

Sampling for spatial discontinuity designs

Florence Kondylis
John Loeser
Jeremy Magruder



DIRECTED BY



FUNDED BY



Sampling for spatial discontinuity designs*

Florence Kondylis John Loeser Jeremy Magruder

March 31, 2026

Abstract

Spatial discontinuity designs are ubiquitous as an identification strategy in the evaluation of place-based policies and infrastructure developments. Yet little guidance exists for sampling and data collection in the context of spatial discontinuity designs. What identifying assumptions are required to estimate causal impacts, and what sampling strategies lend themselves to spatial discontinuity designs? In this chapter, we describe the identification, sampling, and data collection strategies employed by [Jones et al. \(2022\)](#) ([JKLM](#)), who leverage a discontinuity in access to irrigation at the boundaries of hillside irrigation schemes in Rwanda to evaluate their impacts. We discuss challenges unique to sampling and data collection for spatial discontinuity designs, relative to regression discontinuity designs, and how these were addressed in [JKLM](#). Specifically, we discuss the identifying assumptions required to accommodate spatial units of observation that typically occupy discrete areas (e.g., agricultural plots in [JKLM](#)). We then cover the intricacies of matching spatial units to informed respondents (e.g., plots to farmers in [JKLM](#)). Finally, we address the need to account for spillovers across spatial units of observation (e.g., a single farmer managing multiple plots in the context offered by [JKLM](#)). We frame our discussion as a case study in the broader context of sampling and data collection for spatial discontinuity designs in development economics, and discuss alternative identification and sampling strategies used in the literature.

*Kondylis: World Bank, fkondylis@worldbank.org; Loeser: World Bank, jloeser@worldbank.org; Magruder: UC Berkeley, NBER, jmagruder@berkeley.edu. We thank the European Union, the Global Agriculture and Food Security Program (GAISP), the World Bank Rwanda Country Management Unit, the World Bank i2i fund, 3ie, and IGC for generous research funding. The views expressed in this manuscript do not reflect the views of the World Bank. All errors are our own.

1 Identification and estimation with a spatial discontinuity

1.1 What is a spatial discontinuity design?

What is a spatial discontinuity design? When a policy varies discontinuously across a geographic boundary, but other characteristics of individuals do not, average differences in outcomes between individuals on either side of the geographic boundary are attributable to the policy. In a spatial discontinuity design, the impact of the policy is identified as the difference in outcomes between individuals at the boundary who happen to lie on the sides exposed and not exposed to the policy.¹

The spatial discontinuity design in JKLM In JKLM, we observed that access to irrigation (the “policy”) varied discontinuously across gravity fed irrigation canals (the geographic boundary); a photograph of this discontinuity is presented in Figure 1. In Figure 1, plots directly above the canal are not irrigated, and are cultivated predominantly with maize (left of the figure). In contrast, plots directly below the canal, which have access to irrigation are either visibly irrigated and cultivated with horticulture (top right of the figure), or are under preparation, perhaps for a shorter cycle crop than maize (bottom right of the figure). If these plots would have been cultivated with the same crops absent access to irrigation, differences in crop choice between these plots are attributable to the impact of access to irrigation. This is likely the case if the placement of the canal is not affected by determinants of preferred cultivation choices on these plots, as JKLM argue is the case. Consequentially, the impact of access to irrigation on cultivation of horticulture and use of irrigation is simply the differences in these outcomes between plots below the canal, which have access to irrigation, and plots above the canal, which do not.

¹Keele and Titiunik (2015) and Keele et al. (2017) review spatial discontinuity designs, and we follow their discussions closely in this Section 1.1.

Figure 1: Discontinuity in access to irrigation across a gravity-fed canal



Notes: A photograph of plots adjacent to a canal in Karongi 12, one of the hillside irrigation schemes studied in JKLM, is presented in this figure.

