NBER WORKING PAPER SERIES

THE DISTASTE FOR HOUSING DENSITY

Joseph Gyourko
Sean E. McCulloch

Working Paper 33078
http://www.nber.org/papers/w33078

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2024

We thank Braydon Neiszner and Vivian Wang for helpful research assistance. We also appreciate the comments of Hector Blanco (discussant), Gilles Duranton, Fernando Ferreira, Ed Glaeser, Matthew Turner and seminar participants at the University of Wisconsin, the 2024 UEA North American Conference and the Urban Lunch Seminar at Wharton. Naturally, we remain responsible for the final product. Gyourko thanks the Research Sponsor Program of the Zell/Lurie Real Estate Center for financial support, and McCulloch thanks the Population Studies and Training Center at Brown University, which receives funding from the NIH, for training support (T32 HD007338) and for general support (P2C HD041020). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Distaste for Housing Density
Joseph Gyourko and Sean E. McCulloch
NBER Working Paper No. 33078
October 2024
JEL No. R0, R31

ABSTRACT

We characterize the distribution of suburban homeowners' preferences for housing unit density. To measure welfare changes under counterfactual increases in density, we first construct a novel house-level measure of exposure to density and identify its price effects in a boundary discontinuity design. On the borders of municipalities with larger minimum lot sizes, lots are 3,000 ft^2 larger and houses are$40,000 costlier. We exploit the systematic variation in density exposure induced by these discontinuities to estimate price effects. We then connect these estimates to a structural hedonic model of housing choice to retrieve individuals' preferences for density. Overall, we find an average welfare loss among incumbent homeowners from a 1/2 unit per acre increase in density (which is equivalent to a 0.3 standard deviation in density) of about $9,500, with significantly larger losses under counterfactual increases solely from rental units. There is other noteworthy heterogeneity in these preferences, too. Most households have only a moderate preference over density. The median welfare loss is only 55% of the average, implying a long, left tail of those with more extreme aversions to density. This tail disproportionately contains households in affluent, low density neighborhoods. In sum, our results document an important foundation of the demand for density regulation across U.S. suburbs that we hope serves as a valuable input into future research into the considerable costs of that policy.

Joseph Gyourko
University of Pennsylvania
The Wharton School of Business
3733 Spruce Street
Room 448 Dinan Hall
Philadelphia, PA 19104-6302
and NBER
gyourko@wharton.upenn.edu

Sean E. McCulloch
Brown University
Blistein House, 57 Waterman Street
Providence, RI 02906
sean_mcculloch@brown.edu

A data appendix is available at http://www.nber.org/data-appendix/w33078

1. Introduction

Real constant-quality house prices are at all-time highs in many markets, sparking widespread concern about housing affordability.[1] A large literature in economics points to the impact of state and local government restrictions on building activity and housing supply to help explain this phenomenon (Glaeser and Gyourko, 2018). Easing housing unit density limits has become an important policy focus for affordability advocates, with a few jurisdictions changing zoning rules to allow more dense development.[2] That there are not many more communities implementing such an obvious policy response begs the question of why not. A straightforward answer is that local residents do not favor the policy and are able to stop most densification through the political process. Recent survey responses indicate that a majority of Americans prefer less dense housing, and while there are many examples of local opposition to new development on the basis of its supposed negative impacts, the extent of their associated costs has not been quantified.[3]

In this paper, we provide an analysis of the external costs and benefits of housing unit density to local homeowners. By estimating the preference for density at a household level, we can characterize the distribution of preferences, and document heterogeneity across neighborhoods. Our empirical strategy to quantify individuals' preferences for housing density combines a boundary discontinuity design with a structural hedonic model. We build on our previous work (Gyourko and McCulloch, 2023) that documents and analyzes discontinuities in lot size and other housing characteristics across jurisdiction borders. We then construct a novel measure of how exposed a house is to density. Starting with precisely geocoded micro data on the location of each owner-occupied, single-family housing unit in a border area, we then count how many other housing units are within 500 meters of each home based on information in Census block data from the 2010 decennial census. This calculation gives us an explicit house-level measure of density, which can be used to evaluate differences in density within the small geographic areas abutting the borders in our sample.

A discontinuous increase in lot size at a border induces variation in nearby houses' density exposures, as the houses that are in one municipality but are closer to its boundary are more exposed to the size of houses in the other municipality. On average, we find house prices are nearly \$40,000 more expensive and lot sizes are over 3,000ft^2 larger on the side of the border with a more restrictive (i.e., a higher) minimum lot size restriction. This jump in lot size creates variation in density exposure that evolves continuously through the administrative boundary. The

[1]See the plot of the S&P CoreLogic Case-Shiller National Home Price Index in the FRED files at https://fred.stlouisfed.org/series/CSUSHPINSA for the underlying index itself. It shows real, constant-quality prices to have roughly doubled since the middle of the 2010s, and to be well above their pre-global financial crisis peaks.

[2]See Gyourko and McCulloch (2023) for a discussion of the handful of such efforts made across the country. On the political front, even though housing regulation primarily is the responsibility of state and local government, the issue has been 'nationalized' via various policy proposals of the candidates in the current Presidential campaign.

[3]In a survey conducted by Pew Research Center, 57% of Americans said they would prefer to live in a community where "houses are larger and farther apart, but schools, stores and restaurants are several miles away." PewSurvey.

boundary discontinuity design allows us to use this variation in density exposure coming from a known source to estimate the price effects of density exposure while still forming comparisons within a small geographic area.

The empirical strategy is completed by connecting the estimated price effects to a structural hedonic model of housing choice (Rosen, 1974, Bajari and Benkard, 2005). The hedonic model views housing demand as a continuous choice of housing and neighborhood characteristics. In this context, a homeowner's marginal willingness-to-pay (MWTP) is recoverable from the derivative of the estimated hedonic price function. Further welfare analysis can then be conducted with assumptions on the form of homeowners' utility.

Recovering the preference for density from the structural hedonic model reveals a general distaste for density, but one which is far from uniform. Overall, about two-thirds of households have at least a moderate dislike for more density, but a few have an intense dislike and about one-third positively value density. In our baseline specification, the average welfare loss if all homeowners in our sample are exposed to a 1/2 unit per acre increase in density (which is equivalent to a 0.3 standard deviation increase in density) in their surrounding area is about -\$9,500, with the skewness in the distribution evidenced by the fact that the median welfare loss is only about -\$5,200. For context, this counterfactual 1/2 unit per acre increase amounts to adding just under 100 units in the 194 acres of surrounding land within 500 meters of a homeowner's parcel.

Those results are for an increase in density arising from any housing tenure—owned or rented. Using data on the composition of census blocks from the 2010 Census, we are able to provide breakdowns on welfare costs by whether the increase in density arises from additional owner-occupied homes or from renter-occupied units. Stark heterogeneity is evident here, with far higher welfare losses associated with exposure to more rental unit density. The results for the same increase in density arising solely from more owner-occupied homes is similar to that just reported for housing units of any type. In contrast, the mean welfare loss if the density increase is only due to more renter occupied units is more than five times larger, at just over -\$56,000. The median is -\$12,700.

Section 7.C shows that these two key conclusions—that the typical welfare loss is economically, not just statistically, meaningful and that the distaste for rental-related density is far greater than that for owner-occupied related density—are robust to a host of sensitivity analyses including alternative hedonic specifications, different measures of density exposure, and more narrowly defined border areas. These results also hold if we treat each side of a border pair as a distinct market and estimate separate hedonic regressions using only observations on each side of the border. By only forming comparisons between houses within the same municipality, this specification addresses the potential concern that provision of unobserved public goods is driving our estimates. Finally, experimentation with different utility parametrizations (Bishop and Timmins, 2019) indicates that our estimates are conservative in nature.

Our research design and the richness of the CoreLogic microdata also allow us to characterize heterogeneity in preferences across different kinds of neighborhoods. We find larger welfare losses from an increase in housing density in neighborhoods that are both low density and affluent. Households in both the bottom density quartile and the highest income quartile lose an average of about $30,000 from a 1/2 unit per acre increase in total housing unit density, while households in both the highest density quartile and the lowest income quartile lose only $2,000 on average. Estimates of losses from increasing the number of rental units are more extreme. Households in both the bottom density quartile and the highest income quartile lose an average of $263,000 from the equivalent increase in rental unit density, while households in both the highest density quartile and the lowest income quartile lose only $5,000 average on average.

All estimates considered thus far are under a counterfactual in which households *must* continue to live in the (now) higher density neighborhood. In the medium or long run, households with a strong preference against density may adjust their housing consumption by moving to a lower density neighborhood. If they do, allowing households to move at a cost proportional to their house value has only a modest impact on the typical welfare loss associated with increasing owner occupied-related density, but dramatically decreases the magnitude of the average welfare loss from increasing the number of rental units per acre. For example, if households are able to move at a cost equal to 15% of their house value, nearly one-fifth of households in our sample choose to do so, and the average welfare loss from increasing rental unit exposure falls from $56,000 to $9,000. In general, moving at moderate cost limits the influence of the small share of households who suffer extremely high welfare losses under higher rental unit exposure. That said, it is not clear that the minimum cost strategy will be to move. As we discuss below, it could be optimal to fight densification through the local political process, as our results on the nature of the distribution of household-level welfare losses suggest there is a ready coalition to contest and win such a battle.

These estimates of the existing owners' average welfare loss reflect a total willingness-to-pay (WTP) to avoid housing density that may encompass many possible factors. For example, the demand for low housing density may be driven by a preference over neighbors or valuation of the changes in other local amenities such as crime, pollution, traffic flow, noise, etc. The full decomposition of this total WTP into its underlying drivers is an important next step for research on this issue. Given the already long length of this paper, we conclude with a brief analysis that shows how important just one of those potential factors could be. Specifically, we examine the role of preferences over neighbors in mediating our estimate of the distaste for density. We do so by augmenting our hedonic specification with demographic controls for race, income, education, and family structure. Including these controls explains little of the welfare loss associated with increased owner occupied density, but accounts for nearly two-thirds of the loss from an increase in rental density exposure. Thus, neighbor preferences appear to explain some—but not all—of the aversion for housing density, and this factor looks to be be especially important for density arising from rental units.

Finally, it is important to delineate what is and is not accounted for in our welfare calculations. The characterization of the preferences of incumbent local homeowners help us understand a strong demand for the regulation of density, but it is noteworthy that their losses may reflect a strong NIMBYism that does not account for the welfare of non-residents or of overall efficiency. Moreover, the distribution of those preferences helps us recognize why this issue has proven so hard to address though a political system characterized by strong local control of land use. That conclusion begs the question of whether a new approach that limits local control is needed. We hope our findings provide a first step towards examining that question, which clearly will require insights from political economy, not just urban economics. That collaboration across fields will be necessary also is suggested by the highly regressive nature of any compensation scheme that might be introduced to mitigate opposition to densification efforts. In addition, we hope that our results are a stepping stone to develop solutions to the large negative externalities of binding land use regulation at a macro level identified by Hsieh and Moretti (2019), and Duranton and Puga (2023). If, for example, a highly productive labor market area is inefficiently small, the optimal way to change regulatory strictness within that area needs to be informed by results such as ours. Thus, one other direction for future research is to account for the full general equilibrium effects of densification, with our results being an important input into that analysis.

