# The Ubiquity of Tiebout Sorting Across Neighborhood Boundaries

Vikram Maheshri and Kenneth Whaley\*

May 17, 2025

#### 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 in a wide range of observable dimensions – the universe of variables available in the 2020 Census at the Block group level – and six different types of boundaries. We draw two important conclusions from these findings: (1) researchers should implement boundary discontinuity designs with caution because the key identification assumption may not hold except in narrow applications, and (2) even narrowly targeted place-based policies may have much broader impacts if they involve a new administrative boundary. In the case of school zones, where we find the strongest evidence of Tiebout sorting, focusing on the house price capitalization of school quality alone will understate the true cost to access better schools and neighborhoods in US housing markets.

<sup>\*</sup>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

Cities are subdivided into neighborhoods by boundaries that can be natural, physical, political and administrative. In principle, boundaries may partition cities into noticeably distinct neighborhoods characterized by the interactions of households, firms, and local governments. Alternatively, boundaries may simply exist in parallel to the lived experiences of local residents and go otherwise unnoticed. In the former view, a rich literature following from Tiebout (1956) has studied both the process of<sup>1</sup> and the consequences of<sup>2</sup> neighborhood sorting. Fundamental to this literature is the idea that neighborhood boundaries represent more than simply lines on a map, they induce distortions that impact decision makers, making it difficult to attribute differences across boundaries to a single source. Meanwhile, a relatively young literature in empirical social science has embraced the alternative view in hopes of leveraging boundaries for causal inference. <sup>3</sup> Ultimately, it is an empirical question as to which of the two views best describes reality, the answer of which contributes to our fundamental understanding of the distribution of economic activity across space (Clark et al. (2003)).

The extent to which boundaries propagate Tiebout sorting highlights a tension between behavioral responses of economic agents and policy evaluation using boundaries. In spatial econometrics, boundaries represent geographic thresholds where differences in outcomes are caused only by a treatment of interest. So long as unobserved confounders vary smoothly across the boundary space, the models produce estimates that meet a high bar for internal validity (Keele and Titiunik (2015); Butts (2023)). The challenge for researchers is the long-held intuition that adjacent neighborhoods separated by boundaries differ in a variety of ways,<sup>4</sup> and only rough proxies for neighborhood demographics are typically available to rule out Tiebout sorting. Notwithstanding, the rich tapestry of boundaries in the urban landscape lends to research evaluating access to public goods and services, exposure to particular place-based interventions, proximity to private

<sup>&</sup>lt;sup>1</sup>See, for example, Epple et al. (1984, 2003); Bayer et al. (2004); Caetano (2019); Caetano and Maheshri (2023b) <sup>2</sup>See, for example, Black (1999); Bayer et al. (2007); Chetty et al. (2018)

<sup>&</sup>lt;sup>3</sup>The strategy of comparing groups on opposite sides of a boundary to identify the effects of an intervention underlies a broad class of empirical approaches that includes boundary fixed effects, spatial difference-in-differences, and boundary discontinuity designs.

<sup>&</sup>lt;sup>4</sup>The notion of a "wrong side of the railroad tracks" is at least a century old, having found its way into a wedding announcement in the *Daily Inter Ocean* (Chicago) in 1903.

consumption opportunities, residential segregation, and network effects.

It is perhaps surprising then that comprehensive empirical evidence of Tiebout sorting (or the lack thereof) at neighborhood boundaries does not exist despite the broad appeal of spatial models exploiting boundaries for identification. This could be because the estimation of boundary models is data intensive, and the comparative advantage of boundary designs - sharp inference about a particular policy affecting a small area of interest - offers little incentive to conduct comprehensive analysis. This second reason falls under a more general critique of the external validity of boundary designs, as researchers must only rule out sorting and other potential confounders unique to the setting of interest, trading off the license to make broader claims.

In this paper we present a simple theoretical framework and a wealth of evidence to understand how boundaries cause Tiebout sorting and its empirical implications. Our approach begins with a stylized model that illustrates how boundaries distort the location decisions of households, which in turn distort the amenities supplied by private firms and the public goods provided by local governments. The key empirical insight of the model is that these distortions will manifest as discontinuities in the demographic and amenity bundle of neighborhoods and ultimately in neighborhood prices. We then test for the existence of spatial discontinuities in the universe of all publicly available variables from the US Census, analyzing the demographics of Block groups near six very different types of boundaries: historical rail and highway networks, contemporary school district boundaries and attendance zones, county lines, and ZIP code boundaries. Consistent with predictions of the model, we find overwhelming evidence of discontinuities across a wide range of variables and all types of boundaries.

Our model is a simple formalization of the insights of two seminal papers on sorting. Following Hotelling (1929), we start from the notion that certain consumers wish to locate close to certain producers in horizontally differentiated markets. We then incorporate Tiebout (1956) sorting stemming from heterogeneity in public goods supplied across neighborhoods. The basic insight of model is that households of a given type will cluster on the same side of a boundary that provides proximity to their preferred private and public goods. A core prediction is that this clustering will generate a discontinuity at the boundary; the endogenous responses of households, firms, and local

governments create a feedback loop that amplify this discontinuity and hence the causal effects of boundaries on local demographics and amenities. We show that the magnitude of this problem for boundary designs depend on the timing of boundary placement (or treatment assignment to each side) relative to the timing of the explanatory variable of interest. Boundary discontinuity designs may be appropriate to identify the effects of variables that change (and are observed) immediately after the placement (or deletion) of a new boundary. However, as the length of time between boundary placement and the period in which the variable of interest is observed increases, researchers will have to impose stronger and stronger identifying assumptions that are increasingly at odds with the empirical evidence that we present.

Following our model, the econometric framework to test for Tiebout sorting at boundaries simply involves the estimation of a boundary discontinuity design. We implement this test with a dataset comprised of detailed maps of the six aforementioned boundaries throughout the contiguous US and all variables in all Census Block Groups in the US in 2010 and 2020. The average Block Group contains 1395 residents, and there are over 1700 Census variables that describe their demographic characteristics along with features of the housing stock, labor market profiles, and measures of government assistance take-up. Each block group is assigned to a nearby boundary based on distance, and each census variable is taken as an outcome of the model, separately for each boundary type. All types of neighborhood characteristics systematically vary discontinuously at all types of boundaries, with the largest effects found at school zones, where 81% of chracteristics vary discontinuously, and railroad tracks, where 75% of chracteristics vary discontinuously. We also find robust evidence of Tiebout sorting at zip code (63%) and school district (61%) boundaries, and weaker evidence at highway (51%) and county lines (17%). We conduct a series of tests on placebo boundaries to show that our results are not simply statistical artifacts of a high dimensional testing procedure. Additionally, we show that our results are robust to a series of specification tests for sensitivity to bandwidth selection and RD polynomial choice.

