For Online Publication

A Appendix Tables

Table A.1: Sample breakdown over time, by state

State	Census Region	Blocks, Time Avg
California	West	45337
Florida	South	36149
Texas	South	25363
Pennsylvania	Northeast	21620
Massachusetts	Northeast	20114
New York	Northeast	19541
New Jersey	Northeast	15955
Illinois	Midwest	15711
Ohio	Midwest	15601
Virginia	South	12837
Arizona	West	11860
Michigan	Midwest	11444
Washington	West	11037
Connecticut	Northeast	9691
Missouri	Midwest	9324
Maryland	South	9299
North Carolina	South	7366
Georgia	South	7318
Tennessee	South	7149
Indiana	Midwest	6626
Minnesota	Midwest	6042
South Carolina	South	4617
Rhode Island	Northeast	4080
Oklahoma	South	3878
Colorado	West	3868
Alabama	South	3817
Utah	West	3160
Kansas	Midwest	3000
Wisconsin	Midwest	2889
Oregon	West	2052
Delaware	South	1869
District of Columbia	South	1368
Nevada	West	1309
Maine	Northeast	1174
New Hampshire	Northeast	1118
Idaho	West	1025
Mississippi	South	993
Nebraska	Midwest	983
Arkansas	South	907
Iowa	Midwest	659
Kentucky	South	511
New Mexico	West	507

Table A.2: Null Effects Estimated Around Regulatory Borders In Other Stratified Samples

BD Estimates	Block level shares Nonwhite	Black	Asian	Block level shares Nonwhite	Black	Asian
2020 Data	0.008	0.004	0.004	0.026	−0.009	0.001
	(0.009)	(0.005)	(0.006)	(0.019)	(0.010)	(0.006)
2010 Data	−0.019**	−0.013*	−0.002	−0.019	−0.003	−0.011
	(0.009)	(0.007)	(0.004)	(0.015)	(0.015)	(0.010)
1980 Data	−0.015	−0.020	−0.001	0.009	0.005	0.002
	(0.016)	(0.014)	(0.002)	(0.018)	(0.017)	(0.005)
Jurisdiction FE	X	X	X	X	X	X
Treated MLS	6000–11000			11000–22000		
Diff With Compared	0–4000			0–6000		
Total N	359650			132434		

Significance levels: * = 10%; ** = 5%; *** = 1%.

Notes: This table presents outputs of border discontinuity designs over Census blocks b in year t for racial minority m,

$$Share^m_{bt} = \alpha_{j(b)t} + \beta^t \mathbf{1}[Dist_b \geq 0] + \eta^t_- Dist_b + \eta^t_+ Dist_b \cdot \mathbf{1}[Dist_b \geq 0] + \varepsilon_{bt},$$

where each column plots shares for a different group m, "Nonwhite" referring to all residents who are not non-Hispanic white. Separate estimates are made for different years of Census data, and for two samples representing different urban contexts that are detailed further in Section 5.1. In the table, the key density ranges defining the context-specific sample are listed. Point estimates and standard errors are based off of the robust nonparametric procedure in Calonico, Cattaneo and Titiunik (2014). Across specifications, fixed effects are set at the jurisdiction-year level. Standard errors are calculated clustering at the county-year level.

Table A.3: Covariate Balance Over Time, Alternate MLS Samples

BD Estimates	Block level shares Rented Units	>62 y.o.	≤ 18 y.o.	Block level shares Rented Units	>62 y.o.	≤ 18 y.o.
2020 Data	-0.006	0.002	0.002	-0.015	0.028	-0.024^{**}
	(0.018)	(0.007)	(0.004)	(0.019)	(0.028)	(0.011)
2010 Data	-0.005	-0.001	-0.010^{***}	-0.001	0.033^{*}	-0.012
	(0.011)	(0.007)	(0.004)	(0.017)	(0.019)	(0.013)
1980 Data	-0.019	0.001	-0.006	-0.033^{**}	0.004	-0.006
	(0.019)	(0.006)	(0.005)	(0.013)	(0.014)	(0.018)
Jurisdiction FE	X	X	X	X	X	X
Treated MLS		6000–11000			11000–22000	
Diff With Compared		0–4000			0–6000	
Total N		354283			130364	

Significance levels: $^{*}=10\%$; $^{**}=5\%$; $^{***}=1\%$.