A *Relation to the Literature*

This paper contributes to a large literature on land use regulation. Recent research has focused on how regulations such as density limits increase the cost of housing by restricting local market housing supply (Glaeser and Gyourko, 2018, Gyourko and Krimmel, 2022) or decrease aggregate output at the national level by limiting the size of the overall stock in a market (Hsieh and Moretti, 2019, Duranton and Puga, 2023). While there is now significant evidence on the cost of land use regulation at the market or national levels, there is relatively little evidence on land use regulation's benefits, which mostly occur at the local level.[4] Baum-Snow (2023) discusses the role of local land use regulations in driving house price growth in productive U.S. cities, and notes the absence of evidence about how amenities in neighborhoods would be affected if these areas did experience densification. He further comments that quantifying any potential externalities associated with increasing density should be an important input into policies that seek to reduce housing costs through removal of density restrictions. We are able to address this latter concern by providing estimates of the economic cost of density to existing homeowners, as well as by characterizing the nature of that distribution.

[4] An older hedonic literature has interpreted higher house prices from regulations as welfare enhancing. For example, McConnell and Walls (2005) study open space provisions and review reduced form hedonic estimates about the value for proximity to open spaces.

Within the land use regulation literature, several articles exploit variation in zoning around geographic boundaries to study the effects of regulation (e.g., Turner et al. (2014), Kulka et al. (2024), and Song (2024)). Turner et al. (2014)) provides a theoretical framework to understand the price effects of zoning on vacant land sales around a municipal boundary by decomposing zoning's overall welfare effect into three parts: (a) an own lot effect, which reflects the direct cost of a land use constraint on the owner of a parcel; (b) an external effect, which reflects the spatial spillovers of land use constraints on neighboring parcels; and (c) a supply effect, which reflects regulation's impact on the scarcity of developable land. Kulka et al. (2024) present a similar decomposition of the pricing gradient for house rather than land sales.[5] While Kulka et al. (2024) focus on the own lot effect, in this paper we estimate the external effect of density restrictions.

Few other papers have estimated land use regulation's external effects or the spillovers from land use regulation on other nearby parcels, which varies in relation to the boundary. In the framework of Turner et al. (2014), the external effect represents the only channel by which land use regulation can increase welfare. Turner et al. (2014) find a negative external effect which suggests an unambiguous welfare loss from more restrictive land use regulation, but note that their estimates are imprecise.[6] In this paper, we construct an explicit measure of an externality (i.e., exposure to density) that varies in relation to the municipal boundary, and show how the boundary discontinuity generates systematic variation in this measure. We suspect several other neighborhood characteristics of interest to urban economists may follow similar dynamics in relation to municipal boundaries where there are shifts in town policy or provision of public goods.

We use house sales rather than land sales to evaluate residents' welfare changes. An advantage to working with home sales is that the number of observations is much greater and they are not concentrated near the urban fringe. While Turner et al. (2014) exploit the direct connection between land value and welfare that arises in the seminal Alonso-Mills-Muth urban model, the link between house values and welfare is less direct because a house sale is a result of an individual selecting their house based on its physical and neighborhood characteristics. This necessitates a model that can accommodate a consumer maximizing their utility by choosing their house based on its differentiated product characteristics. Hence, we embed a boundary discontinuity design in a structural hedonic model of housing demand (Rosen, 1974, Bajari and Benkard, 2005).

Boundary discontinuity designs have been used to identify individual structural demand

[5]In Kulka et al. (2024), the own lot effect is termed the direct effect and the external effect is termed the neighbor effect.

[6]Another study estimating the external effects from regulation by a different research design is Davidoff et al. (2022), which finds negative price effects from adding accessory dwelling units on neighboring houses' property values following a zoning reform in Vancouver. Like Davidoff et al. (2022), we find evidence of negative price effects from higher density, suggesting a positive external effect from more restrictive zoning that limits density. Blanco and Sportiche (2024) also find negative spillover price effects from higher density housing development bypassing zoning regulation via Massachusetts 40B, but only around larger developments.

parameters in a discrete choice framework (Bayer et al., 2007), but to our knowledge no paper in housing or urban economics has used a boundary discontinuity design to identify demand parameters in a hedonic model. In general, uses of the hedonic model to retrieve demand primitives have been relatively rare in applied work (Greenstone, 2017). We add to the set of articles estimating nonmarginal changes in hedonic models, including Bajari and Kahn (2005), Bishop and Timmins (2019), and notably Diamond and McQuade (2019).[7] The latter apply the methodology from Bajari and Benkard (2005) to estimate homeowners' preferences for proximity to affordable housing projects, identifying price effects through a spatial difference-in-differences strategy. In contexts in which there is a continuous choice of housing characteristics, such as a spatial spillover that varies continuously across administrative boundaries, a hedonic model may be a more natural modeling choice than a discrete choice framework.

Finally, as the discussion of our results above suggests, by using a structural hedonic model to estimate individual preference heterogeneity over housing density, we are able to recover preferences and estimate welfare impacts over density by neighborhood type and demographic category, something that has not been done in previous research.

The plan of the paper is as follows. Section 2 provides a simple model of a border with a minimum lot size restriction. Section 3 describes our data sources and reports summary statistics. Section 4 presents reduced form evidence of discontinuities in lot size and sale price at boundaries where minimum lot size restrictions change. Section 5 details the structural hedonic model used to estimate demand parameters. Section 6 describes the specifications of the hedonic regressions. Section 7 presents estimates of preferences and the welfare change under counterfactual levels of density exposure. Section 8 concludes and discusses directions for future research.

2. Stylized model of a border with a minimum lot size restriction

In this section, we present a stylized model that describes how lot sizes, housing density, and house prices vary around a municipal border where the minimum lot size increases. There is one border between two municipalities in 1-dimensional space x where the boundary is at $x = 0$. The right side of the border is more highly regulated than the left side. On the left, all houses have lot size $\underline{l}$, while on the right a minimum lot size constraint $\overline{l}$ is binding so all houses have lot size $\overline{l} > \underline{l}$. Figure 1 (a) plots houses' lot sizes as a function of distance to the border where lot size jumps discontinuously at the border.

Next, we define a measure for how exposed a house is to neighborhood housing density at each location x, which we term the density exposure $D(x)$. Let R be a constant bounding the spatial extent at which spillovers from housing density are relevant to individual homeowners. A measure of density exposure for a house at location x' is $D(x') = 2R/l_{avg}(x')$, where $l_{avg}(x')$ is the average lotsize for houses within R distance of the house. Equivalently, in this simple model,

[7] Another recent example is Uribe (2022) who exploits notches in housing tax subsidies to retrieve structural parameters in a hedonic model.

$D(x')$ is the number of housing units within R distance of x'. For example, $D(x')$ could be the number of housing units within $R = 500$ meters of the house at location x'. Solving first for $l_{avg}(x')$,

$$l_{avg}(x') = \frac{1}{2R}\int_{x'-R}^{x'+R} l dx = \begin{cases} \underline{l} & \text{if } x' < -R, \\ \bar{l} & \text{if } x' > R, \\ \frac{1}{2R}[\int_{x'-R}^{0} \underline{l} dx + \int_{0}^{x'+R} \bar{l} dx] & \text{if } x' \in (-R, R), \end{cases}$$

where

$$\frac{1}{2R}[\int_{x'-R}^{0} \underline{l} dx + \int_{0}^{x'+R} \bar{l} dx] = \frac{1}{2R}[-(x'-R)\underline{l} + (x'+R)\bar{l}] = \frac{1}{2R}[x'(\bar{l}-\underline{l}) + R(\bar{l}+\underline{l})]$$

Then, the density exposure as a function of x is given by the following,

$$D(x') = \begin{cases} 2R/\underline{l} = \bar{D} & \text{if } x' < -R, \\ 2R/\bar{l} = \underline{D} & \text{if } x' > R, \\ 4R^2/[x'(\bar{l}-\underline{l}) + R(\bar{l}+\underline{l})] & \text{if } x' \in (-R, R). \end{cases}$$

Panel (b) of Figure 1 plots the relationship of density exposure on x. Density exposure is mechanically related to the lot size, and the discontinuity at $x = 0$ induces changes in $D(x)$. Density exposure is flat until $x \in (-R, R)$ where exposure to the other side of the border induces a slope change. Moving from left to right, density exposure declines as houses are located closer to and then further into the municipality with the higher minimum lot size. Additionally, the density exposure can be converted to more common metrics of housing density such as the number of units per acre by dividing by the geographic area within R distance of a location.

Individuals value houses based on a unit's lot size, density exposure, and other property and neighborhood characteristics. Individuals' valuation of housing characteristics are then capitalized into house prices. Assume for exposition that households prefer larger lots and lower neighborhood density. Figure 1 (c) plots a possible relationship between house price and distance to the border. House price jumps discontinuously at the border because of the discontinuous increase in lot size from regulation, but decreasing exposure to density induces a positive relationship between house price and x even before reaching the border.

Figure 1: Stylized border in 1-d space.

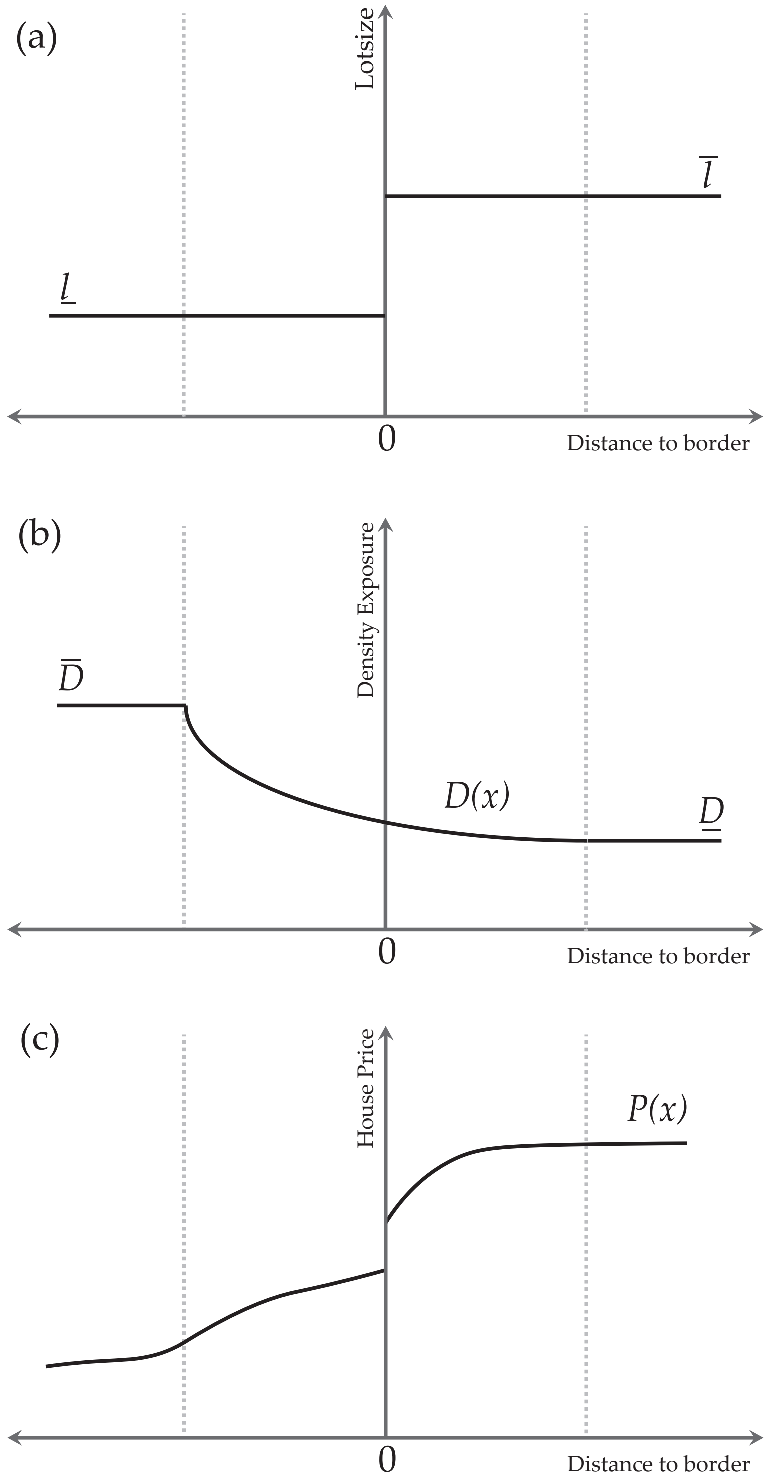

Note: This figure illustrates the relationship between lot size, density exposure, and house prices over space around a municipal boundary with a binding minimum lot size constraint. Panel (a) plots lot size as a flat gradient that increases discontinuously at the boundary entering the more regulated municipality. Panel (b) shows density exposure as a function of the lot size evolving continuously through the municipal boundary. Panel (c) plots the house price gradient if homeowners prefer larger lots and lower density exposure.

