant additive treatment effect, so that $Y_T = Y_C + D\tau_0$, for some constant τ_0 (or, more generally, an exact model for treatment effects under which Y_C could be recovered without error). This requirement is typical of estimators derived from the inversion of hypothesis tests (e.g., Imbens & Rosenbaum, 2005). However, due to its use of sandwich standard errors $SE_s(\cdot)$, our M estimator is robust to some departures from this assumption. Specifically, under Equation 5, $\hat{\tau}$ is consistent for τ_0 solving

$$\mathbb{E}\{e_0(Y_T - D_T\tau_0|R)|Z = 1, R \in \mathcal{W}\} \equiv \mathbb{E}\{e_0(Y_C|R)|R \in \mathcal{W}\},$$

providing such a τ_0 exists. That is, we assume the existence of a constant τ_0 , such that detrended $Y_H = Y - D\tau_0$ is equal to detrended Y_C on average, if not exactly. The simulation study in Section 5.1 bears out this robustness property.

3.5. Postfitting Diagnostics

Once the M estimate for the treatment effect has been found, one inspects the corresponding regression fit for points of high influence. Robust MM regression is helpful here. Besides making influential points easier to see in residual plots, it limits effects of data contamination, as nonconforming influence points are down-weighted as a result of the fitting process. This down-weighting is reflected in “robustness weights,” ranging from 1 for non-discounted observations down to 0 for the most anomalous observations. Plotting robustness weights against residuals may expose opportunities to improve the fit of $\mu_0(R)$, or of the treatment effect model; plotting them against R may expose contaminated subregions of \mathcal{W} that specification testing failed to remove (Maronna et al., 2006).

4. The Effect of AP

Figure 2 plots the LSO study’s primary outcome, GPA in the next term a student was enrolled following his first year (nextGPA), against first-year GPA. In all but 50 of 44,362 cases, being on AP coincided with whether first-year cumulative GPA—the running variable, R —fell below a cutoff. The university in question had three campuses, two having cutoffs of 1.5 and the other having a cutoff of 1.6. To combine data from the three schools, LSO centered each student’s first-year GPA at the appropriate c , making r_i the difference of student i ’s realized first-year GPA and the cutoff at his or her campus. Figure 2 follows LSO in this, displaying these r_i s on its x -axis; it also averages nextGPA values over students with equal first-year GPA, as opposed to plotting students individually.

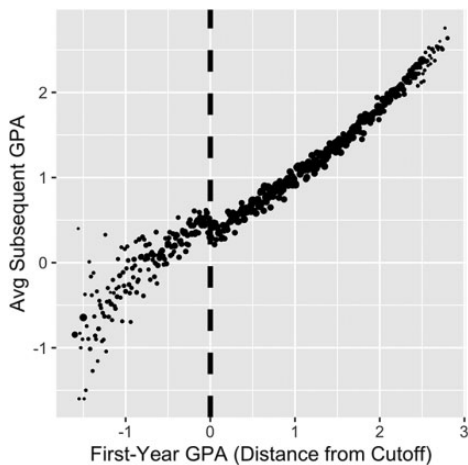


FIGURE 2. *nextGPA* by first-year GPA. Data are plotted with binning by unique value of first-year GPA, vertical coordinates being bin means of *nextGPA* and point sizes being proportional to bin size. GPA = grade point average.

There is both a discontinuity in *nextGPA* values as R crosses 0 and a distinctly nonnull regression relationship on either side of that threshold. How large an AP effect may we infer from these features? How much of the data bear directly in this inference?

4.1. Choosing \mathcal{W} and $\mu_\theta(\cdot)$

The region $\mathcal{W}_{0.5} = \pm 0.5$ grade points include students whose AP status could change if their grades in half their classes changed by a full mark (say from D to C). Simplicity recommends a linear specification for the outcome regression on the forcing variable, and the scatter of Y versus R did not suggest otherwise; so we designated $\mu_\theta(R_i) = \alpha + \beta R_i$ and proceeded to specification checks, as discussed in Section 2.2.

Following LSO, we conducted placebo tests with high-school grade percentile rankings, number of credits attempted in first year of college, first language other than English, birth outside of North America, age at college entry, and which of the university's three campuses the student attended. For the measurement variables, this amounted to fitting ANCOVA models, whereas binary covariates were decomposed as logistic linear in R and Z ; in both cases, subsequent Wald tests of Z 's coefficient used Huber–White standard errors. For $\mathcal{W}_{0.5}$, each (Bonferroni-corrected) p value exceeds .2; downward adjustment of the bandwidth is not indicated.

The McCrary (2008) density test identifies a discontinuity in the running variable at the cut point ($p < .001$). AP is a dubious distinction, and savvy

TABLE 1.

Academic Probation Impact Estimates Using the Method of Section 3 and Variants That Select \mathcal{W} Adaptively, Model the Outcome as a Cubic Function of the Running Variable, or Estimate Intent-to-Treat (ITT) Effects

Specification	Estimate	95% CI	\mathcal{W}	n
Main	.24	[.17, .31]	[0.01, 0.50)	10,014
Adaptive \mathcal{W}	.23	[.18, .27]	[0.01, 1.13)	23,874
Cubic	.24	[.15, .34]	[0.01, 0.50)	10,014
ITT	.24	[.17, .31]	[0.01, 0.50)	10,014

students may try to avoid it. Inspection of the distribution of R reveals an unusual number of students whose first-year GPAs were exactly equal to the AP cutoff, $R = 0$. It would be reasonable to suspect this significant McCrary finding of being an artifact of the discreteness of first-year GPAs, but Frandsen’s (2017) test for manipulation in a discrete running variable likewise detected an anomaly at a wide range of tuning parameter values: $p < .001$ provided that $0 \leq k \leq 0.1$. The finding is further corroborated by the fact that the number of students attempting four or fewer credits was also unusually high in the $R = 0$ subgroup, suggesting that some students dropped courses to dodge AP. In any event, after removing the $R = 0$ subgroup—that is, setting $\mathcal{W} = \{i : |R_i| \in (0, 0.5)\} = \mathcal{W}_{0.5} \setminus \{i : R_i = 0\}$ —the McCrary procedure narrowly avoids rejecting the hypothesis of no manipulation ($p = .15$).

4.2. AP Outcome Analysis

Table 1 gives a set of estimates for the effect of AP, each obtained using the robust procedure of Sections 3.3 and 3.4. The first row of Table 1 gives our main result, with window of analysis $\mathcal{W} = \{i : |R_i| \in (0, 0.5)\}$ and a linear model of $\mathbb{E}(Y_C|R)$, $\mu_0(R) = \alpha + \beta R$, estimating the TOTE based on subjects’ received treatments D . For the best fitting version of the model, robustness weights range from .28 to 1. These weights show little association with R , although the lowest weights occur just above the cut point and near \mathcal{W} ’s edges (Figure 3). The main analysis estimates an average treatment effect of .24, with 95% confidence interval [.17, .31].