These findings have critical implications for a broader class of strategies to identify causal effects with boundaries. Studies that seek to identify the effects of boundaries *per se*, will successfully identify the effects of boundaries on outcomes that are mediated through a broad set of

sorting mechanisms. For example, when boundaries shape historic institutions, the effects may persist today through a variety of channels (Dell (2010); Dell et al. (2018); Cox et al. (2022)). However, when the object of interest is a single attribute that changed historically at a boundary, researchers must soften causal claims regarding contemporary outcomes, even if the data pass preperiod balance tests. Such is the case of property rights (Baragwanath and Bayi (2020)), redlining (Aaronson et al. (2020)), and residential segregation (Monarrez and Schönholzer (2023); Whaley (2024)).

A second strategy involves spatial difference-in-differences assessment of a well-defined policy assigned to one side of an administrative boundary, as with place-based investments (Lu et al. (2019); Albertus (2020); Jia et al. (2021)).<sup>5</sup> Our results suggest that this approach will capture both the intended policy effect plus the knock-on effects of sorting. This implies that stable differences in neighborhoods prior to the treatment (supporting, say, a parallel trends assumption) may not be sufficient to rule out spatial spillovers that affect outcomes of interest after the policy is in place. This is in line with Jardim et al. (2024), who argue that geographically diffuse spillovers of a minimum wage enacted on one side of a boundary affect wages and working hours on the opposite side, which contaminates the estimated impacts of the policy itself. This is not to say that boundary designs are always problematic; indeed, their use can be appropriate to estimate the effect of a single factor when the outcomes of interest are measured soon after boundaries are established (or altered), or soon after a policy is put into place.

The remainder of this paper is organized as follows. In Section 2, we present a simple model of consumer 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 describe the various sources of data that we use to implement this approach. In Section 4 we present our results with an application to house prices at school attendance boundaries that highlights our opposition to treating boundaries with a narrow scope. We find that house prices increase 25-33% on the high quality side of a school zone, a substantially larger estimate than the house price capitalization of school quality alone, and far more indicative of the magnitude and consequences of Tiebout sorting over schools. We

<sup>&</sup>lt;sup>5</sup>Our discussion of spatial difference-in-differences includes spatial differences-in-discontinuies (Butts (2023))

conclude in Section 5.

# 2 Conceptual Framework

We motivate how boundaries generate discontinuities in households' location choices in a simple, 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

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

where the parameter  $\alpha_i$  is drawn from a single peaked distribution over (0, 1). We assume that u' < 0 and u'' > 0. That is,  $\alpha_i$  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

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



Notes: Panel A illustrates the distribution of optimal consumer location choices  $f_X^*$ , given heterogeneous preferences over distance to producers  $y_0^*$  and  $y_1^*$ . Panel B and C illustrate predictions from a model of consumer behavior when a boundary occurs at point B, relative to the initial distribution shown by the dotted line.

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

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

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

All consumers with  $x_i^* < B$  before the introduction of the boundary remain to the left of the boundary and vice versa. 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 this 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.<sup>6</sup> 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. Moreover, these conclusions will persist and likely strengthen if we endogenized the locations of boundaries or producers or if we endogenized the levels of public 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  $\alpha_i$  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,

<sup>&</sup>lt;sup>6</sup>To simplify notation, we assume consumers exactly at the boundary can choose the side of the boundary to which they locate.

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 summarize how each of these mechanisms contributes 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 entirely continuous, the ultimate effect of a boundary on the urban landscape will be discontinuous in the characteristics of residents, characteristics of the amenity bundle, and prices.

We use another diagram to illustrate how this sorting process affects identification of causal effects of either boundaries or neighborhood characteristics affected by them. Consider a boundary that is placed at period  $t_0$ , with  $\vec{x}$  representing a vector of neighborhood characteristics – both observed and potentially unobserved – in periods  $t_1, t_2, \ldots, t_T$ .<sup>7</sup> Solid arrows indicate continuous causal relationships, and dotted arrows indicate discontinuous causal relationships. Suppose  $x^1$  is a variable that the boundary is intended to directly influence (e.g., a local tax rate for an administrative boundary). Then we would expect a discontinuity in  $x^1$  on either side of the boundary. While it is possible that the boundary might be correlated to other neighborhood

<sup>&</sup>lt;sup>7</sup>In principle,  $t_0$  could equal  $t_1$ .

characteristics, it need not generate a discontinuity in those characteristics. However, the sorting of people, firms and amenities in periods  $t_1$  and thereafter imply dynamic linkages between the variables  $x^1, \ldots, x^N$ . Nevertheless, a researcher could estimate the causal effect of an observed  $x_{t_1}^1$ on any observed variable  $x_{t_i}^n$ ,  $n = 1 \ldots N$ , i > 1 using a boundary discontinuity design, as any discontinuity in  $x_{t_i}^n$  at the boundary could only be due to  $x_{t_1}^1$ .



Suppose, however, that  $x^1$  could only be observed starting in period  $t_2$ . We would then re-draw this diagram as



Note that the relationship between the boundary in  $t_0$  and *all* variables in  $t_2$  is discontinuous since these effects are mediated, at least in part, through the unobserved  $x_{t_1}^1$ . In this setting, a boundary discontinuity design would not be valid framework to identify the causal effect of  $x^1$  on *any* future variables.

To mitigate this issue, the period in which the boundary first appears  $(t_0)$  and the period at which the explanatory variable of interest  $x^1$  is first observed, should be as close to each other as possible in time in order to minimizes the scope for endogenous sorting of households and suppliers of amenities to generate discontinuities in confounding variables. For instance, Jia et al. (2021) compare economic activity at the border of two Chinese provinces immediately after one is granted more public sector autonomy, and find evidence for positive effects of pro-growth policies. Alternatively, when the outcome of interest is observed at any point in the future, a boundary discontinuity design will identify the total effect of  $x_{t_1}^1$  on that outcome mediated through all channels, provided the boundary itself does not generate immediate discontinuities in other confounding neighborhood characteristics. Similarly, in a seminal paper in the literature on persistence, Dell (2010) identifies contemporary changes at historic boundaries after restricting the sample to boundary segments where institutional differences are negligible at  $t_0$ . In great detail, Dell (2010) steps through all potential mediating factors to establish a causal path of forced indentured servitude on economic outcomes in Peru centuries later but, insightfully, never fails to attribute this effect to the boundary itself as opposed to any single mediating factor.