Identifying assumptions The key assumption underpinning identification of policy impacts in a spatial discontinuity design is conditional exogeneity.² Under conditional exogeneity, units of observation near the boundary, on either side of the boundary, would have had the same average outcomes absent the policy of interest. In JKLM, conditional exogeneity imposes that, absent access to irrigation below the canal, the same crops would have been cultivated on plots close to the canal on either side. Under conditional exogeneity, differences in outcomes for units of observation near the boundary, on either side of the boundary, reflect the impact of the policy of interest.

The conditional exogeneity assumption is distinct from the “continuity framework”-based identifying assumption typically used in regression discontinuity designs (Cattaneo and Titiunik, 2022). Under the continuity framework, it is assumed that

²Keele et al. (2017) discuss this assumption in depth as “geographic mean independence”.

individuals can be sampled arbitrarily close to the discontinuity. In spatial discontinuity designs, as is generally true of spatial economic data, units of observation occupy discrete space, and it is therefore often not possible to sample units arbitrarily close to the boundary (Keele et al., 2017).

Testing internal validity Standard tests of internal validity for causal inference apply to the conditional exogeneity assumption for spatial discontinuity designs. These include:

- **Balance** tests demonstrate that units on either side of the geographic boundary are observably similar before the introduction of the policy of interest. In JKLM, balance corresponds to the test that characteristics of plots, and their managers, and plot outcomes before the construction of the hillside irrigation scheme, are similar above and below the canal.
- Tests of **no differential attrition** demonstrate that outcomes are equally liked to be observed for units on either side of the boundary. In JKLM, tests of no differential attrition correspond to showing that rentals and purchases of plots, and household non-response, do not differentially induce plot-level attrition above relative to below the canal.
- **Robustness** tests demonstrate that results are unaffected by alternative specifications of the controls conditional on which exogeneity is expected to hold. In JKLM, tests of robustness correspond to showing that results are unaffected by the choice of specifications discussed in Section 1.2.

1.2 Identification and estimation in spatial discontinuity designs in practice

The units of observation in spatial discontinuity designs, and the nature of the boundary, vary substantially across applications; this has led to a range of approaches to estimation, although all build on a version of the conditional exogeneity assumption. We describe estimation motivated by four versions of the conditional exogeneity assumption below, which we link to sampling strategies in our discussion in Section 2. In Section 1.2.1, we describe estimation when exogeneity holds within a short distance of the boundary. We then describe estimation when exogeneity holds conditional on

additional controls for distance to the boundary (Section 1.2.2), or on spatial controls for spatial heterogeneity (Section 1.2.3).

1.2.1 Exogeneity conditional on proximity to discontinuity

The most basic version of conditional exogeneity imposes that, absent the policy of interest, units of observation within some fixed distance of the boundary would have had the same average outcomes. [Keele et al. \(2017\)](#) call this assumption “geographic mean independence”, and note its equivalence to the “local randomization” framework for regression discontinuity designs with distance to the boundary as the running variable ([Cattaneo and Titiunik, 2022](#)).

Application Perhaps the known application of a spatial discontinuity design is [Card and Krueger \(1994\)](#), who compare fast food restaurants on either side of the Pennsylvania-New Jersey border, before and after an increase in the New Jersey minimum wage. After restricting to observations close to the border (that is, by excluding western Pennsylvania), they argue that these restaurants would have experienced the same changes in employment absent the increase in the New Jersey minimum wage.

1.2.2 Exogeneity conditional on distance to discontinuity

In some applications, one may be concerned that outcomes might systematically vary as one moves towards or away from the discontinuity; for example, in [JKLM](#), we note that plots in the command area are systematically at lower elevations, as they lie below the irrigation canal, and therefore elevation is systematically decreasing as one moves towards the command area boundary. This concern motivates controlling for the running variable (along with additional corrections to inference) in the “continuity framework” for regression discontinuity designs; similarly, in spatial discontinuity designs, it is common to control for distance to the boundary, even when restricting to observations very close to the boundary.

