
Covariate Adjustment in Regression Discontinuity Designs

Matias D. Cattaneo, Luke Keele, Rocío Titiunik

CONTENTS

8.1	Introduction	153
8.2	Covariate Adjustment in RD Designs	155
	8.2.1 Overview of the canonical RD design	155
	8.2.2 Efficiency and power improvements	157
	8.2.3 Auxiliary information	159
	8.2.4 Treatment effect heterogeneity	160
	8.2.5 Other parameters of interest and extrapolation	162
8.3	Can Covariates Fix a Broken RD Design?	163
8.4	Recommendations for Practice	164
8.5	Conclusion	165
8.6	Acknowledgments	166
	References	166

8.1 Introduction

In causal inference and program evaluation, a central goal is to learn about the causal effect of a policy or treatment on an outcome of interest. There is a variety of methodological approaches for the identification, estimation, and inference for causal treatment effects, depending on the specific application. In experimental settings, where the treatment assignment mechanism is known, studying treatment effects is relatively straightforward because the methods are based on assumptions that are known to be true by virtue of the treatment assignment rule. In contrast, in observational studies, the assignment mechanism is unknown and researchers are forced to invoke identifying assumptions whose validity is not guaranteed by design. Well-known observational study methods include selection on observables, instrumental variables, difference-in-differences, synthetic controls, and regression discontinuity designs.¹

In experimental and observational studies, baseline covariates generally serve different purposes. In experimental settings covariate adjustment based on pre-intervention measures is often used either for efficiency gains or for the evaluation of treatment effect heterogeneity. On the other hand, in many observational studies, the primary purpose of covariate adjustment is for the identification of causal effects. For example, under selection on observables, the researcher posits that, after conditioning

¹See, for example, [1], [2] and [3] for reviews on causal inference and program evaluation, and [4] for further discussion on the role of natural experiments in observational research.

on a set of pre-intervention covariates, the assignment mechanism is not a function of potential outcomes, thereby mimicking an experimental setting after conditioning on those covariates. Under this assumption, researchers can obtain valid estimates of treatment effects by comparing groups of units whose covariate values are similar. Implementation of covariate adjustment methods in experimental and observational studies can be done in different ways, including via regression methods, inverse probability weighting, and matching. However, the validity of covariate adjustment methods is not automatically guaranteed because it requires additional (usually stronger) assumptions, if valid at all. This is particularly true in the case of the regression discontinuity (RD) design.

Covariate adjustment in RD designs has been employed in a variety of ways and for different purposes in recent years. Because the RD design has a unique feature relative to other non-experimental research strategies, covariate adjustment has distinctive implications for the analysis and interpretation of RD treatment effects. This has led to confusions and misconceptions about the role of auxiliary covariates in RD designs. The main purpose of this chapter is to review the different roles of covariate adjustment in the RD literature, and to offer methodological guidance for its correct use in applications.

To streamline the presentation, we do not provide a comprehensive review of RD designs and methods.² We focus our discussion on the canonical (sharp) RD design, where all units receive an observable score and a treatment is assigned to all units whose value of the score exceeds a known and fixed cutoff, and withheld from units whose score is below it. Under appropriate conditions, the abrupt change in the probability of receiving treatment that occurs at the cutoff justifies comparing units just above the cutoff to units just below it to learn about the causal effect of the treatment. For this canonical RD setup, and employing the taxonomy introduced by [10], we discuss the different roles of covariate adjustment within the two leading conceptual frameworks for the interpretation and analysis of RD designs: the continuity framework and the local randomization framework.

One of the most important roles of baseline covariates in the canonical RD design is for falsification or validation purposes. In the continuity framework for RD analysis, covariates are used to verify that the effect of the treatment on the covariates at the cutoff is zero, offering evidence in support of the assumption that only the probability of treatment status given the score changes discontinuously at the cutoff. In the local randomization approach, covariates are often used to select the window around the cutoff where treated and control units are similar to each other, and where the assumption of as-if random treatment assignment is assumed to hold.

The canonical RD setup is sometimes augmented with additional covariate information for other purposes beside falsification testing, and the role of these additional covariates varies depending on the setting and specific goals. For example, additional covariates may be used for efficiency gains or for the analysis of treatment effect heterogeneity. This use of pre-intervention covariates is motivated by the analogy between randomized experiments and local randomization in RD designs. Alternatively, covariates may be used to construct other treatment effects of interest or for extrapolation purposes (i.e., external validity). This occurs, for instance, when the additional covariates capture other features of the design such as geographic location, evolution over time, or group identity. Finally, sometimes covariates are also used in conjunction with statistical adjustment with the hope of “fixing” an otherwise invalid RD design. See [9] for a comprehensive review and references.

In this chapter we discuss the most common uses that have been proposed for covariate adjustment in the RD literature, including efficiency, heterogeneity, and extrapolation. In each case we discuss whether these methods are best justified from a continuity or a local randomization perspective, and how they methodologically relate to each other. Our main focus is conceptual, and mostly from an identification perspective, although we do discuss implementation issues when appropriate. Covariate adjustment in RD designs can be implemented by either local linear regression, inverse probability weighting, or “matching” methods more generally.

²For recent introductions and practical guides see, for example, [5–9], and references therein.

The rest of the chapter is organized as follows. The next section reviews the different roles of covariate adjustment in RD designs. Then, in [Section 8.3](#), we discuss the important question of whether the use of covariate adjustment such as adding fixed effects for group of units can “restore” the validity of an RD design where the key identifying assumptions do not hold (e.g., settings where units are suspected to have sorted around the cutoff). [Section 8.4](#) offers generic recommendations for practice, and [Section 8.5](#) concludes.

8.2 Covariate Adjustment in RD Designs

For each type of covariate adjustment in RD designs, we discuss specific implications of using covariates. In particular we outline how covariate adjustment can lead to the unintended redefinition of the estimand to a quantity that does not coincide with a parameter of interest. We clarify whether the covariate adjustment applies to the continuity framework, the local randomization framework, or both. While most of the available proposals were introduced within one of these two conceptual frameworks, we also discuss the potential implications that each of the covariate adjustment approaches may have for the other. In this discussion it is important to remember that the core identifying assumptions differ between the two frameworks, and consequently so does the RD parameter of interest. For simplicity, we focus only on sharp RD designs (i.e., settings with perfect treatment compliance or intention-to-treat analysis), but the references given often include discussions of fuzzy, kink, and other related RD setups.