This simple model illustrates what we seek to exploit in our boundary discontinuity design in terms of a discontinuous increase in lot size at a border inducing variation in how exposed those individual homes are to density within a small geographic area. Using comparisons only within a small geographic area limits the confounding effects of unobserved neighborhood characteristics. In the case of density, this exposure is directly related to lot size, although similar reasoning could be applied to other externalities resulting from discontinuous changes. Despite the simplicity of this model, we report below strong visual evidence of similar relationships between these three outcomes and x in the data. For completeness, in Appendix D we present a more formal framework following Turner et al. (2014) and Kulka et al. (2024) that describes the pricing gradient as a function of distance to the border. For the special case in which lot size and density exposure are the only relevant housing characteristics, it is straightforward to show that the former trait encompasses the own lot effect and the latter encompasses the external effect.[8] Furthermore, a Turner et al. (2014)-type model in which housing prices vary by distance to the boundary is equivalent to a hedonic model, but with the variation in housing characteristics also operating through distance to the border. More details are provided in the appendix.

3. Data

Our empirical strategy requires data on land use regulations, maps of administrative boundaries, house prices, and characteristics such as lot and physical structure size, as well as house age. We then construct the measure of density exposure using the housing unit counts of nearby U.S. census blocks. Additionally, in the hedonic regressions we include controls for parcel level measures of slope and elevation, the distance to a variety of other amenities, and neighborhood demographics. We also merge in data on school district reading test scores to control for school quality, which may also change discretely at the border.

A *The Wharton Surveys and Administrative Boundaries*

We use regulation data from both the 2006 and 2018 WRLURI surveys (Gyourko et al., 2008, 2021). Each contains responses from over 2,000 primarily suburban jurisdictions to an array of questions covering the myriad restrictions local governments use. Following the practice in our 2023 working paper (Gyourko and McCulloch, 2023), we focus on density restrictions in the form of minimum lot size restrictions that exist anywhere within the community and how restrictive

[8]In Turner et al. (2014), Kulka et al. (2024) and in this paper, the own lot effect is captured by the discrete jump in housing prices at the boundary, while the external effect is captured by the slope of the pricing gradient around the boundary. In Turner et al. (2014), the supply effect is captured by a parallel vertical shift in the land rent gradient. As we show below, it is the slope of the pricing gradient which drives our welfare calculations. Hence, we interpret our estimates as an external effect. The potential for a local supply effect is discussed in Section 7.C and at the end of Appendix D, but we conclude this is unlikely to play an economically important role in our estmated welfare effects given the patterns in the data.

they are. The range of possible answers varied across the two surveys, which requires us to adopt a standardized set of ranges for density restrictions.

Both surveys asked if there was at least one neighborhood within the community's political boundaries that was in one of the following categories: (a) either no minimum lot size or the most stringent one is less than one-half acre; these are the least strictly regulated places in the sample and possess what we call a "low density restriction index"; (b) those in which the largest minimum lot size ranges from one-half acre to (just under) one acre are moderately regulated and possess a "medium density restriction index" in the analysis below; or (c) those in which the largest minimum is either from 1-2 acres or for 2+ acres; because there were so few communities that reported a 1-2 acre minimum, we group these two categories into a single one for 1+ acre minimums; these places are the most strictly regulated in the sample and are said to possess a "high density restriction index".[9]

Because our research design relies on exploiting variation across administrative boundaries, we restrict our sample to WRLURI survey respondents sharing a border. The U.S. Census provides maps of the administrative boundaries used in our analysis.[10] Along these boundaries, we construct 2 kilometer wide border areas and remove any area covered in water. We then separate each border area into its two administrative sides so that we capture any parcels within 1 kilometer of the boundary on either side. While we begin with fairly large areas around the boundaries to assess the empirical relationships of the outcomes across space, we also present estimates for samples using only closer in parcels.

[9]Each survey question is reproduced in Appendix Figure A1. Density restrictions of some kind exist in almost all WRLURI responding communities, as evidenced by the fact that 94% of all communities responding to the 2018 survey reported having some minimum lot size restriction; 25% reported a 2+ acre restriction somewhere in their jurisdiction. For more detail see Gyourko et al. (2008, 2021)

[10]In the WRLURI data, local governments are typically at the Census Place level, and less frequently at the County Subdivision, or County level. Shapefiles for these geographies and bodies of water were downloaded from the U.S. Census Tiger/Line site at https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html.

Figure 2: Border example from the Cincinnati, OH, metropolitan area

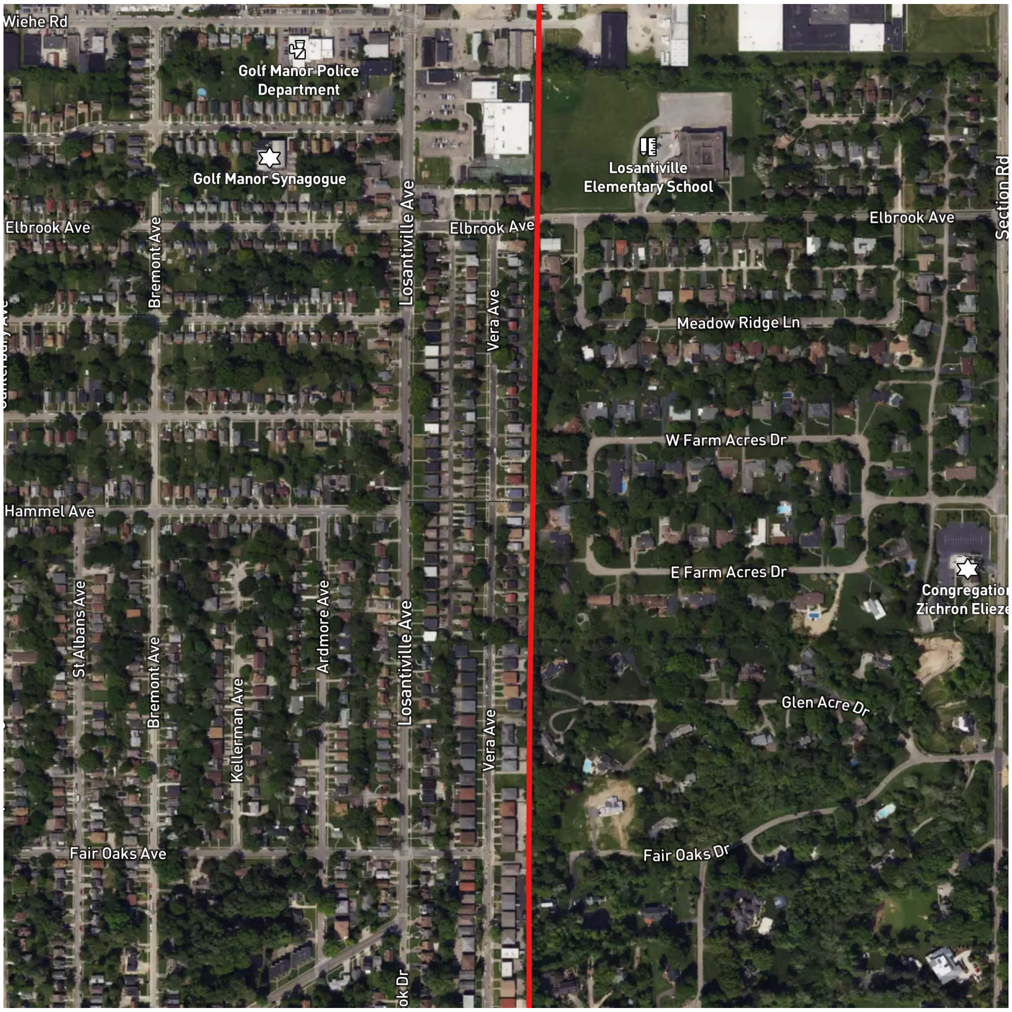

Note: This figure shows a border between two WRLURI-responding municipalities in the Cincinnati, Ohio, metropolitan area. The red line is the boundary between the two municipalities. On the left side of the border is Golf Manor and on the right side is Amberley village. Golf Manor has a low density restriction index while Amberley village has a high density restriction index. Lot and physical structure sizes are immediately larger for houses on the Amberley village side of the border compared to houses in Golf Manor. Basemap by ©OpenStreetMap ©CARTO.

The boundaries between municipalities with measured minimum lot size restrictions form the foundation of our empirical strategy, as is discussed more fully below. As an example, Figure 2 depicts one border pair for the towns of Golf Manor and Amberley village in the Cincinnati (OH) metropolitan area. The red line marks the boundary between the two municipalities. The left side of the border is Golf Manor and the right side is Amberley village. Golf Manor has a low density restriction index and Amberley village has a high density restriction index. Lot and physical structure sizes are immediately larger on the Amberley village side of the border. Municipality-level summary statistics suggest strong evidence of sorting. In terms of demographics, the typical income in the whole of Golf Manor is below the CBSA median at $39,360, while Amberley village is affluent with a typical household income of $145,893. In Amberley, a significantly higher share of households are white and have children.[11] While these statistics suggest that the types of households on either side of the border and their housing preferences may differ, by only forming comparisons within a small border area we can limit confounding of hedonic price estimates coming from correlation between density and unobserved neighborhood quality.

The Wharton surveys contain responses from 1,118 pairs of bordering municipalities across 39 metropolitan areas, with the full distribution reported in the first column of Appendix Table A1. Importantly, not all of these borders represent an area where minimum lot size restrictions and, as a result, lot sizes actually changed at the border. In the analysis sample, we restrict the observations to border areas where there is a change both in the regulatory strictness as measured in the survey and lot size based on the housing transaction microdata. These sample restrictions are detailed below in Section 3.E.

B *Housing and parcel characteristics*

We use parcel-level data from the CoreLogic tax assessment files to measure single family unit density, lot size, sales price, physical structure size and age. These parcels are precisely located by geographic coordinates, allowing us to restrict observations to the parcels within the relevant WRLURI jurisdiction border area. Because of outliers and potentially faulty observations, we winsorize the housing characteristic data at the top (99th) and bottom (1st) percentile. For all border areas in jurisdictions for which regulation data come from the 2018 WRLURI survey, we use housing traits reported as of the 2019 CoreLogic files in our empirical analysis. For those whose regulation data come from the 2006 survey, we exclude all homes built after 2007 in our calculations and analysis. This further reduces the micro samples of homes across all border areas by 3%.

CoreLogic micro data on single-family house sales from 1990-2019 allow us to observe sales prices for virtually all such transactions in each border area. Because relatively few houses sell each year, we construct real 2021 values using the Federal Housing Finance Agency House Price Index (FHFA HPI ®) to adjust nominal values from different sales years.

While the CoreLogic data are known to have good coverage with respect to single-unit homes, there is reason to believe it does not capture all rental units, especially those in structures with multiple units. Hence, we turn to census block data from the 2010 decennial census for information on the total number of units in the relevant area (which include vacant units), as well as a breakdown of the type of unit by whether it is owned or rented (if occupied). These data are used to construct our measure of density exposure discussed below.