Notes: This table presents outputs of border discontinuity designs over Census blocks b in year t for racial minority m,

$$Z_{bt}=\alpha_{j(b)t}+\beta^{t}\mathbf{1}[Dist_b\geq 0]+\eta^{t}_{-}Dist_b+\eta^{t}_{+}Dist_b\cdot\mathbf{1}[Dist_b\geq 0]+\varepsilon_{bt},$$

where Z describes a confounding variable at the block level. Each column plots a separate confounder available in the Census tabulations. Separate estimates are made for different years of Census data, and for two samples representing different urban contexts that are detailed further in Section 5.1. In the table, the key density ranges defining the context-specific sample are listed. Point estimates and standard errors are based off of the robust nonparametric procedure in Calonico, Cattaneo and Titiunik (2014). Across specifications, fixed effects are set at the jurisdiction-year level. Standard errors are calculated clustering at the county-year level.

Table A.4: Comparison of Estimated Effects With Literature

Paper	Policy intervention	Census year	Point estimate	95% confidence interval
Sood and E-S	Minneapolis racial covenants	2020	-0.015	$[-0.04, 0.01]$
Cui and Been	**High density sample estimate**	2010	-0.050	$[-0.082, -0.018]$
Cui and Been	**Large density shift sample estimate**	2010	-0.040	$[-0.068, -0.012]$
AHM	HOLC maps: gap btwn C-B zones	2010	-0.0074	$[-0.023, 0.008]$
M&S	Stacked jurisdiction borders	2010	-0.04	$[-0.042, 0.038]$
Resseger	Stacked SF zones in Boston	2010	-0.0087	$[-0.013, -0.004]$
Sood and E-S	Minneapolis racial covenants	1980	-0.004	$[-0.008, -0.00]$
Cui and Been	**High density sample estimate**	1980	0.003	$[-0.05, 0.06]$
Cui and Been	**Large density shift sample estimate**	1980	-0.048	$[-0.091, -0.005]$
AHM	HOLC maps: gap btwn C-B zones	1980	-0.049	$[-0.126, 0.028]$
M&S	Stacked jurisdiction borders	1990	-0.035	$[-0.03892, -0.0311]$

Notes: Following the discussion in Section 6.1, this Table puts our preferred estimates of racial disparities around lot sizes in 2010 with other border discontinuity estimates of policies in the literature. When abbreviated, “AHM” refers to Aaronson, Hartley and Mazumder (2021) and “M&S” refers to Monarrez and Schönholzer (2023).