8.2.1 Overview of the canonical RD design

We assume that the observed data is (Y_i, D_i, X_i, Z_i) , $i = 1, 2, \dots, n$, where $Y_i = Y_i(0)(1 - D_i) + Y_i(1)D_i$, $Y_i(0)$ is the potential outcome under control, $Y_i(1)$ is the potential outcome under treatment, $D_i = \mathbb{1}(X_i \geq c)$ denotes the treatment assignment rule, X_i denotes the score variable, c denotes the cutoff, and Z_i denotes other pre-determined covariates. Depending on the framework considered, the potential outcomes are taken to be random or non-random, and the data can be seen either as a (random) sample from some large population or as a finite population. The core assumption in the RD design is that the probability of treatment assignment changes abruptly at the cutoff from zero to one, that is, $\mathbb{P}[D_i = 1 | X_i = x] = 0$ if $x < c$ and $\mathbb{P}[D_i = 1 | X_i = x] = 1$ if $x \geq c$.

In the continuity framework, the data is assumed to be a random sample and the potential outcomes are assumed to be random variables. In this context, the most common parameter of interest is the average treatment effect at the cutoff,

$$\tau_{\text{SRD}} = \mathbb{E}[Y_i(1) - Y_i(0) | X_i = c],$$

and the key identifying assumption is the continuity of the conditional expectations of the potential outcomes given the score: $\mathbb{E}[Y_i(1) | X_i = x]$ and $\mathbb{E}[Y_i(0) | X_i = x]$ are continuous in x at c , which leads to the well-known continuity-based identification result [11]:

$$\tau_{\text{SRD}} = \lim_{x \downarrow c} \mathbb{E}[Y_i | X_i = x] - \lim_{x \uparrow c} \mathbb{E}[Y_i | X_i = x].$$

Many methods are available for estimation, inference, and validation of RD designs within the continuity framework. The most common approach is to use (non-parametric) local polynomial methods to approximate the two regression functions near the cutoff. The implementation requires choosing a bandwidth around the cutoff, and then fitting two polynomial regressions of the outcome on the score – one above and the other below the cutoff – using only observations with scores within

the chosen region around the cutoff as determined by the bandwidth. The bandwidth is typically chosen in an optimal and data-driven fashion to ensure transparency and to avoid specification searching. The most common strategy is to choose a bandwidth that minimizes the asymptotic mean squared error (MSE) of the RD point estimator (see [12], and references therein).

Although the implementation of local polynomial methods reduces to fitting two local, lower-order polynomials using least-squares methods, the interpretation and properties are different from the typical parametric least squares framework. Local polynomials are meant to provide a non-parametric approximation to the regression functions of the potential outcomes given the score. Because these functions are fundamentally unknown, the local polynomial approximation has an unavoidable misspecification error. When the bandwidth is chosen to be MSE-optimal, this approximation error appears in the standard distributional approximation and renders the conventional least squares confidence intervals invalid because conventional inference assumes that the approximation error is zero [13] developed valid confidence intervals based on a novel distributional approximation for a test statistic in which the RD point estimator is adjusted using bias correction and also the standard error is modified to account for the additional variability introduced by the bias estimation step. These robust bias-corrected confidence intervals lead to valid inferences when using the MSE-optimal bandwidth, and remain valid for other bandwidth choices, in addition to having several other optimal properties [12].

The canonical implementation of local polynomial estimation and inference in RD designs uses only the outcome and the score, without additional covariates. We discuss below how the local polynomial approach can be augmented to include pre-determined covariates in the two polynomial fits. In addition we discuss how the local polynomial estimation and inference methods are adapted to incorporate covariates for other purposes such as heterogeneity analysis or extrapolation.

In the local randomization framework for RD analysis, the parameter of interest is defined differently because the core identifying assumptions are different from those in the continuity framework [10, 14]. Because in the local randomization approach the central assumption is that the RD design creates conditions that resemble a randomized experiment in a neighborhood near the cutoff, the parameter is defined within this neighborhood—as opposed to at the cutoff as in the continuity approach. The local randomization parameter can be generically written as

$$\tau_{\text{SLR}} = \frac{1}{N_{\mathcal{W}}} \sum_{i: X_i \in \mathcal{W}} \mathbb{E}_{\mathcal{W}}[Y_i(1) - Y_i(0)]$$

where $\mathcal{W} = [c - w, c + w]$ is a symmetric window around the cutoff, $N_{\mathcal{W}}$ is the number of units with $X_i \in \mathcal{W}$, $\mathbb{E}_{\mathcal{W}}$ denotes an expectation computed conditionally for all units with $X_i \in \mathcal{W}$, and two conditions are assumed to hold within \mathcal{W} : (i) the joint distribution of scores in \mathcal{W} is known; and (ii) the potential outcomes are unaffected by the score X_i in \mathcal{W} . Under these local randomization assumptions, estimation and inference can proceed using standard methods for the analysis of experiments, including Fisherian methods where the potential outcomes are assumed non-random and inferences are finite-sample exact, and Neyman or super-population methods where inferences are based on large-sample approximations with either non-random or random potential outcomes. In general the parameter of interest can be written as

$$\tau_{\text{SLR}} = \frac{1}{N_{\mathcal{W}}} \sum_{i: X_i \in \mathcal{W}} \mathbb{E}_{\mathcal{W}} \left[\frac{T_i Y_i}{\mathbb{P}_{\mathcal{W}}[T_i = 1]} \right] - \frac{1}{N_{\mathcal{W}}} \sum_{i: X_i \in \mathcal{W}} \mathbb{E}_{\mathcal{W}} \left[\frac{(1 - T_i) Y_i}{1 - \mathbb{P}_{\mathcal{W}}[T_i = 1]} \right],$$

where $\mathbb{P}_{\mathcal{W}}$ denotes a probability computed conditionally for all units with $X_i \in \mathcal{W}$, which indicates a natural estimator by virtue of the known assignment mechanism within \mathcal{W} and the other assumptions imposed.

