

Geographic Difference-in-Discontinuities

Kyle Butts*

SEPTEMBER 14, 2021

A recent econometric literature has critiqued the use of regression discontinuities where administrative borders serves as the ‘cutoff’. Identification in this context is difficult since multiple treatments can change at the cutoff and individuals can easily sort on either side of the border. This note extends the difference-in-discontinuities framework discussed in [Grembi et al. \(2016\)](#) to a geographic setting. The paper formalizes the identifying assumptions in this context which will allow for the removal of time-invariant sorting and compound-treatments similar to the difference-in-differences methodology.

Keywords: difference-in-discontinuities, spatial econometrics, sorting, causal inference

JEL Classification Number: C01, R15, R58

1 — Introduction

An increasingly popular estimation strategy involves using administrative borders as cutoffs in a regression discontinuity (RD) setting where the ‘running variable’ is the distance to the border. The purpose of using observations close to the border ‘cutoff’ is to try and better match treated and control units based on unobservable characteristics. Identification using the standard regression discontinuity continuity assumption is problematic because many laws and institutions change discontinuously (i.e. compound

*University of Colorado Boulder, Economics Department (kyle.butts@colorado.edu)

treatment) at the border cutoff and people chose to sort on either side of the borders (i.e. sorting around cutoff), leading to important differences between units in close geographic proximity even in the counterfactual world without treatment.¹

The intuition of the difference-in-discontinuities design is very similar to the difference-in-differences design. A pre-treatment regression discontinuity identifies time-invariant effects of other laws as well as the discontinuity in outcomes due to time-invariant sorting. A post-treatment regression discontinuity identifies those previous two discontinuities plus the one caused by the treatment of interest. The difference between the two identifies the treatment effect. In this note, I extend the difference-in-discontinuities identification strategy formalized in [Grembi et al. \(2016\)](#) to the context of geographic regression discontinuities and discuss the particular identifying assumptions needed for the above identification sketch to be true when using a geographic regression discontinuity.

I contribute to the econometric literature on regression discontinuity in three ways. First, I contribute to the nascent literature formalizing the difference-in-discontinuities identification strategy ([Grembi et al., 2016](#); [Galindo-Silva et al., 2021](#); [Millán-Quijano, 2020](#)). The results of the previous papers only consider the case of compound treatment where multiple treatment occurs at a cutoff. This paper formalizes the effectiveness of using difference-in-discontinuities to address the problem of sorting around the cutoff.

The second contribution is that I extend on the work of [Keele and Titiunik \(2015\)](#) who formalize identification with geographic-border based RDs into the *geographic* difference-in-discontinuities setting. The authors raise the problem of sorting on either side of the boundary as well as multiple laws changing discontinuously at the boundary and propose stringent assumptions to avoid these problems in the cross-section. This paper uses the difference-in-discontinuities methodology which provide a solution to these problems under arguably less stringent assumptions by leveraging the panel nature of data to estimate time-invariant sorting and effects of other policy/institution changes.

1. Identification through randomization local to the cutoff does not make sense in the geographic context because that would require people to randomly be located on either side of the border.

Last, I show that that estimation of difference-in-discontinuities with panel data can be done by running regression discontinuity on outcomes that have been first-differenced. This allows for the use of modern advancements in estimation and inference of regression discontinuities.² In cases where panel data is not available, then the local regression framework proposed in [Grembi et al. \(2016\)](#) can be used.

2 — Theory

2.1. Traditional RD Identification

Before introducing the difference-in-discontinuities method, I first review geographic regression discontinuity to highlight potential pitfalls in the method. I consider the standard context of an independent and identically distributed random cross-sectional sample of units $i \in \{1, \dots, n\}$. There is a running variable D_i for observation i and treatment turns on when the running variable is greater than the cutoff which without loss of generality we normalize to zero. Therefore treatment status is defined by $T_i = \mathbf{1}(D_i \geq 0)$. In the context of geographic regression discontinuity, D_i is the measure of the distance to a border of a treated area with distances are positive within the treated area and negative outside.³ The observed outcome is given by

$$y_i = f(D_i) + \tau(D_i)\mathbf{1}(D_i \geq 0) + \underbrace{X_i\beta + u_i}_{\equiv \varepsilon_i}.$$

The function $f(D)$ summarizes location-specific variables that affect outcomes. For example, one side of the border may be closer to a city and therefore $f(D)$ would represent

2. See [Cattaneo et al. \(2019\)](#) and [Cattaneo et al. \(2019\)](#) for an overview of modern techniques. The formulation using first-differences is practically useful as estimation can be done using the suite of regression discontinuity packages found at <https://rdpackages.github.io/>.