B Appendix Exhibits

Figure B.1: Border discontinuities over 1980 Census data

(a) Outcome: Block-level racial minority shares

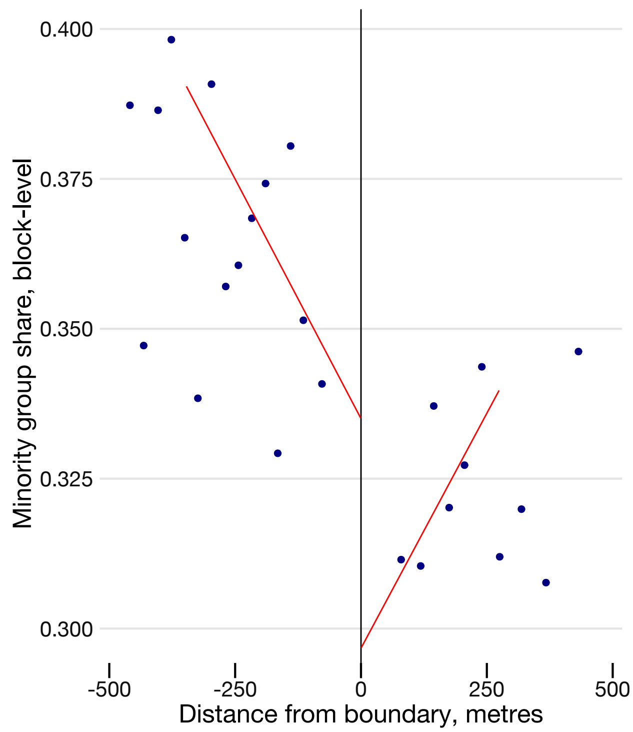

(b) Outcome: Block-level shares of rented housing units

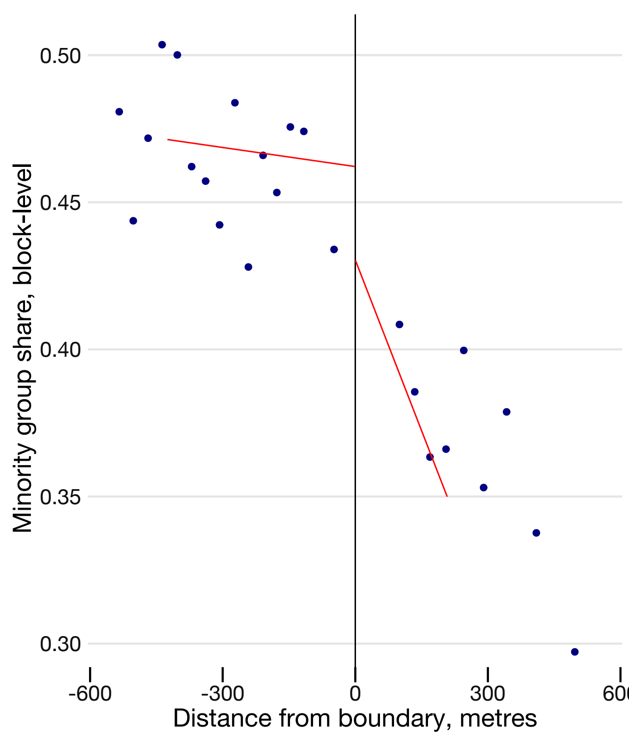

Sources: Calculations from 2020 NHGIS Tables (Manson et al. (2021)) and CoreLogic Tax Records.