[11] Demographic estimates from the *American Community Survey* 2016-2020 estimates, `https://data.census.gov/profile/`.

a Elevation and slope

Additionally, we augment our housing trait information with additional controls for the elevation and slope at each house's location. To measure the local topography, we download digital elevation maps (DEMs) from NASA. Specifically, we use maps derived from the Shuttle Radar Topography Mission (SRTM) Global 1-arc second data collection.[12] These data report the elevation for most of the Earth's surface at a spatial resolution of about 30 meters. We also construct maps of the local terrain's slope by translating these elevation maps using GDAL's slope function.[13] We then assign elevation and slope values according to each parcel's location in these maps using the precise geographic coordinates of parcels reported by CoreLogic.

C *Constructing a measure of density exposure*

To create a measure of every individual house's exposure to housing density, around each house's location in the relevant border area, we initially count the number of housing units within 500 meters. We then convert this count to the number of units per acre by dividing by the area of a circle with a 500 meter radius ($\pi 500^2$ square meters) and then translate that to acres. Figure 3 shows the construction of the density exposure for one house very close to the Golf Manor-Amberley border. The red line again represents the boundary between the two municipalities. The blue dot is the location reported by CoreLogic for this particular house. The blue circle is centered at that house and has a 500 meter radius. The 2- or 3-digit numbers in black are the counts of total housing units in each census block. The location of a block is given by a single set of GPS coordinates at its interior point as reported by the Census. If the number is within the blue circle, all units within the block are considered within the 500 meter radius; if the number is outside the blue circle, none of the units count towards the density exposure measure, no matter how close the block coordinates are to the boundary. After repeating this process on every house within 1 kilometer of our sample borders, each house is endowed with a new housing characteristic: the number of housing units per acre in its immediate 500 meter vicinity. In addition, we construct analogous measures that can differentiate between owner versus renter occupied units using the same counts by housing tenure.

[12] Description and downloading of the SRTM data is available at `https://lpdaac.usgs.gov/products/srtmgl1v003/`.

[13] GDAL is a translator library for raster and vector geospatial data. More detail can be found here `https://gdal.org/`.

Figure 3: A measure of a border parcel's density exposure

Note: This figure shows the construction of a density exposure measure for one parcel close to the Golf Manor-Amberley border in the Cincinnati, OH metropolitan area. The blue dot is the location reported by CoreLogic for the parcel of interest and the blue circle shows the area within 500 meters of that parcel. The black numbers are counts of the total number of housing units in a given census block plotted at its interior point as reported by the U.S. Census. To construct a measure of density we total the counts of all census blocks with internal points within the blue circle and divide by the circle's geographic area ($\pi 500^2$ square meters). We then convert from square meters to acres. We apply this process to every single family parcel in the border areas in our sample. Basemap by ©OpenStreetMap ©CARTO.

Figure 4 reports the cumulative probability distribution and summary statistics of the density exposure variable among all houses within 1 kilometer of bordering WRLURI municipalities. The mean (median) number of total units per acre, which included owner-occupied, renter-occupied and vacant units, is 3.4 (3.0) with a standard deviation of 2.5. The interquartile range for our density measure runs from 1.8 units per acre to 4.3 units per acre. Note that there are typically some renter-occupied units in most border areas, although it is quite rare for them to be the dominant tenure type. Some vacancy is common, too.

Figure 4: Distribution of density exposures

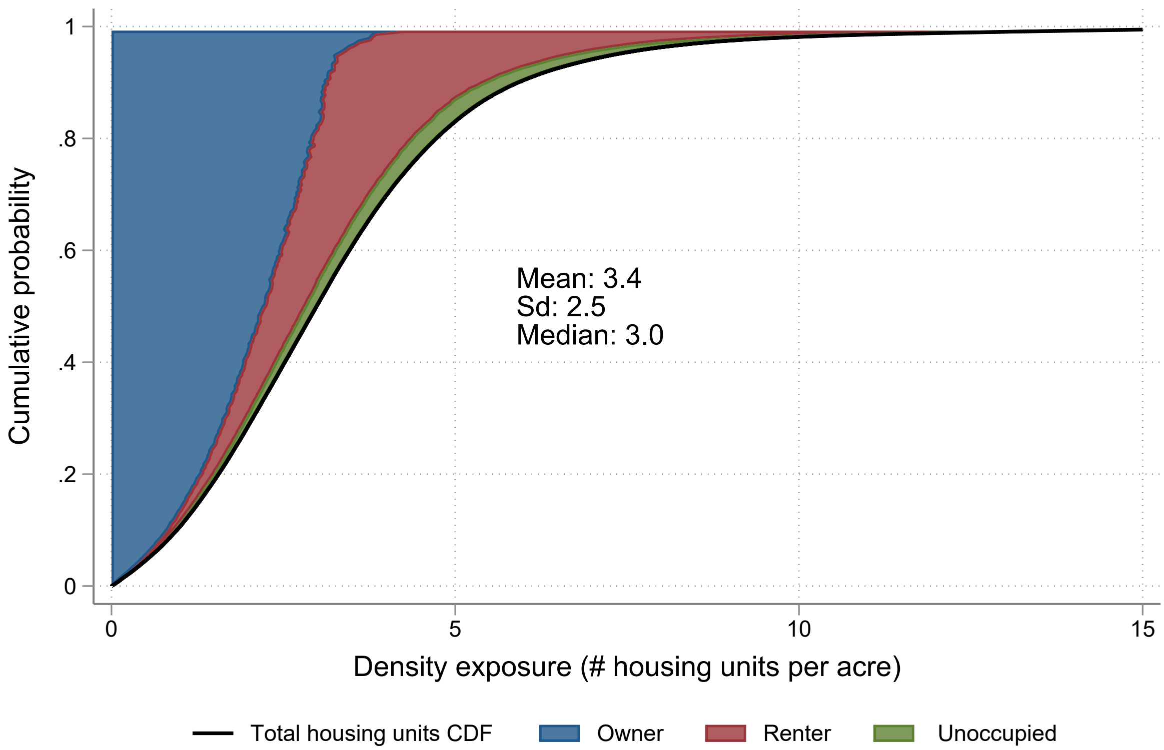

Note: This figure plots the empirical CDF of the total density exposure measure constructed following the process outlined in Section 3.C. The x-axis is the density exposure measured by the number of housing units per acre within 500 meters of a house. The y-axis is the cumulative proportion of observations occurring in each density exposure bin. The colored areas show the typical share of housing units by tenure status across the total housing unit density exposure CDF. The average number of owner occupied, renter occupied, and unoccupied housing units per acre were calculated over 200 bins equally spaced across the y-axis measured by the horizontal distance in each shaded region. A density exposure of 1.81 units per acre and 4.34 units per acre represent the 25th and 75th percentiles respectively among all parcels within 1 kilometer of bordering WRLURI municipalities.

For reference, returning to the Golf Manor-Amberley border pair example, the typical house toward the interior of Golf Manor (houses 250 to 500 meters away from the border) has a density exposure of 5.4 units per acre, which is at the 87th percentile of the overall distribution. For houses located closer to the border, the density exposure drops—the average within 250 meters of the border is 4.2 units per acre, which is at the 73rd percentile of the distribution. Finally 250 to 500 meters away from the border on the Amberley side, density exposure falls to 1.57 units per acre or the 21st percentile. Thus, moving from the interior of Golf Manor to the interior of Amberley Village reduces the average density exposure from the 87th to the 21st percentile of the overall distribution. While the Golf Manor-Amberley border pair is an example with a particularly large lot size discontinuity, we find similar qualitative patterns when pooling all the border areas where density restrictions changed. These results are described further in Section 4.

D *Additional controls*

We also control for a number of other factors in our specifications. Some, such as distance to local amenities that also vary with distance to the border, are potential confounders of our density preference estimate. Others such as school quality often change discretely at the administrative border, but do not vary with distance from the border. Still others such as the racial or demographic composition of the area are included to help us better understand how one's preference for density might be mediated by a preference over neighbors.

a Distance to local amenities

We always include controls for the distance to various local amenities in the hedonic regressions. We expect density exposure to vary across space as houses approach the boundary. If other local amenities also vary systematically relative to the boundary, these variables could confound our price effect estimates. For example, if more highly regulated municipalities tend to have more public parks, we could be capturing the value of proximity to parks rather than the value of lower density exposure *per se*. In response, we include controls for distance to the nearest green space, highway, body of water, public and private school, and the center of the CBSA's central business district.

Shapefiles for these geographic features were retrieved from the following sources:

- Green spaces are collected from the OpenStreetMap's API by selecting 'parks' features.[14]
- Highways are also collected from the OpenStreetMap's API by selecting 'motorway', 'trunk' and 'primary' features.
- Bodies of water are from the Census 2019 TIGER/line shapefiles for water area.[15]
- Locations of public and private schools are defined by the Common Core Data (CCD) for the Homeland Infrastructure Foundation-Level Data (HIFLD).[16]
- CBSA centers are based on the coordinate location from searching the central city of the metropolitan area on Google Maps.

For every parcel in our border areas, we calculate the distance to the nearest shape for each of the categories listed above.

[14]OpenStreetMap features were retrieved using the Python package OSMnx. More detail can be found at `https://github.com/gboeing/osmnx`.

[15]Water shapefiles were retrieved from the U.S. Census' FTP site at `https://www2.census.gov/geo/tiger/TIGER2019/AREAWATER/`.

[16]These data can be downloaded from HIFLD's open data website at `https://hifld-geoplatform.opendata.arcgis.com/search?collection=Dataset&groupIds=f16c582f00184cb094affff556fe57ee`.

b School district test scores

Because school district boundaries often coincide with administrative boundaries, school quality is a potential confounder of the impact of regulation. Access to high quality schools likely affects housing characteristics and could be correlated with regulation itself. To address this issue, we include controls for school quality by matching school district test score data from the Stanford Education Data Archive (SEDA) to each side of a border.[17] Specifically, we measure a border side's school quality by the standardized average 3rd grade reading score in its respective school district. By construction, every parcel on the same side of the border is associated with the same school district.

We discretize our school quality measure as follows. Each border area is assigned its jurisdiction's 3rd grade reading score for 2018 or the most recent year available. We then assign each border area to one of three categories: (a) 'Lowest School Quality', which is comprised of those with scores in the bottom quartile of the sample distribution; (b) 'Average School Quality', which is comprised of those with scores in the interquartile range of the distribution; or (c) 'Highest School Quality', which is comprised of those with scores in the top quartile of the distribution.

c Demographics

Finally, we collect data on local demographics from 2010 U.S. Census block and block group data. Specifically, we collect counts on race, age, marital status, and presence of children in nearby census block and we retrieve the median household income and share of the population over 25 with 4+ years of college education in each parcel's block group. For demographic variables available at the census block level, we construct exposure measures similar to the density exposure measure. For example, we construct race exposure variables for the share of nearby households that are Black or Hispanic. For the Black exposure measure, we sum the counts of people identifying as Black in census blocks within 500 meters of a parcel and divide by the total population in these census blocks. For each house in our sample this process creates a measure of the share Black within 500 meters. We perform an identical process for the share Hispanic, middle aged, married, and with children.