3. [Keele and Titiunik \(2015\)](#) discuss the relative advantages and disadvantages of using a single measure of distance versus a two-dimensional running variable. The difference-in-discontinuity method can be extended into the two-dimensional framework easily, but data will usually render the two-dimensional case implausible.

the effect of proximity to a city has on the outcome variable. Also, amenities and labor markets change across space and $f(D)$ summarizes the effect these have on the outcome at distance D from the border. On the other hand, ε_i represents the unobserved and *potentially* observed individual-specific characteristics that affect outcome variable. The quantity $Y_i(0) = f(D_i) + \varepsilon_i$ determines the outcome variable in the absence of treatment T and $\tau(D_i) = Y_i(1) - Y_i(0)$ is the average treatment effect at distance D_i .

Identification of the treatment effect relies on the assumption that location-specific and individual-specific characteristics don't change discontinuously at the border. These are formalized in assumption (1).

Assumption 1 (RD).

- (i) The functions $f(D)$ and $\tau(D)$ are continuous at the cutoff, $D = 0$.
- (ii) $\mathbb{E}[\varepsilon_i | D_i = D]$ is continuous at the cutoff, $D = 0$.

Part (i) of assumption (RD) says that the effect of the running variable on the outcome with and without treatment is continuous at the cutoff and part (ii) says that the effect of other *potentially* unobserved individual-specific variables on the outcome are continuous at the cutoff. In the context of geographic discontinuities, a discontinuity in $f(D)$ could arise from multiple policies changing at the border and a discontinuity in $\mathbb{E}[\varepsilon_i | D_i = D]$ could arise from sorting across the border based along some characteristics that affect the outcome (Keele and Titiunik, 2015). These two problems are quite common in geographic RDs and represent a central threat to identification of treatment effects.

If the two continuity assumptions are satisfied, observations in the control area close to the border identify the limiting value of $f(0)$ and observations in the treated area close to the border identify the limiting value of $f(0) + \tau(0)$.⁴ The difference between these two limits identify $\tau(0)$.

4. Since researchers have no observations exactly on the cutoff, RD estimation requires extrapolation from observations near the border. There are many ways to extrapolate from observed data to the limits at the cutoff, but a local linear or quadratic polynomial regression are the most common. See Cattaneo et al. (2019) for discussions of estimation.

We define for a variable z , the left and right limits at the cutoff as $z^+ \equiv \lim_{D_i \rightarrow 0^+} z_i$ and $z^- \equiv \lim_{D_i \rightarrow 0^-} z_i$. With assumption (RD), it is easy to show that the regression discontinuity estimate identifies the treatment effect, i.e. $\tau = y^+ - y^-$.⁵

Theorem 1 (RD Identification). Under assumption (RD), $\tau(0) = y^+ - y^-$

Proof. First note that

$$y^+ = \lim_{D_i \rightarrow 0^+} y_i = f^+ + (\tau(D_i))^+ + \varepsilon^+, \text{ and } y^- \equiv \lim_{D_i \rightarrow 0^-} y_i = f^- + \varepsilon^-.$$

$$\begin{aligned} y^+ - y^- &= (f^+ + (\tau(D_i))^+ + \varepsilon^+) - (f^- + \varepsilon^-) \\ &= \tau(0) + f^+ - f^- + \varepsilon^+ - \varepsilon^- = \tau(0), \end{aligned}$$

where the second equality comes from the continuity of $\tau(D)$ and the last equality comes from continuity of $f(D)$ and part (ii) of assumption (RD). \square

2.2. Difference-in-Discontinuities Identification

Now we turn to the panel setting where we observe outcomes before and after implementation of the treatment of interest, $t \in \{0, 1\}$. I assume that we observe panel data such that $\{(y_{i0}, y_{i1}, D_i)\}_{i=1}^n$ are independent and identically distributed. In this setting we are able to view discontinuities in outcomes before treatment occurs which we will subtract from the discontinuity after treatment occurs to remove problems of time-invariant sorting and multiple policy changes at the border.