Table 1’s next three rows relax each of the main model’s specifications. The row labeled “Adaptive \mathcal{W} ” reports the results using the wider, adaptively chosen window $\mathcal{W}_a = \{i : |R_i| \in (0, 1.13)\}$. The “Cubic” row allows for a cubic relationship between first-year GPA and subsequent GPA, with $\mu_0(R) = \alpha + \beta_1 R + \beta_2 R^2 + \beta_3 R^3$. This specification performed well in simulations (Section 5), suggesting that the warnings in Gelman and Imbens (2019) against higher order global polynomials may not apply to analyses that use robust fitters. Finally, the “ITT” row gives an “intent to treat” analysis, ignoring the difference between students’ actual probation and what we would have expected




FIGURE 3. Robustness weights from robust MM estimation of the model that $\mathbb{E}(Y_C|R = r)$ is linear in r while $Y_T = Y_C + \tau_0 D_T$, with τ_0 set at $\hat{\tau} = 0.24$.

TABLE 2.
The Effect of Academic Probation From Our Main Analysis Compared With Permutation and Ordinary Least Squares Analyses

Method	Estimate	95% CI	\mathcal{W}	n
Local Linear	.24	[.19, .28]	[0.00, 1.25)	26,647
Limitless	.24	[.17, .31]	[0.01, 0.50)	10,014
Local Permutation	.10	[.04, .15]	[0.01, 0.19)	3,766

based on their GPAs. According to all four analyses, AP gave a modest benefit over this range.

4.3. Comparison With Selected Alternatives

For comparison purposes, we reanalyzed the LSO data using two alternative methods: local linear regression (e.g., Imbens & Kalyanaraman, 2012), which targets the difference of limits of regression functions, and the permutational method of Cattaneo et al. (2015), which does not require a limit-based interpretation. The three sets of results are in Table 2.

The local linear approach used the widest window, including observations with $R < 1.25$, and the local permutation approach used the smallest window, including only observations with $R < 0.19$. The effect estimates from our method and local linear regression largely agree, whereas the local permutation

approach finds a smaller effect, with a confidence interval excluding the other two point estimates.

The data sample for the local linear approach differed from ours in two ways. First, since the goal of local linear analysis is to estimate regression functions at the cutoff, it makes little sense to discard observations with $R = c$, despite counterindications from the McCrary and Frandsen tests (Section 4.1). Second, the Imbens and Kalyanaraman (2012) bandwidth is based on nonparametric estimates of the curvature of the mean function $\mathbb{E}(Y|R = r)$ rather than on covariates. We computed this bandwidth using Dimmery’s (2013) implementation in R, using the “rectangular” kernel option to facilitate comparisons across methods. The resulting \mathcal{W} is the widest shown in Table 2—too wide, from viewpoints either of Section 3.1, or of local permutation analysis. For example, Section 3.2’s placebo tests reject comparability of detrended covariate residuals when applied with this \mathcal{W} ($p = .046$).

Local linear effect estimation resembles our method, in that both require analysts to specify and fit models for Y_C and τ . However, whereas ours calls for robust M estimation, the local linear method uses weighted least squares—when the kernel is rectangular, as in our example, this reduces to OLS within the chosen window. Confidence intervals are of the Wald type—that is, $\hat{\tau} \pm z_{\alpha/2} SE(\hat{\tau})$, where $z_{\alpha/2}$ is an appropriate normal or t distribution quantile—rather than inversions of a family of hypothesis tests. Recent elaborations and extensions include those of Calonico, Cattaneo, and Titiunik (2014), Imbens and Wager (2019), and Kolesár and Rothe (2018).

Similar to our approach, the permutation-based procedure of Cattaneo et al. (2015) uses covariates to select a window of analysis. However, its covariate balance tests do not adjust for R , instead seeking a \mathcal{W} over which $X \perp Z$ is not rejected. In the LSO case, \mathcal{W}_b is rejected as long as $b \geq 0.19$. Recall that our R -adjusted check found no fault with bandwidths as large as 1.13. (In both cases, we tested each \mathcal{W} at level $\alpha = .15$, addressing multiplicity of covariates using the Bonferroni method.)

Within the chosen window, the permutational approach estimates effects under the assumption of ignorability of treatment assignment, $Z \perp Y_C$. Failure of this assumption may explain differences between the permutation-based estimate of the AP effect and estimates from the other two methods shown in Table 2. A correlation between nextGPA and R —possible even in regions in which covariate balance cannot be rejected—would bias a positive effect toward zero. The Bayesian method of Li et al. (2015), which begins from a similar ignorability assumption, nevertheless models the relationship between R and Y within the chosen region, to guard against the assumption’s failure. Doing so in the LSO data set yields a similar point estimate as does the permutational approach but with a wider confidence interval that includes the estimates from our and the local linear approach.

TABLE 3.

Empirical Bias and 95% Confidence Interval Coverage (%) and Width for the Analyses of 5,000 Simulated Data Sets Using Either Permutation Tests, Limitless, or Local Ordinary Least Squares (OLS) Methods

n	Effect	Error	Permutation			“Limitless”			Local OLS		
			Bias	Cover	Width	Bias	Cover.	Width	Bias	Cover.	Width
50	0	$\mathcal{N}(0, 1)$.37	64	.91	-.00	93	1.75	-.00	93	1.69
	0	t_3	.37	50	.74	.01	94	1.41	-.00	94	1.66
	t_3	t_3	.37	65	.95	.01	93	1.80	.01	93	2.04
250	0	$\mathcal{N}(0, 1)$.37	3	.39	-.00	95	0.77	-.00	95	0.75
	0	t_3	.37	0	.29	-.00	95	0.57	-.00	95	0.74
	t_3	t_3	.37	3	.38	-.00	95	0.73	-.00	95	0.91
2,500	0	$\mathcal{N}(0, 1)$.38	0	.12	.00	95	0.24	.00	95	0.24
	0	t_3	.37	0	.09	.00	96	0.17	.00	95	0.23
	t_3	t_3	.37	0	.11	.00	94	0.22	.00	95	0.29

Note. The average treatment effect was zero in all conditions; in six conditions the effect was uniformly zero, and in three it was distributed as t_3 .

5. Simulation Studies

5.1. Point and Interval Estimates for Three RDD Methods

Our first simulation study compares the performance—bias and confidence interval coverage and width—of our “limitless” method to local-OLS and local permutation methods. Across all simulation runs, the running variable R was generated as $\text{Uniform}(-0.5, 0.5)$, and control potential outcomes were generated as $Y_C = 0.75R + \varepsilon$, where the .75 slope was chosen to approximately match the estimated slope from the LSO study. Within this framework, we varied three factors: (a) sample size, (b) the distribution of regression error ε , and (c) the treatment effect. We considered three sample sizes: $n = 50, 250,$ and $2,500$. Regression errors were distributed as either Normal or Student’s t with 3 df ; to mimic the LSO data, we forced the errors to have a standard deviation of .75. Finally, the treatment effect was either exactly zero—so $Y_T = Y_C$ —or was generated randomly, as $Y_T = Y_C + \eta$, where $(\sqrt{3}/0.75)\eta \sim t_3$ (so η was drawn from the same distribution as was ε). Each simulation scenario was run 5,000 times.

