Regression Discontinuity Designs*

Matias D. Cattaneo† Rocio Titiunik‡

January 11, 2024

The regression discontinuity (RD) design is a research strategy employed to study the causal effect of a treatment (e.g., intervention or policy) on an outcome of interest using observational data. The design was first introduced by Thistlethwaite and Campbell (1960) to analyze the effect of public recognition of students' achievement on their attitudes towards education, where the receipt of public recognition was given only to students who scored above a cutoff on a qualifying test. RD designs identify causal effects by exploiting the abrupt change in treatment assignment induced by a discontinuous allocation rule; for example, in the original application, only students who scored above a cutoff on a qualifying test were eligible to receive the treatment. Cattaneo and Titiunik (2022) give a comprehensive review of the RD literature, and recent practical introductions to RD methods can be found in Cattaneo, Idrobo and Titiunik (2020, 2024) and Cattaneo, Keele and Titiunik (2023). Popular general-purpose statistical software implementing RD methods is available at `https://rdpackages.github.io/`.

The RD design belongs to the class of non-experimental methods (e.g., selection-on-observables, instrumental variables, difference-in-differences), but it is often referred to as a quasi-experimental method because it can resemble a randomized control experiment. RD designs are also often described as "local" randomized experiments (Lee and Lemieux, 2010), and thus viewed as natural experiments (Titiunik, 2021). Due to its popularity and wide applicability, there are many variants of the RD design, but at their core they all share the same three distinctive elements: a score, a cutoff, and a discontinuous (as a function of the score) treatment assignment rule that relates the score and cutoff. These elements differentiate RD designs from all other non-experimental methods: each unit in the study receives a score, and the treatment is only offered to those units whose scores are equal to or larger than the cutoff, but not offered to the units whose scores are smaller than the cutoff. Mathematically, suppose there are n units in the study, indexed by $i = 1, 2, \dots, n$, and denote the score of each unit by X_i, then the RD treatment assignment rule is $T_i = \mathbb{1}(X_i \geq c)$, where $\mathbb{1}(A)$ is the indicator function that is equal to one when A is true and zero otherwise, and c is the known (to the researcher at least) fixed cutoff.

*Entry prepared for the 2024 Encyclopedia of Experimental Social Science, Elgar Encyclopedias in the Social Sciences series. Cattaneo and Titiunik gratefully acknowledge financial support from the National Science Foundation (SES-2241575).

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

‡Department of Politics, Princeton University.

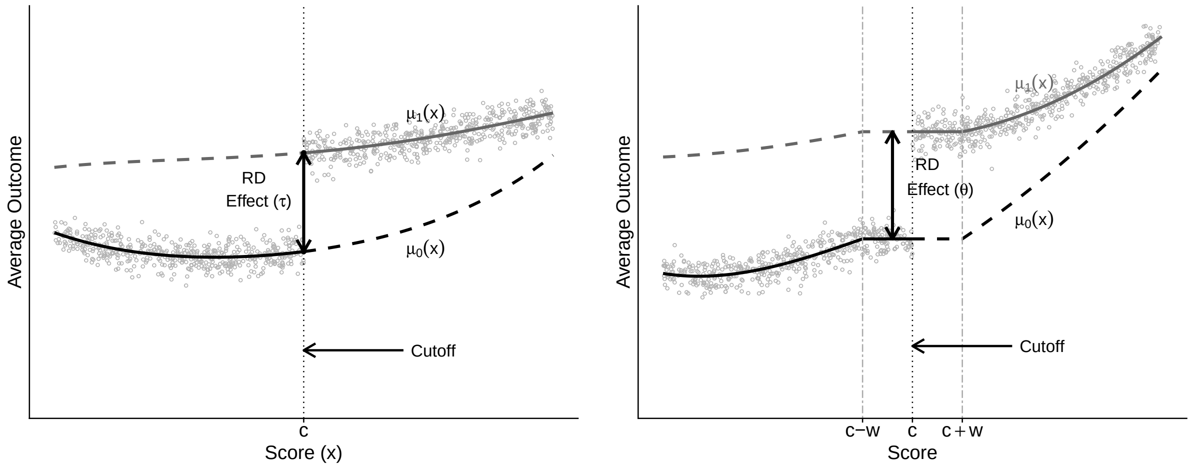

(a) *Continuity-Based Framework* **(b)** *Local Randomization Framework*

Figure 1: *Schematic Illustration of RD design*

Note: $\mu_1(x)$ and $\mu_0(x)$ are the average outcomes under treatment and control, respectively, seen as functions of the score. Dashed lines represent unobserved outcomes, solid lines represent observed outcomes.