Figure B.2: Relative sizes of stratified samples

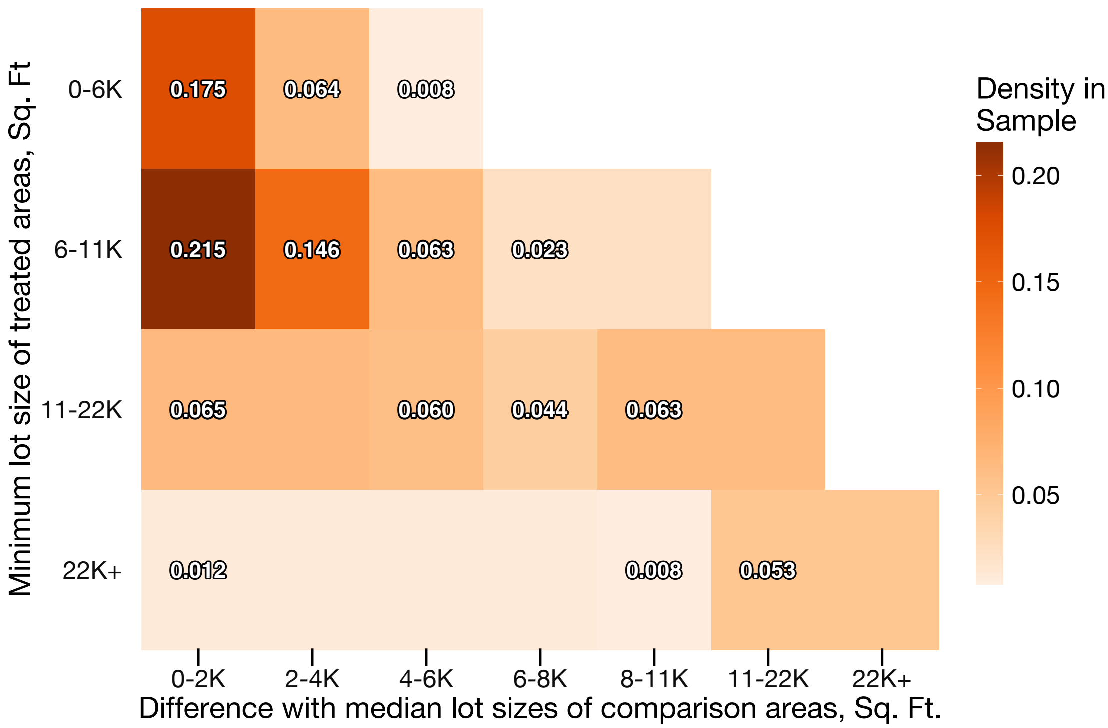

Notes: This figure presents a heatmap plotting a histogram of detected lot size discontinuities, used in our analysis sample. Each cell in the heatmap plots the share of all blocks used in the 2010 Census data satisfying two criteria. First, the sample includes blocks belonging to a minimum lot size district whose values are in one of four ranges on the Y axis. Second, the surrounding urban context of the lot size regulated blocks have higher densities that differ within the ranges on the X axis. When cells span multiple X axis columns, that means the underlying sample includes comparison blocks in the union of the ranges. Only discontinuities detected within jurisdiction boundaries, using the algorithm of Section 3, are included.
Sources: Calculations from 2020 NHGIS Tables (Manson et al. (2021)) and CoreLogic Tax Records.

Figure B.3: Heatmap of Black American residential disparities across minimum lot size areas