The results are displayed in Table 3. With a linear data generating model and a symmetric window, the bias for the local permutation approach will generally be equal to the product of the slope and the bandwidth; in our scenario, its bias was approximately $0.75 \times 0.5 \approx 0.37$ across simulation runs. The coverage of permutation confidence intervals decreased with sample size. The limitless and local OLS methods were approximately unbiased, and 95% confidence intervals achieved approximately nominal coverage for $n = 250$ or $2,500$, and




FIGURE 4. Data-generating models for the polynomial simulation.

undercovered for $n = 50$. Notably, random treatment effect heterogeneity did not affect bias or coverage.

Across the board, the local permutation method gave the smallest confidence intervals; however, this came at the expense of coverage. Our limitless RD method tended to have equal or slightly narrower interval widths than the local OLS approach, with greater advantage when ε was distributed as t_3 than when ε was normally distributed.

5.2. Polynomial Regression

When Y may not be linear in R , flexibility in the $\mu_\theta(R)$ function takes on added importance. We ran an additional simulation to explore the potential of robust polynomial regression to mitigate influence, as discussed in Section 3.4, while adding flexibility to the specification of the Y_C on R regression. We compared limitless RD analysis, with μ_θ a polynomial in R with degree 1, 2, 3, 4, or 5 to analogous estimates from OLS. In the OLS regressions, we followed the advice of Lee and Lemieux (2010, p. 318) and included interactions between the R polynomial and Z . Finally, we compared these methods to local-linear regression with the triangular kernel and the bandwidth of Imbens and Kalyanaraman (2012). The OLS and limitless methods used the entire range of data. We simulated data sets of size $n = 500$ by drawing R and ε from Uniform $(-1, 1)$ and t_3 distributions, respectively, then adding ε to one of the three functions of R shown in Figure 4 to form Y_C .

Table 4 displays the results. For the linear data generating model, all estimates were unbiased, while root mean squared errors (RMSEs) were lowest for the limitless method. For the nonlinear data generating models, the limitless estimators using linear and quadratic specifications had substantial bias, and bias was much lower for higher order polynomial specifications. In contrast, OLS estimators of all polynomial degrees were heavily biased. The local linear model does not employ higher order polynomials. It fared better than OLS, having similar bias but higher RMSE than limitless with higher order polynomials.

TABLE 4.
Results From 5,000 Simulations of Polynomial Specifications for RDD Analysis, Using Limitless, OLS, or Local Linear Regression

DGM	Measure	Limitless					OLS					Local Linear
		Polynomial Degree					Polynomial Degree					
		1	2	3	4	5	1	2	3	4	5	
Linear	Bias	0.0	0.0	.0	.0	.0	0.0	-0.0	0.0	0.3	-1.7	.0
	RMSE	0.2	0.2	.3	.3	.4	0.3	1.1	4.7	22	106	.5
Anti-Sym	Bias	-0.6	-0.6	-0.0	-0.0	.1	-0.6	1.7	1.8	-9.0	-9.4	-0.0
	RMSE	0.7	0.7	.3	.3	.4	0.7	2.0	5.0	24	106	.5
Sine	Bias	1.2	1.2	.2	.2	.0	1.2	-2.6	-2.2	1.8	0.2	.1
	RMSE	1.2	1.2	.3	.3	.4	1.2	2.9	5.2	21	103	.5

Note. Data generating models were as depicted in Figure 4, with t_3 errors; sample size for all runs was 500; there was no treatment effect. RMSE = root mean squared error.

OLS and limitless estimation sharply diverge in the quality of their point estimates, with the OLS estimates’ RMSEs exceeding those of comparable robust M estimates by factors exceeding 200. As Gelman and Imbens (2019) would predict, OLS estimates’ RMSEs increased sharply with each increment of polynomial degree, whatever the form of the data generating model. In marked contrast, under nonlinear data generating models, higher degree polynomial terms increased the accuracy of the robust, limitless method; under linear data generating models, including higher degree terms imposed little penalty.

6. Discussion

Beginning with Thistlethwaite and Campbell (1960), the dominant mode of RDD analysis has built upon ANCOVA models. A modern variant instead targets parameters defined in terms of limits as $r \rightarrow c$, such as the LATE. However, in RDDs with discrete running variables such as LSO’s, neither $\lim_{r \downarrow 0} \mathbb{E}(Y|R = r)$ nor $\lim_{r \uparrow 0} \mathbb{E}(Y|R = r)$ exists, except perhaps as $\mathbb{E}(Y|R = -.01)$ or $\mathbb{E}(Y|R = 0)$. A separate embarrassment for limit-based modeling of RDDs occurs if a donut-shaped \mathcal{W} is necessary to address potential manipulation of the running variable, as we found to occur with the LSO data.

An alternative approach (e.g., Cattaneo, Frandsen, & Titiunik, 2015; Li et al., 2015) takes the “local randomization” heuristic more literally, analyzing data in a small region around the cutoff as if it were from a randomized experiment. However, this approach assumes that potential outcomes are independent of R in a window around the cutoff. That assumption is plausible of neither the Hurricane Maria example nor the LSO case study. In both settings, it is necessary

to acknowledge and model the R – Y relationship in order to set the stage for a credible claim of independence. The method of Cattaneo et al. (2015) performed poorly in our simulations (Section 5); its discrepant estimate of LSO’s AP effect (Table 2) embodies systematic error.

In contrast, this article’s RDD analysis framework links ANCOVA and local randomization heuristics. Residual ignorability (Equation 5) assumes that the component of Y_C that depends on R may be modeled and removed, leaving residuals $e_{\theta_0}(\mathbf{Y}|\mathbf{R})$ that are independent of Z . Like the local randomization approach, it targets the TOTE or ATE within \mathcal{W} , as opposed to a difference of limits, and accommodates discreteness in R and donut designs. In the special case of residual ignorability models with $e_{\theta}(\mathbf{Y}_C|\mathbf{R}) \equiv \mathbf{Y}_C$ in Equation 5, it reduces to the local randomization method. Like the limit-based approach, it models and accounts for the correlation between R and Y . Under certain modeling and fitting choices, it returns the classical ANCOVA estimate (Section 2.1).

The method of this article improves upon each of these approaches by using robust M estimation to adjust for R . For analysis of potentially imperfect RDDs, we see this as a necessity. For instance, covariate balance tests will necessarily be underpowered to detect imbalance in a small fraction of the sample, so the proper bandwidth b will be uncertain. Likewise, if the initial sample includes subjects who manipulated their recorded R values, then the use of donut-shaped \mathcal{W} may remove some, but not all such subjects. Robust M estimation retains consistency under scenarios such as these, with moderate amounts of contamination (He, 1991; Yohai & Zamar, 1997), whereas OLS does not.

If a large fraction of the dataset violates Equation 5, even robust M estimators can be misled. Thus, these methods should be used in addition to, rather than instead of, preliminary specification checks.