Application: Linear control for distance An example application of this approach is in [JKLM](#), in which we control linearly for distance to the boundary and its interaction with an indicator for lying within the boundary in our “RDD” specification. This control is motivated by the observation that, even within a small

distance of the canal, plots above the canal will be systematically located higher on the hillside and this may be associated with differences in agronomic characteristics; as these differences are expected to be continuous with respect to distance to the canal, controlling linearly for distance can partly account for such differences.

Application: Robustness to flexible controls for distance In another application, Dell (2010) estimate the impacts of a historical system of forced labor, using a spatial discontinuity in its imposition, on long-run household welfare. They make use of part of the boundary which deviated from the Andean precipice, along which they demonstrate balance by showing that districts on either side of the boundary were observably similar before the imposition of the system of forced labor. To test robustness, they implement specifications that either control for polynomials ranging from degree one through four in distance to the boundary, either interacted or not interacted with an indicator for lying within the boundary, or control for polynomials ranging from degree one through four in latitude and longitude.

1.2.3 Exogeneity conditional on spatial fixed effects

In practice, when units of observation occupy discrete space and as distances become small, the notion of distance to the boundary as a running variable may lose some meaning; in the application of Dube et al. (2016), for instance, all units are located at the boundary, and in that sense have a distance to the boundary of 0. One common approach in this case is to apply “spatial fixed effects” (Magruder, 2012), that is to estimate impacts by taking differences in outcomes between spatially proximate units on either side of the boundary. Spatial fixed effects are motivated by a natural extension of the basic version of conditional exogeneity: absent the policy of interest, *nearby* units of observation within some fixed distance of the boundary would have had the same average outcomes. A number of closely related implementations of spatial fixed effects have emerged in practice, which may vary in appropriateness by the nature of the discontinuity.

Application: Pairwise matching Dube et al. (2016) estimate the impacts of minimum wages using discontinuities in changes in minimum wages across state borders. They pair counties that share a border, but are in different states, and control for county-pair fixed effects. Estimated impacts with pairwise matching are differences

in outcomes between units of observation (counties) on either side of the boundary (state borders), among paired observations (pairs of counties that share a border).

Application: Boundary segment fixed effects Dell (2010) divide their boundary into four equal-length segments, and control for boundary segment fixed effects. Estimated impacts with boundary segment fixed effects are differences in outcomes between units of observation (districts) on either side of the boundary, among units of observation closest to the same boundary segment.

Application: Spatial differencing Magruder (2013) estimate the impacts of minimum wages in Indonesia using discontinuities in changes in minimum wages across Province (and sometimes sub-Province) borders. Before running regressions of outcomes on treatment (minimum wages), they spatially demean outcomes and treatment by subtracting off their average across units of observation within some fixed distance. Estimated impacts with spatial differencing are differences in outcomes between units of observation on either side of the boundary; however, spatial differencing places a higher weight on units with many observations on the opposite side of the boundary within a given distance, in a manner similar to controlling linearly for the running variable in a one-dimensional regression discontinuity design.

Application: Spatial k -nearest neighbors differencing In JKLM, in our “SFE” specification, we first spatially demeaned outcomes and treatment (that a plot is located in the command area) by subtracting off their average across the five nearest plots. Estimated impacts with spatial k -nearest neighbors differencing are differences in outcomes between units of observation (plots) on either side of the boundary, with a higher weight placed on units i with many observations on the opposite side of the boundary either among i ’s k -nearest neighbors, or that have i as one of their k -nearest neighbors.

2 Sampling spatially around a spatial discontinuity

Analogy to regression discontinuity designs In the canonical regression discontinuity design, it is imagined that with a larger-and-larger sample, the sample used to estimate the effect of the discontinuity will be progressively shrunk to an

arbitrarily small bandwidth around the discontinuity. Consequentially, when sampling for a regression discontinuity design, it is optimal for internal validity to sample units as close as possible to the discontinuity (Cattaneo et al., 2019). Such sampling is possible, and desirable, when the location of units relative to the discontinuity is known, but key outcomes of interest require costly surveys to observe (Cattaneo and Titiunik, 2022).

