Boundaries Generate Discontinuities in the Urban Landscape

Vikram Maheshri and Kenneth Whaley*

June 1, 2024

Abstract

Neighborhood boundaries are often determined by physical topography, transportation networks or the administration of public goods (e.g., school attendance zones). We present a simple model of boundaries that predicts discontinuities in household demographics, the supply of amenities and home prices at physical and administrative boundaries. We take these predictions to the data and find abundant evidence of discontinuities across a wide range of observable dimensions – the universe of variables that are available in the 2020 Census at the Block Group level – and five different types of boundaries. We draw two important implications from these findings: even narrowly targeted place based policies may have much broader impacts if they involve a new administrative boundary, and researchers should implement boundary discontinuity designs with caution, as the key identification assumption is may not hold except in narrow applications.

*Vikram Maheshri (vmaheshri@uh.edu): Department of Economics, University of Houston. Kenneth Whaley (kennethwhaley@usf.edu): Department of Economics, University of South Florida. We thank various seminar and conference participants for their helpful feedback. All errors are our own.

1 Introduction

For as long as there have been cities, there have been neighborhoods. For as long as there have been neighborhoods, households, firms, and governments have sorted between them. A rich literature following from Tiebout [1956] has studied both the process of[1] and the consequences of[2] neighborhood sorting. Fundamental to this literature is the idea that neighborhood boundaries matter. They represent more than simply lines on a map, and to the extent that they induce distortions that impact decisionmakers, we should observe different outcomes across boundaries. As such, boundaries will shape both where people live and how amenities are distributed spatially.

In this paper, we develop a simple model that clarifies how boundaries may distort the location decisions of households, which in turn will distort the amenities that are supplied by private firms and the public goods that are provided by local governments. Importantly, our model suggests that these distortions will manifest as discontinuities in the demographic and amenity bundle of neighborhoods and in neighborhood prices. We test for the existence of spatial discontinuities in the universe of all publicly available variables from the 2020 US Census, analyzing block groups near five very different types of borders: historical rail and highway networks, contemporary school district boundaries and attendance zones, and county lines.

We find overwhelming evidence of discontinuities across a wide range of variables and all types of boundaries. The demographic characteristics of residents, features of the housing stock, labor market profiles, and government assistance take-up all systematically vary dis-

[1] See, for example, Epple et al. [1984, 2003], Bayer et al. [2004], Caetano [2019], Caetano and Maheshri [2023]

[2] See, for example, Black [1999], Bayer et al. [2007], Chetty et al. [2018]

continuously at boundaries. The net effect of this is ultimately capitalized into prices: we estimate average price differentials of 7-22% across these boundaries. We conduct a series of placebo tests to ensure that these differentials correspond to boundary effects as opposed to other confounders, and that these boundary effects are not merely statistical artifacts of our estimation and testing procedure.

Our results imply that researchers should be cautious when attempting to use historical boundaries to identify the effects of contemporary neighborhood characteristics, e.g., socioeconomic compositions, home prices, crime, or other urban features in a boundary discontinuity design (BDD). The key identifying assumption that is common to such studies is that the effects of boundaries are mediated through a single characteristic. Our findings suggest that such identifying assumptions are unlikely to hold, as boundaries affect a wide range of neighborhood characteristics. We conjecture that this is likely due to the fact that households and suppliers of amenities continuously sort, so changes in the features of neighborhoods will beget further changes in features of neighborhoods in a self-reinforcing feedback loop. The upshot is that BDDs may be more appropriate to estimate short-run effects before the sorting process can unfold. It is common practice to support the identifying assumptions of a BDD by showing that potentially confounding observable characteristics do not vary discontinuously at neighborhood boundaries. We provide evidence that such demonstrations are probably statistically under powered when the analysis covers smaller geographies, e.g., a single city, county or state.

Our analysis also underscores the importance for researchers to broaden the scope of their assessments of the impacts of historical boundaries. For instance, a large literature has shown that historical redlining practices in US cities still affects residential segregation today

(Zenou and Boccard [2000], Gale [2021]). Our findings indicate that these studies almost certainly understate the scope of the effects of redlining, as these efforts likely affected the urban landscape in a profound way.

To model the boundary effects we draw from seminal papers that leverage quasi-random assignment of individuals and households near boundaries such as Black [1999] and Dell [2010]. The key innovation in our analysis is that we explore whether *every* neighborhood attribute varies at the boundary, in contrast to a standard approach of exploiting arguably exogenous variation in a single variable at the boundary. Finally, this paper is related to the empirical literature that leverages historical boundaries for identification. A growing literature that expands on the findings of Ananat [2011] to show that historical rail placement subdivides cities in a way that creates more neighborhood options, segregation, and worse outcomes for non-whites (Chyn et al. [2022], Cox et al. [2022]). In the context of our findings, historic railroad boundaries likely contribute to mobility and crime outcomes through local differentials in the entire bundle of neighborhood characteristics, not simply a low dimensional proxy of racial segregation.

The remainder of this paper is organized as follows. In Section 2, we present a simple model of consumers location decisions, and we characterize how they are affected by physical and administrative boundaries. In Section 3, we present our empirical approach to estimate discontinuities in a large set of variables with no *a priori* spatial ordering, and in Section 4, we describe the various sources of historical and contemporary data that we use to implement this approach. In Section 5, we present our results. We discuss these results in the context of other research in Section 6 before concluding in Section 7.

2 Conceptual Framework

We motivate how boundaries generate discontinuities in households' location choices in an extremely simplified, illustrative model. Boundaries are modeled as either physical distortions that increase the distance between points on opposite sides of the boundary, or mechanisms that allow for different levels of public goods to be supplied (or both). For simplicity, we assume that these boundaries are exogenously located along with two producers who are exogenously located at points $0 < y_0 \leq y_1 < 1$ on the unit interval, and a unit mass of consumers, indexed by i, who endogenously locate at points $x_i \in [0, 1]$. There is also a public good that is supplied exogenously at a level of g.

Let $d_i^j = |x_i - y_j|$ be the distance between consumer i and producer j. Then consumer i's utility is given by

$$U(x_i) = \alpha_i u\left(d_i^0\right) + (1 - \alpha_i) u\left(d_i^1\right) + v(g) \tag{1}$$

where the parameter α_i is drawn from a single peaked distribution over $(0, 1)$. We assume that $u' < 0$ and $u'' > 0$. That is, α_i represents consumer i's relative preference for producer 0 to producer 1, and all consumers prefer locating closer to producers (with a diminishing loss in marginal utility in distance to producers). We also assume that $v' > 0$.