# 3 Empirical Approach

### 3.1 Local Linear Regression Model

To test for these predicted discontinuities, we propose 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; 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 small, equal length pieces. We denote these as boundary segments indexed by b. For each boundary segment, we consider a set of nearby neighborhoods denoted as  $J_b$ . We locate the nearest neighborhoods in  $J_b$  on either side of the boundary and refer to them as index neighborhoods. For each neighborhood characteristic C, boundary segment 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 the distance from j to the boundary, and for each j on the low side, we define  $d_{bj}$  to be -1 times the distance from j to the boundary. We then estimate the following regression:

$$C_{bj} = \delta_C H_{Cbj} + f_-(d_{Cbj}) \times 1 (d_{Cbj} < 0) + f_+(d_{Cbj}) \times 1 (d_{Cbj} > 0) + \epsilon_{Cbj}$$
(6)

where  $f_{-}(\cdot)$  and  $f_{+}(\cdot)$  are flexible functions of the distance to the boundary and  $\epsilon_{Cbj}$  is an error term. The parameter  $\delta_{C}$  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,  $\delta_{C}$  will be identified and can be estimated by least squares. We implement the discontinuity model with bandwidth selection procedures formalized by Imbens and Kalyanaraman (2012), Calonico et al. (2014), and Calonico et al. (2017).

#### **Placebo Validation Exercise**

We consider the following placebo exercise to validate our empirical strategy. To simulate placebo boundaries, we shift every neighborhood in a boundary segment by 0.5 miles to the left or right in the two dimensional boundary space with the direction of the shift determined randomly. That is, we reshuffle the position of each neighborhood  $j \in J_b$  by  $\xi_{Cb} = \{-0.5, 0.5\}$  with equal probability. For each characteristic C we estimate the placebo regression

$$C_{bj} = \delta_C^p H_{Cbj} + f_-^p \left( d_{Cbj} + \xi_{Cb} \right) \times 1 \left( d_{Cbj} + \xi_{Cb} < 0 \right) + f_+^p \left( d_{Cbj} + \xi_{Cb} \right) \times 1 \left( d_{Cbj} + \xi_{Cb} > 0 \right) + \epsilon_{Cbj}^p$$
(7)

which can be understood as an analog to equation (6) in which the location of each boundary segment has been randomly assigned across the map containing neighborhoods  $j \in J_b$ . If our identifying assumption is satisfied, then we would expect our estimate of  $\delta_C^p$  to be zero.

### 3.2 Data

To estimate the boundary discontinuities  $\delta_C$ , we construct a dataset that is comprised of Census Block groups geospatially merged to nearby boundaries. 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 area near each boundary segment along more than one thousand dimensions. We describe both the boundary network data and neighborhood variables in further detail.

#### 3.2.1 Boundary Networks

We analyze three types of boundaries: transportation networks, educational boundaries, and administrative boundaries. Transportation networks form physical boundaries as they deform the urban landscape, 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.<sup>8</sup> We measure the historical US rail network using the Atack (2013) 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.<sup>9</sup>

<sup>&</sup>lt;sup>8</sup>American Association of Railroads, Chronology of America's Freight Railroads. https://www.aar.org/chronology-of-americas-freight-railroads/.

 $<sup>^{9}</sup>$ Our analysis omits segments of interstate < 0.5 mile long. Additionally, state managed highways are not present in the shape file.

School district and school attendance zone boundaries are constructed using shapefiles 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. Unlike school districts, there is scant nationally representative spatial data for school attendance zones. NHGIS hosts the 2009-2010 shapefiles created by the NSF funded SABINS project. The analysis of school attendance boundaries is restricted to elementary schools, which are defined as those enrolling third grade students (and none above ninth grade). Because the best available data for US school boundaries are from the 2009-2010 school year, we estimate this model with Census Block group data from 2010.

We also consider administrative boundaries in the form of 2020 county lines and ZIP code boundaries. A large share of county boundaries exist in remote areas with few nearby Block groups, however, our sample is restricted to boundary segments with descriptive Census data on both sides. Such a restriction results in smaller samples for county line regressions, despite the near 100% geographic coverage. ZIP code shapefiles also have extensive geographic coverage, but they contrast with county lines in that the data restriction is far less binding.

#### 3.2.2 Sample Description

We provide summary statistics of our data in Appendix Table 1, where we compare selected neighborhood characteristics in each boundary network. All Block groups are within one mile of each boundary type, with the sample restricted to boundaries with at least two block groups on either side. For five of the six boundaries we use Block group data from the 2020 Decennial Census and the 2016-2020 American Community Survey (ACS). For school zones we use 2010 Decennial Census Block group data. For 2020, there are over 1,700 variables characterizing the people, housing, and amenities for each Block group in the sample contiguous 48 states (with over 2,000 for 2010). Tens of 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). On average, Block group population means range from 1,674 to 1,859 compared to a population average of 1,776. The variation in home values and income across boundary types reflect the exclusion of boundaries

in rural areas that do not meet the sampling restriction. Incomes for the sample (\$77,768) are relatively higher than US national averages (\$67,521) at the time.