In simulated RDDs of moderate size, our estimates were unbiased and our confidence intervals were typically narrower than those from an OLS-based approach, while achieving nominal coverage. Further simulations found robust M estimation to be compatible with the use of global cubic and quartic polynomials to accommodate nonlinear, imperfectly modeled relationships between R and Y_C , in marked contrast to methods using OLS to adjust for trend.

Appendix

A. Large-Sample Randomization Inference for Regression Discontinuity Designs (RDDs)

In Section 3 and following, the conditioning is with respect to

$$\mathcal{F}_n^* = \sigma[e_{\theta_0}(\mathbf{Y}_C|\mathbf{R}), \mathbf{D}_T, \{(Y_{Ti}, Y_{Ci}, D_{Ti}, R_i)\mathcal{I}[R_i \notin \mathcal{W}]\}_{i=1}^n], \quad (\text{A1})$$

a sigma field bearing information about Y_C s but not Y_T s. In contrast to sigma fields used in Fisherian randomization inference, it does not carry information about the sizes of the treatment and control groups (within $\{i : R_i \in \mathcal{W}\}$). In Equation A1, conditioning on full information about subjects i whose R values fell outside \mathcal{W} is a formal reflection of those subjects' removal from the analytic sample.

As indicated in Section 1.1, we assume noninterference, the model that a subject's response may depend on his but not also on other subjects' treatment assignments (Cox, 1958; Rubin, 1978); for fuzzy RDDs, we also assume the exclusion restriction and monotonicity (of (D_C, D_T) ; see Section 2.1.1). In varying degree, the sections that follow place additional assumptions bounding Y_T , Y_C , or transformations of them.

The Residual Ignorability condition of Section 3.1 assumes $\hat{\theta}$ to be determined in a "sufficiently regular" fashion, noting that bounded influence MM estimation meets this requirement. A weaker regularity condition than boundedness of the influence function is that the fitter's influence function $\text{IF}_{\hat{\theta}}(w, (\theta, \eta))$, where (θ, η) denotes the full parameter and $w = (y, d, r)$, must satisfy: $\mathbb{E}[\text{IF}_{\hat{\theta}}(W, (\theta, \eta))] = 0$ for a unique $\theta = \theta_0$; each (θ, η) solving $\mathbb{E}[\text{IF}_{\hat{\theta}}(W, (\theta, \eta))] = 0$ makes $\mathbb{E}[\text{IF}_{\hat{\theta}}(W, (\theta, \eta))\text{IF}_{\hat{\theta}}(W, (\theta, \eta))']$ finite.

A.1. Distributional Approximation for $\hat{\theta} - \theta_0$

Recall Section 3.1 assumes that for the true parameter (θ_0, η_0) , $\mathbb{E}[\text{IF}_{\hat{\theta}}(W, (\theta_0, \eta_0))] = 0$ and $\Sigma = \mathbb{E}[\text{IF}_{\hat{\theta}}(W, (\theta_0, \eta_0))\text{IF}_{\hat{\theta}}(W, (\theta_0, \eta_0))']$ is finite. As noted by Stefanski and Boos (2002, section 3), these entail that $n^{-1/2}\sum_{i=1}^n \text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) \xrightarrow{d} N(0, \Sigma)$ and $n^{1/2}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, \Sigma)$, and if fitting is done by MM estimation also that sandwich estimates are consistent, for Σ . This argument applies also if $\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0))$ is the influence function of $(\hat{\theta}, \hat{\gamma})$, as opposed only to $\hat{\theta}$, where $\hat{\gamma}$ is a Z coefficient as discussed in Section 3.4.

For inference conditioned on \mathcal{F}_n^* as in Equation A1, we require an approximation to the distribution of $\sum_{i=1}^n \text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) - \mathbb{E}[\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) | \mathcal{F}_n^*]$. For methods recommended in this article, it suffices to consider the case that $\text{IF}_{\hat{\theta}}(w, (\theta, \eta))$ is bounded. Write

$$\mathcal{G}_n = \sigma\left(\{W_i\}_{i=1}^n; \left\{ \left(e_{\theta_0}(\mathbf{Y}_C | \mathbf{R}), \mathbf{D}_{T_i}, (Y_{T_i}, Y_{C_i}, D_{T_i}, R_i) \mathcal{I}[R_i \notin \mathcal{W}] \right) \right\}_{i=n+1}^{\infty} \right),$$

so that $\mathcal{F}_n^* \subseteq \mathcal{G}_m$ for all n and m , while $\sum_{i=1}^n \text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) - \mathbb{E}[\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) | \mathcal{F}_n^*]$ is adapted to filtration $(\mathcal{G}_n : n)$. Fixing a vector t of the same dimension as $\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0))$, and writing $M_n = \sum_{i=1}^n t \text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) - t \mathbb{E}[\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) | \mathcal{F}_n^*]$, we see that $(M_n : n)$ is a martingale. If $\text{IF}_{\hat{\theta}}(W_n, (\theta_0, \eta_0))t$ is \mathcal{F}_n^* measurable then $M_n = 0$ a.s. for all n and is asymptotically

Normal(0, 0²). Otherwise $\mathbb{E}\{\text{Var}[t\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0))|\mathcal{F}_n^*]\} > 0$ and $\sum \text{Var}[t\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0))|\mathcal{F}_n^*] = O_P(n)$. The Lindeberg condition follows by dominated convergence, since $\text{IF}_{\hat{\theta}}(w, (\theta, \eta))$ is bounded. Thus, $\{n^{1/2}M_n\}$ is asymptotically Normal by Lévy's martingale central limit theorem. Because t was arbitrary, it follows that $n^{-1/2}\sum_{i=1}^n t\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0)) - \mathbb{E}[t\text{IF}_{\hat{\theta}}(W_i, (\theta_0, \eta_0))|\mathcal{F}_n^*]$ converges in distribution to a Normal distribution with mean 0 and variance $\text{Var}(t\text{IF}_{\hat{\theta}}(W_1, (\theta_0, \eta_0))|\mathcal{F}_1^*)$. Accordingly, $n^{1/2}(\hat{\theta}, \hat{\gamma}) - (\theta_0, \gamma_0)$ is asymptotically MVN with covariance $\text{Cov}[\text{IF}_{\hat{\theta}}(W_1, (\theta_0, \eta_0))|\mathcal{F}_1^*]$.

Sandwich covariance estimates converge to Σ , not $\text{Cov}[\text{IF}_{\hat{\theta}}(W_1, (\theta_0, \eta_0))|\mathcal{F}_1^*]$; but since $t\text{Cov}[\text{IF}_{\hat{\theta}}(W_1, (\theta_0, \eta_0))|\mathcal{F}_1^*]t' \leq t\text{Cov}[\text{IF}_{\hat{\theta}}(W_1, (\theta_0, \eta_0))]t'$ for each t , under H , we have $\hat{\gamma}_H/\text{SE}_s(\hat{\gamma}_H) \rightarrow^d \text{Normal}(0, \nu)$ for some $\nu \leq 1$ (where “SE_s” indicates a sandwich estimate).