The window \mathcal{W} where the local randomization assumptions are assumed to hold is unknown in most applications. The recommended approach to choose this window in practice is based on

balance tests performed on one or more pre-determined covariates [14]. The main idea behind this method is that the treatment effect on any pre-determined covariate is known to be zero, and thus a non-zero effect is evidence that the local randomization assumptions do not hold. Assuming that the covariate is correlated with the score outside of \mathcal{W} but not inside, a test of the hypothesis that the mean (or other feature of the distribution) of the pre-determined covariate is the same in the treated and control groups should be rejected in all windows larger than \mathcal{W} , and not rejected in \mathcal{W} or any window contained in it. The method thus consists of performing balance tests in nested (often symmetric) windows, starting with the largest possible window and continuing until the hypothesis of covariate balance fails to be rejected in one window and in all windows contained in it.

The window \mathcal{W} in the local randomization approach is analogous to the bandwidth in the continuity-based local polynomial analysis. An important distinction, however, is that the bandwidth can be chosen in a data-driven and optimal way using only data on the outcome and the score, while choosing the local randomization window \mathcal{W} requires data on auxiliary pre-determined covariates. In this sense, the local polynomial analysis can be fully implemented without additional covariates, while the local randomization analysis typically cannot—unless \mathcal{W} is known or chosen in an arbitrary manner. For further details on RD estimation and inference under the continuity-based and local randomization approaches, we refer the reader to [5–9], and references therein.

8.2.2 Efficiency and power improvements

In the analysis of randomized experiments, it is well known that including pre-intervention covariates in the estimation of treatment effects can increase precision and statistical power, as discussed in standard textbooks on the analysis of experiments [1, 2, 15]. Covariate adjustment in experimental analysis can be implemented both before and after randomization has occurred. Adjustment before randomization is typically done via blocking or stratification, where units are first separated into groups according to their values of one or more pre-treatment covariates, and the treatment is then assigned randomly within groups. These methods are effective at increasing precision, but unfortunately they are not available for RD designs because the assignment of treatment is not under the control of the researcher, who typically receives the data after the treatment has already been assigned.³

One popular ex-post method of covariate adjustment for experiments transforms the outcome to subtract from it a prior measure (also known as “pre-test outcome” or simply a “pre-test”) that may come from a baseline survey or pilot study collected before the treatment was assigned. For example, instead of estimating the average treatment effect on the outcome, the researcher estimates the average treatment effect on the transformed outcome. The adjusted estimator is consistent for the average treatment effect, and leads to efficiency gains when the pre-test outcome is highly predictive of the post-treatment outcome. This idea can be easily implemented using modern non-parametric or machine learning methods.

Another popular method for ex-post covariate adjustment in experiments, and one of the most widely used, is to add pre-treatment covariates to a linear regression fit between the outcome and the treatment indicator. This regression-based covariate adjustment typically leads to efficiency gains when the added covariates are correlated with the outcome, and has been studied in detail in various contexts (see, for example, [2], and references therein). An interesting finding in this literature is that a regression model that incorporates the pre-intervention covariates and their interaction with the treatment variable is more efficient than a simple covariate adjusted estimator excluding the interaction term, unless the treatment and control groups have equal number of observations or the covariates are uncorrelated with the individual treatment effect.

³In recent years there has been interest in designing ex-ante surveys for participants in RD designs. At the moment, the RD literature only has methodological developments for simple power calculations and survey sample design ([16]), but no work is available taking into account covariates for how to use ex-ante survey information to increase efficiency in the ex-post analysis.

These ideas can be applied to the analysis of RD designs with suitable modifications. In the local randomization RD approach, the particular properties of covariate adjustment follow directly from applying the results from the literature on experiments to the window around the cutoff where the local randomization assumptions are assumed to hold. Depending on the assumptions imposed (Fisherian, Neymann, or super-population), and the specific estimation and inference methods considered, the efficiency gains may be more or less important.

Furthermore, regression-based covariate adjustment can also be used in the local randomization approach to relax some of the underlying assumptions. For example, [10] model the potential outcomes as polynomials of the running variable, to allow for a more flexible functional form of the potential outcomes inside the window \mathcal{W} where the local randomization assumptions are assumed to hold. This model can be extended directly to include pre-determined covariates in the polynomial model, in addition to the running variable.

In the continuity-based framework, applying the same arguments is not immediate because, although the local polynomial RD estimator is based on linear regressions above and below the cutoff, these regressions are non-parametric and the average treatment effect is estimated by extrapolating to the cutoff point. The properties of regression-based covariate adjustment in this context were first studied by [17], who show that the inclusion of covariates other than the score in a local polynomial analysis can lead to asymptotic efficiency gains, if carefully implemented. The authors also show that covariate adjustment can result in unintended changes to the parameter that is being estimated depending on how the covariates are introduced in the estimation, a topic we revisit in [Section 8.3](#).

The standard local linear estimator of the RD treatment is implemented by running the weighted least squares regression of Y_i on a constant, T_i , X_i , and $X_i T_i$ using only units with scores inside the chosen bandwidth, $X_i \in [c - h, c + h]$, and applying weights based on some kernel function. This leads to a point estimator, which can be interpreted as MSE-optimal if the bandwidth employed is chosen to minimize the mean squared error of the RD point estimator. This estimator includes only the score X_i , and thus it is said to be “unadjusted” because it does not incorporate any additional pre-treatment characteristics of the units. As discussed above, under standard continuity assumptions, the unadjusted local linear estimator $\hat{\tau}_{\text{SRD}}$ is consistent for the continuity-based RD treatment effect τ_{SRD} , and robust bias-corrected inference can be employed as is now standard in the literature.

In practice, it is common for researchers to adjust the RD estimator by augmenting the local polynomial specification with the additional pre-determined covariates Z_i . [17] show that augmenting the local polynomial specification by adding covariates in an additive-separable, linear-in-parameters way that imposes the same common coefficient for treated and control groups leads to a covariate-adjusted RD estimator that remains consistent for the canonical sharp RD treatment effect τ_{SRD} , while offering a reduction in its variance in large samples (see also [18]). The key required restriction, of course, is that the covariates are pre-intervention, and that they have non-zero correlation with the outcome. More specifically, whenever the effect of the additional pre-intervention covariates on the potential outcomes near the cutoff is (approximately) the same in the control and treatment groups, augmenting the specification with these covariates can lead to efficiency gains. Furthermore, the authors also show that other approaches commonly used in the applied RD literature for regression-based covariate adjustment can lead to inconsistent estimators of the parameter of interest τ_{SRD} .