Variables	Full Data	Railroad	Highway	School District	School Zone	County Line	ZIP Code
Home Value	304,807.22 (249299.8)	298,811.16 (257542.2)	340,305.52 (275536.2)	329,255.35 (263742.7)	307,648.80 (211625.8)	340,140.41 (266735.6)	331,725.58 (271314.8)
Income	$77,768.43 \\ (40300.6)$	$74,\!642.87 \\ (39464.3)$	$79,372.11 \\ (40816.8)$	85,446.41 (42854.5)	$74,\!813.60 \\ (38145.6)$	87,853.63 (43871.7)	81,567.33 (41879.6)
Share White	0.60 (0.302)	$0.59 \\ (0.305)$	$\begin{array}{c} 0.53 \ (0.305) \end{array}$	0.61 (0.300)	$0.62 \\ (0.298)$	$0.63 \\ (0.291)$	$\begin{array}{c} 0.58 \\ (0.304) \end{array}$
Share Latino	0.19	0.19	0.21	0.17	0.18	0.14	0.19
	(0.235)	(0.239)	(0.245)	(0.227)	(0.227)	(0.180)	(0.237)
Share Black	$0.12 \\ (0.201)$	$0.12 \\ (0.207)$	0.14 (0.217)	$0.11 \\ (0.194)$	$\begin{array}{c} 0.13 \\ (0.213) \end{array}$	$\begin{array}{c} 0.13 \\ (0.219) \end{array}$	$0.12 \\ (0.201)$
Rooms / Unit	5.82 (1.374)	5.65 (1.316)	5.61 (1.399)	5.97 (1.471)	5.78 (1.272)	6.01 (1.562)	5.85 (1.447)
Age of Housing	$52.96 \\ (159.1)$	55.73 (143.3)	56.00 (152.9)	60.05 (184.4)	35.33 (17.98)	$60.69 \\ (190.6)$	59.06 (181.3)
Observations	65639	56815	39875	33685	35604	14266	45765

Table 1: Block Group Summary Statistics

Notes: Means and standard deviations (in parentheses) for the full sample of Census Block groups that lie within 1 mile of any boundary (column (2)), and the six subsamples of Block groups for each corresponding boundary network.) Because school zone boundaries reflect the 2009-2010 school year, Census data for this column come from the 2010 Census. For all other boundaries the data are from the 2020 census. Income and home value measures displayed are in 2020 dollars.

# 4 Results

# 4.1 Tiebout Sorting By Boundary Type

Our primary empirical exercise in this paper is to estimate the share of Census variables that exhibit statistically significant discontinuities, and the average magnitude of the identified changes for each boundary type. In order to facilitate comparison of these estimates across neighborhood characteristics and boundary types, we normalize the standard deviation of all dependent variables to 1. Table 2 highlights the fact that a substantial share of neighborhood characteristics vary at all types of boundaries we study, consistent with widespread Tiebout sorting.

We rank the results in Panel A of Table 2 by the fraction of statistically significant effects column (1). For each boundary type the fraction is substantially greater than 5%, with neighborhoods

	Frac. Significant Effects		Avg. Signifi	cant Effect Size
	All Vars.	Selected Vars.	All Vars.	Selected Vars.
	(1)	(2)	(3)	(4)
Panel A: Actual Boundaries				
School Attendance Zone	0.810	0.841	0.403	0.406
Railroad	0.756	0.783	0.264	0.270
Zip Code	0.635	0.683	0.219	0.219
School District	0.610	0.643	0.262	0.304
Highway	0.508	0.531	0.307	0.311
County	0.167	0.178	0.471	0.499
Panel B: Placebo Boundaries				
School Attendance Zone	0.040	0.036	0.086	0.075
Railroad	0.046	0.045	-0.012	-0.007
Zip Code	0.048	0.048	-0.001	0.006
School District	0.050	0.050	-0.045	-0.019
Highway	0.053	0.047	-0.037	-0.003
County	0.039	0.043	-0.027	-0.025

Table 2: Summary of Results

Notes: In columns (1) and (3), we present statistics for  $\delta_C$  for the entire sample of neighborhood characteristics. In columns (2) 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. Each row summarizes the results of independent regressions taking each Census descriptive variable as an outcome, grouped by the particular boundary used in each set of regressions. Statistical significance is reported at the 95% level.

varying at remarkably high rates given the large number of characteristics we test. We report that 81% of characteristics change discontinuously at school attendance zone boundaries, followed by 76% at railroad tracks, 64% at ZIP code boundaries, and 61% at school district boundaries. High-ways have less of an influence than railroads, but affect over 50% of neighborhood characteristics. The weakest prevalence of boundary effects occur at county lines, and in Appendix Awe provide further analysis of the boundary geography and location, which in part explain heterogeneity in the estimated effects.

Column (3) compares the average estimated effect size for statistically significant variables at each boundary type. The average magnitude ranges from 0.22 to 0.48 standard deviations from the mean. The larger discontinuities that we find for school zones reflect greater household sorting across school zones than other boundary types. Interestingly, county lines yield the fewest significant effects but the largest average effect size. This may indicate that when a discontinuity exists at a county boundary, it tends to be large. To ensure that these results are not driven by a large number of Census variables being uninformative, in columns (2) 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 analysis. Nevertheless, we find a similar prevalence and size of significant effects as we did in the unrestricted sample of columns.

The results of our placebo tests are summarized in Panel B of Table 2. All entries in columns (1) and (2) are close to 5%, which is what we would expect when considering significance at the 95% level. Meanwhile, columns (3) and (4) shows that even the few statistically significant placebo effects that we do find are not economically significant. To further support our identification of widespread boundary effects, we break down the distributions of estimates (including placebo estimates), and present sensitivity tests for RD bandwidth and polynomial choices in Appendix B.

### 4.2 Undetected Tiebout Sorting in Small Samples

Our finding of robust, widespread boundary discontinuities invites caution to practitioners seeking to implement regression discontinuity designs and other approaches that exploit geographical boundaries for identification. In the case of boundary or geographic discontinuity designs, researchers seek 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 that boundary ( $H_{Cbj} < 0$ ). The central identifying 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 this assumption is to demonstrate that other potential confounders  $C'_{bj}$  that are observed by 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. We discuss interpretation of estimates obtained by using this research design and other boundary methods that rely on similar assumptions about the data generating process.

Internal validity in boundary discontinuity designs may stem from the focus on a single city,

metropolitan area or even state. We speculate that balance tests of discontinuities in neighborhood characteristics are likely to be under-powered in such settings. To support this, we replicate our main analysis 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 3. School zones are where we find the most prevalent effects, and we select these 9 CBSAs because they have the largest share of total Census population residing within school attendance zones available in the GIS data. In this sense, these results should be seen as conservative since they are the most statistically powered subsamples. 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 only 10-25% of neighborhood characteristics depending on city.

		(1)	(2)
City (CBSA)	Observations	Fraction of	Fraction of
		Statistically	Statistically
		Significant Effects	Significant Effects,
		(All)	Selected Variables
All US (Pooled)	53966	0.810	0.841
Miami	2203	0.101	0.109
Philadelphia	1598	0.088	0.087
Minneapolis	1371	0.228	0.237
Houston	1344	0.223	0.234
Tampa	1141	0.184	0.188
Atlanta	1104	0.371	0.397
Riverside	1010	0.287	0.288
Washington, DC	958	0.163	0.163
Denver	855	0.118	0.102

Table 3: School Attendance Zone Boundary Discontinuities by City

Notes: In column (1) we present statistics 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. Selected cities are chosen based on data coverage, ie the share of MSA population residing in a school zone available in GIS shapefiles. Sample cities are sorted in the table by total number of observations. Statistical significance is reported at the 95% level.