Sampling for spatial discontinuity designs Similarly, in spatial discontinuity designs, differences between units are most credibly attributable to the boundary for units very close to the boundary. Common specifications discussed in Section 1.2 involved restricting to observations within small distances of the boundary, and observations further from the boundary do not contribute to estimates of treatment effects in such specifications. Under a fixed sample size, sampling observations within a small distance of the boundary therefore maximizes power.

In many cases, no sampling is needed; data may be available for all potential units of observation, as in Dube et al. (2016) for US counties, or when remotely sensed outcomes are of interest.

In other contexts, sampling may be possible from administrative data or a census. For instance, Card and Krueger (1994) sample *all* fast food restaurants from four chains within their discontinuity sample (eastern Pennsylvania and New Jersey), using white-pages telephone listings. However, in many contexts, such a “census” may not be available.

Constraints on sampling for spatial discontinuity designs However, in many important contexts, fully censusing units of observation and identifying their locations may not be practical. This was our case in JKLM and we discuss our context in Section 2.1. In a related example, Lowes et al. (2017) oversample communities in which individuals are disproportionately likely to have ancestry in their discontinuity sample, in a context (Kananga) in which no census of ancestral origins is available and conducting such a census would have been prohibitively costly. In general, the appropriate approach will vary across contexts.

Roadmap In this section, we summarize the approach to sampling plots, and surveying their associated plot managers, near the boundary of the command areas of

the hillside irrigation schemes in JKLM. In addition, we discuss the generalization of the approach to sampling for spatial discontinuity designs more broadly, including when particular steps are and are not necessary.

- **Context** Section 2.1 describes the setting of JKLM, and how it motivated spatial sampling.
- **Step 1: The boundary** Section 2.2.1 describes administrative data on the spatial discontinuity.
- **Step 2: Sampling space** Section 2.2.2 describes the process of sampling GPS coordinates in the neighborhood of the spatial discontinuity.
- **Step 3: From space to a listing** Section 2.2.3 describes the process of interpretably describing each sampled GPS coordinate and identifying its appropriate respondent.
- **Step 4: From a listing to households** Section 2.2.4 describes the process of sampling from this list to construct a survey sample of households and descriptions of their associated GPS coordinates.
- **Step 5: Surveying households** Section 2.2.5 describes the process of surveying this sample.
- **Testing exogeneity** Section 2.3 describes the process of testing exogeneity.
- **Interpreting estimates** Section 2.4 describes how the employed sampling strategy shapes the interpretation of estimates under the approaches to estimation described in Section 1.2.

2.1 Context and spatial sampling

In JKLM, the ideal discontinuity sample around the boundary of hillside irrigation schemes was exceptionally local. Each command area, the area of the scheme in which farmers are able to access irrigation from the canal, was between 1 and 3 square kilometers. The hillside irrigation schemes lay in areas with substantial geographic heterogeneity, reflected in the photograph of one scheme in Figure 2: the schemes were built into rolling hills, with canals that were far from neither the tops of the hills nor the marshlands in the valley.

Figure 2: A hillside irrigation scheme




Notes: Appendix Figure A1 of [JKLM](#), a photograph of Karongi 12 hillside irrigation scheme.

While censuses of households were available at the village-level, rural Rwanda is a context where many households have plots in villages that are not where they live ([Kondylis, 2008](#)). As a result, sampling individuals in villages located at the boundary would result in a low probability of sampling individuals with plots within a small distance from the boundary.

Alternatively, a census of households with land in command area was available. Unfortunately, this data did not identify the proximity of households' plots to the boundary. And perhaps more crucially, this data did not identify households with plots outside the command area.

This context motivated our “spatial sampling” approach described in Section 2.2: broadly, we directly sampled plots located near the boundary, and followed up with their managers.

Figure 3: Implementing spatial sampling



Notes: The steps of our spatial sampling procedure in JKLM are visualized in this figure along a segment of the command area boundary, from the boundary shapefile in 3a to the sample plot maps along that segment in 3f. Additional details are in Section 2.2.