E *Summary statistics*

We define two distinct subsamples of the borders of contiguous WRLURI municipalities that are used in the analyses. First, for the regression discontinuity analysis, we restrict the estimation sample to borders where the density restrictions actually changed. For over half of all bordering municipalities, we do not find a change in the reported density restriction index. This occurs when both municipalities report the same density restriction index (e.g., both municipalities could

[17] These data may be downloaded at `http://purl.stanford.edu/db586ns4974` from the Stanford Education Data Archive (Version 4.1), Reardon, S.F., et. al.

report low density restrictions). At these borders, we do not expect to find a discontinuity in lot size because there was no change in regulatory strictness.[18] Ordering our border pairs from lower to higher density restriction index, we only include pairs with one of the following density restriction index patterns: (1) Low to Medium Density Restriction Indexes; (2) Low to High Density Restriction Indexes; or (3) Medium to High Density Restriction Indexes. Additionally, we drop borders that do not have at least 100 single-family houses within 1 kilometer with non-missing sales price information, as well as borders that follow a highway or river.[19] We call the subsample resulting from these restrictions "ΔDRI" because there is a measured change in the level of density restriction index at these borders.

Second, for the hedonic regression analysis we further restrict the ΔDRI border areas to those where we find a measurable increase in lot size as one moves across the border from left to right (i.e., from the lower Density Restriction Index community to the higher Density Restriction Index community). Specifically, we restrict the sample of borders in ΔDRI to those where we find a regression discontinuity estimate for lot size greater than $+500$ ft^2. While on average we find significant discontinuities in lot size and sale price pooling together all the ΔDRI borders, it is not the case that every border with a reported change in density restriction also has a change in lot size at the border. This could be because density restrictions are nonbinding in that area, for example if developers would have built large lots regardless of the regulation, or because the survey instrument may reflect a density restriction somewhere else in the municipality. We call the subsample of borders where we find a measurable discontinuity in lot size ΔLot. To estimate a hedonic regression with this subsample, we must have sufficient parcels with sale price information. In that regard, we restrict the estimation sample to border areas that have at least 25 observations on both sides of the boundary, in addition to having at least 100 observations in total.

Table 1 provides summary statistics on our underlying data by the subsample of borders. The first column provides information on all potential observations from any jurisdiction that responded to either of the WRLURI surveys, with parcel level data from Corelogic included on any home within 1,000 meters of the border. This clearly is a suburban-dominated sample, as the typical lot size is over one-quarter acre, with constant \$ 2021 price just over \$468,000. This is reinforced by the fact that there are only 24 central cities of CBSAs in these data, and they only constitute 47 distinct border pairings with suburban jurisdictions. The other 4/5ths of the pairs compare a border area of one suburban jurisdiction with that of another suburban jurisdiction. Finally, there are millions of housing observations within the communities on either side of the 1,118 borders.

[18] In the left column of Appendix Figure A3 we show plots for this set of borders and do not find evidence of discontinuities in lot size or house prices, nor systematic variation in density exposure around these borders.

[19] Specifically, we drop any border where 25% or more of the parcels within 100 meters of the border are also within 100 meters of either a highway or body of water.

The remaining columns restrict the sample in various ways, with the final one containing the hedonic analysis sample. The second column imposes the restriction that there be a change in the density restriction index (labeled ΔDRI in the table) at the border. This cuts the number of borders and municipalities by more than half, but still contains well over 850,000 housing observations with non-missing sales prices. Relatively little else changes significantly in terms of house quality, or demographics such as race. The third column further restricts the border depth to 500 meters. This reduces the sample size substantially, but does not affect the averages very much. The final column restricts to borders with an increase in lot size (ΔLot). This yields the final sample used in the hedonic analysis below that still contains over one-quarter million house price observations from 217 borders and 325 jurisdictions, which reflects the fact that some municipalities have multiple borders with different communities.

Comparing housing traits across the samples, most characteristics are fairly similar with the exception of the rental unit density. Restricting observations to areas with binding minimum lot sizes and sufficient single family housing sales skews the sample toward areas that are more single family. That said, suburban, predominantly single family neighborhoods are the primary neighborhood type of interest concerning density restrictions, and the other observable characteristics do not suggest our results are otherwise based on a highly selected sample.

In the final hedonic sample, the 217 border pairs (implying 434 sides) are distributed across the 32 metropolitan areas with nonzero border pairs listed in column 3 of Appendix Table A1. The Chicago metro area makes up the highest percentage of border pairs at over 15%, Cincinnati contains nearly 10% and Detroit 9%. The other two-thirds of the pairs come from a wide array of other markets. The interquartile range for the number of observations in a border area for this depth runs from 476 to 1,458 parcels.

Note also that density exposure differs materially depending upon the underlying set of housing units being counted. For total units (row 2 of Table 1), there are 2.75 housing units per acre on average in the exposure area. There are 194 acres in our exposure circles, which implies about 534 housing units in the circle with a 500 meter radius (row 2 of Table 1).[20] For owner-occupied units, the density exposure in the regression sample is just under 2 units per acre (row 3 of Table 1). There are only 0.64 rental units in the exposure area on average (row 4 of Table 1).

[20]The area of circle is given by $\pi 500^2$ square meters, which amounts to 786,148 square meters or 8,462,018 square feet. Dividing that number by the 43,560 square feet in an acre yields approximately 194 acres.

Table 1: Summary statistics

	(1) All < 1,000m	(2) < 1,000m, ΔDRI	(3) < 500m, ΔDRI	(4) < 500m, ΔLot
Lot size (ft^2)	12,574	12,566	12,549	12,527
	(17,453)	(15,808)	(15,797)	(14,147)
Density exposure, total units per acre	3.31	3.02	2.97	2.75
	(2.21)	(1.86)	(1.84)	(1.61)
Density exposure, owner occ. units per acre	2.09	2.02	2.00	1.96
	(1.09)	(1.00)	(1.00)	(0.97)
Density exposure, renter occ. units per acre	0.98	0.79	0.77	0.64
	(1.32)	(1.06)	(1.03)	(0.88)
Sale price ($ 2021)	$468,798	$459,657	$458,361	$424,224
	($449,452)	($440,324)	($436,339)	($369,507)
Living area (ft^2)	1,864	1,908	1,908	2,019
	(864)	(867)	(861)	(896)
House age (years)	55	52	51	48
	(26.3)	(25.8)	(25.4)	(25.3)
Share black < 500m	0.11	0.10	0.10	0.08
	(0.20)	(0.18)	(0.18)	(0.14)
Household income ($ 2021)	$91,148	$93,453	$93,289	$98,196
	($43,524)	($43,276)	($42,987)	($43,638)
Num. parcels	2,989,361	1,360,558	719,181	263,340
Num. parcels (sale price)	1,869,071	856,294	454,035	263,340
Num. borders	1,118	425	425	217
Num. municipalities	1,052	560	558	325

Note: This table reports means and standard deviations of housing characteristics for single family parcels near the borders of adjacent WRLURI municipalities. Col. (1) reports summary statistics for all parcels within 1 kilometer of any border between WRLURI municipalities. Col. (2) reports on parcels within 1 kilometer of borders between municipalities where the reported Density Restriction Index increased (ΔDRI) after removing borders following a major highway or a river. Col. (3) reports on parcels in a similar sample but shortening the border area depth to 500 meters. Col. (4) shows summary statistics on the parcels used in the hedonic regressions, which represent parcels with sale price information in border areas where there is a measured discontinuous increase in lotsize (ΔLot). Housing traits are winsorized at the 1% level.

4. Discontinuities and density exposure at municipal boundaries

Figure 5: Distance gradients in ΔDRI borders

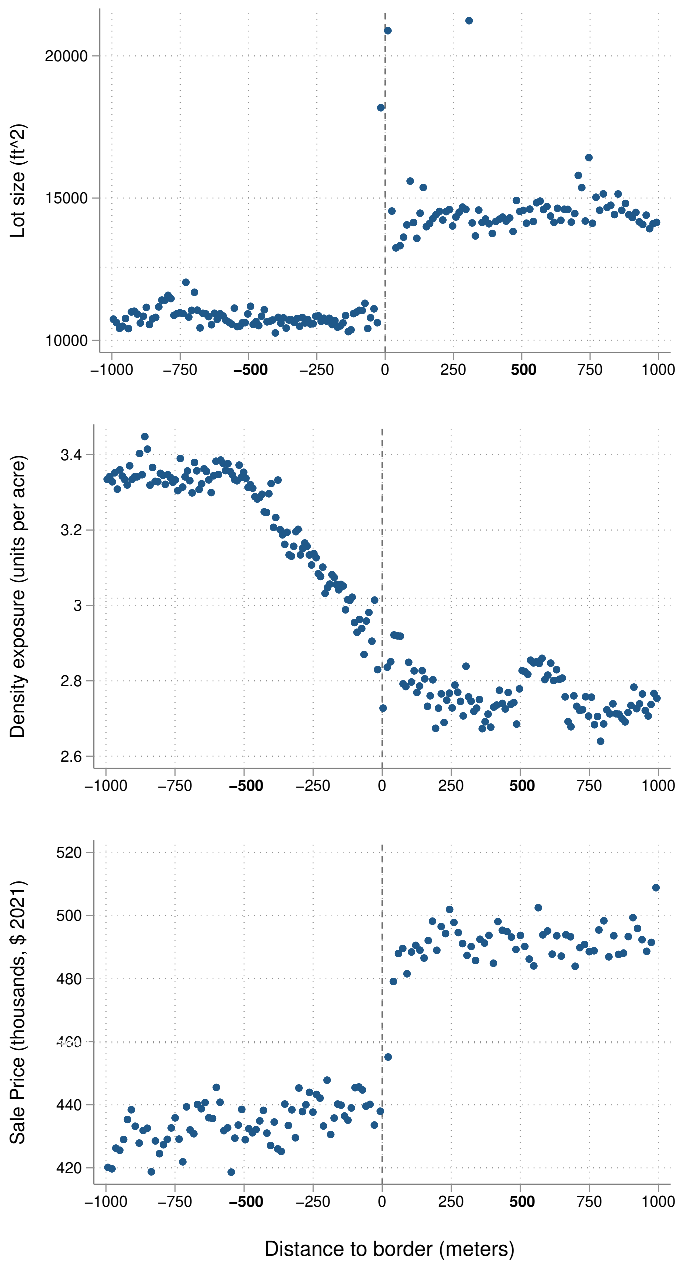

Note: This figure presents binned scatter plots of the three outcomes lot size, density exposure, and sale price on distance to the municipal boundary pooling together the border areas in the ΔDRI sample. The plotted relationships between the outcomes and distance to the boundary adjust for border pair specific fixed effects following Cattaneo et al. (2024a,b). The top panel is the plot for lot size, the middle panel is for density exposure, and the final panel is for sale price. The first two panels are calculated using 1,360,558 parcels spread throughout 425 border areas. The final panel is calculating using the 856,294 parcels with sale price information.

Figure 5 depicts binned scatter plots of the data on lot size, density exposure and sale price against distance to the relevant municipal boundary using the ΔDRI subsample of borders where there was a change in density restrction. The top panel shows that lot size jumps discretely at the border by about 3,000ft^2. Density exposure (middle panel) starts to drop in the lower Density Restriction Index community well before the border, as expected. Finally, the third panel shows a sharp increase in sales price of over $40,000 at the border. Empirically, our data are quite consistent with the predictions of our stylized model of the border discussed above in Section 2.

Table 2 presents regression discontinuity estimates by bandwidth for lot size and sales price comparable to the top and bottom panels in Figure 5. The estimates almost always are highly statistically significant, with the exception of sales price within 50 meters of the border. Using smaller border area depths provides better comparison groups because the homes are all very close to the border. However, there is an inevitable tradeoff in statistical precision because effective sample sizes become much smaller.[21] More generally, the impact on both lot size and sales price decreases as the bandwidth narrows, but the changes remain economically meaningful for samples of homes within 175 meters of the border.