The results in Table 3 suggest researchers should be cautious using covariate balance tests to validate boundary research designs. This also applies to dynamic applications that require similarly

strict assumptions as standard panel data models. Difference-in-discontinuities, for example, is amendable to pre-existing differences in covariate levels so long as changes in treated and control units are parallel (Grembi et al. (2016); Butts (2023). However, the proximity of treated and control units to the same boundary creates potential for contamination. Any sorting that occurs during the treatment window, because of pre-existing differences in neighborhood characteristics, will lead to false conclusions about parallel trends and in turn, bias estimates for treatment effects.

The upshot is that in settings with small pre-existing differences, researchers can make a case for minimal short-run sorting concerns. However, in the periods following policy dictated by a boundary (or establishment of a new administrative boundary), sorting changes the composition of both treated and control neighborhoods. Further, the effect is increasing in the magnitude of the policy change, posing a substantive threat to boundary designs (Jardim et al. (2024)). Our results imply the potential for false conclusions regarding the presence of potential sorting confounders.

### 4.3 Applications: House prices, Schools, and Segregation

The Tiebout sorting we find at attendance boundaries is consistent with the idea that schools facilitate residential segregation across broader housing markets (Reardon (2016); Monarrez (2023)). Cutler et al. (1999) and many others point to house prices, which capitalize amenities like neighborhood race and income composition, as a reinforcing mechanism for segregation. In that vein, we turn the analysis to Census home values at school attendance boundaries to estimate the price of the entire bundle of amenities near high quality schools. In doing so we broaden the literature with equity implications beyond traditional hedonic estimates of school quality in the US (Black (1999); Bayer et al. (2007)) and the UK (Gibbons et al. (2013)). It is the full housing cost premium for access to the school quality, social networks, and neighborhood amenities that amplifies the inequality of decentralized racism predicted by Cutler et al. (1999).

Table 4 presents results from the model taking Census home values as the outcome of our school zone boundary discontinuity. In the upper panel of column (1) we estimate the difference in price level to be \$53,211 for the sample including all block groups within the optimal bandwidth 0.44 miles from a school zone boundary. Moving from left to right in the table we assess the potential for

overlapping boundaries by excluding settings where the school boundary overlaps with a district, railroad, or highway segment. The largest effects are found where school zones do not overlap with any other boundary, and the pattern of the results is consistent with our headline findings of weaker Tiebout effects in settings where highways are present. Removing those boundaries from the sample in columns 4 and 5 increases the discontinuity by over \$15,000 relative to column (1).

	All	No District Effect	No RR Effect	No HWY Effect	No Effects 2-4
	(1)	(2)	(3)	(4)	(5)
Price Level					
Estimate	$53211.4^{*}$	48953.9	$57567.3^{*}$	68503.9**	69223.3*
SE	(23059.2)	(26184.0)	(27918.4)	(25421.7)	(33840.9)
Log(Price)					
Estimate	$0.264^{**}$	$0.251^{*}$	$0.256^{*}$	$0.329^{***}$	$0.329^{*}$
SE	(0.0894)	(0.0993)	(0.106)	(0.0983)	(0.128)
Boundary Fixed Effects					
Estimate	$0.104^{***}$	$0.0895^{***}$	$0.107^{***}$	$0.105^{***}$	$0.105^{**}$
SE	(0.0241)	(0.0268)	(0.0299)	(0.0273)	(0.0350)
N	14223	9954	9395	10558	5704
Bandwidth (Miles)	0.44	0.44	0.44	0.44	0.44
RD Polynomial	3	3	3	3	3

Table 4: Hedonic Estimates at School Zone Boundaries

We take the natural log of house prices and report the results in the middle panel of Table 4, which illustrates a discontinuity ranging from 25-33%. This figure approximates the percentage increase in full on board housing costs to reside on the higher quality side of the school boundary, reflecting house price capitalization of the full slate of Tiebout sorting effects in Table 2. Naturally, the result is markedly larger than estimates for the capitalization of test-scores alone.<sup>10</sup>

Testing for the effects of overlapping boundaries like railroads, highways, or even rivers is a sensitivity check for the influence of fixed characteristics in the urban space. The potential for bias in this large confounder is evidenced by 5704/14223 = 40% of block groups near school zones considered unexposed to other boundaries. In the bottom panel, we present results using boundary fixed effects that account for the presence of local confounders common to both sides of

<sup>&</sup>lt;sup>10</sup>Hedonic estimates of school quality begin with Black (1999) findings of a WTP of 2.5 percent for a 5 percent increase in test scores. Bayer et al. (2007)estimates a much smaller WTP, less than 1 percent, for a 5 percent increases in school performance. Dhar and Ross (2012) similarly test for the capitalization of school district quality, finding a 10% increase in house prices for a one standard deviation increase in district test scores.

a given boundary. Consistent with Dhar and Ross (2012), the inclusion of boundary fixed effects absorbs a substantial portion of the identifying variation, yielding more consistent estimates across samples. However, this potentially interesting variation may better explain the lack of integration of higher mobility neighborhoods. If the price premium shown in the upper two panels of Table 4 is prohibitive, it is no wonder why endogenous household location choices have little responsibility for the decline in US segregation today (Caetano and Maheshri (2023a)).

# 5 Conclusion

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, educational and administrative boundaries. These discontinuities are sizable, 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 under-powered 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.

As an illustrative example we unpack the price house price capitalization of Tiebout sorting at school attendance boundaries, and find discrete price changes close to five times the magnitude of estimates for the capitalization of school quality alone. The large price differential reflects broad changes in the bundle of amenities, including school quality and social networks that influence long-run labor market outcomes. While school quality equalization is worthwhile to pursue, the sorting we find along a number of different margins (at various boundary types) suggests that such a policy addresses only a portion of location choices that influence lifetime outcomes.