A.2. Equivalence of Tests From z Coefficient and From Mean Difference of Partial Residuals

Let $f_0(r) = 1, f_1(r) = \mathcal{I}[r < 0] - \mathbb{E}\mathcal{I}[R < 0] = z - p_z$, and $f_j(\cdot), j = 2, \dots, J$, be functions defining a design matrix for regression on Z and R ; denote coefficients of this regression α, γ, β , where α and γ are scalar coefficients for $f_0(r), f_1(r)$ and β is a scalar ($J = 2$) or column vector ($J > 2$) multiplier for row vectors $\vec{f}(r) = [f_2(r) \dots f_J(r)]$. We demonstrate that when α, γ, β and perhaps other components of θ are fitted simultaneously, in an MM estimation type regression procedure making their sampling variability $O_P(n^{-1/2})$, then $\hat{\gamma}$ differs only by $o_P(n^{-1/2})$ from a constant multiple of $\overline{e_{\hat{\theta}}(y_H|r)}_{z=1} - \overline{e_{\hat{\theta}}(y_H|r)}_{z=0}$, the difference in means of partial residuals.

Denote parameters other than γ , that is, α, β , a scale parameter s and perhaps others, collectively by η , so that $\theta = (\gamma, \eta)$. Fixing a (strong) hypothesis H , write w for (y_H, r, z) and $e_{(\gamma, \eta)}(w)$ for $\Psi\{(y_H - \alpha - \gamma f_1(r) - \vec{f}(r)\beta)/s\}$ such that the estimating functions are $\Psi_j(w, \gamma, \eta) = e_{(\gamma, \eta)}(w)f_j(r), j \leq J$. The column equation $\mathbb{E}\Psi(W, \gamma_H, \eta_H) = 0, \Psi(w, \gamma, \eta) = [\Psi_1(w, \gamma, \eta), \dots, \Psi_J(w, \gamma, \eta)]^t$, implicitly defines γ_H , and defines or contributes to an implicit definition of θ_H . It's permitted that there be additional estimating equations, not involving γ , that contribute to the implicit definition of θ : in robust regression, equations defining the preliminary scale estimate; an equation $\mathbb{E}Z = p_z$ defining p_z . Note that the factor of Ψ_1 involving $z, f_1(r) = z - p_z$, has been centered on p_z . This facilitates the argument by ensuring that $\sum_i \Psi_1(W_i, \gamma, \eta) = np_z(1 - p_z) \left\{ \overline{e_{(\gamma, \eta)}(W)}_{Z=1} - \overline{e_{(\gamma, \eta)}(W)}_{Z=0} \right\}$ but does not affect the value of γ or its estimates $\hat{\gamma}$, due to the inclusion via Ψ_0 of an intercept term. Note that $e_{\hat{\theta}}(y_H|r) = e_{(\gamma, \eta)}(y_H|r) = e_{(0, \eta)}(w)$.

Finally, assume that with these estimating equations and $\mathcal{L}(W)$ jointly the estimating function is asymptotically linear,

$$\begin{aligned}
 & n^{-1} \sum_{i=1}^n \Psi(W_i, \hat{\gamma}_n, \hat{\eta}_n), \\
 & = n^{-1} \sum_{i=1}^n \Psi(W_i, \gamma_H, \eta_H) + \dot{\Psi}_H [\hat{\gamma}_n - \gamma_H, \hat{\eta}_n - \eta_H]' + o_P(n^{-1/2}), \quad (A2)
 \end{aligned}$$

for consistent estimators $(\hat{\gamma}_n, \hat{\eta}_n)$ of (γ_H, η_H) , where $\dot{\Psi}_H$ is a matrix with rows $\nabla_{\gamma, \eta} \{\mathbb{E}\Psi_j(W, \gamma, \eta)\}$, as evaluated at $(\gamma, \eta) = (\gamma_H, \eta_H)$. Assume the fitting procedure generates consistent estimators of (γ_H, η_H) and $\dot{\Psi}_H$, denoted $(\hat{\gamma}, \hat{\eta})$ and $\hat{\dot{\Psi}}$, with either $\sum_{i=1}^n \Psi(W_i, \hat{\gamma}, \hat{\eta}) = 0$ or at least $|\sum_{i=1}^n \Psi(W_i, \hat{\gamma}, \hat{\eta})|_2 = o_P(n^{-1/2})$.

For robust regression embodying a residual transformation $\psi(\cdot)$ that redescends, consistency and Equation A2 follow if Ψ is twice differentiable in the parameters, with derivatives of $\Psi(W, \gamma, \beta, \sigma)$ bounded by a function $K(W)$ with finite expectation. This can be assumed of robust regression using the bisquare or lqq ψ functions, as in R's robustbase package or Stata's mmregress, if Y_H and $(f_j(R) : j)$ have finite second moments. If Ψ is the estimating function of an ordinary or generalized linear model, or robust regression with Huber loss, suitable conditions for consistency and Equation A2 are given in, for example, He and Shao (2000).

If H is true then $\gamma_H = 0$, and either of $(\hat{\gamma}, \hat{\eta})$ and $(0, \hat{\eta})$ is consistent for $(\gamma_H, \eta_H) = (0, \eta_H)$. Under H , then, we can apply Equation A2 to either of these estimator sequences, with the consequence that

$$\begin{aligned}
 & n^{-1} \sum_{i=1}^n \Psi_1(W_i, \hat{\gamma}, \hat{\eta}) - n^{-1} \sum_{i=1}^n \Psi_1(W_i, 0, \hat{\eta}), \\
 & = \left[\frac{\partial}{\partial \gamma} \mathbb{E}\Psi_1(W, \gamma, \eta) \Big|_{(0, \eta_H)} \right] \hat{\gamma} + o_P(n^{-1/2}), \\
 & = \dot{\Psi}_{1\gamma} \hat{\gamma} + o_P(n^{-1/2}).
 \end{aligned}$$

Since $n^{-1} |\sum_{i=1}^n \Psi(W_i, \hat{\gamma}, \hat{\eta})|_2 = o_P(n^{-1/2})$, this means that up to differences of order $o_P(n^{-1/2})$, $n^{-1} \sum_{i=1}^n \Psi_1(W_i, 0, \hat{\eta}) \approx \dot{\Psi}_{1\gamma} \hat{\gamma}$. But $n^{-1} \sum_{i=1}^n \Psi_1(W_i, 0, \hat{\eta}) = \bar{z}(1 - \bar{z}) \left\{ \overline{e_{(0, \hat{\eta})}(W)}_{Z=1} - \overline{e_{(0, \hat{\eta})}(W)}_{Z=0} \right\}$, so

$$n^{1/2} \left| \frac{\bar{z}(1 - \bar{z})}{\hat{\dot{\Psi}}_{1\gamma}} \cdot \left[\overline{e_{(0, \hat{\eta})}(W)}_{Z=1} - \overline{e_{(0, \hat{\eta})}(W)}_{Z=0} \right] - \hat{\gamma} \right| \xrightarrow{P} 0.$$