More recently, [19] proposed a high-dimensional implementation of the covariate-adjusted local polynomial regression approach of [17], which allows for selecting a subset of pre-intervention covariates out of a potentially large covariate pool. Their method is based on penalization techniques (Lasso) and can lead to further efficiency improvements in practice. More generally, combining modern high-dimensional and machine learning methods can be a fruitful avenue for further efficiency improvements via covariate adjustment in RD settings. For example, non-parametric or machine learning adjustments can be applied directly to the outcome variable to then employ the residualized outcomes in the subsequent local polynomial analysis, a procedure that can lead to further efficiency and power improvements.

In sum, our discussion highlights several important issues. First, standard methods for covariate adjustment in the analysis of experiments can in principle be applied to RD designs with suitable modifications depending on the specific conceptual framework used. These methods and their properties apply immediately to RD estimation and inference in a local randomization framework, within the window \mathcal{W} where the local randomization assumption is assumed to hold. The extension to continuity-based RD analysis is less immediate, as it requires considering the properties of the specific techniques at boundary points in a non-parametric local polynomial regression setup. However, there is a critical conceptual point: when using covariates for efficiency gains in RD designs, adjusting for those covariates should not change the RD point estimate. Regardless of the methods considered, covariate adjustment should not change the parameter of interest, and therefore unadjusted and adjusted point estimators should be similar in applications.

Second, pre-intervention covariates can also affect tuning parameter selection and implementation of the RD design more broadly. For example, in the local randomization framework, covariates can be used to select the window where the key assumptions hold and to relax those assumptions, while in the continuity-based framework they can be used to further refine bandwidth selection for local polynomial estimation and robust bias-corrected inference.

Finally, while augmenting RD analysis with covariates to increase efficiency is a principled goal, researchers should be careful to avoid using covariate adjustment for specification searches. In empirical studies, researchers should always report unadjusted RD results first, and covariate-adjusted results second. Covariate adjustment can be implemented in many different ways, and different adjustments might lead to different results and conclusions. Instead of trying multiple ways of covariate adjustment at the time of analysis, researchers should pre-register their preferred adjustment method, and perform only that specification as opposed to trying multiple regression models. We return to this important point when discussing recommendations for practice in [Section 8.4](#).

8.2.3 Auxiliary information

In the context of RD designs, covariate adjustment has also been proposed to address (i) missing data and measurement error, and (ii) to incorporate prior information via Bayesian methods.

A common concern in RD designs is that the RD running variable may be measured with error [20], and several approaches employing covariates have been recently proposed to address this issue. Generally speaking, key ideas from the imputation literature can be applied with appropriate modifications to RD designs. In the continuity-based framework, [21] study the case where the RD score is measured with error and this error changes across different groups of observations, and the researcher has access to auxiliary data that can be used to learn about the group-specific measurement error. The general procedure is to replace the mismeasured running variable by a “corrected” running variable of the same dimension that is estimated using the auxiliary covariates. Similarly, [22] study a fuzzy RD design where the running variable is measured with error, but the researcher has access to an additional covariate – the true running variable – that is observed for a sample of treated individuals. With this auxiliary data, the RD treatment effect at the cutoff can be identified. The procedure thus leads to heterogeneity tests that can be directly employed in empirical applications to assess whether the treatment affects different populations differently. Finally, in the local randomization framework, conventional methods can be used based on the standard missing-at-random assumption, where the missingness indicator and the potential outcomes are assumed to be independent conditional on pre-determined covariates. This assumption and associated methods can be applied directly within the window \mathcal{W} to justify the use of imputation approaches.

The use of covariates has also been proposed to incorporate information via Bayesian approaches to RD analysis – which fall in a separate category, different from both local randomization and continuity-based RD approaches. [23] propose an infinite-mixture Bayesian model that allows the density of the outcome variable to depend flexibly on covariates. At a minimum, these covariates

include the running variable and the treatment indicator, but other covariates can be added to the specification. A somewhat more general use of pre-determined covariates in Bayesian approaches to RD estimation is to use them to inform priors. Here, researchers sometimes use pre-treatment information rather than specific covariates. For example, in their study of the effect of prescribing statins to patients with high cardiovascular disease risk scores on their future LDL cholesterol levels, [24] use data on previous experimental studies to formulate informative priors in a Bayesian approach to estimate RD effects. Finally, in a related approach, [25] use principal stratification in a Bayesian framework to estimate complier average treatment effects in fuzzy RD designs. They model the imperfect compliance near the cutoff using a discrete confounder variable that captures different types of subjects: compliers, never-takers, and always-takers. Their setup leads to a mixture model averaged over the unknown subject types. The compliers' potential outcomes are modeled as a function of observed pre-determined covariates and the running variable, plus an idiosyncratic shock that follows a t-distribution. Given a prior distribution on the type probabilities, the posterior distribution of the complier average treatment effect is obtained by Markov chain Monte Carlo (MCMC) methods.

These types of methods have not been widely adopted in practice, but they provide useful examples of the different roles that covariate adjustment methods can have in RD applications. Importantly, in all these methods, the goal is ultimately to identify, estimate and conduct inference on the canonical (sharp and other) RD treatment effects, that is, without changing the parameter of interest. In the upcoming sections, we discuss covariate adjustment methods that do change the estimand and therefore require careful interpretation.

8.2.4 Treatment effect heterogeneity

Another important use of covariates for RD analysis is to assess whether the treatment has different effects for different subgroups of units, where the subgroups are defined in terms of observed values of pre-determined covariates. The use of covariates to explore heterogeneity has a long tradition in the analysis of both experimental and non-experimental data, as researchers are frequently interested in assessing the effects of the treatment for different subpopulations. From a broader perspective, this analysis is equivalent to estimating treatment effects conditional on pre-treatment characteristics, which can typically be implemented using standard non-parametric or machine learning methods.

In RD designs, and regardless of the conceptual framework employed, the analysis of treatment effect heterogeneity necessarily changes the parameter of interest. For example, in the continuity-based framework, the parameter of interest is no longer τ_{SRD} , but rather

$$\tau_{\text{CSR D}}(z) = \mathbb{E}[Y_i(1) - Y_i(0) | X_i = c, Z_i = z],$$

where Z_i denotes the additional pre-intervention covariate. The parameter $\tau_{\text{CSR D}}(z)$ corresponds to the (conditional, sharp) RD treatment effect at the cutoff for the subpopulation with covariates $Z_i = z$. Naturally, under appropriate assumptions, the canonical sharp RD treatment effect τ_{SRD} becomes a weighted average of the conditional RD treatment effects, with weights related to the conditional distribution of $Z_i | X_i$ near the cutoff. Similar ideas apply to the local randomization framework, depending on the specific assumptions invoked.