Consumer i chooses x_i to maximize the objective in equation (1). The first order condition implies

$$\frac{\alpha_i}{(1 - \alpha_i)} \frac{u'(d_i^0)}{u'(d_i^1)} = 1 \tag{2}$$

in equilibrium. This has a familiar interpretation, as the left hand side of equation (2) is the marginal rate of substitution between the two producers. The right hand side of the equation

Figure 1: Spatial Equilibrium

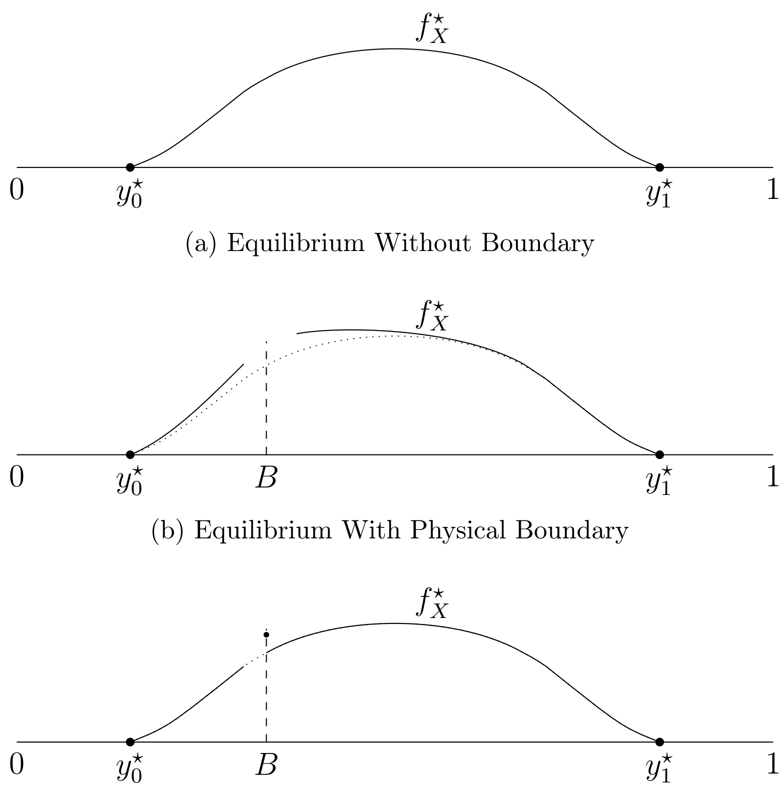

(a) Equilibrium Without Boundary

(b) Equilibrium With Physical Boundary

(c) Equilibrium With Administrative Boundary

corresponds to the price ratio if we understand distances to producers to be effective prices since $-\frac{\partial d_i^1}{\partial x_i} / \frac{\partial d_i^0}{\partial x_i} = 1$. No consumer will ever locate outside of the interval $[y_0, y_1]$, as they would be strictly better off moving into the interval. Hence we can illustrate the spatial equilibrium in the top panel of Figure 1.

We now introduce an exogenous boundary at some point $B \in (y_0, y_1)$. A basic characteristic of many boundaries such as highways or rivers is that they distort the physical environment. We capture this by modeling the distance between any two points on opposite sides of B increased by $\beta \geq 0$. Consumer i's first order condition is now

$$\frac{\alpha_i}{(1-\alpha_i)} \frac{u'\left(d_i^0\right)}{u'\left(d_i^1+\beta\right)} = 1 \qquad x_i < B \tag{3}$$

$$\frac{\alpha_i}{(1-\alpha_i)} \frac{u'\left(d_i^0+\beta\right)}{u'\left(d_i^1\right)} = 1 \qquad x_i > B \tag{4}$$

It follows that all consumers with $x_i^\star < B$ before the introduction of the boundary remain to the left of the boundary and vice versa. Moreover, the distortion has the effect of shifting the mass of consumers away from the boundary with a greater shift for those consumers who are located farthest from the boundary. We illustrate the effects of a physical boundary on spatial equilibrium in panel (b) of Figure 1. In general, physical boundaries generate discontinuities in the locations of consumers.

A second characteristic of many boundaries such as school zones or political borders is that they allow for public goods to be differentiated.[3] We model this by specifying $g(x) = g_0$ at all points $x \leq B$ and $g(x) = g_1$ at all points $x > B$ where $g_0 < g_1$ without loss of generality. This has the effect of shifting a mass of consumers just to the left of B across the boundary. We illustrate the effects of an administrative boundary on spatial equilibrium in panel (c) of Figure 1. In general, administrative boundaries also generate discontinuities in the locations of consumers. Of course, many boundaries are both physical and administrative. For example, political boundaries may coincide with rivers or school district boundaries may coincide with roadways. This does not affect the qualitative conclusions of our analysis. We should note that these conclusions are likely to persist – or even strengthen – if we endogenized the locations of boundaries or producers or if we endogenized the levels of public

[3] To simplify notation, we assume consumers exactly at the boundary can choose the side of the boundary to which they locate.

goods since this would increase incentives for sorting.

If consumers belonged to different demographic groups, and these groups had systematically different tastes for the producers (the distributions of α_i differed across groups) or different tastes for public goods ($v(\cdot)$ differed across groups) then this simple analysis would yield further insights. Because boundaries would generate discontinuities in the locations of both groups of consumers, then we would generically observe discontinuities in demographic compositions across boundaries. Moreover, if producers adjusted their outputs to cater to their clientele, then this would imply discontinuities in the amenities supplied across boundaries. Similarly, to the extent that governments respond to the preferences of their constituencies, then this would imply discontinuities in the public goods supplied across boundaries. Finally, as amenities and public goods are capitalized into home prices, this may affect the demographics of new consumers.

We can summarize each of these mechanisms contribute to the positive feedback loops shown in the following diagram:

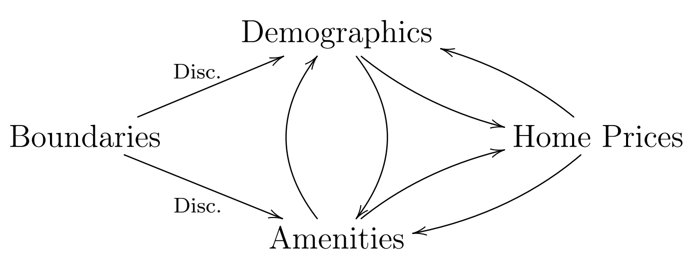

Boundaries generate discontinuities in demographics and amenities. Demographics and amenities are then co-determined with home prices through the sorting of households (consumers) and adjustments made by the suppliers of private amenities (producers) and public amenities (governments). Even if this sorting process is continuous, the ultimate effect of a boundary on the urban landscape will be discontinuities in the characteristics of residents, characteristics of the amenity bundle, and prices.

3 Empirical Approach

In order to test for these predicted discontinuities, we employ a scalable approach to estimate discontinuous boundary effects on a large set of variables. We observe a boundary network (e.g., the interstate highway network) as a series of curves in space, and we observe characteristics (e.g., population demographics or house prices) at a set of discrete points in space, which, in an abuse of nomenclature, we refer to as neighborhoods. We index neighborhoods with j.