# References

- Aaronson, D., Hartley, D.A., Mazumder, B., 2020. The effects of the 1930s holc'redlining'maps.
- Albertus, M., 2020. Land reform and civil conflict: Theory and evidence from peru. American Journal of Political Science 64, 256–274.
- Atack, J., 2013. On the use of geographic information systems in economic history: The american transportation revolution revisited. The Journal of Economic History 73, 313–338.
- Baragwanath, K., Bayi, E., 2020. Collective property rights reduce deforestation in the brazilian amazon. Proceedings of the National Academy of Sciences 117, 20495–20502.
- Bayer, P., Ferreira, F., McMillan, R., 2007. A unified framework for measuring preferences for schools and neighborhoods. Journal of Political Economy 115, 588–638.
- Bayer, P., McMillan, R., Rueben, K., 2004. An equilibrium model of sorting in an urban housing market. Technical Report. National Bureau of Economic Research.
- Black, S., 1999. Do better schools matter? parental valuation of elementary education. The Quarterly Journal of Economics 114, 577–599.
- Butts, K., 2023. Geographic difference-in-discontinuities. Applied Economics Letters 30, 615–619.
- Caetano, G., 2019. Neighborhood sorting and the value of public school quality. Journal of Urban Economics 114, 103193.
- Caetano, G., Maheshri, V., 2023a. Explaining recent trends in us school segregation. Journal of Labor Economics 41, 175–203.
- Caetano, G., Maheshri, V., 2023b. A unified empirical framework to study neighborhood segregation. Working Paper .
- Calonico, S., Cattaneo, M.D., Farrell, M.H., Titiunik, R., 2017. rdrobust: Software for regressiondiscontinuity designs. The Stata Journal 17, 372–404.

- Calonico, S., Cattaneo, M.D., Titiunik, R., 2014. Robust nonparametric confidence intervals for regression-discontinuity designs. Econometrica 82, 2295–2326.
- Chetty, R., Friedman, J.N., Hendren, N., Jones, M.R., Porter, S.R., 2018. The opportunity atlas: Mapping the childhood roots of social mobility. Technical Report. National Bureau of Economic Research.
- Clark, G.L., Feldman, M.P., Gertler, M.S., 2003. The Oxford handbook of economic geography. Oxford University Press.
- Cox, R., Cunningham, J.P., Ortega, A., Whaley, K., 2022. Black lives: The high cost of segregation. University of Southern California Working Paper .
- Cutler, D.M., Glaeser, E.L., Vigdor, J.L., 1999. The rise and decline of the american ghetto. Journal of political economy 107, 455–506.
- Dell, M., 2010. The persistent effects of peru's mining mita. Econometrica 78, 1863–1903.
- Dell, M., Lane, N., Querubin, P., 2018. The historical state, local collective action, and economic development in vietnam. Econometrica 86, 2083–2121.
- Dhar, P., Ross, S.L., 2012. School district quality and property values: Examining differences along school district boundaries. Journal of Urban Economics 71, 18–25.
- Epple, D., Filimon, R., Romer, T., 1984. Equilibrium among local jurisdictions: toward an integrated treatment of voting and residential choice. Journal of Public Economics 24, 281–308.
- Epple, D., Romer, T., Sieg, H., 2003. Interjurisdictional sorting and majority rule: An empirical analysis. Econometrica 69, 1437–1465.
- Gibbons, S., Machin, S., Silva, O., 2013. Valuing school quality using boundary discontinuities. Journal of Urban Economics 75, 15–28.
- Grembi, V., Nannicini, T., Troiano, U., 2016. Do fiscal rules matter? American Economic Journal: Applied Economics , 1–30.

Hotelling, H., 1929. Stability in competition. The economic journal 39, 41–57.

- Imbens, G., Kalyanaraman, K., 2012. Optimal bandwidth choice for the regression discontinuity estimator. The Review of economic studies 79, 933–959.
- Jardim, E., Long, M.C., Plotnick, R., Vigdor, J., Wiles, E., 2024. Local minimum wage laws, boundary discontinuity methods, and policy spillovers. Journal of Public Economics 234, 105131.
- Jia, J., Liang, X., Ma, G., 2021. Political hierarchy and regional economic development: Evidence from a spatial discontinuity in china. Journal of public economics 194, 104352.
- Keele, L.J., Titiunik, R., 2015. Geographic boundaries as regression discontinuities. Political Analysis 23, 127–155.
- Lu, Y., Wang, J., Zhu, L., 2019. Place-based policies, creation, and agglomeration economies: Evidence from chinaâs economic zone program. American Economic Journal: Economic Policy 11, 325–360.
- Monarrez, T., Schönholzer, D., 2023. Dividing lines: racial segregation across local government boundaries. Journal of Economic Literature 61, 863–887.
- Monarrez, T.E., 2023. School attendance boundaries and the segregation of public schools in the united states. American Economic Journal: Applied Economics 15, 210–237.
- Reardon, S.F., 2016. School segregation and racial academic achievement gaps. RSF: The Russell Sage Foundation Journal of the Social Sciences 2, 34–57.
- Tiebout, C., 1956. A pure theory of local expenditures. The Journal of Political Economy 64, 416–424.
- Whaley, K., 2024. Residential segregation at physical neighborhood boundaries. Journal of Economics, Race, and Policy 7, 141–153.

# **Online Appendix**

# A Data and Sampling

The approach to constructing our data is the same for each of the six boundaries. Original source data shapefiles are loaded into ArcGIS, and each boundary is cut into segments at maximum two miles in length. Block groups are assigned to nearby boundary segments using the latitude and longitude coordinates of each Block group population center. For each boundary type, we allow Block groups to be assigned to multiple boundary segments. Although Block groups are discrete points across space, the data proxy for a continuous distribution of amenities near each boundary. In that sense, a single Block group that is 0.75 miles from one boundary and 0.1 miles from another can provide valuable information about each. The sample is restricted to Block groups less than a mile from a boundary, and boundary segments must have data on both sides to be included.

### A.1 Physical Boundaries

The data for railroad and highway networks is comprehensive and covers the entire United States. The sum total of railroad mileage is greater than highway mileage, and the arrangement of both networks may differ in urban and rural areas. As an example, Figure 2 contains a comparison of railroad and highway coverage in a large urban area (Atlanta, GA) and a suburan/rural county (Anderson, SC) in the same region. Evident from the upper panel are transportation networks serving as major arteries in the city structure, with substantial clustering of neighborhoods near both railroads and highways. Economic activity in rural areas is more likely to be clustered near railroads, as shown in the the lower panel of Figure 2.