When the covariates Z_i are few and discrete, the simplest strategy to explore heterogeneity is to conduct the RD analysis for each subgroup defined by $Z_i = z$, for each value z , separately. For example, in a medical experiment, researchers may be interested in separately estimating treatment effects for patients in different age groups, e.g., patients between 45 and 64 versus patients 65 and above. In such cases separate analyses by subgroups have the advantage of being entirely non-parametric and involving no additional assumptions. One disadvantage is that the number of observations in each subgroup is lower than in the overall analysis, which reduces statistical power. Whether this is a limitation will depend on the number of observations and other features

of the data generating process in each particular application. Practically, this type of analysis can also be implemented by generating indicator variables for each category of the pre-intervention covariates, and then conducting estimation and inference using a fully saturated, interacted local polynomial regression model; if the indicators cover all possible subgroups, this strategy is equivalent to estimating effects separately by subgroup. This approach is also valid with respect to estimation and inference, and can be deployed in both the continuity-based and local randomization frameworks. However, if the covariates are multi-valued or continuous, using an interactive model imposes additional parametric assumptions, local to the cutoff, and the validity of inferences will depend on the validity of these assumptions.

If the number of covariates is large and researchers wish to explore heterogeneity in multiple dimensions, estimating separate average treatment effects by subgroups may be infeasible or impractical. In this case, researchers are usually interested in exploring heterogeneity in a large number of covariate partitions without specifying subgroups a priori. Standard non-parametric or machine learning methods can be useful in this situation, as they allow researchers to learn the relevant subgroups from the data. Machine learning methods in the continuity-based RD framework require modifications and have begun to be explored only recently. In particular, [26] explores heterogeneity of the RD effect in subpopulations defined by levels of pre-determined covariates, creating a tree where each leaf contains an RD effect estimated on an independent sample. The approach assumes a parametric q -th order polynomial in each leaf of the tree; under this parametric assumption and additional regularity assumptions, the method leads to subgroup point estimates and standard errors that can be used for inference.

A different continuity-based approach to examine RD treatment effect heterogeneity on different subpopulations is proposed by [27]. The authors impose stronger continuity conditions than in the standard RD design, requiring (in the sharp RD case) continuity of the expectation of the potential outcomes conditional on both the running variable and the additional covariates, and continuity of the conditional distribution of the additional covariates given the running variable. Under these conditions, they derive conditional average treatment effects given the covariates, and propose tests for three hypothesis: (i) that the treatment is beneficial for at least some subpopulations defined by values of the pre-determined covariates, (ii) that the treatment has any impact on at least some subpopulations, and (iii) that the treatment effect is heterogeneous across all subpopulations.

[27] define the null and alternative hypotheses by conditional moment inequalities given both the running variable and the additional covariates, and convert these conditional moment inequalities into an infinite number of unconditional-on-covariates moment inequalities (that is, inequalities that are conditional only on the running variable and not on the additional covariates). Once the null and the alternative hypotheses are re-defined as instrumented moment conditions, these are estimated with local linear polynomials. Under regularity conditions, the asymptotic distribution of the local polynomial estimators of the moment conditions can then be used to derive the distribution of test statistics under the null hypothesis. In a follow-up paper, [28] employ similar methods to develop a monotonicity test to assess whether a conditional local average treatment effect in a sharp RD design or a conditional local average treatment effect for compliers in a fuzzy RD design has a monotonic relationship with an observed pre-determined covariate – that is, the null hypothesis is that the conditional local average treatment effect, seen as a function of a covariate z , is non-decreasing in z .

Covariate adjustments for heterogeneity analysis within the local randomization framework may be more challenging due to limited sample sizes. Since the local randomization framework is generically more appropriate for very small windows \mathcal{W} around the cutoff, and these windows typically contain only a few observations, conducting estimation and inference for subsets of observations according to $Z_i = z$ can be difficult in practice, leading to limited statistical power.

In sum, there are several important issues when using covariate adjustment for the analysis of heterogeneous RD treatment effects. First, focusing on treatment effect heterogeneity requires that researchers redefine the parameter of interest. Second, estimating different RD treatment effects for

different subgroups defined by covariate values can be done in a non-parametric way by estimating fully saturated models – or equivalently, by analyzing each sub-group separately. When the covariate dimensionality is too high and sub-group analysis is not possible, machine learning methods may offer an attractive strategy to discover relevant subgroups. Finally, augmented models that include interactions between the RD treatment and non-binary covariates require additional parametric assumptions for identification and inference.

8.2.5 Other parameters of interest and extrapolation

Covariate adjustment has also been proposed in RD designs to identify additional parameters of interest. This includes cases where covariates are used to define new parameters for subpopulations, similar to the case of heterogeneity analysis discussed above, and cases where covariates are used to extrapolate treatment effects for score values far from the cutoff.

In some RD designs, researchers have access to special covariates, and in exploring heterogeneity along those covariates, they redefine the parameter of interest. For example, it is common in practice to see treatments that follow the RD assignment rule but use different cutoffs for different subpopulations of units, that is, $T_i = \mathbb{1}(X_i \geq c_j)$ for $j = 1, 2, \dots, J$. In this case, the cutoffs c_1, c_2, \dots, c_J naturally define subpopulations of units, where each subpopulation is defined by the cutoff that those units are exposed to [29] introduced the Multi-cutoff RD design, and investigated how multiple cutoffs can be used to provide more information about RD effects. In particular the heterogeneity of the average treatment effect across cutoffs can be explored directly by performing RD analysis separately for each subgroup of units exposed to each cutoff value, as discussed in the prior section. The authors also discuss the role of the commonly used normalizing-and-pooling strategy, which leads to a causal interpretation of the canonical RD treatment effect as a weighted average of the cutoff-specific RD treatment effects under additional assumptions. See [30] for an empirical application.