Table 2: Regression discontinuity estimates by border area depth (bandwidth)

	1000m	500m	425m	300m	175m	50m
Lot size (ft^2)	3,686***	3,454***	3,312***	2,937***	2,684***	3,662***
	(462)	(442)	(447)	(446)	(547)	(1,050)
Sale price ($ 2021)	45,809***	39,891***	38,621***	34,785***	28,848***	11,085
	(8,107)	(7,167)	(6,821)	(6,021)	(5,951)	(11,243)
Num obs	1,463,632	742,978	629,957	437,785	244,666	53,475
Num obs (sale price)	919,858	468,908	398,096	277,524	155,249	33,150
Num borders	425	425	425	425	425	414

*Note: This table presents regression discontinuity estimates at various choices of border area depth (bandwidth) using local linear regression with a uniform kernel. Like in Figure 5, these specifications adjust for border pair specific fixed effects. Each cell represents an RD estimate using only parcels within the distance to the boundary displayed at the top of that column. In the first row, the outcome is lot size. In the second row, the outcome is sale price. Standard errors, clustered by border area, are reported in parentheses. Significance stars * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.*

While we presented the regression discontinuity results using the ΔDRI border sample to show that minimum lot size restrictions are often binding, we can further refine our sample to the border areas where we confirm there is a lot size discontinuity (ΔLot) because we also observe

[21] In our application, 50 meters is a very small geographic distance. For the typical lot size in our sample, the length of just one parcel is roughly 30 meters (1/4 acres $\sim$ 1012 meters^2 and $\sqrt{1012 \text{ meters}^2} \sim$ 30 meters) and municipal boundaries also typically sit on roads. Taken together, a depth of 50 meters will often only accommodate one house on either side of a boundary. Hence, we are skeptical of estimates based on such a narrow range because of relatively small sample sizes and the possibility that homes right on the border just are different from the average in the broader border area.

lot size at these boundaries. In Appendix Figure A3, we compare the three outcome gradients in borders with no change in minimum lot size restrictions versus the ΔDRI borders (Figure 5 above) versus the ΔLot borders. The left column of this figure serves as a useful placebo test. In borders without a change in density restriction, we do not observe discontinuities in lot size or house prices, and there is no systematic downward gradient in the density exposure. The next two columns compare the regression discontinuity sample ΔDRI to the borders used in hedonic estimation, ΔLot. The final column shows the relationships in Figure 5 are made even sharper by restricting to the borders where we actually observe a discontinuous increase in lot size. Finally, Appendix Figure A4 shows the analogous plots using only the observations used in the estimation sample for the hedonic regressions, i.e. restricting to houses with sales price information within 500 meters of the boundary.[22]

5. Structural hedonic model of housing choice

A *Review of equilibrium in the hedonic model*

To describe equilibrium in a hedonic model, we follow Greenstone's (2017) recent review of the hedonic approach. In Rosen (1974)'s classic model, a differentiated good is described by a vector of product characteristics, $\boldsymbol{z} = (z_1, z_2, \ldots, z_K)$. For a house, these characteristics may consist of structural attributes such as the lot size and the number of bedrooms or neighborhood qualities such as distance to schools, open spaces, or the density of nearby housing. The market price for a house j is a function of the product characteristics:

$$p_j = p_j(z_{j,1}, z_{j,2}, \ldots, z_{j,K}) = p_j(\boldsymbol{z}).$$

The equilibrium matchings of consumers and producers leads to the observed function between house prices and characteristics, termed the hedonic price schedule. Consumers' utility depends on their house's characteristics and consumption of a numeraire good, c.

$$u(c, \boldsymbol{z}) \;\; s.t. \;\; w - p(\boldsymbol{z}) - c = 0.$$

Maximization of the utility function naturally implies that consumers choose characteristics to satisfy

$$\frac{\frac{\partial u}{\partial z_{j,k}}}{\frac{\partial u}{\partial c}} = \frac{\partial p}{\partial z_{j,k}}. \tag{1}$$

Thus, a consumer's marginal willingness to pay (MWTP) for characteristic $z_{j,k}$ is related to the partial derivative of the hedonic price function at her equilibrium product choice.

[22] We perform several other visual checks of the validity of our research design. In Appendix Figure A5, we plot alternative measures of the density exposure on distance to the municipal boundary (middle panel in Figure 5). All density exposure measures have a shape similar to the prediction of the simple model in Section 2 that is kinked at -500 meters and declines continuously through the municipal boundary. In Appendix Figure A6, we separate the sample by whether the change in density restriction was low to high, low to medium, or medium to high. In all cases we find similar relationships over the three outcomes. Hence, we do not think that pooling these three subgroups is affecting our estimation.

It is useful to reformulate the consumer's utility function in terms of a bid (indifference) curve. To do so, substitute the budget constraint into the utility function and fix utility, $\tilde{u} = u(w - p(\boldsymbol{z}), \boldsymbol{z})$. By inverting this equation to solve for the price, we can define the bid curve $\theta(\boldsymbol{z}; w, \tilde{u})$. Then, holding the other housing characteristics fixed, $\theta(z_k; \boldsymbol{z^*_{-k}}, w, \tilde{u})$ expresses the largest amount a consumer would pay for alternative values of z_k at the same level of utility. In equilibrium, individuals maximize utility by choosing the bundle of product characteristics at which their bid curve is tangent to the hedonic price schedule. Heterogeneity in individuals' bid curves stemming from different preferences over product characteristics leads to the different housing purchases observed along the hedonic price schedule.

Figure 6 plots the hedonic price schedule considering density exposure as a negative product characteristic ($z_k = D$) and a bid curve for one consumer type. A homeowner has purchased a house at D^* for the price P_0. In the original equilibrium at D^*, the resident's bid curve function θ is tangent to the hedonic price schedule. Consider a consumer of the same type who has not yet purchased her house. For this resident to achieve the same utility at a higher density exposure D', she would need to pay the below equilibrium price P_2 allowing her to afford a high enough consumption of the numeraire good c to offset losses from the change in density exposure.

Figure 6: Monetary measure of the welfare change.

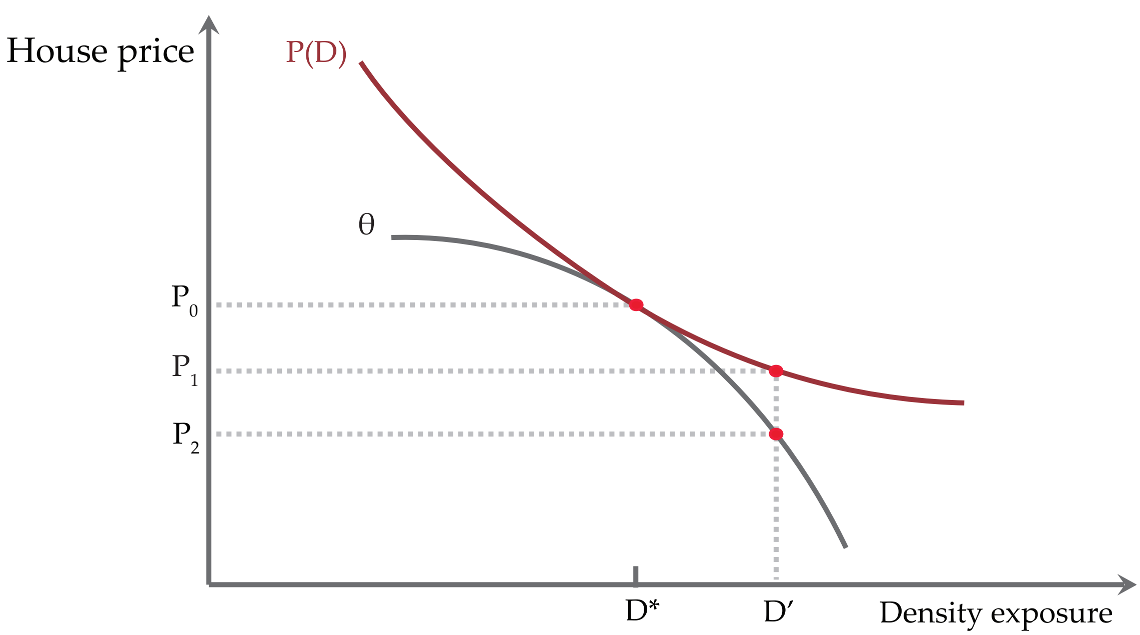

Note: Diagram of the hedonic price schedule as in Rosen (1974). The red line plots housing prices as a function of the density exposure. For one consumer type, we plot her bid curve θ and its tangency point in equilibrium at the density exposure D^*. The partial derivative of $p(D)$ at D^* characterizes that consumer's Marginal Willingness to Pay (MWTP) for density exposure. If θ is known, the welfare change under counterfactual levels of density exposure can be estimated.

If we know consumers' bid curves, we can calculate monetary measures of the welfare change under counterfactual choices of D. Consider the same consumer who owns a house at D^* being suddenly forced to live under the higher density exposure D'. Because the consumer now owns their house, she loses welfare from two channels: the price loss from moving down the hedonic price schedule and the utility change from having to consume a different level of density exposure.

$$\Delta p = P_1 - P_0, \quad \Delta u = P_2 - P_0. \tag{2}$$

The total welfare loss in monetary terms is then the sum of these two components: ΔWelfare $= \Delta p + \Delta u$. Because $\frac{\partial \theta}{\partial w} = 1$, an alternative way to consider ΔWelfare is as the compensation necessary for the homeowner to accept living under the higher density exposure.

B ***Recovering demand parameters***

In the model above, Equation (1) relates each individual's MWTP with the derivative of the hedonic price function at the chosen equilibrium bundle. However, in order to conduct nonlocal welfare analysis, we need to characterize individuals' entire bid curves. Bajari and Benkard (2005) suggest making a parametric assumption on individuals' utility functions. Then for each individual, her entire bid curve is recoverable from the one product choice made in equilibrium. In particular, they use a utility function that is log linear in the product characteristics.

Following those authors, let $j \in J$ be a house consisting of $z_{j,k}$ housing characteristics and ξ_j unobservable characteristics. Let $\bar{D}$ be the highest density exposure in the market and set $z_{j,1} = \bar{D} - D_j$ the inverted density exposure.[23] p_j is the price of house j. w_i is the wealth of individual i. For individual i that considers density exposure a disamenity, assume:

$$\begin{aligned} u_{ij} = \beta_{i,D} \log(\bar{D} - D_j) + \beta_{i,2} \log(z_{j,2}) + ... + \beta_{i,k} \log(z_{j,k}) + \beta_{i,\xi} \log(\xi_j) + c \\ \text{s.t. } p_j + c \leq w_i \end{aligned} \tag{3}$$

Then individual i's taste parameter for density $\beta_{i,D}$ can be solved by

$$\beta_{i,D} = -(\bar{D} - D_j) \frac{\partial p}{\partial D}(D_j). \tag{4}$$

For the majority of individuals in our data we observe a negative price derivative, indicating density exposure is a negative product characteristic. However, there is significant heterogeneity in individuals' preferences for density based on our estimates. For individuals that instead consider density exposure an amenity assume:

$$\begin{aligned} u_{ij} = \beta_{i,D} \log(D_j) + \beta_{i,2} \log(z_{j,2}) + ... + \beta_{i,k} \log(z_{j,k}) + \beta_{i,\xi} \log(\xi_j) + c \\ \text{s.t. } p_j + c \leq w_i \end{aligned} \tag{5}$$

[23] Flipping the axis of density exposure $\bar{D} - D_j$ ensures concavity of u_{ij} when density exposure is a disamenity. Concavity of the utility function is necessary for equilibrium in the hedonic model. In Appendix B, we show the bid curve using a log linear utility function with a negative product characteristic where the characteristic is not flipped. The resulting bid curve is convex and there is no equilibrium along the hedonic price schedule.

Then $\beta_{i,D}$ can be solved by

$$\beta_{i,D} = D_j \frac{\partial p}{\partial D}(D_j). \tag{6}$$

We can therefore estimate individuals' structural demand parameters $\beta_{i,D}$ from the data using the estimated price derivative with respect to density exposure and the observed density exposure D_j.