#### Figure 2: Railroads and Highways



Notes: In each panel, the left Figure includes only railroad lines, while the right Figure includes both railroad and highway lines. Railroad lines are colored solid black and highway lines are outlined in grey.

### A.2 School Boundaries

There is an extensive body of research on the formation and consequences of school boundaries. There is substantial heterogeneity in the political economy behind how school district and attendance zone boundaries are drawn by state. For the two school boundary types, we try to use the broadest definition possible. The school district sample includes unified districts that do not solely serve charter schools. School attendance zones in the sample include those that contain third grade students and are intended to have open enrollment. Open enrollment school zones allow students with the catchment area to choose from two or more schools that share the same boundary. School district maps can vary substantially by state, and are a classic example of variation in the political economy that has historically shaped school districts. On one extreme are school districts drawn synonomously with county lines, covering large spatial areas and governed by a single school board. Such design is prominent is southern states. By contrast, other counties are fragmented with smaller school districts each managed by an independent school board. In the latter case school district lines may coincide with municipal boundaries at a lower geography than the county. As an example, Figure 3 contains two counties (shaded areas) with similar populations in 2020. On the left is Mecklenberg County, population 1.13 million, in which Charlotte is center city. On the right is Cuyahoga County, population 1.23 million, where Cleveland is the center city. As of 2020, Cuyahoga county is served by over 25 school districts.

Figure 3: School District Heterogeneity



Notes: Mecklenburg County, NC (4.1) has a population of 1.13 million in 2020, in comparison to the 1.23 million residents of Cuyahoga County, OH (4.2). The upper panel illustrates the single unified Charlotte-Mecklenburg school district in contrast to the many districts serving the Cleveland-Cuyahoga County area.

# A.3 Administrative Boundaries

Data for county lines and ZIP codes cover nearly entire lower 48 geographic area. As Figure 4 shows, counties have broader administrative geography with boundary sections in both urban and remote areas. There are over 3,000 US counties, and across county lines are predicted changes in tax rates levied on property and sales, municipal service quality, and school quality. There is substantial

heterogeneity in county population sizes, with the average US county being approximately 100,000 people.



Figure 4: Example: Greater Miami MSA

Notes: Census Block group locations are overlaid with spatial data for county lines and ZIP code boundaries in the Greater Miami metropolitan area, including Dade, Broward, and Palm Beach counties.

As of the 2020 census there are 41,500 zip codes containing an average 8,340 people in each. As shown in the second panel of Figure 4, the spatial correlation between ZIP code density and population density combined with a contiguous ZIP code map yields a substantially larger analysis sample of ZIP code boundaries. The primary purpose of ZIP codes is to delineate federal postal routes, in theory an otherwise meaningless distinction that in practice is used to facilitate other locally provided services. Public school catchments, police and fire precincts, county lines and other municipal boundaries have the potential to overlap with ZIP codes to produce statistically significant amenity differentials in our analysis.

# A.4 Real and Placebo Boundaries

The location choice of residents and the endogenous production of heterogeneous neighborhood amenities tie the data describing block group population centers to boundary segments in the spatial data. It is the location choices, not the endogeneity of amenities and neighborhood demographics, that result in the observed boundary effects. If the statistical correlation between neighborhood amenities remained constant but the location of neighborhoods were randomly assigned, there would be a much weaker boundary effect, particularly for physical boundaries.

To implement our placebo test, we map neighborhoods to the same boundary segment but randomize the position of population centers relative to the threshold. In doing so we preserve the statistical correlation between neighborhood amenities to simulate a world absent the specific type of spatial segregation that generates boundary effects. Alternatively, if consumers (producers) have within-neighborhood preferences for proximity to a specific producer (consumer) type, but are indifferent about the characteristics of nearby neighborhoods, boundaries will have no effect on the spatial distribution of amenities.

# **B** Supplementary Analysis

### **B.1** Supplement : Main Results

For each boundary type, we present histograms of all estimated boundary effects in Figure 5 overlaid with histograms of those effects 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.

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 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. We estimate fewer statistically significant effects at county boundaries, but a substantial fraction of statistically significant effects at ZIP code boundaries. This is likely due to the fact that the majority of county borders are in rural areas that offer fewer observations for estimation, a drawback that does not apply to ZIP code boundaries.





Note: Statistical significance is determined at the 95% level. Each histogram contains the frequency of statistically significant and insignificant effects arranged by the magnitude of the boundary effect (x-axis). For comparison across estimates, the standard deviation of all dependent variables is normalized to 1.

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 6. 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.<sup>11</sup> (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.

<sup>&</sup>lt;sup>11</sup>The slight skew to the right is likely due to our selection procedure for the "high side."



#### Figure 6: Placebo Boundary Effects

Note: Statistical significance is determined at the 95% level. Each histogram contains the frequency of statistically significant and insignificant effects arranged by the magnitude of the boundary effect (x-axis). For comparison across estimates, the standard deviation of all dependent variables is normalized to 1.Detailed description of the placebo boundary exercise is in Section 3.

In this section we stress test the results for robustness to specification choices in the boundary discontinuity design. We focus on school attendance zones, which produce the most widespread amenity effects in our analysis and have been a focus in the literature for quite some time. In a setting where house prices capitalize changes in school quality and endogenous amenities at the boundary, bandwidth selection involves choosing the distance from the boundary used to restrict the estimation sample. Further, modeling the change in prices relative to distance to the boundary involves choosing the order of polynomial used to fit the regression. Through this analysis we illustrate how results of boundary - hedonic price regressions are sensitive to research design, and discuss the limitations of causal claims in the absence of high frequency data for both prices and amenities.

# **B.2** Sensitivity Analysis

As standard with regression discontinuity designs we stress test our results across two specification margins: the choice of RD polynomial and bandwidth. Section B.2.1 presents results modeling each outcome as a linear function of distance to the boundary (as opposed to the cubic polynomial functional form in the main results). Section B.2.2 contains a table that shows our results are robust to ad hoc choices of bandwidth around each threshold.

#### **B.2.1** Polynomial Choice

	Frac. Significant Effects		Avg. Signific	cant Effect Size
	All Vars.	Selected Vars.	All Vars.	Selected Vars.
	(1)	(2)	(3)	(4)
Panel A: Actual Boundaries				
School Attendance Zone	0.902	0.948	0.339	0.359
Railroad	0.829	0.871	0.194	0.209
Zip Code	0.752	0.793	0.144	0.151
School District	0.721	0.754	0.191	0.203
Highway	0.609	0.656	0.158	0.165
County	0.278	0.289	0.135	0.148

Table 5: Summary of Results

Notes: In columns (1) and (3), we present statistics for  $\hat{\delta}_C$  for the entire sample of neighborhood characteristics. In columns (2) 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. Each row summarizes the results of independent regressions taking each Census descriptive variable as an outcome, grouped by the particular boundary used in each set of regressions. Statistical significance is reported at the 95% level.