Standard approaches to estimate boundary effects require researchers to know which side of a boundary is treated and which side of a boundary is untreated. These approaches then can then identify treatment effects in a regression discontinuity framework where the running variable is distance to boundary (untreated neighborhoods are usually assigned negative distances, and treated neighborhoods are usually assigned positive distances). In our setting, we do not have *a priori* treated and untreated sides of boundaries, as our goal is simply to identify discontinuities and estimate their magnitudes. Moreover, we seek to estimate these effects on vast, highly intersecting boundary networks that span the entire United States. For these reasons, we must modify the standard approach.

First, for a given boundary network, we divide all boundaries into smaller sections of equal length. We denote these as boundary sections, and we index them with b. For each boundary section, we consider a set of nearby neighborhoods, which we denote as J_b. We locate the nearest neighborhoods on either side of boundary section b in J_b and refer to them as index neighborhoods. For each neighborhood characteristic C, boundary section b, and neighborhood $j \in J_b$, we construct a dummy $H_{Cbj} = 1$ for all neighborhoods that are on the same side of the boundary as the index neighborhood with a higher value of C (i.e., the high side) and $H_{Cbj} = 0$ otherwise.

For each $j \in J_b$ on the high side, we define d_{Cbj} to be equal to the distance from j to the boundary, and for each j on the low side, we define d_{bj} to be equal to -1 times the distance from j to the boundary. We then estimate the following regression:

$$C_{bj} = \delta_C H_{Cbj} + f_- \left(d_{Cbj}\right) \times 1 \left(d_{Cbj} < 0\right) + f_+ \left(d_{Cbj}\right) \times 1 \left(d_{Cbj} > 0\right) + \epsilon_{Cbj} \tag{5}$$

where $f_-(\cdot)$ and $f_+(\cdot)$ are flexible functions of the distance to the boundary and ϵ_{Cbj} is an error term. The parameter δ_C corresponds to the boundary effect for characteristic C. Under the assumption that the unobservable determinants of C_{bj} vary continuously at the $d_{Cbj} = 0$ threshold, δ_C will be identified and can be estimated by least squares.

Placebo Validation Exercise

We consider the following placebo exercise to validate our empirical strategy. For each characteristic C and boundary section b, we randomly draw ξ_{Cb} from XXX. We then estimate the placebo regression

$$C_{bj} = \delta_C^p H_{Cbj} + f_-^p \left(d_{Cbj}\right) \times 1 \left(d_{Cbj} < \xi_{Cb}\right) + f_+^p \left(d_{Cbj}\right) \times 1 \left(d_{Cbj} > \xi_{Cb}\right) + \epsilon_{Cbj}^p \tag{6}$$

which can be understood as an analog to equation (5) in which each boundary section has been perturbed by ϵ_{Cb} towards the high side. If our identifying assumption is satisfied, then we would expect our estimate of δ_C^p to be equal to zero.

4 Data

To estimate the boundary discontinuities δ_C, we construct a dataset that is comprised of Census Block groups that are geospatially merged to the nearest segments of each boundary type. Latitude and longitude estimates for the center of population of each Block group is provided by the National Historical Geographic Information System (NHGIS), and we use ARCGIS software to map population centers to 2 mile segments of each boundary type. The rich set of publicly available data at the Census Block group level allows us to describe the spatial area surrounding each boundary segment along more than one thousand dimensions. We describe both the boundary network data and neighborhood variables in further detail.

4.1 Boundary Networks

We analyze three classes of boundaries: transportation networks, educational boundaries, and political boundaries. Transportation networks form physical boundaries as they deform the urban landscape. They are costly to cross and often delineate distinct neighborhoods. We consider historical rail and highway networks. The US railway network peaked at 254,000 miles of track in the early twentieth century, and today it is comprised of approximately 160,500 miles of track.[4] We measure the historical US rail network using the Atack [2013]

[4] American Association of Railroads, *Chronology of America's Freight Railroads. https://www.aar.org/chronology-of-americas-freight-railroads/.*

historic GIS transportation database created from the New Century Atlas maps published in 1911. This includes all passenger and freight rail lines that were in operation circa 2011. The Interstate Highway System stretches nearly 50,000 and is part of a larger network that includes state highways. For the analysis we employ a digitized map of only interstate highways made publicly available by the US Department of Transportation as of 2020.[5] Figures 4 and 5 lay bare the spatial features of the transportation data, and illustrates the relative density of the US railroad network compared to the interstate highway network.

School district and school attendance zone boundaries are analyzed using shape files provided by NHGIS. The National Center for Education Statistics (NCES) conducts an annual update of school district boundaries dating back to 1995, and we obtain boundaries from the 2020 update. Figure 6 shows that our school district sample is nationally representative and reflects the political economy of different states. In Southern states it is common for school districts to span the entire county, with an extreme example being Florida where counties and districts cover expansive areas. In Northern states like New Jersey and Ohio, counties are divided into cities and small townships that manage school districts. Unlike school districts, there is scant nationally representative spatial data for school attendance zones. NHGIS hosts the 2009-2010 shape file created by the NSF funded SABINS project. Figure 7 shows the national coverage of our attendance boundaries, which we restrict to elementary schools.[6]

Finally, we consider political boundaries in the form of county lines. The publicly available shape file is hosted by NHGIS and is the Census Tigerline 2020 US County map.

[5] Our analysis omits segments of interstate < 0.5 mile long. Additionally, state managed highways are not present in the shape file.

[6] For our purposes, a school is defined as an elementary school if it enrolls third grade students.

Boundaries for all counties in the contiguous US are shown in Figure 8. While the county shape file has near full national coverage (over 3000 counties), a large share of county boundaries exist in remote areas. We mitigate this sampling issue in two ways for all boundaries. By imposing that each boundary strip be no longer than 2 miles, we restrict the distance that two Block groups in remote areas can be mapped to the same boundary. Furthermore, our research design requires both sides of each boundary segment to have a sufficicent number of Block groups within one mile. In practice, this means that most rural county boundaries are not used for estimation. In Section 4.3, we present summary statistics for the Block groups near each transportation network and compare the samples to the universe of Census Block groups.

4.2 Neighborhood Characteristics

Block groups are small Census-designated geographies with between 600 and 3,000 residents. We obtain data at the Block group level from the 2020 Decennial Census and the 2016-2020 American Community Survey (ACS). The two sources provide over 1,700 variables that characterize the people, housing, and amenities for each Block group in the sample contiguous 48 states.

Our analysis relies on a spatial assignment of Block groups to (specific sides of) boundaries. NHGIS provides estimates of the latitude and longitude for each Block group center of population, which we use as a proxy for the location of each neighborhood. Each boundary network is divided into two mile segments, and neighborhoods within a one-mile radius of each segment are mapped to that boundary. This yields cross-sections of neighborhood level characteristics for each of our five boundary networks.

4.3 Sample Description

In Table 1, we compare the Block groups (neighborhoods) in each boundary network cross-section to the universe of Block groups in the contiguous US. The neighborhoods in columns 2-5 are within one mile of each boundary type, with the sample restricted to boundaries with at least two block groups on either side. Millions of people live within a mile of our sample boundaries with the most being near railroads (98.7 million) and the least living near county lines (25.9 million). Neighborhood population means range from 1,674 to 1,859 compared to a population average of 1,776.