2.2 Implementing spatial sampling

2.2.1 Administrative data on boundary

To sample plots within a small distance of the boundary, we began with a shapefile of the command area boundary, an example of which is presented in Figure 3a.

2.2.2 Sampling GPS coordinates

In a second step, we used the boundary shapefile to construct three sampling areas in each of the hillside irrigation sites with a discontinuity: the area within 50 meters of the boundary both inside and outside the command area, and the remainder of

the command area. We then placed a grid of GPS coordinates at 2 meter resolution over the site, and sampled coordinates from each of the three areas (with an equal number of points sampled within 50 meters of the boundary inside and outside the command area).

An example of our sample of coordinates is presented in Figure 3b. We observe coordinates both outside the command area (on the bottom right) and coordinates inside the command area (on the top left). All coordinates outside the command area lie within 50 meters of the boundary, while some additional coordinates were sampled more than 50 meters from the boundary inside the command area.

2.2.3 From GPS coordinates to listing

In a third step, we provided enumerators with GPS devices with the coordinates of each sampled point (uploaded to the GPS devices as `.gpx` files), and asked enumerators to visit each sampled point; while visiting points, enumerators used the “record” feature of their GPS device, which allowed verification that enumerators visited the intended points. At each point that fell on cultivable land (80.1%), enumerators worked with a local informant (often farmers working nearby) to identify the cultivator responsible for the plot (name, village of residence, and phone number), and describe the plot with sufficient detail that the cultivator would be able to identify the exact plot described. Names of identified cultivators were then fuzzily matched, with the fuzzy matching confirmed through phone calls with village leaders.

An example of the output of this listing is visualized in Figure 3c; each of the points sampled in Section 2.2.2 (which, in this case, all happened to fall on cultivable land) is now colored based on the identity of the cultivator, with points belonging to a common cultivator enclosed in that cultivator’s color. Many cultivators had multiple points sampled, which in some cases lay on multiple distinct plots.

2.2.4 From listing to survey sample

In a fourth step, we randomly sampled one point among those each cultivator was responsible for. An example output of this sampling is visualized in Figure 3d, in which only the sampled points are left darkened. Some cultivators have no sampled points in Figure 3d, as their sampled point lay in another area of the hillside irrigation scheme.

The description of the plot at the sampled point, which we refer to as the “sample plot”, along with the identity of the cultivator, comprise our survey sample. Note that points, rather than plots (of which we did not have a comprehensive census), were sampled – as a consequence, very large plots were oversampled, which we return to in Section 2.4.

2.2.5 Surveying

In a final step, we surveyed cultivators on their sample plot, and also on their household and their other plots. Fuzzy matching the initial names and confirming through phone calls with village leaders was particularly important, and we were able to reach 90.2% (1,695 households) of our initial sample. Of these 1,695, we successfully mapped the sample plot for 91.6% (1,552).³ An example of the mapped sample plots, and their associated sampled points, is presented in Figure 3e, with the plot maps overlaid with satellite imagery in Figure 3f.

2.3 Testing exogeneity

Our goal of sampling plots close to the boundary in Section 2.2 was to estimate the impacts of access to irrigation, by comparing plots just inside and just outside the command area, using the methods described in Section 1.2. These estimates recover the causal effect of access to irrigation under exogeneity conditional on proximity to the discontinuity: absent access to irrigation, outcomes would have been the same on plots just inside and just outside the command area.

In JKLM, we implement a number of tests of exogeneity: balance of household characteristics, balance of plot characteristics, balance of remotely sensed outcomes on sampled plots, and balance of attrition in follow up surveys. These are tests of exogeneity conditional on our sampling strategy. Such tests may therefore reject balance if either outcomes would have discontinuously varied at the boundary absent access to irrigation, or alternatively if our sampling strategy differentially sampled households or plots discontinuously at the boundary.