#### B.2.2 Distance Bandwidth

As in traditional RD settings, the primary hurdle in spatial RD design is determination of the bandwidth, defined by the distance from the boundary on either side that determines which data points to be included in the estimation sample. A common practice is to repeatedly estimate the sharp RD model, reducing the sample each iteration by shrinking the bandwidth. We follow Imbens and Kalyanaraman (2012) bandwidth selection process that minimizes an empirical approximation of the mean squared error. Using the Calonico et al. (2017) RDRobust package implements Imbens and Kalyanaraman (2012) optimal bandwidth selection for each regression outcome in the main analysis.

	Frac. Significant Effects		Avg. Signifi	cant Effect Size
	All Vars.	Selected Vars.	All Vars.	Selected Vars.
	(1)	(2)	(3)	(4)
Panel A: 1 Mile Bandwidth				
School Attendance Zone	0.896	0.942	0.325	0.343
Railroad	0.812	0.849	0.186	0.194
Zip Code	0.744	0.781	0.140	0.148
School District	0.712	0.745	0.184	0.191
Highway	0.606	0.634	0.136	0.311
County	0.191	0.196	0.114	0.119
Panel B: 0.5 Mile Bandwidth				
School Attendance Zone	0.040	0.036	0.086	0.075
Railroad	0.046	0.045	-0.012	-0.007
Zip Code	0.048	0.048	-0.001	0.006
School District	0.050	0.050	-0.045	-0.019
Highway	0.053	0.047	-0.037	-0.003
County	0.039	0.043	-0.027	-0.025

Table 6: Summary of Results

Notes: In columns (1) and (3), we present statistics for  $\delta_C$  for the entire sample of neighborhood characteristics. In columns (2) 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. Each row summarizes the results of independent regressions taking each Census descriptive variable as an outcome, grouped by the particular boundary used in each set of regressions. Statistical significance is reported at the 95% level.

### **B.3** Supplement : Hedonic Regressions and Boundary Fixed Effects

Figure 7 includes raw data scatter plots of census block group home prices for owner-occupied housing. Prices are displayed in logs on the y-axis, and the x-axis arranges block groups by distance to the boundary. For each boundary segment, the median house price is computed for each side, and block groups on the higher income side take positive values. Each panel includes identical underlying data but vary by the polynomial degree chosen for the line of best fit. In a sharp RD design, the point estimate is the difference between two regression functions evaluated at the boundary, represented by the vertical difference the two fit lines. Comparing the panels illustrates how the potential for misspecification if non-linearities in the house price gradient go unmodeled.







Notes: Scatterplots of raw data for Census Block groups arranged by distance to a school attendance zone boundary. Each Figure varies only by the polynomial selected to model the local regression on each side of the boundary. Point estimates in a hedonic model of prices are computed as the vertical distance between the two regression lines as they intersect the boundary.

Non-parametric estimation of house price changes at school boundaries captures the general equilibrium effect of better quality schools, neighborhood demographics that change with school quality, and other amenities correlated with neighborhood demographics. This contrasts with using school boundaries and house prices to estimate household valuation of school quality while conditioning on demographic and amenity differences. On one hand, it is important to know the value that households place on educational investments. On the other hand, the overall general equilibrium effect is informative as to the amount households must pay to access better quality schools. In places without open enrollment, i.e. school attendance zones are binding, households must be willing to pay for both higher school quality and the bundle of neighborhood demographics and amenities that accompany better schools.

Better educated families with higher incomes may sort into high quality school zones for a variety of reasons. Thus, the demand for housing near high quality schools also reflects the fact that better educated neighbors with higher incomes are positive neighborhood amenities, and other amenities near such neighbors are also likely to be of higher quality. Households choosing where to live may find these neighborhoods desirable over and above the quality of their schools. Our results in the main body of the paper suggest that households must be willing to pay for both the higher school quality and the price capitalization of other amenities. Column 2 of Table 7 shows that the average premium is approximately 26%, or \$51,000 as shown in Column 5.

Boundary fixed effects are a way to de-mean house prices in block groups relative to neighbors in the same boundary segment. In a hedonic model this approach compares residual variation in prices to residual variation in an observable neighborhood characteristic to obtain a hedonic valuation of that particular amenity. Incorporating boundary fixed effects into a local linear RD involves two stages : regressing house prices on a set of indicators for each boundary segment, then fitting the RD model with residuals from the first stage. Estimates for the hedonic RD with boundary fixed effects are found in columns 3 and 6 of Table 7.

	(1) Log	(2) Log	(3) Log	(4) Level	(5) Level	(6) Level
Boundary Effect	0.206***	$0.264^{**}$	$0.104^{***}$	45870.6***	$51117.0^{*}$	$26528.9^{***}$
SE	(0.0399)	(0.0894)	(0.0241)	(11180.5)	(23106.1)	(6896.2)
N	14482	14482	14223	14074	14074	13767
Polynomial	1	3	3	1	3	3
Controls	None	None	Boundary FE	None	None	Boundary FE

Table 7: Hedonic Estimates at School Zone Boundaries

\* p < 0.05, \*\* p < 0.01, \*\*\* p < 0.001

Outcome: Log(Price)	All	No District Effect	No RR Effect	No HWY Effect	No Effects 2-4
Sharp RD	(1)	(2)	(3)	(4)	(5)
Estimate	$53211.4^*$	48953.9	57567.3*	$68503.9 \setminus **$	69223.3 <sup>*</sup>
SE	(23059.2)	(26184.0)	(27918.4)	(25421.7)	(33840.9)
Boundary Fixed Effects RD	(6)	(7)	(8)	(9)	(10)
Estimate	$27309.2^{***}$	$25429.3^{***}$	34503.7 <sup>***</sup>	$29258.2^{***}$	$34296.9^{***}$
SE	(6835.9)	(7664.5)	(8585.0)	(7595.4)	(9863.1)
N	14223	9954	9395	10558	5704
Polynomial	3	3	3	3	3

# Table 8: Hedonic Estimates at School Zone Boundaries