A.3. Differences of Residual Means Involving Parameters Estimated During Detrending

Fix $H : Y_T = Y_C + \tau_0 D_T$ and write

$$Y_{HC} = Y_T - D_T \tau_0 \quad \text{and} \quad Y_H = ZY_{HC} + (1 - Z)Y_C, \quad (A3)$$

so that for each i , Y_{HCi} is the y_C value that would be reconstructed from data (y_{Ti}, d_{Ti}) under H , if (y_{Ti}, d_{Ti}) rather than y_{Ci} were observed, whereas \mathbf{y}_H is the reconstruction of \mathbf{y}_C according to H based on those data that were actually observed. Thus, $\mathbf{Y}_{HC} \equiv \mathbf{Y}_C$ if H is true, but not otherwise. Let $\bar{\theta} = (\bar{\alpha}, \bar{\beta}, \bar{s}, \dots)$ describe a solution of a system of estimating equations, including a subsystem $\mathbb{E}\psi\{(Y_H - \vec{f}(R)(\alpha, \beta))/s\} \vec{f}(R) = \mathbf{0}$, some $\vec{f}(\cdot) = [f_0(\cdot) f_1(\cdot) \dots f_k(\cdot)]$ and some $\psi(\cdot)$. (For n vectors \mathbf{u} , \mathbf{v} , and \mathbf{x} , “[$\mathbf{u}\mathbf{v}\mathbf{x}$]” denotes the $n \times 3$ matrix with columns \mathbf{u} , \mathbf{v} , and \mathbf{x} .) In the method described and exemplified in Section 3 and subsequent parts of this article, $e_\theta(y_H|r) = \psi\{(y_H - \vec{f}(r)(\alpha, \beta))/s\}$; however, for the arguments presented in this appendix, it is permitted that θ be estimated using an unrelated $\psi(\cdot)$ function, prior to and separately from the residualization $(y, r) \mapsto e_\theta(y|r)$ figuring in tests of H .

Given modest regularity conditions on the distribution of $(Y_H, R) = (Zy_{HC} + (1 - Z)y_C, R)$ (cf. Equation A3) and the transformation $e_\theta(\cdot|\cdot)$, one has

$$\left\{ \overline{e_\theta(Y_H|R)}_{Z=1} - \overline{e_\theta(Y_H|R)}_{Z=0} \right\} - \left\{ \overline{e_{\bar{\theta}}(Y_H|R)}_{Z=1} - \overline{e_{\bar{\theta}}(Y_H|R)}_{Z=0} \right\} = \nabla_\theta [\mathbb{E}\{e_\theta(Y_H|R)|Z=1\} - \mathbb{E}\{e_\theta(Y_H|R)|Z=0\}]_{\theta=\bar{\theta}} (\hat{\theta} - \bar{\theta})^t + o_p(n^{-1/2}), \quad (\text{A4})$$

where $\bar{\theta} = \mathbb{E}(\hat{\theta}(\mathbf{Y}_H, \mathbf{R}))$; see Proposition 1 below. Approximation A4 justifies the robustness to non-additivity claimed in Section 3.4. When $\hat{\theta}$ is $n^{1/2}$ consistent and asymptotically Normal, approximation A4 warrants the use of Huber–White estimates of $\text{Var}\left\{ \overline{e_\theta(Y_H|R)}_{Z=1} - \overline{e_\theta(Y_H|R)}_{Z=0} \right\}$, in turn providing a basis for large-sample t tests.

Approximation A4 holds for randomized and quasiexperimental designs alike, but in randomized controlled trials (RCTs), the bracketed difference of expectations vanishes under H (Hansen & Bowers, 2009; Lin, 2013), and the product at right of Equation A4 vanishes as well. Up to $o_p(n^{-1})$, the variances $\text{Var}\left\{ \overline{e_\theta(Y_H|R)}_{Z=1} - \overline{e_\theta(Y_H|R)}_{Z=0} \right\}$ and $\text{Var}\left\{ \overline{e_{\bar{\theta}}(Y_H|R)}_{Z=1} - \overline{e_{\bar{\theta}}(Y_H|R)}_{Z=0} \right\}$ are the same, the latter being no greater than the expected value of the squared standard error of two-sample inference with unequal variances, as applied to $e_{\bar{\theta}}(y_H|r)$. In practice, that specific standard error is unavailable, $\bar{\theta}$ being unknown; but its ratio with $\text{SE}_u\left\{ \overline{e_\theta(Y_H|R)}_{Z=1} - \overline{e_\theta(Y_H|R)}_{Z=0} \right\} \Big|_{\theta=\hat{\theta}}$, the unequal variances two-sample standard error as figured with substitution of $\hat{\theta}$ for θ , tends in probability to 1, by Slutsky’s Lemma; this suffices for the limiting null distribution of the Studentized t statistic to be standard Normal. So in an RCT, there is no need for any explicit acknowledgment of sampling variability in $\hat{\theta}$.

But this argument does not extend to RDDs. Rather, in an RDD, the expected values compared in Equation A4 coincide only for a single value, ordinarily 0, of

θ 's slope component. The differential at right of Equation A4 is generally non-zero; it must be permitted to make a contribution. Fortunately, any ordinary standard error attaching to the z coefficient of the regression of \mathbf{Y}_H on $\mathbf{1}$, \mathbf{R} and \mathbf{Z} will do so; this includes Huber–White standard errors.

Let $\theta \mapsto e_\theta(y|y) = e_\theta(w)$ be continuously differentiable (for each $w = (y, r)$). By the mean value theorem,

$$\begin{aligned} & \left\{ \left[\overline{e_{\hat{\theta}}(w)}_{z=1} - \overline{e_{\bar{\theta}}(w)}_{z=1} \right] - \left[\overline{e_{\hat{\theta}}(w)}_{z=0} - \overline{e_{\bar{\theta}}(w)}_{z=0} \right] \right\}, \\ & = \nabla_\theta \left[\overline{e_\theta(w)}_{z=1} - \overline{e_\theta(w)}_{z=0} \right]_{\theta = \theta^*} \cdot (\hat{\theta} - \bar{\theta}), \end{aligned} \quad (\text{A5})$$

some θ^* on the line segment connecting $\hat{\theta}$ and $\bar{\theta}$. Of course $\theta^* \xrightarrow{P} \bar{\theta}$ if $\hat{\theta} \rightarrow^P \bar{\theta}$.

Let there be a compact neighborhood Θ of $\bar{\theta}$ and an accompanying envelope function $k_\Theta(\cdot)$, that is, $|\nabla_\theta e_\theta(w)| \leq k_\Theta(w)$, all $\theta \in \Theta$ and all w , that is, integrable, $\mathbb{E}k_\Theta(W) < \infty$. With this assumption, the uniform strong law (e.g., Ferguson, 1996, chapter 16) entails that if $\theta^* \xrightarrow{P} \bar{\theta}$ then

$$\overline{(\nabla_\theta e_\theta(w)|_{\theta=\theta^*})}_{Z=z} \xrightarrow{P} \mathbb{E}[\nabla_\theta e_\theta(W)|_{\theta=\bar{\theta}}|Z = z], \quad z = 0 \text{ or } 1. \quad (\text{A6})$$

Moreover, since there is an integrable envelope function, dominated convergence and the mean value theorem combine to warrant interchanging the expected value and differentiation operations at right of Equation A6.