In most applications, the score is correlated with the outcome of interest, which means that treated and control units are systematically different: a comparison of all units below the cutoff (control units) to all units above the cutoff (treated units) cannot identify a causal treatment effect. For example, in the original application, students with high test scores may have more optimistic attitudes towards education regardless of whether they receive the treatment. The RD design works around this selection problem by leveraging the identifying assumption that any systematic differences between treated and control units in the absence of the treatment disappear once attention is restricted to units whose scores are "near" the cutoff: a causal treatment effect can be recovered by comparing (i) the outcomes of treated units with scores barely above the cutoff to (ii) the outcomes of control units with scores barely below the cutoff. While units with extreme values of the score will be systematically different, the RD design only focuses on units whose scores are sufficiently close to the cutoff determining treatment assignment so that they are comparable with each other in all respects, except for the fact that those with score barely above the cutoff are assigned to treatment while those with score barely below the cutoff are not.

The most basic RD design is called the Sharp RD design, and corresponds to the case when compliance with the treatment assignment is perfect: all units assigned to the treatment condition actually receive the treatment, and no units assigned to the control condition receive the treatment. In this case, the defining feature of the design is that the probability of receiving the treatment jumps from exactly zero to exactly one at the cutoff. There are two ways to formalize the assumption of comparability near the cutoff in the Sharp RD design. In the continuity-based approach, first introduced by Hahn, Todd and van der Klaauw (2001), comparability takes the form of continuity at the cutoff, and relates to the idea of extrapolation using information from units whose scores are located near the cutoff. The continuity-based approach is illustrated in Figure 1: the average outcomes under treatment and control, seen as functions of the score and represented by $\mu_1(x)$ and $\mu_0(x)$, respectively, are assumed to be continuous on the score at the cutoff. This assumption implies that the average outcome for units with score values near the cutoff would have been roughly the same in the absence of the treatment, regardless of whether their scores were below or above the cutoff. Thus, under these continuity assumptions, a causal average effect of the treatment at the cutoff, represented by τ in Figure 1(a), is identifiable as the vertical difference between the average treated outcome at the cutoff and the average control outcome at the cutoff.

In the local randomization approach, the comparability of units near the cutoff is formalized by assuming that, for units with scores near the cutoff, the assignment of the treatment resembles the assignment that one would have seen in a randomized experiment that had assigned the treatment randomly among the subpopulation of units with scores near the cutoff. This analogy between the RD design and a randomized experiment was first noted by Thistlethwaite and Campbell (1960), then heuristically discussed by Lee (2008) and Lee and Lemieux (2010), and later formalized by Cattaneo, Frandsen and Titiunik (2015) and Cattaneo, Titiunik and Vazquez-Bare (2017). The local randomization assumption is illustrated in Figure 1(b), where the average outcome functions under treatment ($\mu_1(x)$) and control ($\mu_0(x)$) become constant for all values of the score in the window $[c-w, c+w]$ around the cutoff, where w is the half-width of the so-called local randomization window. The flatness of both average outcomes indicates that there is no relationship between these outcomes and the score inside the window, and therefore units with scores above the cutoff in this window are comparable to units with scores below the cutoff in this window. The local randomization causal RD treatment effect is the vertical distance between the average outcome under treatment and the average outcome under control inside the window, represented by θ in Figure 1(b).

Methods for estimation and inference differ depending on whether the continuity-based approach or the local randomization approach is adopted. In the continuity-based approach, the most common implementation uses local polynomial regression methods (Fan and Gijbels, 1996). This non-parametric method approximates the two functions $\mu_1(x)$ and $\mu_0(x)$ in Figure 1(a) separately for observations above and below the cutoff, fitting in each case a least squares regression of the observed outcome on a polynomial of the score of order p, and keeping the intercept of each fit. The point estimate is obtained by taking the difference between these two intercepts. The bandwidth that determines which observations are included in the fit can be chosen in different ways, depending on the optimality criterion used (Calonico, Cattaneo and Farrell, 2020). If the goal is to obtain an optimal point estimator, a common method is to choose the value of h that minimizes an approximation to the mean squared error (MSE) of the local polynomial point estimator. The MSE-optimal bandwidth balances the variance-bias trade-off: as h shrinks, the error or bias in the polynomial approximation decreases, but the variance increases because fewer observations are included, and vice versa when h increases. Alternatively, if the goal is to obtain optimal inference, the bandwidth can be chosen to minimize the coverage error of the confidence intervals based on the RD point estimator. See Cattaneo and Vazquez-Bare (2016) for more discussion on tuning parameter selection in RD designs.