Potential outcomes for individual i at time t are now given by

$$y_{it} = f_t(D_i) + \gamma(D_i)\mathbf{1}(D_i \geq 0) + \tau(D_i)\mathbf{1}(D_i \geq 0)\mathbf{1}(t = 1) + \varepsilon_{it}, \quad (1)$$

where $\gamma(D)$ represent the time-invariant discontinuity at the cutoff which could be due to time-invariant sorting and/or the effects of other policies that change at the border; the untreated location-specific component, $f_t(D)$, can vary across periods; and $\tau(D)$ remains the treatment effect of interest.

5. See Theorem 1 in Hahn et al. (2001).

The assumptions necessary to identify the treatment effect $\tau(0)$ requires the traditional RD assumptions to hold in both periods.

Assumption 2 (Diff-in-Disc).

- (i) The function $f_t(D)$ is continuous at the cutoff, $D = 0$, for $t \in \{0, 1\}$ and the function $\tau(D)$ is continuous at $D = 0$,
- (ii) $\mathbb{E}[\varepsilon_{it}|D_i = D]$ is continuous at the cutoff, $D = 0$, for $t \in \{0, 1\}$.

These assumptions warrant a bit of discussion. Note that the continuity assumption on f_0 does not require there to be no other policies that change at the border by $t = 0$ because their effects are included in $\gamma(D)$ which is allowed to be discontinuous, but rather that outcomes are continuous in the counterfactual absence of all policies. This is an improvement over traditional geographic RDs as compound treatment is often a big concern (Keele and Titiunik, 2015). Similarly, the continuity assumption of $\mathbb{E}[\varepsilon_{i0}|D_i = D]$ does not rule out time-invariant sorting as this too is captured in $\gamma(D)$.⁶

In the post period, part (ii) of assumption (Diff-in-Disc) requires two things conceptually. First, $f_1(D)$ being continuous requires that no other policies turn on between periods $t = 0$ and 1 that would cause a discontinuity at the border in the absence of the treatment of interest. For example, this would fail if a location passes a population threshold that triggers a policy change of interest as well as multiple other federal policy changes (Eggers et al., 2018). Second, it requires that the effects of previous policies were already fully developed in period 0. If the effects of other policies change over time, then the changes in effects over time would not be absorbed in $\gamma(D)$ and would cause a discontinuity in $f_1(D)$ that would be mistaken as the treatment effect. In the post period, part (ii) of assumption (Diff-in-Disc) requires that no additional sorting can occur between 0 and 1, whether that be from the treatment or lagged sorting of other previous treatments.

6. This is similar to the difference-in-differences method in that we can allow for differences in initial levels as long as they remain constant across periods.

To help with estimation of the treatment effect, we can reformulate our potential outcomes in a first-difference model,

$$(y_1 - y_0) = (f_1(D_i) - f_0(D_i)) + \tau(D_i)\mathbf{1}(D_i \geq 0) + (\varepsilon_{i1} - \varepsilon_{i0}),$$

where $\gamma(D_i)$ cancel out because it is time-invariant.

Theorem 2 (Diff-in-Disc Identification). Under assumption (Diff-in-Disc), $\tau(0) = (y_1 - y_0)^+ - (y_1 - y_0)^-$.

Proof.

$$\begin{aligned} (y_1 - y_0)^+ - (y_1 - y_0)^- &= \tau(D)^+ + (f_1 - f_0)^+ + (\varepsilon_{i1} - \varepsilon_{i0})^+ - ((f_1 - f_0)^- + (\varepsilon_{i1} - \varepsilon_{i0})^-) \\ &= \tau(0) + (f_1^+ - f_1^-) + (f_0^+ - f_0^-) + (\varepsilon_1^+ - \varepsilon_1^-) + (\varepsilon_0^+ - \varepsilon_0^-) \\ &= \tau(0), \end{aligned}$$

where the second equality comes from continuity of $\tau(D)$ and the last equality comes from the two continuity assumptions (Diff-in-Disc) and (Diff-in-Disc). \square

The above theorem says that so long as sorting and other policies are fully observed in the pre-period, a regression discontinuity estimated on a first-differenced outcome will identify the treatment effect. This theorem is similar to [Grembi et al. \(2016\)](#). However, their context assumes that the only difference is a second policy that occurs in period $t = 0$ whereas the model used here allows for both compound treatment and sorting.