Using these estimates and the parametric assumption, we can now recover individuals' whole bid curves. For an individual i, rewrite u_{ij} in terms of her bid curve. First, plug in the budget constraint and fix the level of utility:

$$\begin{aligned}\tilde{u}_{ij} = \beta_{i,D}\log(\bar{D} - D_j) + \beta_{i,2}\log(z_{j,2}) + ... + \beta_{i,k}\log(z_{j,k}) + \beta_{i,\xi}\log(\xi_j) \\ + w_i - p\end{aligned} \tag{7}$$

Then, holding all other product characteristics constant and inverting the equation for the price, we get individual i's bid curve along D

$$\begin{aligned}\theta_i(D; \boldsymbol{z^*_{-k}}, w, \tilde{u}) = \beta_{i,D}\log(\bar{D} - D_j) + \beta_{i,2}\log(z^*_{j,2}) + ... + \beta_{i,k}\log(z^*_{j,k}) + \beta_{i,\xi}\log(\xi^*_j) \\ + w_i - \tilde{u}_{ij}\end{aligned} \tag{8}$$

And, we can consider changes along the bid curve at different levels of density exposure. Consider the change from D to D'.

$$\begin{aligned}\Delta\theta_i &= \theta_i(D'; \boldsymbol{z^*_{-k}}, w, \tilde{u}) - \theta_i(D; \boldsymbol{z^*_{-k}}, w, \tilde{u}) \\ &= \beta_{i,D}\log(\bar{D} - D'_j) - \beta_{i,D}\log(\bar{D} - D_j) = \beta_{i,D}\log(\frac{\bar{D} - D'_j}{\bar{D} - D_j})\end{aligned} \tag{9}$$

$\Delta\theta_i$ is the utility change Δu from having to live at a higher density exposure, or the utility loss in Equation (2). By similar arguments for individuals that consider density exposure an amenity, the change along the bid curve is

$$\Delta\theta_i = \beta_{i,D}\log(\frac{D'_j}{D_j}). \tag{10}$$

We can now express the welfare change for individual i in terms that can be estimated from the data. For density exposure as a disamenity,

$$\begin{aligned}\Delta\text{Welfare}_i &= \Delta p + \Delta\theta_i = \Delta p + \beta_{i,D}\log(\frac{\bar{D} - D'_j}{\bar{D} - D_j}) \\ &= \Delta p - (\bar{D} - D_j)\frac{\partial p}{\partial D}(D_j)\log(\frac{\bar{D} - D'_j}{\bar{D} - D_j}).\end{aligned} \tag{11}$$

The terms D_j, D'_j and $\bar{D}$ are observed directly, while Δp and $\frac{\partial p}{\partial D}(D_j)$ can be estimated by hedonic regression. Lastly, the welfare change if density exposure is considered an amenity is given by

$$\begin{aligned}\Delta\text{Welfare}_i &= \Delta p + \Delta\theta_i = \Delta p + \beta_{i,D}\log(\frac{D'_j}{D_j}) \\ &= \Delta p + D_j\frac{\partial p}{\partial D}(D_j)\log(\frac{D'_j}{D_j})\end{aligned} \tag{12}$$

in which, again, all the terms are either directly observed or can be estimated by the data.

6. Hedonic regressions

We estimate the relationship between density exposure and house price for each border area in our sample under many different specifications. Thus, we treat each combined border area as a distinct market and run a separate hedonic regression for it. We then repeat this process varying the set of controls included in the regressions. This strategy lets us control flexibly for potential confounding variables and characterize heterogeneity across neighborhoods by allowing for different relationships between the price, density exposure, and relevant covariates in each neighborhood.

In particular, we consider the following hedonic pricing model. The market price is a function of the density exposure using the observations in a market l. We regress the price $p_{j,l}$ of house j in market l on the density exposure D_j and a set of covariates X_j. In general form,

$$p_{j,l} = f_l(D_j) + X_j'\theta + \epsilon_{j,l} \text{ for } l \in \{1, 2, \ldots, L\}. \tag{13}$$

where l is a border area and $f_l(D_j)$ is a nonlinear function of the density exposure. The controls included in X_j are described in more detail below.

A key assumption of the hedonic model is that comparisons are made in a market that satisfies the law of one price. This is appropriate if the prices of identical homes on different sides of the border area in (say) Figure 2 are the same, *ceteris paribus*. Given the small geographic size of each border area, we believe this is a reasonable assumption.[24]

As our primary specification, we consider a semiparametric model following Robinson (1988) where $f_l(D_j)$ is estimated nonparametrically by local linear kernel regression. This model does not impose any functional form assumptions on the the relationship between density exposure and price while assuming the remaining controls enter linearly.[25] A semiparametric model allows relatively flexible estimation of the relationship of interest while avoiding issues related to the curse of dimensionality in nonparametric models.[26]

After fitting these regressions, $\hat{f}_l(d)$ is used to estimate the price derivative $\frac{\widehat{\partial p}}{\partial D_j}$ for each market and each value of D_j in the data. Identification of the price derivatives relies on a conditional independence assumption: within a market, conditional on X_j, density exposure is unrelated to the error term. In our hedonic regressions, comparisons are always made within a small

[24] However, to check for the possibility that market or supply conditions change at the border, for robustness we also estimated results using separate hedonic specifications for each side of the border. Essentially, this treats each side of the border as a separate market and price comparisons are only made between houses within the same municipality. Those findings are not materially different from our baseline estimates reported in Table 3 and are discussed further in Section 7.C.

[25] In Robinson (1988)'s 2-step procedure, control and outcome variables are first residualized on density exposure nonparametrically. These residuals are then used to estimate the linear component of the model by OLS. After subtracting the linear component due to the other housing characteristics, $f_l(D_j)$ can be estimated by standard local polynomial regression.

[26] As a robustness check, in the appendix we re-estimate all the results under two completely parametric specifications, a simple linear model and a model in which all variables are log transformed.

geographic area and, as is discussed more fully in the next section, we include specifications with a large number of controls to help mitigate potential omitted variable bias.

We also consider specifications in which the density exposure related to renter-occupied housing units is considered separately from density exposure arising from owner-occupied units. Let $D_{j,rent}$ be the rental units per acre in the immediate area of house j and $D_{j,own}$ the owner-occupied units per acre. In this case, we include both terms but treat $D_{j,rent}$ as the nonparametric object of interest while $D_{j,own}$ is treated as a linear term. The augmented specification is

$$p_{j,l} = f_l(D_{j,rent}) + \gamma_{own} D_{j,own} + X_j'\theta + \epsilon_{j,l} \text{ for } l \in \{1, 2, \ldots, L\}. \tag{14}$$

The price derivatives are then estimated by $\frac{\partial \hat{f}_l}{\partial D}(D_{j,rent})$ for the renter occupied density exposure and $\hat{\gamma}_{own}$ for the owner occupied exposure.

A *Recovering preference estimates*

We estimate individual MWTPs, taste parameters, and welfare changes by plugging in the empirical price derivative estimates from the hedonic regressions into the consumer's first order condition (1), and equations (4), (6), (11) and (12). In our dataset, there is an individual i who purchased each house j in our sample to maximize her utility. For this individual i, her MWTP is estimated by the derivative of the price function at D_j.

$$\widehat{\text{MWTP}}_i = \frac{\partial \hat{p}}{\partial D}(D_j) = \frac{\partial \hat{f}}{\partial D}(D_j).$$

where $f(D)$ is the function in equation (13) estimated by hedonic regression.

For the taste parameters, if the price derivative is negative at D_j, implying she considers density a disamenity, we estimate her individual preference parameter $\beta_{i,D}$ by

$$\hat{\beta}_{i,D} = -(\bar{D} - D_j)\frac{\partial \hat{p}}{\partial D}(D_j) = -(\bar{D} - D_j)\frac{\partial \hat{f}}{\partial D}(D_j).$$

If the price derivative is positive, we estimate $\beta_{i,D}$ by

$$\hat{\beta}_{i,D} = D_j\frac{\partial \hat{p}}{\partial D}(D_j) = D_j\frac{\partial \hat{f}}{\partial D}(D_j).$$

With these estimates we can then characterize the aggregate distribution of preferences for housing density in our sample.

Using the preference parameters, we estimate welfare changes for individuals under counterfactual levels of density exposure D'. If the price derivative is negative,

$$\begin{aligned}
\Delta\widehat{\text{Welfare}}_i &= \hat{\Delta p} + \hat{\beta}_{i,D}\log\left(\frac{\bar{D} - D'}{\bar{D} - D_j}\right) \\
&= \hat{\Delta p} - (\bar{D} - D_j)\frac{\partial \hat{p}}{\partial D}(D_j)\log\left(\frac{\bar{D} - D'}{\bar{D} - D_j}\right) \\
&= \hat{f}(D') - \hat{f}(D_j) - (\bar{D} - D_j)\frac{\partial \hat{f}}{\partial D}(D_j)\log\left(\frac{\bar{D} - D'}{\bar{D} - D_j}\right).
\end{aligned}$$

And, if the price derivative is positive,

$$\Delta \widehat{\text{Welfare}}_i = \hat{f}(D') - \hat{f}(D_j) + D_j \frac{\partial \hat{f}}{\partial D}(D_j) \log(\frac{D'}{D_j}).$$

We also retrieve estimates of the MWTPs, preference parameters, and welfare changes by owner versus renter occupied housing density. These can be solved in exactly the same way using the price derivatives estimated by Equation (14).

7. Results

A *Empirical hedonic price schedule*

In this section, we first return to the example of the Golf Manor-Amberley Village border to illustrate how we empirically estimate a hedonic price schedule for one market. We then present aggregate results on preferences and welfare changes using all the border areas.

Figure 7: Hedonic price schedule in Golf Manor-Amberley border area, Cincinnati, OH

Note: This figure plots the relationship between sale price and the density exposure in the Golf Manor-Amberley border area after controlling for observable characteristics. The red line is the fit line from a nonparametric local polynomial regression of the (nonparametrically) residualized outcome on the inverted density exposure. Covariate adjustment was carried out following Robinson (1988). Results are plotted after adjusting for: a house's lot size, living area, and age; the parcel's elevation and slope; the average 3rd grade reading score in its school district; and the distance to the nearest body of water, park, highway, public school, private school, and to the city center. Bandwidths for the local polynomial regressions in Robinson (1988)'s 2-step procedure were selected to minimize the integrated mean square error using Calonico et al. (2019), and all local polynomial regressions used a Gaussian kernel. The shaded area plots the bias-corrected confidence interval.

Figure 7 plots the relationship between sale price and the density exposure in the Golf Manor-Amberley Village border area after adjusting for numerous house-level covariates identified in the notes to the figure using only observations with this one border area.[27] The red line plots the local polynomial fit following Robinson (1988); the shaded region plots the bias corrected confidence interval following Calonico et al. (2019). The estimated derivatives at each point of the red line is our estimate of homeowners' MWTPs used to estimate individual taste parameters $\hat{\beta}_{i,D}$ following equations (4) and (6). Using the intuition from Rosen (1974), the negative slope of the pricing function indicates the typical homeowner in this border area dislikes higher housing density. After recovering $\hat{\beta}_{i,D}$, we can then estimate welfare changes for these homeowners under counterfactual levels of density exposure.

We repeat this process for each border area, estimating a separate hedonic price schedule (e.g., the red line in Figure 7) and its derivative at each homeowners' density exposure. These derivatives are then used to estimate individuals' taste parameters and welfare changes for different levels of density exposure according to the equations in Section 6.A.