The construction of confidence intervals based on local polynomials is affected by the error in the local polynomial approximation, that is, by how different the least squares fit is from the actual functional form of the average outcomes near the cutoff. If an MSE-optimal bandwidth is used, the approximation error affects the Gaussian approximation upon which conventional confidence intervals are based. A method to build valid confidence intervals in the presence of misspecification error is the robust bias correction approach proposed by Calonico, Cattaneo and Titiunik (2014). This method corrects the point estimate by subtracting an estimate of the approximation error, and simultaneously increases the variance to account for the additional variability introduced in the error correction estimation step. A practical discussion of continuity-based methods for RD estimation and inference is given by Cattaneo, Idrobo and Titiunik (2020). This approach to identification, estimation, and inference is the default methodology in empirical work.

In the local randomization approach, estimation and inference is instead implemented by following methods commonly used for the analysis of experiments (Rosenbaum, 2010). Implementation usually proceeds in two steps. In the first step, researchers use predetermined covariates (that is, variables not affected by the treatment) to choose a window near the cutoff such that the covariate distribution of the treated units in that window is similar to the covariate distribution of the control units in that window. If such a window exists, the average treatment effect in the window is estimated by the difference in means between treated and control units in that window, as is normally done in randomized experiments. For inference, both finite-sample and large-sample methods can be used. Finite-sample methods assume that units in the window were assigned to treatment with a known distribution, and use that distribution in combination with the so-called Fisherian null hypothesis to enumerate all the possible values that the test statistic could have had, which are then used to build an exact p-value. In contrast, large-sample methods rely on the usual Gaussian approximations to standard test statistics. For a practical discussion of these methods, see Section 2 in Cattaneo, Idrobo and Titiunik (2024). These methods are often used as complementary robustness checks in settings where the continuity-based approach is valid, but they are sometimes used in isolation when the canonical continuity RD assumptions are not satisfied. See Cattaneo, Keele and Titiunik (2023) for more practical discussion, and Hyytinen, Meriläinen, Saarimaa, Toivanen and Tukiainen (2018) for an interesting empirical application juxtaposing both approaches.

The continuity and local randomization assumptions are not implied by the RD treatment assignment rule. Instead, they must be invoked, and in cases where units have the ability to change their score to obtain the treatment condition that they most desire, these assumptions may not hold. Such active manipulation of the score by the units of analysis is the most important threat to the validity of any RD design. For this reason, it is common for researchers to include validation and falsification tests as part of their RD analysis, which include (i) studying the effects of the treatment on pre-determined covariates and placebo outcomes, both of which should be unaffected by the treatment, (ii) checking that the score density near the cutoff is continuous, and (iii) repeating estimation and inference using an artificial cutoff value and for different values of tuning parameters. The implementation of these falsification analyses differs depending on whether a continuity-based or a local randomization approach is used; see, respectively, Section 5 in Cattaneo, Idrobo and Titiunik (2020) and Section 2.3 in Cattaneo, Idrobo and Titiunik (2024).

The canonical Sharp RD design has been extended in multiple directions. The Fuzzy RD design arises when compliance with the treatment assignment is imperfect, so that some of the units assigned to the treatment condition fail to actually receive the treatment, and some units assigned to the control condition take the treatment anyway. In this case, the probability of receiving treatment continues to change discontinuously at the cutoff but, in contrast to the Sharp RD design, the change in the probability of receiving treatment at the cutoff is less than one. Local randomization and continuity-based approaches for RD estimation and inference with non-compliance parallel the instrumental variables approaches commonly used in experimental and non-experimental settings. For practical discussions, see Section 3 in Cattaneo, Idrobo and Titiunik (2024), and Cattaneo, Keele and Titiunik (2023).

Multi-dimensional RD designs are another common extension of the canonical RD setup. They arise when either the score has more than one dimension, or there is more than one cutoff determining treatment assignment. In the Multi-Score RD design, the score that determines treatment assignment has two or more components, a generalization that was developed by Papay, Willett and Murnane (2011) and Reardon and Robinson (2012), among others. For example, a scholarship may be given to students on the basis of two exams, so that only students who score above a qualifying cutoff in both a mathematics and an English exam are eligible to receive the scholarship. An important special case of the Multi-Score RD design is the Geographic RD design (Dell, 2010; Keele and Titiunik, 2015), where units reside in a geographic region that is split into two adjacent areas by a boundary, and the treatment is offered only to units residing in one of those areas. In this case, the score is a two-dimensional vector that captures the geographic coordinates of every unit in space, and the boundary between the treated and control areas collects all the points at which the treatment assignment changes abruptly from zero to one. The presence of a boundary separating treatment and control areas in two dimensions implies that many treatment effects of interest can be studied, including treatment effects at a specific boundary point, as well as treatment effects averaged along the entire boundary. See Section 4 in Cattaneo, Idrobo and Titiunik (2024) for more discussion.