Strengthening the consistency assumption on $\hat{\theta}$ to root- n consistency, $|\hat{\theta} - \bar{\theta}|_2 = O_P(n^{-1/2})$, it now follows that the difference of the right-hand side of Equation A5 and

$$\nabla_\theta \{ \mathbb{E}[e_\theta(W)|Z = 1] - \mathbb{E}[e_\theta(W)|Z = 0] \}_{\theta=\bar{\theta}} \cdot (\hat{\theta} - \bar{\theta}), \quad (\text{A7})$$

is $o_P(n^{-1/2})$. This suffices for Equation A4. To summarize:

Proposition 1 Under (1)–(3) below, $n^{1/2}\{(\text{A7}) - \text{RHS of (A5)}\} \xrightarrow{P} 0$ as $n \uparrow \infty$, and in consequence, Equation A4 holds.

1. for each (y, r) , $\nabla_\theta e_\theta(y|r)$ exists and is continuous in θ ;
2. for some compact Θ , open $S \subseteq \Theta$ with $\bar{\theta} \in S$, and $k_\Theta(\cdot, \cdot)$ with $\mathbb{E}k_\Theta(Y_H, R) < \infty$, $|\nabla_\theta e_\theta(y|r)| \leq k_\Theta(y, r)$ for all (y, r) ; and
3. $|\hat{\theta} - \bar{\theta}|_2 = O_P(n^{-1/2})$.

The argument can be generalized to cover residual transformations that are Lipschitz but not continuously differentiable in θ , but the generalization is not needed for ordinary generalized linear model fitters nor for common bounded-influence alternatives (e.g., robust regression with bisquare psi function).

Authors' Note

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors.

Acknowledgments

The authors thank Susan Dynarski, Xinzhou Guo, Xuming He, Guido Imbens, Brian Junker, Justin McCrary, Walter Mebane, Kerby Shedden, Jeff Smith, Rocío Titiunik, the participants in the University of Michigan Causal Inference in Education Research Seminar, and three anonymous reviewers for helpful suggestions. They also thank Jeffrey Howard and Alexis Santos-Lozada for sharing nonpublic research replication materials and John E. Bellquist for editing.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research and/or authorship of this article: This research was supported by the Institute of Education Sciences, U.S. Department of Education (R305B1000012), the U.S. National Science Foundation (SES 0753164), and an NICHD center grant (R24 HD041028).

Notes

1. For recent methods addressing bias when R is rounded see, for example, Dong (2015) and Kolesár and Rothe (2018). In those cases, there is a continuous running variable, say R^* , that is unobserved, while observed $R = f(R^*)$ for some $f(\cdot)$ that may be unknown or noninvertible; then the local average treatment effect may be defined in terms of limits of realizations of R^* . In contrast, in the Lindo, Sanders, and Oreopoulos's example, R is discrete *by definition*.
2. *Design of Experiments* (1935, Chapter 1) takes sample sizes to be fixed, but Fisher's example of the purple flowers (e.g., Little, 1989; Upton, 1992) demonstrates his view that random N_0 and N_1 should be treated as fixed after conditioning upon them. For quite different assumptions supporting permutation tests in RDDs, see Canay and Kamat (2017).
3. \mathbf{A}^* carries full information about \mathbf{E}_C but none on \mathbf{E}_T . \mathbf{A}^* would not be ancillary to targets of the form $\mathbb{E}(Z|\varepsilon)$ or $\mathbb{E}(E_C|\varepsilon)$, some event ε .

References

- Acosta, R. J., & Irizarry, R. A. (2018). *Post-hurricane vital statistics expose fragility of Puerto Rico's health system*. BioRxiv Preprint, <https://doi.org/10.1101/407874>.
- Agresti, A. (2011). Score and pseudo-score confidence intervals for categorical data analysis. *Statistics in Biopharmaceutical Research*, 3, 163–172.
- Agresti, A., & Gottard, A. (2005). Comment: Randomized confidence intervals and the mid- p approach. *Statistical Science*, 20, 367–371.

- Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996, June). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, *91*, 444–455.
- Angrist, J. D., & Lavy, V. (1999). Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *The Quarterly Journal of Economics*, *114*, 533–575.
- Baiocchi, M., Cheng, J., & Small, D. S. (2014). Instrumental variable methods for causal inference. *Statistics in Medicine*, *33*, 2297–2340.
- Barreca, A. I., Guldi, M., Lindo, J. M., & Waddell, G. R. (2011). Saving babies? Revisiting the effect of very low birth weight classification. *The Quarterly Journal of Economics*, *126*, 2117–2123.
- Bell, R. M., & McCaffrey, D. F. (2002). Bias reduction in standard errors for linear regression with multi-stage samples. *Survey Methodology*, *28*, 169–182.
- Berk, R. A., & Rauma, D. (1983). Capitalizing on nonrandom assignment to treatments: A regression-discontinuity evaluation of a crime-control program. *Journal of the American Statistical Association*, *78*, 21–27.
- Bloom, H. S. (1984). Accounting for no-shows in experimental evaluation designs. *Evaluation Review*, *8*, 225.
- Boos, D. D. (1992). On generalized score tests. *The American Statistician*, *46*, 327–333.
- Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, *82*, 2295–2326.
- Canay, I. A., & Kamat, V. (2017). Approximate permutation tests and induced order statistics in the regression discontinuity design. *The Review of Economic Studies*, *85*, 1577–1608.
- Cattaneo, M. D., Frandsen, B. R., & Titiunik, R. (2015). Randomization inference in the regression discontinuity design: An application to party advantages in the US senate. *Journal of Causal Inference*, *3*, 1–24.
- Cook, R. D., & Weisberg, S. (1982). *Residuals and influence in regression*. New York, NY: Chapman & Hall.
- Cox, D. R. (1958). *The planning of experiments*. New York, NY: John Wiley.
- Cox, D. R. (2006). *Principles of statistical inference*. Cambridge, MA: Cambridge University Press.
- Dimmery, D. (2013). *RDD: Regression discontinuity estimation* (R package version 0.54) [Computer software manual]. Retrieved from <http://CRAN.R-project.org/package=rdd>
- Dong, Y. (2015). Regression discontinuity applications with rounding errors in the running variable. *Journal of Applied Econometrics*, *30*, 422–446.
- Eggers, A., Fowler, A., Hainmueller, J., Hall, A. B., & Snyder, J. M. (2015). On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races. *American Journal of Political Science*, *59*, 259–274.
- Ferguson, T. S. (1996). *A course in large sample theory*. London, England: Chapman & Hall.
- Fisher, R. A. (1935). *Design of experiments*. Edinburgh, England: Oliver and Boyd.
- Frandsen, B. R. (2017). Party bias in union representation elections: Testing for manipulation in the regression discontinuity design when the running variable is discrete. In M. D. Cattaneo & J. C. Escanciano (Eds.), *Regression discontinuity designs: Theory and applications* (pp. 281–315). Bingley, England: Emerald.