Differences between the neighborhoods in each boundary network sub-sample and the full sample reflect the fact that our sample restrictions exclude rural areas. This implies higher house prices and less White populations in all sub-samples than in the full sample. Urban incomes are higher, and areas with denser county lines are more likely to be in the Northeast. However, the disamenity of living very close to physical boundaries implies lower incomes near railroad and highway boundaries.

Because our analysis compares block groups within boundary type, differences in means across samples would only matter if the goal was to compare empirical results by boundary type for inference about behavioral differences. Our analysis does not rest on determining which boundary produces the strongest behavioral response.

5 Results

Figure 2 presents histograms of the estimates of boundary effects (δ_C) for all neighborhood characterisics and five different boundary types. For each bounday type, we present a his-

Table 1: Summary Statistics

	Full Data	Railroad	Interstate	School District	School Zone	County Line
Variables	(1)	(2)	(3)	(4)	(5)	(6)
Population	1,775.72	1,673.50	1,752.93	1,738.72	1,859.47	1,738.77
	(862.0)	(786.3)	(806.1)	(798.1)	(909.6)	(822.7)
House Prices	309,455.81	318,913.88	356,345.96	342,514.65	318,225.34	361,239.23
	(261229.1)	(282986.5)	(293976.6)	(280384.3)	(237131.8)	(292144.2)
Income	79,729.77	75,243.12	79,682.70	85,568.20	82,652.66	89,051.29
	(40723.4)	(40989.6)	(41770.9)	(43935.7)	(42067.4)	(45393.4)
Public Asst.	0.11	0.13	0.12	0.11	0.11	0.10
	(0.121)	(0.134)	(0.130)	(0.123)	(0.121)	(0.117)
Share White	0.61	0.56	0.51	0.57	0.55	0.59
	(0.304)	(0.308)	(0.305)	(0.303)	(0.293)	(0.299)
Share Latino	0.18	0.21	0.22	0.19	0.21	0.15
	(0.233)	(0.247)	(0.250)	(0.238)	(0.242)	(0.194)
Share Black	0.11	0.13	0.15	0.12	0.13	0.15
	(0.198)	(0.215)	(0.224)	(0.205)	(0.207)	(0.233)
Own Home	0.31	0.39	0.39	0.33	0.33	0.33
	(0.241)	(0.249)	(0.258)	(0.250)	(0.247)	(0.257)
Single Family	0.74	0.68	0.67	0.72	0.74	0.69
	(0.266)	(0.283)	(0.300)	(0.287)	(0.278)	(0.311)
Rooms / Unit	5.91	5.61	5.57	5.91	5.91	5.94
	(1.378)	(1.325)	(1.385)	(1.479)	(1.483)	(1.586)
Observations	190900	56444	41507	30286	34509	12956
Pop. Estimate	306M	98.7M	74.4M	59.1M	63.5M	25.9M

Notes: Means and standard deviations (in parentheses) for the full sample of 2020 Census Block groups (column (1)) and the five subsamples of Block groups for each corresponding boundary network (columns (2)-(6)).

togram of all $\hat{\delta}_C$ and we overlay a histogram of all $\hat{\delta}_C$ that are significant at the 95% level. In order to facilitate comparison across estimates, we normalize the standard deviation of all dependent variables to 1.[7]

For the two physical boundaries, railroads and highways, all statistically significant boundary effects are positive, which supports our selection procedure for H_{Cbj} (the "high side"). Moreover, the vast majority of economically significant effects (defined as $\hat{\delta}_C > 0.05$) are statistically significant. The range of effects for highways and railroads is similar, though railroad effects tend to be slightly larger than highway effects.

For the two administrative educational boundaries, we estimate even larger effects than for the physical boundaries. Once again, all statistically significant boundary effects are positive, and the vast majority of economically significant effects are statistically significant. Meanwhile, for the administrative political boundary, we estimate fewer statistically significant effects, though the qualitative patterns from other boundaries are reproduced. This is possibly due to the fact that the majority of county borders are in rural areas that offer fewer observations for estimation.

[7]For all results in the main text, we model the relationship between each outcome and distance to the boundary as a third order polynomial function (i.e., f_+ and f_- are cubic polynomials).

Appendix Figure 9 suggests that our estimates of δ_C could vary based on modelling assumption. As a specification test we replicated all results assuming f_+ and f_- were linear, which we summarize in Appendix Table 5. A comparison with our main results in Table 2, shows that our cubic model is more conservative in the sense that we estimate fewer statistically significant outcomes.

Figure 2: Boundary Effects

(a) Railroads

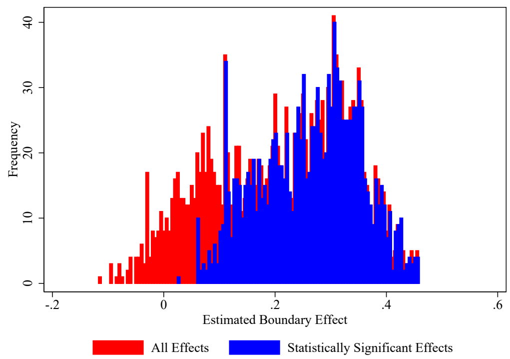

(b) Highways

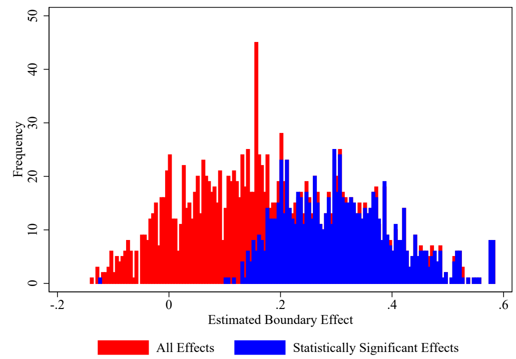

(c) School Districts

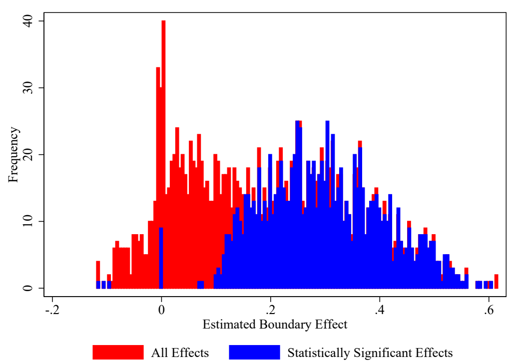

(d) School Attendance Zones

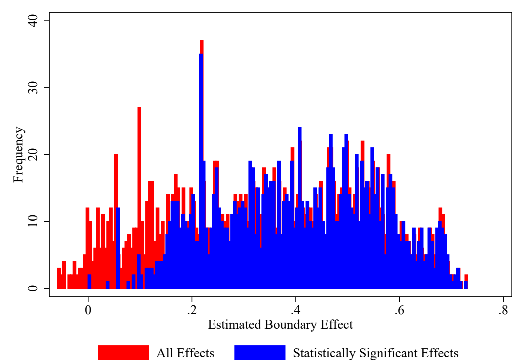

(e) Counties

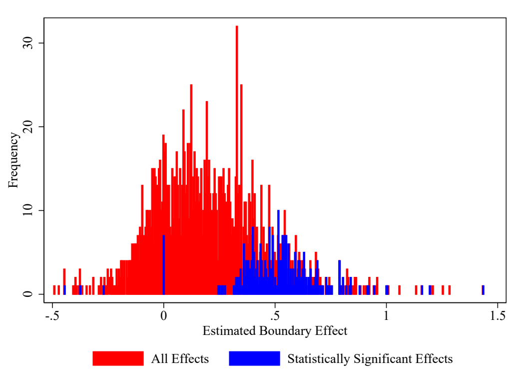

Note: Statistical significance is at the 95% level.