³Attrition was often due to data corruption during transfer from the GPS devices. Modern CAPI software enables plot mapping with the tablet used for surveying, which can be used with a bluetooth GPS receiver. This simplifies plot mapping protocols, and has eliminated data corruption as an issue.

2.4 Interpreting estimates

Under exogeneity conditional on proximity to the discontinuity, what weighted average of treatment effects do comparisons recover under this sampling strategy? Our sampling strategy, as mentioned in Section 2.2.4, oversamples large plots. To interpret this oversampling, it is helpful to consider two limiting cases. First, suppose that each household holds a vanishingly small share of total area: in this case, our procedure would identify exactly one point per cultivator, cultivators would be sampled with probability proportional to their total area managed, and plots within cultivator would be sampled with probability proportional to their area. Second, suppose that we sampled sufficiently densely that at least one point was sampled from each household: in this case, all cultivators would have been sampled, but plots within cultivator would be sampled with probability proportional to their area. In JKLM, we averaged 4.2 points per cultivator and identified multiple points for 63% of cultivators, suggesting that we sampled sufficiently densely (or that there were sufficiently few total cultivators) that in practice we were closer to sampling all cultivators. However, in general, this procedure will sample cultivators with probability increasing in their total area managed, and will sample plots within cultivator proportional to their area.

As a consequence, estimated impacts on sample plots using this sampling strategy reflect average effects on land within household,⁴ with households weighted by their sampling probability. In JKLM, for instance, our discussion above would suggest that estimates reflect average impacts across households, with plots weighted proportionally to their area within household.

References

- Card, David and Alan B Krueger**, “Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania,” *The American Economic Review*, 1994, 84 (4), 772–793.
- Cattaneo, Matias D and Rocio Titiunik**, “Regression discontinuity designs,” *Annual Review of Economics*, 2022, 14 (1), 821–851.

⁴The interpretation of these estimated impacts, distinct from the weighted average of comparisons they reflect, is shaped by across plot spillovers, which we discuss at length in JKLM.

- , —, and **Gonzalo Vazquez-Bare**, “Power calculations for regression-discontinuity designs,” *The Stata Journal*, 2019, *19* (1), 210–245.
- Dell, Melissa**, “The persistent effects of Peru’s mining mita,” *Econometrica*, 2010, *78* (6), 1863–1903.
- Dube, Arindrajit, T William Lester, and Michael Reich**, “Minimum wage shocks, employment flows, and labor market frictions,” *Journal of Labor Economics*, 2016, *34* (3), 663–704.
- Jones, Maria, Florence Kondylis, John Loeser, and Jeremy Magruder**, “Factor market failures and the adoption of irrigation in rwanda,” *American Economic Review*, 2022, *112* (7), 2316–2352.
- Keele, Luke J and Rocio Titiunik**, “Geographic boundaries as regression discontinuities,” *Political Analysis*, 2015, *23* (1), 127–155.
- Keele, Luke, Scott Lorch, Molly Passarella, Dylan Small, and Rocio Titiunik**, “An Overview of Geographically Discontinuous Treatment Assignments with an Application to Children’s Health Insurance,” *Regression discontinuity designs: Theory and applications*, 2017, pp. 147–194.
- Kondylis, Florence**, “Agricultural outputs and conflict displacement: Evidence from a policy intervention in Rwanda,” *Economic Development and Cultural Change*, 2008, *57* (1), 31–66.
- Lowes, Sara, Nathan Nunn, James A Robinson, and Jonathan L Weigel**, “The evolution of culture and institutions: Evidence from the Kuba Kingdom,” *Econometrica*, 2017, *85* (4), 1065–1091.
- Magruder, Jeremy R**, “High unemployment yet few small firms: The role of centralized bargaining in South Africa,” *American Economic Journal: Applied Economics*, 2012, *4* (3), 138–166.
- , “Can minimum wages cause a big push? Evidence from Indonesia,” *Journal of Development Economics*, 2013, *100* (1), 48–62.

IGC

theigc.org