In some cases researchers may prefer to use the information provided by the multiple cutoffs in combination with additional assumptions to define new parameters of interest. For example, [31] studies a Multi-Cutoff RD setup where each individual is exposed to a particular cutoff and a particular treatment dosage that changes as a function of that cutoff. Under additional assumptions, the author uses the changes in treatment dosage applied to individuals at various levels of the running variable to explore average treatment effects for counterfactual policies that set the treatment dosage to levels not observed in the data. [32] combine a Multi-Cutoff RD design with additional assumptions to define the average treatment effect at a particular cutoff for a subpopulation of units exposed to a different cutoff value. Both approaches lead to extrapolation of RD treatment effects – that is, to parameters that capture the average effect of the treatment at values of the score different from the original cutoff to which units were exposed.

Another common setup involves an RD design observed in two periods of time, where the treatment of interest is confounded by an additional treatment in one period but not the other [33] call this setup a differences in discontinuities design. In the first period, there are two treatments that are assigned with the same RD rule, that is, the probability of receiving each treatment goes from zero to one at the same cutoff. In the second period, only one treatment changes at the cutoff, but the other treatment – which is the treatment of interest—is not active. The authors propose an estimator that takes the difference between the first-period RD effect and the second-period RD effect. Whether the limit of this estimator is the standard RD parameter will depend on the particular assumptions imposed. Under the usual parallel trends assumption of difference-in-differences designs, the estimand will remain unchanged. However, researchers could also make different assumptions about the change of the regression functions before and after, and redefine the parameters accordingly.

When multiple cutoffs or time periods are not available, extrapolation of treatment effects can be done by using auxiliary pre-determined covariates. In this setting, with additional assumptions, researchers can learn about the treatment effect for units whose scores are not in the immediate neighborhood of the cutoff. For example, [34] propose a conditional independence assumption to

study the effect of the treatment for observations far from the cutoff. This is formalized by assuming that, conditional on the covariates, the potential outcomes are mean-independent of the running variable. This “selection on observables” assumption (together with the standard common support assumption) immediately allows for extrapolation of the RD treatment effect, since it is assumed to hold for the entire support of the running variable. In this case covariates allow for the identification of new parameters that capture the effect of the treatment at different values or regions of the running variable.

A similar conditional independence assumption has been used by in the context of a geographic or multi-score RD designs. In a geographic RD design, a treatment is assigned to units in a geographic area and withheld from units in an adjacent geographic area. In this setup, units can be thought of as having a score (such as a latitude-longitude pair) that uniquely defines their geographic location and allows the researcher to calculate their distance to any point on the boundary between the treated and control areas. The assignment of the treatment can then be viewed as a deterministic function of this score, and the probability of receiving treatment jumps discontinuously at the border that separates the treated and control areas. As discussed by [35], seen in this way, the geographic RD design can be analyzed with the same continuity-based and local randomization tools of standard, one-dimensional RD designs [36] consider a geographic RD application where units appear to choose their location on either side of the boundary in a strategic way. They propose a conditional independence assumption according to which, for each point on the boundary, the potential outcomes are independent of the treatment assignment conditional on a set of covariates for units located in a neighborhood of that point (see also [37] for an extension to multiple dimensions). An important difference between their assumption and the conditional independence assumption proposed by [34] is that the former imposes conditional independence in a neighborhood of the cutoff and thus allows for extrapolation only within a neighborhood, while the latter imposes (mean) conditional independence along the entire support of the running variable and thus allows for extrapolation in the entire support of the running variable. Despite the differences, in both cases covariates are used for identification purposes, to define new parameters of interest that capture the average treatment effect at values of the running variable that are different from the cutoff.

Another use of auxiliary covariates for extrapolation purposes is proposed by [38], who augment the usual RD design with an exogenous measure of the outcome variable. Under the assumption that the exogenous outcome data can be used to consistently estimate the regression function of the actual outcome on the score in the absence of the treatment, treatment effects can be extrapolated to values of the score other than the actual cutoff used for assignment.

In sum, additional covariates can be used in the RD design to define new parameters of interest, including parameters capturing treatment effects for units with score values different from the original cutoff. These covariates are sometimes part of the original RD design, as in the case of multiple cutoffs or multiple time periods, while in other settings they are external to the design and must be collected separately, as in the case of auxiliary unit characteristics or exogenous outcomes.

8.3 Can Covariates Fix a Broken RD Design?

In a randomized experiment the fact that the treatment is randomly assigned implies that the treatment variable D_i is independent of both potential outcomes $(Y_i(1), Y_i(0))$ and predetermined covariates Z_i . This means that, in experimental settings, we have independence between potential outcomes and treatment assignment, $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i$, and also conditional independence between potential outcomes and treatment assignment given the covariates, $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i | Z_i$. At the same time, it is well known that the assumption of conditional independence given covariates alone, $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i | Z_i$, does not imply $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i$. In observational studies researchers

often assume that $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i | Z_i$ in order to identify treatment effects. This assumption directly justifies covariate adjustment via regression models, inverse probability weighting, and other matching methods.

This observation has motivated some researchers to employ covariate adjustment to “fix” invalid RD designs where the identifying assumptions are suspected not to hold. In these cases, it is common to encounter RD analyses that include fixed effects or other covariates via additive linear adjustments in (local) polynomial regression estimation. In the methodological RD literature, adjustment for covariate imbalances has been proposed using multi-step non-parametric regression [39] and inverse probability weighting [40].

In the continuity-based framework, covariate adjustment for imbalances in RD designs is done within a framework where $\mathbb{E}[Y_i(t) | X_i = x]$ is assumed to be discontinuous at the cutoff but $\mathbb{E}[Y_i(t) | X_i = x, Z_i = z]$, seen as a function of x , is assumed continuous in x , for $t = 0, 1$. These assumptions combined imply that the conditional distribution of $Z_i | X_i$ must be discontinuous at the cutoff, which in turn implies that canonical average RD treatment effects are necessarily not identifiable in general. Rather, in the absence of additional strong assumptions, the estimand that can be identified in this case is a weighted average of treatment effects where the weights depend on the conditional distributions of $Z_i | X_i$ for control and treatment units, which are different from each other [17, 39]. As a consequence, in imbalanced RD designs, canonical RD treatment effects are not identifiable in general, and by implication covariate adjustment cannot fix a broken RD design. In the local randomization framework, the situation is analogous.