In order to ensure that our results are not simply a statistical artifact of our estimation and testing procedure, we reproduce these histograms using estimated effects from our placebo validation exercise in Figure 3. Three observations are immediate for all boundary types: (1) The overwhelming majority of placebo effects are not statistically significantly different than zero. (2) The distributions of placebo effects are roughly symmetric around zero.[8] (3) The magnitudes of estimated placebo effects are substantially smaller than the magnitudes of the estimated boundary effects. All three observations strongly support our identification strategy.

[8] The slight skew to the right is likely due to our selection procedure for the "high side."

Figure 3: Placebo Boundary Effects

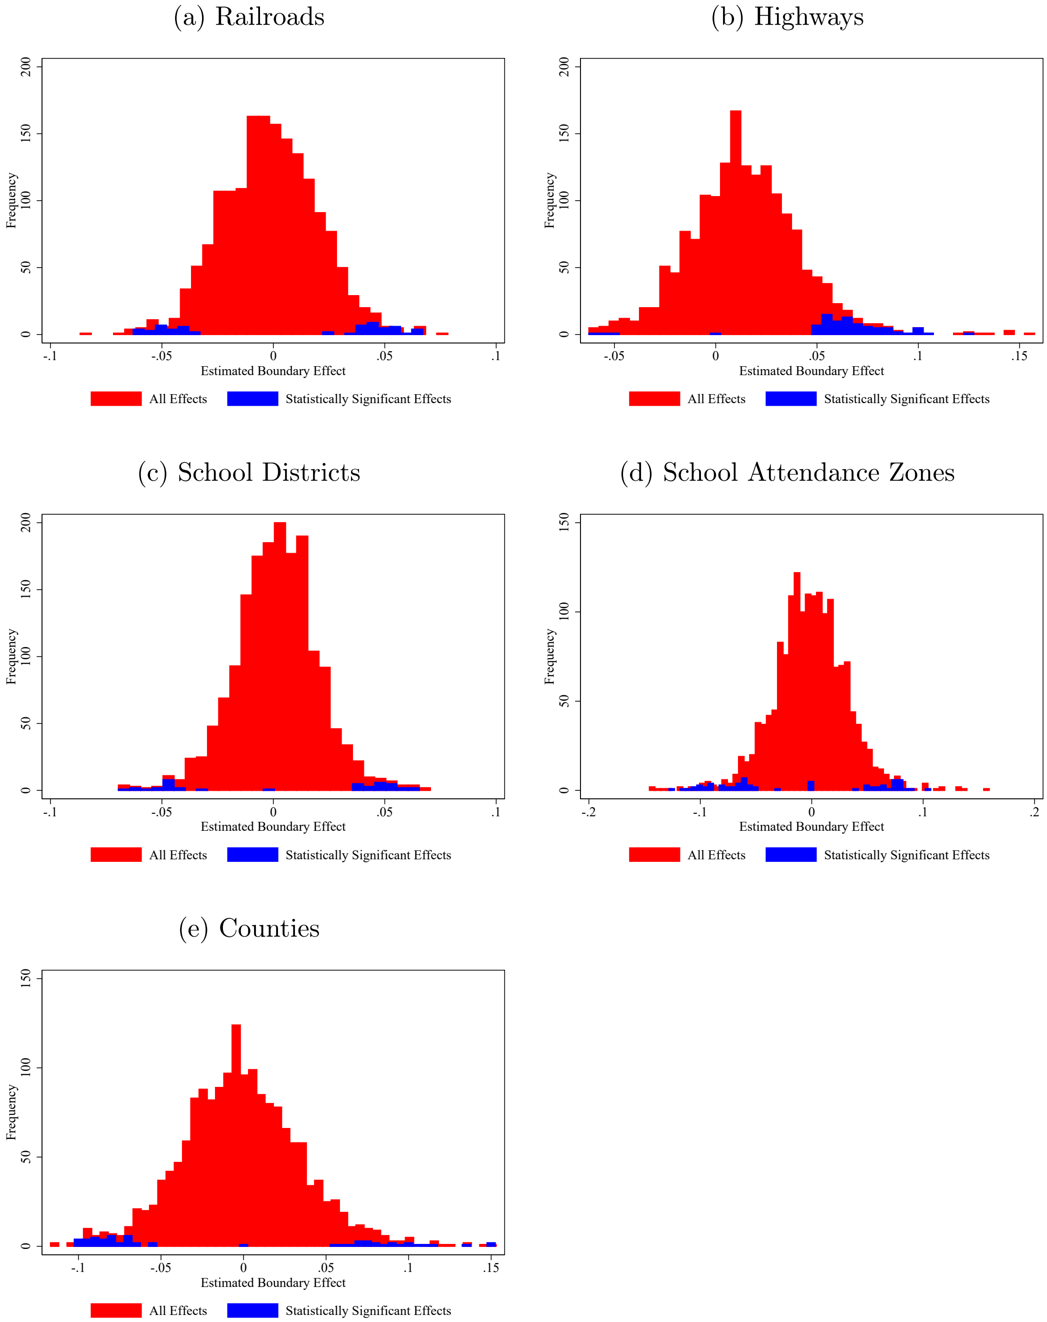

Note: Statistical significance is at the 95% level.

Table 2: Summary of Results

Boundary Type	(1) Frac. Statistically Significant Effects (All)	(2) Avg. Statistically Significant Effect Size (All)	(3) Frac. Statistically Significant Effects (Selected Variables)	(4) Avg. Statistically Significant Effect Size (Selected Variables)
Railroad	0.756	0.264	0.783	0.271
Highway	0.508	0.301	0.531	0.306
School District	0.610	0.262	0.643	0.304
School Attendance Zone	0.810	0.403	0.841	0.405
County	0.167	0.471	0.178	0.501

Notes: In columns (1) and (2), we present stastics for $\hat{\delta}_C$ for the entire sample of neighborhood characteristics. In columns (3) and (4), we present statistics for $\hat{\delta}_C$ for the subsample of non-negative neighborhood characteristics that have fewer than 10% missing observations and no extreme outliers. Statistical significance is reported at the 95% level.