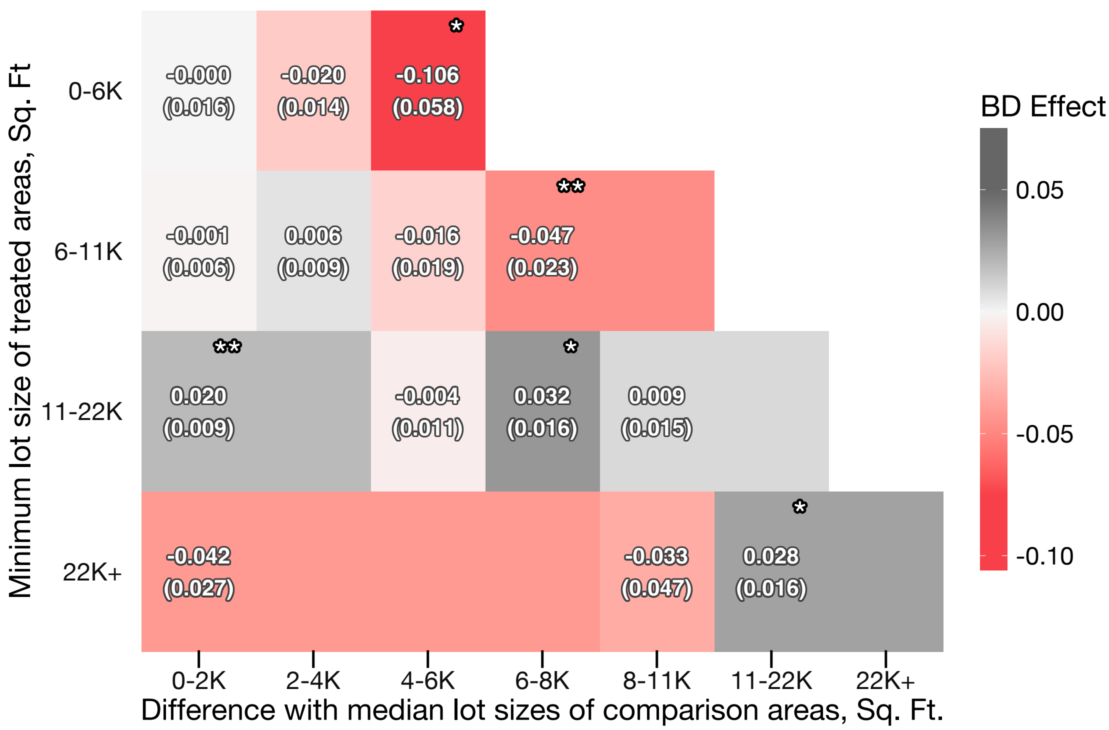

Notes: This figure presents border discontinuity effects on the 2020 Census block-level Black share, estimated across stratified samples of lot size discontinuities. Each cell in the heatmap plots border discontinuity estimates, using the specification in 4.2, for one sample defined off of two criteria. First, the sample includes blocks belonging to a minimum lot size district whose values are in one of four ranges on the Y axis. Second, the surrounding urban context of the lot size regulated blocks have higher densities that differ within the ranges on the X axis. When cells span multiple X axis columns, that means the underlying sample includes comparison blocks in the union of the ranges. Only discontinuities detected within jurisdiction boundaries, using the algorithm of Section 3, are included. Standard errors clustered at the county level to account for spatial autocorrelation.
Significance levels for each sample are reported as follows: * = 10%; ** = 5%; *** = 1%.
Sources: Calculations from 2020 NHGIS Tables (Manson et al. (2021)) and CoreLogic Tax Records.

C Further details on lot size automated procedure

Local support vector machine problem. In each of the interior cells, we take all geocoded lots within the cell's borders. Based on the minimum lot size we want to measure, the lots are classified into two classes $Y \in \{1, -1\}$. We want to find a linear boundary over two variables, the longitude and latitude, which is the separating hyperplane for if a lot belongs in either class. The optimization problem is to choose the linear boundary such that the lots are correctly classified as much as possible.

The objective function solved by the SVM classifier is pick the support vector of the boundary (w, b) and misclassification distances ξ, then solve

$$\min_{w,b,\xi} \frac{1}{2} w^T w + C \sum_{i=1}^{n} \xi_i$$
$$s.t. y_i(w^T x_i + b) \geq 1 - \xi_i,$$
$$\xi_i \geq 0.$$

When minimizing the objective, the SVM classifier is trading off accurate prediction — in which case $\text{sgn}(y_i(w^T x_i + b)) \geq 0$ and can be normalized to be ≥ 1 — and minimizing deviations in distance between misclassified lots and the boundary. The strength of this tradeoff is governed by the hyperparameter C.

We implement the SVM procedure using the `scikit-learn` package in Python, which speeds up the procedure by solving the dual problem to the above primal problem. The dual problem involves minimizing a quadratic form subject to linear constraints.

K-nearest neighbor procedure. A common enough edge case is that in the same interior tile, detection of a linear boundary was successful for multiple lot sizes with bunching behavior. The multiple detection could correspond to actual zoning district borders that are irregular in that area, but without further adjustment this leads to overlap issues with regions. Where classified treated areas intersect, a block could be identified as belonging to two separate minimum lot size districts at once.

For these tiles, we refine boundaries using a K-nearest neighbor algorithm. Figure C walks through this exercise for a sample cell and geocoded property data within it. For this interior cell, the SVM procedure detected regulatory border segments for three separate lot sizes. The implied lot size areas detected are shown on the left graphic, where overlapping between the detected areas is evident.

To set up the KNN algorithm, we classify all properties in the cell into $N+1$ classes, where N is the number of detected lot sizes. The first class is all properties lower than the least detected lot size, visualized as white diamonds in the figure. Then, indexing the detected lot sizes in ascending order $\underline{\ell}_1, \ldots, \underline{\ell}_N$, properties are classified as class $k = 2, \ldots, N$ if their lot sizes are in the interval $[\underline{\ell}_{k-1}, \underline{\ell}_k]$. All lots above $\underline{\ell}_N$ are classified as class $N + 1$.

With the target classes set up, we run the KNN algorithm by predicting class at interior points in the cell based on the nearest k_{mult} lots. We only run this prediction at a certain resolution, which is determined by the radius parameter r_{mult}. We standardize distances within the cell, at which point the larger r_{mult} is the larger the gap between points where we predict using the KNN classifer.

The figure on the right side of Figure C shows the output of the KNN procedure for the cell. First, note that lots below any detected lot size are now in their own class. Then, the regions where each lot size minimum is predicted to apply are the union of predictions at local points, expanded by taking a square buffer around each.