2.3. Estimation

In cases of panel data, formulating the above identification result in terms of first differences is advantageous. Since $(y_1 - y_0)^+ - (y_1 - y_0)^-$ is a standard RD estimate of the difference between the right and left limits, this unlocks the wide set of econometric tools used in RD estimation including local polynomial regression, data-driven bandwidth selection, and bias-corrected inference. [Cattaneo et al. \(2019\)](#) and [Cattaneo et al. \(2199\)](#) provide a literature review of the modern RD literature and include a set of R and Stata programs containing powerful estimation tools.

In the non-panel case, estimation can proceed in a local-polynomial regression framework as proposed by [Grembi et al. \(2016\)](#). They recommend running the following regression using observations within a small interval around $D_i = 0$:

$$Y_{it} = \delta_0 + \delta_1 * D_i + \mathbf{1}(D_i \geq 0) (\gamma_0 + \gamma_1 D_i) + \mathbf{1}(t = 1) (\alpha_0 + \alpha_1 * D_i + \mathbf{1}(D_i \geq 0) (\beta_0 + \beta_1 D_i)) + \eta_{it}.$$

From standard regression results, β_0 , will be the difference-in-discontinuities estimate. This estimation strategy, however, does not as easily allow for the use of modern bias-robust estimators.

3 — Concluding Remarks

This paper extended the difference-in-discontinuities framework proposed by [Grembi et al. \(2016\)](#) into the context of geographic discontinuities. This setting faces the same problem of compound treatment that other RD contexts exhibit and since individuals can sort across the border, this context provides additional difficulties. This paper formalizes the necessary assumptions in the *geographic* context in order to identify the treatment effect of a policy. Moreover, in the presence of panel data, this paper proposes improved estimation techniques by recasting the estimator as a regression discontinuity estimator on first-differenced data.

References

- Cattaneo, Matias D., Nicolas Idrobo, and Rocio Titiunik.** 2019. “A Practical Introduction to Regression Discontinuity Designs: Foundations.” *arXiv:1911.09511 [econ, stat]*. [10.1017/9781108684606](https://arxiv.org/abs/1911.09511), arXiv: 1911.09511.
- Cattaneo, Matias D, Nicolas Idrobo, and Rocio Titiunik.**, “A Practical Introduction to Regression Discontinuity Designs: Volume II.” 106.

- Eggers, Andrew C., Ronny Freier, Veronica Grembi, and Tommaso Nannicini.** 2018. “Regression Discontinuity Designs Based on Population Thresholds: Pitfalls and Solutions: RDD BASED ON POPULATION THRESHOLDS.” *American Journal of Political Science* 62 (1): 210–229. [10.1111/ajps.12332](https://doi.org/10.1111/ajps.12332).
- Galindo-Silva, Hector, Nibene Habib Some, and Guy Tchuente.** 2021. “Fuzzy Difference-in-Discontinuities: Identification Theory and Application to the Affordable Care Act.” *arXiv:1812.06537 [econ]*, <http://arxiv.org/abs/1812.06537>, arXiv: 1812.06537.
- Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano.** 2016. “Do Fiscal Rules Matter?” *American Economic Journal: Applied Economics* 8 (3): 1–30. [10.1257/app.20150076](https://doi.org/10.1257/app.20150076).
- Hahn, Jinyong, Petra Todd, and Wilbert Klaauw.** 2001. “Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design.” *Econometrica* 69 (1): 201–209. [10.1111/1468-0262.00183](https://doi.org/10.1111/1468-0262.00183).
- Keele, Luke J., and Rocío Titiunik.** 2015. “Geographic Boundaries as Regression Discontinuities.” *Political Analysis* 23 (1): 127–155. [10.1093/pan/mpu014](https://doi.org/10.1093/pan/mpu014).
- Millán-Quijano, Jaime.** 2020. “Fuzzy difference in discontinuities.” *Applied Economics Letters* 27 (19): 1552–1555. [10.1080/13504851.2019.1696930](https://doi.org/10.1080/13504851.2019.1696930).