In conclusion, for invalid RD designs where important covariates are imbalanced at the cutoff, covariate adjustment cannot restore identification of the canonical RD treatment effect. At best, covariate adjustment can identify some weighted average of the conditional treatment effects, $\mathbb{E}[Y_i(1) | X_i = x, Z_i = z] - \mathbb{E}[Y_i(0) | X_i = x, Z_i = z]$, with weights determined by the properties of the two distinct conditional distributions of $Z_i | X_i$ below and above the cutoff. Unfortunately, in most applications, this covariate-adjusted estimand is not a parameter of interest and requires strong, additional assumptions on the underlying data generating process to deliver a useful parameter of interest. In other words, the main advantages of canonical RD designs in terms of identification are lost once the key identifying assumptions fail, and cannot be restored by simply using covariate adjustment methods without additional untestable assumptions.

8.4 Recommendations for Practice

When employed correctly, baseline covariates can be useful for the analysis and interpretation of RD designs. When pre-intervention covariates are available, they can be used for two main purposes without affecting the main RD identification strategy: (i) to improve efficiency and power, and (ii) to define new parameters of interest. However, analysts should keep in mind that canonical RD designs do not necessitate covariate adjustments, and therefore researchers should always report unadjusted RD treatment effect estimates and associated unadjusted robust bias-corrected inference methods.

Covariate adjustment for efficiency and power improvement is a natural way of incorporating baseline covariates in the RD analysis. This can be done by using regression adjustments, or via other well-known approaches from the experimental literature such as outcome adjustments, inverse probability weighting, or matching methods. Importantly, because the treatment effect of interest is the same with and without covariate adjustment, researchers should always report both unadjusted and covariate-adjusted RD results. The main and only goal of this practice is to improve efficiency and power (e.g., shorten confidence intervals), but the RD point estimate should remain unaffected, which should be borne out by the empirical analysis. Furthermore, it is important to guard against p-hacking in this context. Researchers should be cognizant of the dangers of multiple hypothesis

testing when incorporating covariate adjustment in the RD analysis. One possible solution to address this concern is pre-registration of the analysis, which requires practitioners to declare ex-ante the covariate adjustment approach to be used in the subsequent RD analysis in order to avoid ex-post specification searching.

Pre-intervention covariates can also be employed for identification of other useful RD treatment effects. In [Section 8.2](#), we discussed different examples such as treatment effect heterogeneity, multi-cutoff, geographic and multi-score designs, and extrapolation. These approaches offer relatively principled ways of incorporating covariates into an RD analysis. In the case of treatment effect heterogeneity, no substantial additional assumptions are needed, particularly in cases where the covariates are discrete and low-dimensional, in which case subset analysis is a natural way to proceed empirically. In the other cases, additional assumptions are needed in order to exploit covariate adjustment for identification of other meaningful RD treatment effects. These covariate adjustment practices are reasonable, and can be used in applications whenever the necessary additional assumptions are clearly stated and judged to be plausible. Again, it is important to guard against model specification searching, just like when employing pre-intervention covariates for efficiency and power improvements.

Finally, an important misconception among some practitioners and methodologists is that covariates can be used to restore identification or “fix” an RD design where observations just above the cutoff are very different from observations just below it – in other words, a design where there is covariate imbalance at or near the cutoff and the RD assumptions are not supported by the empirical evidence. However, as we discussed in [Section 8.3](#), covariate adjustment in such cases requires strong additional assumptions, and cannot in general recover canonical RD treatment effects. In the absence of additional assumptions, adjusting for imbalanced covariates will at best recover other estimands at the cutoff that are unlikely to be of substantive interest in applications.

8.5 Conclusion

We provided a conceptual overview of the different approaches for covariate adjustment in RD designs. We discussed benign approaches based on pre-intervention covariates, including methods for efficiency and power improvements, heterogeneity analysis, and extrapolation. Those methods are principled and generally valid under additional reasonable assumptions on the data generating process. However, we also highlighted a natural tension between incorporating covariates in an RD analysis and issues related to p-hacking and specification searching. As a consequence, researchers employing covariate adjustment in RD designs should always be cognizant of issues related to model/covariate selection when implementing those methods in practice. Last but not least, we addressed the common misconception that covariate adjustment can fix an otherwise invalid RD design, and discussed how those methods can at best recover other estimands that may not be of interest.

In conclusion, covariate adjustment can be a useful additional tool for the analysis and interpretation of RD designs when implemented carefully and in a principled way. In canonical RD settings it should never replace, only complement, the basic RD analysis based on the score and outcome variables alone. In other settings, covariate adjustment can offer additional insights in terms of heterogeneous or other treatment effects of interest, under additional design-specific assumptions.

8.6 Acknowledgments

We thank our collaborators Sebastian Calonico, Max Farrell, Michael Jansson, Xinwei Ma, Gonzalo Vazquez-Bare, and Jose Zubizarreta for stimulating discussions on the topic of this chapter. Cattaneo and Titiunik gratefully acknowledge financial support from the National Science Foundation (SES-2019432), and Cattaneo gratefully acknowledge financial support from the National Institutes of Health (R01 GM072611-16).

References

- [1] Paul R. Rosenbaum. *Design of Observational Studies*. Springer-Verlag, 2010.
- [2] Guido W Imbens and Donald B Rubin. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015.
- [3] Alberto Abadie and Matias D. Cattaneo. Econometric methods for program evaluation. *Annual Review of Economics*, 10:465–503, 2018.
- [4] Rocio Titiunik. Natural experiments. In J. N. Druckman and D. P. Gree, editors, *Advances in Experimental Political Science*, chapter 6, pages 103–129. Cambridge University Press, 2021.
- [5] Matias D. Cattaneo, Rocio Titiunik, and Gonzalo Vazquez-Bare. The regression discontinuity design. In L. Curini and R. J. Franzese, editors, *Handbook of Research Methods in Political Science and International Relations*, chapter 44, pages 835–857. Sage Publications, 2020.
- [6] Matias D. Cattaneo, Nicolás Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press, 2020.
- [7] Matias D. Cattaneo, Nicolás Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Extensions*. Cambridge University Press (to appear), 2023.
- [8] Matias D. Cattaneo, Luke Keele, and Rocio Titiunik. A guide to regression discontinuity designs in medical applications. Working Paper, 2022.
- [9] Matias D. Cattaneo and Rocio Titiunik. Regression discontinuity designs. *Annual Review of Economics*, 14:821–851, 2022.
- [10] Matias D. Cattaneo, Rocio Titiunik, and Gonzalo Vazquez-Bare. Comparing inference approaches for rd designs: A reexamination of the effect of head start on child mortality. *Journal of Policy Analysis and Management*, 36(3):643–681, 2017.
- [11] Jinyong Hahn, Petra Todd, and Wilbert van der Klaauw. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209, 2001.
- [12] Sebastian Calonico, Matias D. Cattaneo, and Max H. Farrell. Optimal bandwidth choice for robust bias corrected inference in regression discontinuity designs. *Econometrics Journal*, 23(2):192–210, 2020.
- [13] Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014.