As the example shows, the output lot boundaries may look more irregular or even disjoint. To ensure robustness of the KNN procedure, we apply the same filter as the SVM procedure: KNN predictions with a misclassifcation rate above $\overline{m}_{err}$ are dropped as well from the analysis sample.

Figure C.1: Illustration of additional KNN procedure

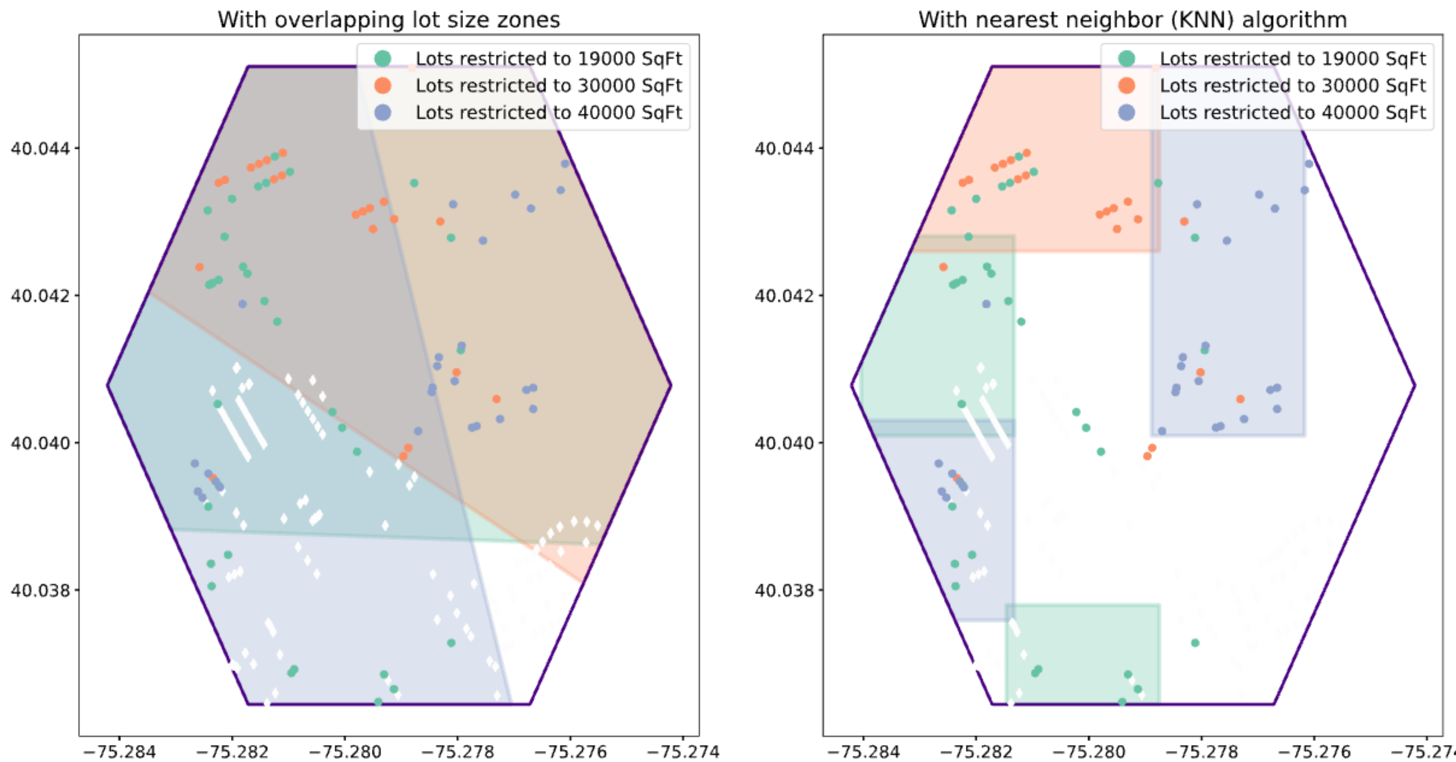

Sources: Calculations from CoreLogic Tax Records.