We summarize these results in Table 2. The fraction of statistically significant effects in column (1) is substantially greater than 5%, which supports our theory that boundaries should generate systematically generate discontinuities in neighborhood characteristics. From column (2), we see that these effect sizes are of similar size for railroads and highways, slightly larger for school districts and counties, and roughly 50% larger for school attendance zones. This may capture a greater degree of household sorting across school zones than other boundaries.

To ensure that these results are not driven by a large number of Census variables being uninformative, in columns (3) and (4) we consider only a subset of variables for which fewer than 10% of observations are missing, all entries are numeric, no observations accidentally take on negative values, and no variables contain extreme outliers (a coefficient of variation over 2000). This eliminates roughly 20% of all neighborhood characteristics from our analy-

Table 3: Summary of Placebo Tests

Boundary Type	(1) Fraction of Statistically Significant Placebo Effects (All)	(2) Average Statistically Significant Placebo Effect Size (All)
Railroad	0.039	0.045
Highway	0.053	0.065
School District	0.026	0.035
School Attendance Zone	0.043	0.020
County	0.038	0.010

Notes: In columns (1) and (2), we present stastics for $\hat{\delta}_C$ for the entire sample of neighborhood characteristics. In columns (3) and (4), we present statistics for $\hat{\delta}_C$ for the subsample of non-negative neighborhood characteristics that have fewer than 10% missing observations and no extreme outliers. Statistical significance is reported at the 95% level.

sis. Nevertheless, we find a similar prevalence and size of significant effects as we did in the unrestricted sample of columns.

Finally, we summarize our placebo tests in columns (1) and (2) of Table 3. All entries in column (1) are close to 5%, which is what we would expect when considering significance at the 95% level. (Importantly, all entries in columns (1) and (4) of 2 are much larger than 5%.) Meanwhile, column (2) shows that even the few statistically significant placebo effects that we do find are not economically significant.

6 Discussion

Our finding of robust, widespread boundary discontinuities invites caution to practitioners seeking to implement regression discontinuity designs at geographical boundaries, which are commonly known as boundary discontinuity designs (BDD) or geographic regression discontinuity (GRD) designs. Using the notation above, in such a design, a researcher hopes to estimate the effect of a treatment on an outcome C_{bj} in which treated units j lie to one side of a boundary b ($H_{Cbj} > 0$) and untreated units lie to the other side of a boundary ($H_{Cbj} < 0$). The primary idenfiying assumption in such a design is that the average treatment effect should be continuous as we approach the boundary from either side (Keele and Titiunik [2015]). A common test of the this assumption is to demonstrate that other potential confounders C'_{bj} that are observed to the researcher do not vary discontinuously at the boundary, with the implication being that any discontinuity that is estimated at the boundary can then be attributed to treatment.

Our results suggest that such an assumption is unlikely to be satisfied in the practice. Yet, this is seemingly at odds with the covariate balance tests that often accompany such designs (e.g., Bayer et al. [2007], Gibbons et al. [2013]). We argue that this is due to the fact that tests of discontinuities in neighborhood characteristics are likely to be underpowered when the sample is a single city, metropolitan area or even state. To make this point explicitly, we replicate our anlaysis separately for 9 large Core-based statistical areas (CBSA) in the United States using school attendance areas as boundaries and summarize our results in Table 4.[9] While the estimated effect sizes for each CBSA are similar to those of the nation as a whole, the precision of these estimates is dramatically smaller. Instead of finding statistically significant discontinuities in over 80% of neighborhood characteristics, we find statistically significant discontinuities in 10-25% of neighborhood characteristics depending

[9] We select these 9 CBSAs because they have the largest number of recorded school attendance boundaries and the largest number of observations that can be used for estimation of boundary discontinuities. In this sense, these results should be seen as conservative since they are the most stastically powered subsamples.

Table 4: School Attendance Zone Boundary Discontinuities by City

City (CBSA)	(1) Fraction of Statistically Significant Effects (All)	(3) Fraction of Statistically Significant Effects, Selected Variables
All US	0.810	0.841
Atlanta	0.192	0.179
Washington, DC	0.126	0.124
Denver	0.107	0.106
Houston	0.187	0.192
Miami	0.242	0.222
Minneapolis	0.152	0.146
Philadelphia	0.192	0.183
Riverside	0.185	0.193
Tampa	0.162	0.161

Notes: In column (1) we present stastics for $\hat{\delta}_C$ for the entire sample of neighborhood characteristics, and in column (2) we present statistics for $\hat{\delta}_C$ for the subsample of non-negative neighborhood characteristics that have fewer than 10% missing observations and no extreme outliers. Statistical significance is reported at the 95% level.

on city. Hence, researchers should be careful when relying on such tests of (the lack of) discontinuities in potential confounders to validate their research designs.

We should note that our findings do not imply that boundary discontinuity designs are never appropriate. Instead, they indicate that these designs are better suited for estimating certain classes of treatment effects. The fact that we find discontinuities in such a broad set of neighborhood characteristics suggests some (endogenous) relationship between these characteristics. As alluded to in Section 2, an attractive basis for this relationship is some endogenous process of sorting of households and suppliers of amenities. Of course, such a process takes time to unfold. As a result, boundary discontinuity designs may be more successful at identifying short-run treatment effects where the outcome is observed fairly soon after the introduction (or removal) of a boundary.

7 Conclusion

Boundaries are ubiquitous and unavoidable. Physical structures, both natural and man-made, distort the urban landscape. So too do administrative boundaries that allow for the differentiation of public goods. In this paper, we present a simple model that yields the prediction that these distortions will manifest as discontinuities in neighborhood characteristics across boundaries of many types. We then show that a comprehensive set of neighborhood characteristics – the universe of publicly available characteristics in the decennial Census – exhibit discontinuities at a broad set of physical, administrative and political boundaries. These discontinuities are sizeable, systematic, and not merely statistical artifacts of how spatial data are collected.

Given these findings, we argue that the popular boundary discontinuity design should be applied with caution as its core identifying assumption may not hold in certain settings, and a standard validation exercise of this assumption is, in practice, probably underpowered to draw a meaningful conclusion. This yields an insight that should be taken to heart by both policymakers and researchers. Although the short-run effects of boundaries may be narrow, the long-run effects of boundaries are likely to be broad in scope, even if the treatment induced by the boundary is very narrow. Shifting a school attendance boundary has the scope to affect far more than educational outcomes; adding a highway will affect neighborhoods in far more profound ways than changing traffic patterns; past institutional boundaries such as redlines that are no longer in effect may still generate dramatic discontinuities in the present day. Policymakers would be wise to consider these knock on effects when assessing if and where to place boundaries. And researcher should perhaps trade-off the hope of using boundaries for the sharp identification of narrow treatment effects for the prospect of using boundaries to explain a broader set of spatial phenomena.

References

Elizabeth Oltmans Ananat. The wrong side (s) of the tracks: The causal effects of racial segregation on urban poverty and inequality. *American Economic Journal: Applied Economics*, 3(2):34–66, 2011.