- [14] Matias D. Cattaneo, Brigham Frandsen, and Rocio Titiunik. Randomization inference in the regression discontinuity design: An application to party advantages in the u.s. senate. *Journal of Causal Inference*, 3(1):1–24, 2015.
- [15] Donald P. Green Green and Alan Gerber. *Field Experiments: Design, Analysis, and Interpretation*. W. W. Norton & Company, 2012.
- [16] Matias D. Cattaneo, Rocio Titiunik, and Gonzalo Vazquez-Bare. Power calculations for regression discontinuity designs. *Stata Journal*, 19(1):210–245, 2019.
- [17] Sebastian Calonico, Matias D. Cattaneo, Max H. Farrell, and Rocio Titiunik. Regression discontinuity designs using covariates. *Review of Economics and Statistics*, 101(3):442–451, 2019.
- [18] Jun Ma and Zhengfei Yu. Empirical likelihood covariate adjustment for regression discontinuity designs. arXiv:2008.09263, 2022.
- [19] Yoichi Arai, Taisuke Otsu, and Myung Hwan Seo. Regression discontinuity design with potentially many covariates. arXiv:2109.08351, 2021.
- [20] Zhuan Pei and Yi Shen. The devil is in the tails: Regression discontinuity design with measurement error in the assignment variable. In Matias D. Cattaneo and Juan Carlos Escanciano, editors, *Regression Discontinuity Designs: Theory and Applications (Advances in Econometrics, volume 38)*, pages 455–502. Emerald Group Publishing, 2017.
- [21] Otavio Bartalotti, Quentin Brummet, and Steven Dieterle. A correction for regression discontinuity designs with group-specific mismeasurement of the running variable. *Journal of Business & Economic Statistics*, 39(3):833–848, 2021.
- [22] Laurent Davezies and Thomas Le Barbanchon. Regression discontinuity design with continuous measurement error in the running variable. *Journal of Econometrics*, 200(2):260–281, 2017.
- [23] George Karabatsos and Stephen G Walker. A bayesian nonparametric causal model for regression discontinuity designs. In R. Mitra and P. Müller, editors, *Nonparametric Bayesian Inference in Biostatistics*, pages 403–421, 2015.
- [24] Sara Geneletti, Aidan G O’Keeffe, Linda D Sharples, Sylvia Richardson, and Gianluca Baio. Bayesian regression discontinuity designs: Incorporating clinical knowledge in the causal analysis of primary care data. *Statistics in Medicine*, 34(15):2334–2352, 2015.
- [25] Siddhartha Chib and Liana Jacobi. Bayesian fuzzy regression discontinuity analysis and returns to compulsory schooling. *Journal of Applied Econometrics*, 31(6):1026–1047, 2016.
- [26] Agoston Reguly. Heterogeneous treatment effects in regression discontinuity designs. *arXiv preprint arXiv:2106.11640*, 2021.
- [27] Yu-Chin Hsu and Shu Shen. Testing treatment effect heterogeneity in regression discontinuity designs. *Journal of Econometrics*, 208(2):468–486, 2019.
- [28] Yu-Chin Hsu and Shu Shen. Testing monotonicity of conditional treatment effects under regression discontinuity designs. *Journal of Applied Econometrics*, 36(3):346–366, 2021.
- [29] Matias D. Cattaneo, Luke Keele, Rocio Titiunik, and Gonzalo Vazquez-Bare. Interpreting regression discontinuity designs with multiple cutoffs. *Journal of Politics*, 78(4):1229–1248, 2016.

- [30] Yasin Kursat Onder and Mrittika Shamsuddin. Heterogeneous treatment under regression discontinuity design: Application to female high school enrolment. *Oxford Bulletin of Economics and Statistics*, 81(4):744–767, 2019.
- [31] Marinho Bertanha. Regression discontinuity design with many thresholds. *Journal of Econometrics*, 218(1):216–241, 2020.
- [32] Matias D Cattaneo, Luke Keele, Rocio Titiunik, and Gonzalo Vazquez-Bare. Extrapolating treatment effects in multi-cutoff regression discontinuity designs. *Journal of the American Statistical Association*, 116(536):1941–1952, 2021.
- [33] Veronica Grembi, Tommaso Nannicini, and Ugo Troiano. Do fiscal rules matter? *American Economic Journal: Applied Economics*, 8(3):1–30, 2016.
- [34] Joshua D Angrist and Miikka Rokkanen. Wanna get away? regression discontinuity estimation of exam school effects away from the cutoff. *Journal of the American Statistical Association*, 110(512):1331–1344, 2015.
- [35] Luke J Keele and Rocio Titiunik. Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1):127–155, 2015.
- [36] Luke Keele, Rocio Titiunik, and Jose R Zubizarreta. Enhancing a geographic regression discontinuity design through matching to estimate the effect of ballot initiatives on voter turnout. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 178(1):223–239, 2015.
- [37] Juan D Diaz and Jose R Zubizarreta. Complex discontinuity designs using covariates for policy impact evaluation. *Annals of Applied Statistics*, forthcoming, 2022.
- [38] Coady Wing and Thomas D Cook. Strengthening the regression discontinuity design using additional design elements: A within-study comparison. *Journal of Policy Analysis and Management*, 32(4):853–877, 2013.
- [39] Markus Frolich and Martin Huber. Including covariates in the regression discontinuity design. *Journal of Business & Economic Statistics*, 37(4):736–748, 2019.
- [40] Sida Peng and Yang Ning. Regression discontinuity design under self-selection. In *International Conference on Artificial Intelligence and Statistics*, pages 118–126. PMLR, 2021.