The Multiple-Cutoff RD design is another popular extension that arises when multiple subpopulations of units are assigned to treatment according to an RD assignment rule, but each subpopulation uses a different cutoff value for assignment. This is common, for example, when social programs are assigned on the basis of need, and socio-economic conditions change by geographic region. A formalization of the Multiple-Cutoff RD design was given in Cattaneo et al. (2016), and a practical discussion can be found in Section 4 in Cattaneo, Idrobo and Titiunik (2024).

Yet another generalization of the RD design is known as the Regression Kink Design, which was developed by Card, Lee, Pei and Weber (2015). In this case, different levels of the treatment are given on the basis of a score according to a policy formula that changes slope at known cutoff values (known as kink points). Just like in the RD design we expect the level of the average outcome to change abruptly at the cutoff, in a Regression Kink design the expectation is that the first derivative of the observed outcomes will change abruptly at the kink point. A common example is tax rates that increase at certain cutoffs of baseline income. This and other generalizations, as well as a thorough discussion of the extensive RD methodological literature, is reviewed by Cattaneo and Titiunik (2022).

References

Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2020), "Optimal Bandwidth Choice for Robust Bias Corrected Inference in Regression Discontinuity Designs," *Econometrics Journal*, 23, 192–210.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014), "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82, 2295–2326.

Card, D., Lee, D. S., Pei, Z., and Weber, A. (2015), "Inference on Causal Effects in a Generalized Regression Kink Design," *Econometrica*, 83, 2453–2483.

Cattaneo, M. D., Frandsen, B., and Titiunik, R. (2015), "Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate," *Journal of Causal Inference*, 3, 1–24.

Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2020), *A Practical Introduction to Regression Discontinuity Designs: Foundations*, Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge University Press.

——— (2024), *A Practical Introduction to Regression Discontinuity Designs: Extensions*, Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge University Press.

Cattaneo, M. D., Keele, L., and Titiunik, R. (2023), "A guide to regression discontinuity designs in medical applications," *Statistics in Medicine*, 42, 4484–4513.

Cattaneo, M. D., Keele, L., Titiunik, R., and Vazquez-Bare, G. (2016), "Interpreting Regression Discontinuity Designs with Multiple Cutoffs," *Journal of Politics*, 78, 1229–1248.

Cattaneo, M. D., and Titiunik, R. (2022), "Regression discontinuity designs," *Annual Review of Economics*, 14, 821–851.

Cattaneo, M. D., Titiunik, R., and Vazquez-Bare, G. (2017), "Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality," *Journal of Policy Analysis and Management*, 36, 643–681.

Cattaneo, M. D., and Vazquez-Bare, G. (2016), "The Choice of Neighborhood in Regression Discontinuity Designs," *Observational Studies*, 2, 134–146.

Dell, M. (2010), "The persistent effects of Peru's mining mita," *Econometrica*, 78, 1863–1903.

Fan, J., and Gijbels, I. (1996), *Local Polynomial Modelling and Its Applications*, Vol. 66 of *Monographs on Statistics and Applied Probability.*, CRC Press.

Hahn, J., Todd, P., and van der Klaauw, W. (2001), "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69, 201–209.

Hyytinen, A., Meriläinen, J., Saarimaa, T., Toivanen, O., and Tukiainen, J. (2018), "When Does Regression Discontinuity Design Work? Evidence from Random Election Outcomes," *Quantitative Economics*, 9, 1019–1051.

Keele, L. J., and Titiunik, R. (2015), "Geographic Boundaries as Regression Discontinuities," *Political Analysis*, 23, 127–155.

Lee, D. S. (2008), "Randomized Experiments from Non-random Selection in U.S. House Elections," *Journal of Econometrics*, 142, 675–697.

Lee, D. S., and Lemieux, T. (2010), "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 48, 281–355.

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

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

Rosenbaum, P. R. (2010), *Design of Observational Studies*, New York: Springer.

Thistlethwaite, D. L., and Campbell, D. T. (1960), "Regression-Discontinuity Analysis: An Alternative to the Ex-Post Facto Experiment," *Journal of Educational Psychology*, 51, 309–317.

Titiunik, R. (2021), "Natural Experiments," in *Advances in Experimental Political Science*, eds. J. N. Druckman and D. P. Green, chapter 6, Cambridge University Press, pp. 103–129.