- Freedman, D. A. (2006). On the so-called “Huber sandwich estimator” and “robust standard errors.” *The American Statistician*, *60*, 299–302.
- Gelman, A., & Imbens, G. W. (2019). Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business and Economic Statistics*, *37*, 447–456.
- Hansen, B. B., & Bowers, J. (2009). Attributing effects to a cluster randomized get-out-the-vote campaign. *Journal of the American Statistical Association*, *104*, 873–885.
- Hansen, B. B., & Sales, A. C. (2015). Comments on “Observational Studies,” by William G. Cochran. *Observational Studies*, *1*, 184–193.
- He, X. (1991). A local breakdown property of robust tests in linear regression. *Journal of Multivariate Analysis*, *38*, 294–305.
- He, X., & Shao, Q.-M. (2000). On parameters of increasing dimensions. *Journal of Multivariate Analysis*, *73*, 120–135.
- Huber, P. J. (1964). Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, *35*, 73–101.
- Huber, P. J. (1967). The behavior of maximum likelihood estimates under nonstandard conditions. *Proceedings of the fifth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, University of California Press, Berkeley, CA, pp. 221–233.
- Imbens, G. W., & Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *The Review of Economic Studies*, *79*, 933–959.
- Imbens, G. W., & Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, *142*, 615–635.
- Imbens, G. W., & Rosenbaum, P. R. (2005). Robust, accurate confidence intervals with a weak instrument: Quarter of birth and education. *Journal of the Royal Statistical Society, Series A: Statistics in Society*, *168*, 109–126.
- Imbens, G., & Wager, S. (2019). Optimized regression discontinuity designs. *Review of Economics and Statistics* *101*, 264–278.
- Kolesár, M., & Rothe, C. (2018). Inference in regression discontinuity designs with a discrete running variable. *American Economic Review*, *108*, 2277–2304.
- Lee, D. S. (2008). Randomized experiments from non-random selection in US House elections. *Journal of Econometrics*, *142*, 675–697.
- Lee, D. S., & Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, *48*, 281–355.
- Li, F., Mattei, A., & Mealli, F. (2015). Evaluating the causal effect of university grants on student dropout: Evidence from a regression discontinuity design using principal stratification. *The Annals of Applied Statistics*, *9*, 1906–1931.
- Lin, W. (2013). Agnostic notes on regression adjustments to experimental data: Re-examining Freedman’s critique. *The Annals of Applied Statistics*, *7*, 295–318.
- Lindo, J. M., Sanders, N. J., & Oreopoulos, P. (2010). Ability, gender, and performance standards: Evidence from academic probation. *American Economic Journal: Applied Economics*, *2*, 95–117.
- Little, R. J. A. (1989). Testing the equality of two independent binomial proportions. *The American Statistician*, *43*, 283–288.
- Long, J. S., & Ervin, L. H. (2000). Using heteroscedasticity consistent standard errors in the linear regression model. *The American Statistician*, *54*, 217–224.

- MacKinnon, J. G., & White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29, 305–325.
- Maronna, R. A., Martin, D., & Yohai, V. (2006). *Robust statistics*. Hoboken, NJ: John Wiley.
- Mattei, A., & Mealli, F. (2016). Regression discontinuity designs as local randomized experiments. *Observational Studies*, 2, 156–173.
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142, 698–714.
- Neyman, J. (1990). On the application of probability theory to agricultural experiments. Essay on principles. Section 9 (trans. D. M. Dabrowska and T. P. Speed). *Statistical Science*, 5, 463–480). (Original work published 1923)
- Pustejovsky, J. E., & Tipton, E. (2018). Small-sample methods for cluster-robust variance estimation and hypothesis testing in fixed effects models. *Journal of Business & Economic Statistics*, 36, 672–683.
- Rivera, R., & Rolke, W. (2018). Estimating the death toll of Hurricane Maria. *Significance*, 15, 8–9.
- Rosenbaum, P. R. (1996, June). Identification of causal effects using instrumental variables: Comment. *Journal of the American Statistical Association*, 91, 465–468.
- Rosenbaum, P. R. (2008). Testing hypotheses in order. *Biometrika*, 95, 248–252.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 41–55.
- Rousseeuw, P., Croux, C., Todorov, V., Ruckstuhl, A., Salibian-Barrera, M., Verbeke, T., . . . Maechler, M. (2015). *robustbase: Basic robust statistics* (R package version 0.92-5) [Computer software manual]. Retrieved from <http://CRAN.R-project.org/package=robustbase>
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66, 688.
- Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 6, 34–58.
- Rubin, D. B. (2007). The design versus the analysis of observational studies for causal effects: Parallels with the design of randomized trials. *Statistics in Medicine*, 26, 20–36.
- Sales, A. C., Hansen, B. B., & Rowan, B. (2018). Rebar: Reinforcing a matching estimator with predictions from high-dimensional covariates. *Journal of Educational and Behavioral Statistics*, 43, 3–31.
- Salibian-Barrera, M., & Yohai, V. J. (2006). A fast algorithm for s-regression estimates. *Journal of Computational and Graphical Statistics*, 15, 414–427.
- Santos-Lozada, A. R., & Howard, J. T. (2018). Use of death counts from vital statistics to calculate excess deaths in Puerto Rico following hurricane Maria. *Journal of the American Medical Association*, 320, 1491–1493.
- Stefanski, L. A. (1991). A note on high-breakdown estimators. *Statistics & Probability Letters*, 11, 353–358.
- Stefanski, L. A., & Boos, D. D. (2002). The calculus of M-estimation. *The American Statistician*, 56, 29–38.

- Thistlethwaite, D. L., & Campbell, D. T. (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology*, *51*, 309.
- Upton, G. J. (1992). Fisher's exact test. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, *155*, 395–402.
- Verardi, V., & Croux, C. (2009). Robust regression in Stata. *The Stata Journal*, *9*, 439–453.
- Wong, V. C., & Wing, C. (2016). The regression discontinuity design and the social corruption of quantitative indicators. *Observational Studies*, *2*, 183–209.
- Yohai, V. J., & Zamar, R. H. (1997). Optimal locally robust M-estimates of regression. *Journal of Statistical Planning and Inference*, *64*, 309–323.

Authors

ADAM C. SALES is the director of Statistics, Measurement and Research Design Techniques in Education Research consulting service in the College of Education, University of Texas, Austin, TX 78712; email: asales@utexas.edu. His research interests include principal stratification, machine learning, and causal inference methods incorporating administrative or computer log data.

BEN B. HANSEN is an associate professor in the Department of Statistics and a research associate in the Survey Research Center at the University of Michigan, Ann Arbor, MI 48104; email: bbh@umich.edu. His research interests include matching, sensitivity analysis and distribution-free methods for observational studies and experiments, and applications in education and health.

Manuscript received May 11, 2018

Revision received April 19, 2019

Accepted August 2, 2019