Jeremy Atack. On the use of geographic information systems in economic history: The american transportation revolution revisited. *The Journal of Economic History*, 73(2): 313–338, 2013.

P. Bayer, R. McMillan, and K. Rueben. An equilibrium model of sorting in an urban housing market. Technical report, National Bureau of Economic Research, 2004.

P. Bayer, F. Ferreira, and R. McMillan. A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638, 2007.

S.E. Black. Do better schools matter? parental valuation of elementary education. *The Quarterly Journal of Economics*, 114(2):577–599, 1999.

Gregorio Caetano. Neighborhood sorting and the value of public school quality. *Journal of Urban Economics*, 114:103193, 2019.

Gregorio Caetano and Vikram Maheshri. A unified empirical framework to study neighborhood segregation. *Working Paper*, 2023.

Raj Chetty, John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter. The opportunity atlas: Mapping the childhood roots of social mobility. Technical report, National Bureau of Economic Research, 2018.

Eric Chyn, Kareem Haggag, and Bryan A Stuart. The effects of racial segregation on intergenerational mobility: Evidence from historical railroad placement. Technical report, National Bureau of Economic Research, 2022.

Robynn Cox, Jamein P Cunningham, Alberto Ortega, and Kenneth Whaley. Black lives: The high cost of segregation. *University of Southern California Working Paper*, 2022.

Melissa Dell. The persistent effects of peru's mining mita. *Econometrica*, 78(6):1863–1903, 2010.

D. Epple, T. Romer, and H. Sieg. Interjurisdictional sorting and majority rule: An empirical analysis. *Econometrica*, 69(6):1437–1465, 2003.

Dennis Epple, Radu Filimon, and Thomas Romer. Equilibrium among local jurisdictions: toward an integrated treatment of voting and residential choice. *Journal of Public Economics*, 24(3):281–308, August 1984.

William G Gale. Public finance and racism. *National Tax Journal*, 74(4):953–974, 2021.

Stephen Gibbons, Stephen Machin, and Olmo Silva. Valuing school quality using boundary discontinuities. *Journal of Urban Economics*, 75:15–28, 2013.

Luke J Keele and Rocio Titiunik. Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1):127–155, 2015.

Charles Tiebout. A pure theory of local expenditures. *The Journal of Political Economy*, 64 (5):416–424, 1956.

Yves Zenou and Nicolas Boccard. Racial discrimination and redlining in cities. *Journal of Urban economics*, 48(2):260–285, 2000.

Online Appendix

A Data and Analysis

A.1 Data and Sample Construction

Figure 4: RailRoads

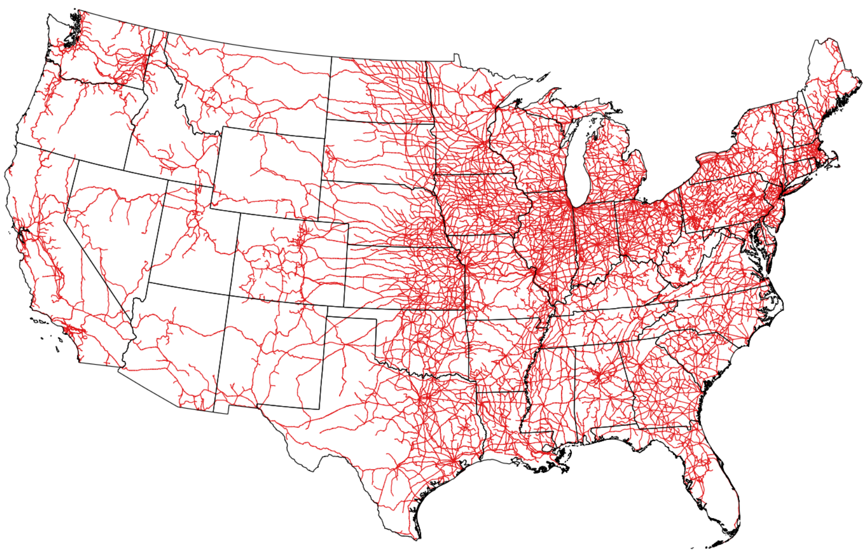

Figure 5: Highways

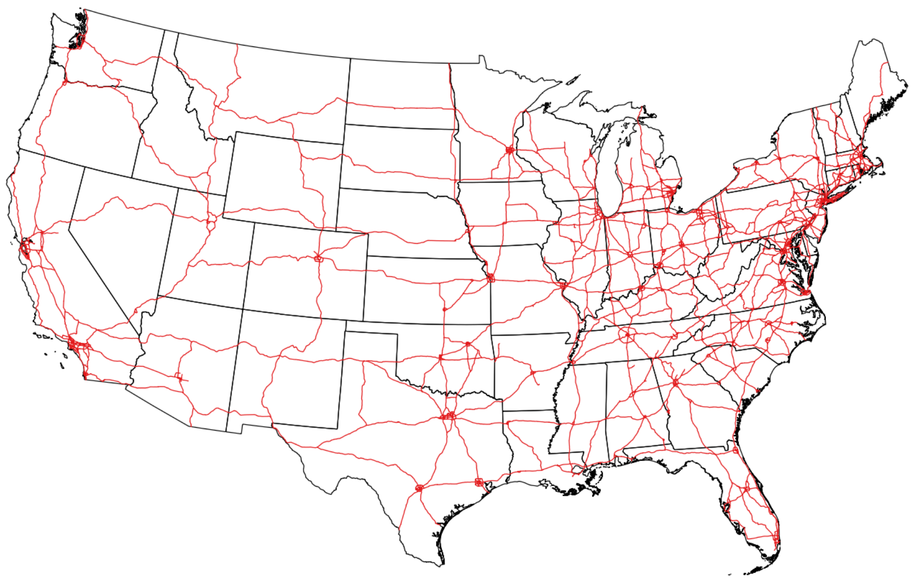

Figure 6: School Districts

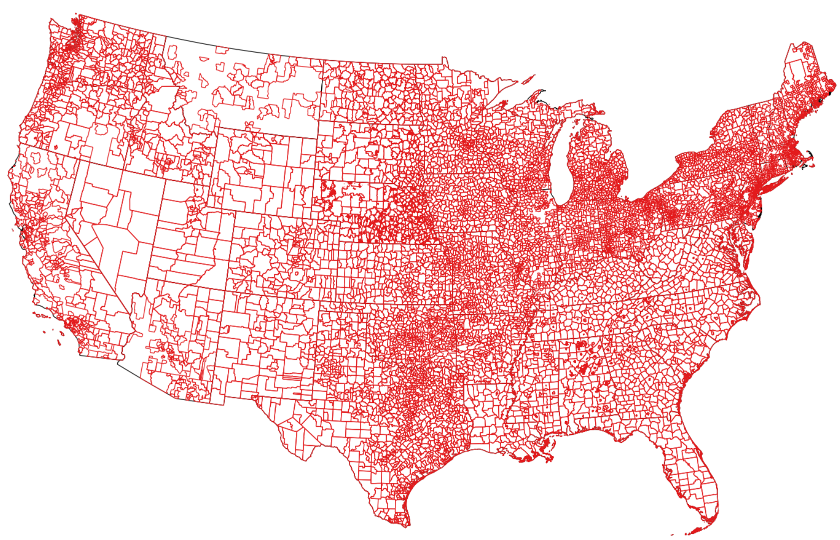

Figure 7: School Attendance Zones

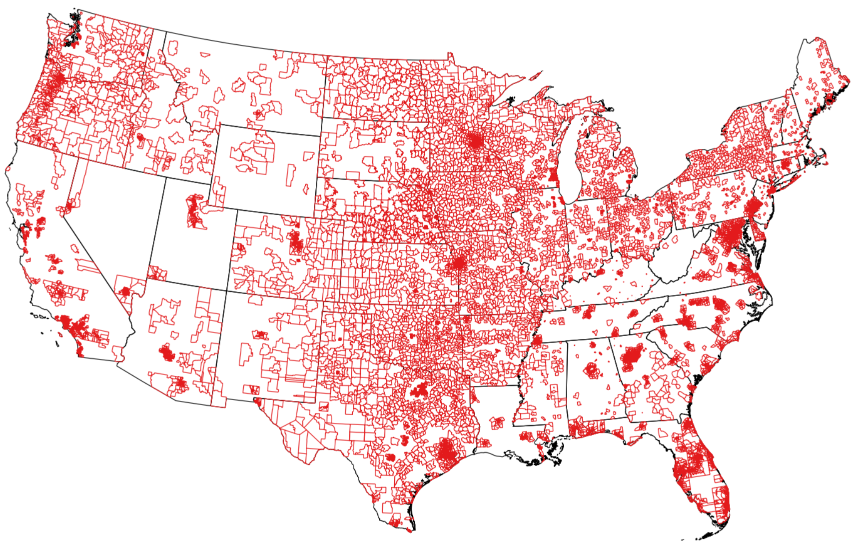

Figure 8: Counties

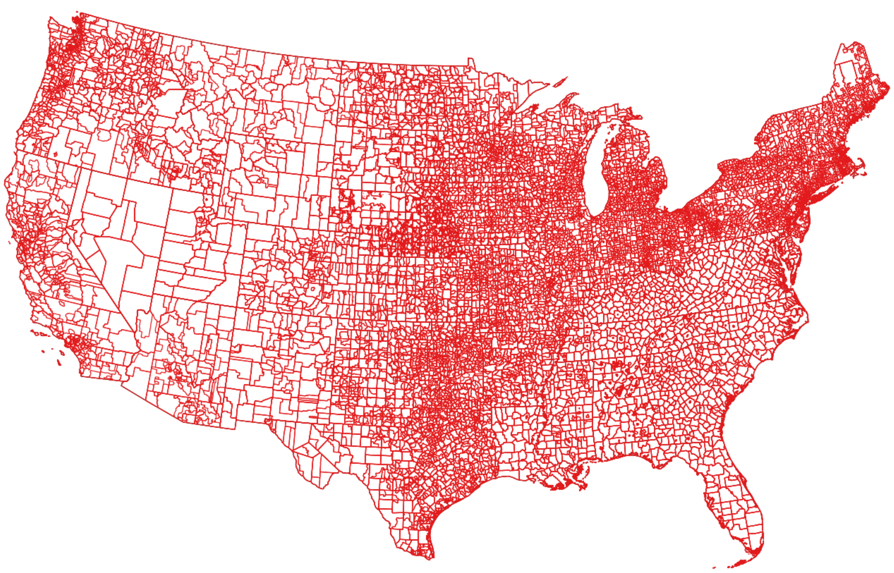

A.2 Robustness of Model Specification

Figure 9: Mean Home Values by Distance to the Boundary, 2020 Census

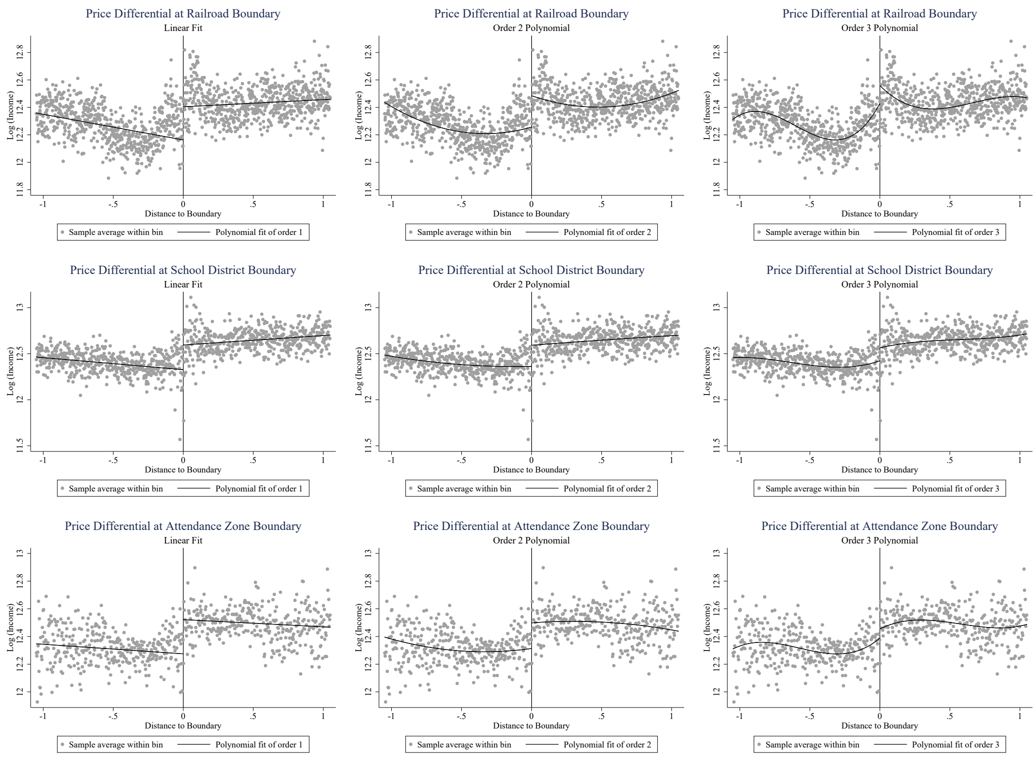

Table 5: Summary of Results

Boundary Type	(1) Fraction of Statistically Significant Effects (All)	(2) Average Statistically Significant Effect Size (All)	(3) Fraction of Statistically Significant Effects, Selected Variables	(4) Average Statistically Significant Effect Size, Selected Variables
Railroad	0.835	0.228	0.873	0.232
Highway	0.623	0.223	0.646	0.222
School District	0.731	0.251	0.759	0.256
School Attendance Zone	0.899	0.375	0.930	0.377
County	0.268	0.271	0.283	0.294