B *Welfare counterfactuals*

Table 3 reports our estimates of the mean welfare change based on results from the semiparametric specification discussed above if every homeowner in our sample were exposed to a 1/2 unit per acre increase in density exposure. This is about a 0.3 standard deviation increase in density for our hedonic regression sample.[28] Welfare impacts are reported by row for the three different exposures discussed above: to total housing units, to owner-occupied housing units (holding constant renter-occupied exposure), and to renter-occupied housing units (holding constant owner-occupied exposure).

The table reports results from our baseline specification that includes hedonic controls for housing traits, local geography (slope and elevation), distance to local amenities and our reading test score proxy for school quality. Thus, this model controls for house and neighborhood quality differences that might confound our estimate of the preference for density.[29] Mean welfare changes are always negative and highly statistically significant. For example, $9,462 is required on average to compensate an existing owner for a 1/2 unit per acre increase in density, corresponding to nearly a 100 unit increase within a circle with a 500 meter radius centered

[27]Figure A7 also recreates binned scatter plots analogous to the gradients in Figure 5 using only observations within 500 meters of the Golf Manor-Amberley Village border. For this border area, the relationships between each outcome and the distance to the border follow patterns similar to those seen above in Figure 5, although they are estimated with less precision because they only use observations in this one border area.

[28]See column 2 of Table 4 above: 0.5/1.61=0.31, which is a bit less than a one-fifth increase in the sample mean density (0.5/2.75=0.18). Finally, results for the two other specifications estimated–the linear and log models are included in Appendix Table A2.

[29]These controls work as expected. For example, not controlling for distance to amenities and school quality moderately increases the estimated average welfare loss associated with a 1/2 unit per acre increase in All Housing Units from -$9,500 to -$11,000. That is as we expected: if house values in the less regulated, low density community border area are being bid up because they also happen to be closer to attractive amenities such as (say) nice parks, the price estimate will be biased up because of the uncontrolled for amenity.

on her home. The second and third rows of the table highlight how different the welfare loss is if that increase arises from more owner-occupied versus rental units. Not surprisingly, the mean welfare loss for an increase in owner-occupied units is close to that for All Housing Units, as owner-occupied units are the dominant type of housing in our suburban sample. However, the average required compensation is 5-6 times larger if the exposure increase is due to added renter-occupied units.[30] In our specific case, the heightened distaste for density could arise because a roughly 100-apartment unit increase is likely to involve a high rise complex that some people may view as fundamentally changing the nature of the area.[31]

Table 3: Average ΔWelfare estimates by housing tenure

Tenure	Baseline
All housing units	-\$9,462
	(\$1,457)
Owner occupied	-\$10,366
	(\$3,053)
Renter occupied	-\$56,392
	(\$12,552)
Housing traits	Y
Elevation/slope	Y
Distance to amenities	Y
School test score	Y

Note: This table presents estimates of the average change for individual homeowners from increasing their density exposure by 1/2 unit per acre by housing tenure. All regression estimates were retrieved separately by border area using 263,340 houses spread throughout the 217 border areas. Standard errors are calculated by cluster resampled bootstrap. The 217 border areas were resampled with replacement for 1,000 bootstrap replications.

C *Robustness Analysis*

In this subsection, we show that our key conclusions are robust to a variety of different assumptions. We begin by estimating different specifications of the hedonic price function. We also explore the implications of using smaller border depths and smaller circles for measuring density exposure. Next, we show results assuming each side of the border is a separate market and then discuss the extent to which hyperlocal supply effects could be affecting our results. Finally, we discuss estimates under alternative parametrization assumptions to homeowners' utility functions.

[30]Note that the result for All housing units is slightly outside the range for Owner occupied and Renter occupied. Slightly different samples explain that. That is, the sample for All housing units is not the same as for the sum of Owner-occupied and Renter occupied. It is larger because it includes vacant units. An increase in vacant units appears to be modestly less distasteful than an increase in an occupied unit.

[31]This general pattern of results is not dependent upon the specific functional form assumption we make on the underlying hedonic estimation. However, we do find larger absolute welfare losses on increased rental density exposure with non-linear specifications (i.e., the log and semiparametric ones) compared to a simple linear hedonic. See the discussion in Section 7.C directly below.

Appendix Table A2 reports results from two alternative hedonic specifications. In both the linear and log models, our key conclusions still hold. That is, we always find an economically meaningful average welfare loss and substantially larger losses associated with increased renter occupied density. A key difference between the fully linear hedonic model's results and those from our preferred semiparametric specification is in the smaller absolute welfare loss from a increase in density due to more renter-occupied units.[32] The linear model's estimated welfare loss strictly from more apartments is less than half that we reported from the semiparametric model (-$26,736 versus -$56,392). Note that this is not the case with the log model. Its estimated loss from a purely renter-occupied increase in density is a much larger -$107,475. While there is a large range in the rental unit density estimates across different functional forms, our preferred semiparametric specification, which makes no functional form assumptions on the object of interest and is consistent with the recommended best practice that suggests nonlinear flexible estimation of the hedonic price function (Bishop et al., 2020), still does not result in an estimate larger in magnitude than the standard log-log parametric specification.

The first column of Appendix Table A3 then shows that these key conclusions are robust to assuming smaller border depths (using the baseline semiparametric specification). Those results are based on border depths of only 250 meters (as opposed to 500 meters in our baseline). The results in column 1 are similar to, but not identical, to our baseline results in Table 3. The welfare loss associated with a general increase in density is smaller, at about -$8,000 while the estimates for purely owner-occupied and purely renter-occupied increases in density are each greater. Thus, estimation using shorter border depths certainly does not yield findings suggestive that there is no distaste for density, and it confirms a strong dislike of rental unit-generated density.

The second column of Appendix Table A3 provides welfare loss estimates based on density exposure being created using circles with a smaller 250 meter radius. Here, we find qualitatively similar results but the estimates are smaller in magnitude. Note that the standard errors also shrink, implying a similar confidence level in a negative average effect. This decline in magnitudes is at least partially due to the use of smaller circles to measure the density exposure changes the underlying counterfactual. In our baseline model with a 500 meter radius, the area of the circle was about 785,000 square meters or about 194 acres. Hence, a 1/2 unit increase per acre implies about 97 more housing units. With the smaller radius of 250 meters, the area of the circle is only 49 acres, which is nearly three-quarters smaller than the area of the circle with a 500 square meter radius. For 49 acres, a 1/2 unit per acre increase in housing amounts to about 24 units. That is barely one-quarter of the unit increase in our baseline estimation. Hence, it is not surprising to find smaller absolute impacts. Even so, the pattern of sharply higher welfare losses arising from renter-occupied density relative to owner-occupied density still is apparent.

For completeness, we also estimate average welfare losses using only the single family units

[32] The differences in estimated welfare costs from more total housing units and for only more owner occupied units are relatively minor.

from CoreLogic. An advantage of this measure is that it is constructed using more granular location data; however, it does not account for rental units. The average losses reported in the third column of Appendix Table A3 are similar in size to reports for density increases associated solely with more owner-occupied units in Table 3. That consistency is comforting, but we much prefer our main sample using census block data that allows us to break out impacts for owner-occupied versus renter-occupied increases in density, as that heterogeneity clearly is important.

Our next robustness test estimated separate hedonic specifications for each side of a border area. Conceptually, this is akin to treating each side as a distinct housing market. Even if one does not literally believe they are separate housing markets, this robustness analysis helps control for the potential impact of unobserved publicly provided services, such as as trash collection, police services, etc., by only using observations in a single municipality to estimate the underlying hedonic specification.[33] The fourth column of Appendix Table A3 reports results analogous those in Table 3. Note that they are not materially different from our baseline findings. Statistically speaking, one cannot reject the null that the coefficients in column 4 of Appendix Table A3 are the same as those in Table 3. The heterogeneity by tenure is readily apparent, too, with the distaste for renter-occupied density still at least five times larger than that for owner occupied density. The modestly lower estimates for All housing units and owner occupied units might be consistent with an impact of omitted locally produced services and amenities in our baseline specification, but even if so, the influence cannot be very large. Moreover, a reduction in magnitude is not observed in the moderately higher renter occupied density estimate. In sum, treating the border area as two separate markets does not change our key conclusions in any material way.

Another potential concern is that our results are being driven by hyperlocal supply effects rather than preferences over amenities, either because of the discontinuity in housing types crossing the border or from supply conditions decaying across space similar to how a preference over proximity to a disamenity may decay across space. The former is tested by the previous robustness check noted above, which estimated hedonic pricing functions using comparisons only between houses in the same municipality within 500 meters of the border.

Regarding the latter issue, the literature (Hartley, 2014, Anenberg and Kung, 2014, Pennington, 2021, Li, 2022, Asquith et al., 2023) notes that if supply effects decay across space in a small geographic area, then it is very difficult to disentangle disamenity effects from supply effects. However, we argue that the patterns in our results are hard to rationalize as coming primarily from a supply effect, but readily explained when interpreted as preference estimates. That conclusion is based on a rationale similar to that in Hartley (2014), who decomposes the local price effects of home foreclosures by placing assumptions on the segmentation between the single family and multifamily housing markets. In our case, the intuition is as follows: if the observed price effects on density are primarily from supply effects, an increase in the single

[33]This sample is slightly smaller at 178 versus 217 border areas because we further restrict to those areas where each side contains at least 100 usable observations for our hedonic models.

family housing stock (all else constant) should lower single family housing prices more than a similar quantity increase in the multifamily stock because additional single family stock is more directly substitutable with single family houses. In Equation (14) above, we consider an equal quantity increase in renter occupied housing, holding owner occupied quantity constant and *vice versa*. But, in our willingness-to-pay estimates for housing density in Table 3, we find 5 times larger welfare losses from increases in rental unit density than in owner occupied density.

Another way to look at this is to let our willingness-to-pay estimates by tenure be driven in part by preferences over disamenities and in part by a supply effect,

$$\Delta\text{Welfare}_{own} = Disamenity_{own} + Supply_{own} = -\$10{,}366$$

$$\Delta\text{Welfare}_{rent} = Disamenity_{rent} + Supply_{rent} = -\$56{,}392,$$

in which the monetary amounts are based on the average welfare change estimates in Table 3. Per the argument above, a weak assumption is that $Supply_{own} \leq Supply_{rent}$. That is, the reduction in single family house prices from adding owner occupied stock, which is more directly substitutable with single family housing, is the same or larger in magnitude than the reduction from adding the same quantity of renter occupied stock. The most extreme case is to assume supply effects explain all of the owner occupied welfare change. Even if so, supply effects can explain at most $-\$10{,}366$ / $-\ \$56{,}392 = 18\%$ of the rental welfare change. While we can not rule out supply effects totally, it is hard to imagine that supply effects drive 100% of the owner occupied estimate and just 18% of the renter occupied estimate.[34] In sum, while we can not completely rule out hyperlocal supply effects like those advanced in the papers referenced above, the patterns in our willingness-to-pay estimates are not consistent with a large supply effect, but are well explained as a reflection of preferences and sorting across neighborhoods.[35]

[34] Additional support that our estimates primarily capture a preference estimate emerges when we look at the patterns in average willingness-to-pay splitting the sample across different neighborhoods (shown in Table 4 below). For example, this same pattern by housing tenure is even more pronounced in neighborhoods that are predominantly white. For those areas, the average owner occupied welfare loss is even smaller in magnitude while the rental occupied welfare loss is even larger (-\$4,748 versus -\$70,896). Based on the same heuristic argument, this would imply supply effects can explain at most 7% of the rental welfare change estimate in these neighborhoods. On the other hand, it it would be very easy to explain this pattern as a reflection of differing preferences across demographic groups based on differential rates of homeownership. For example, data from the St. Louis Federal Reserve Bank's FRED files in the fourth quarter of 2020 shows the homeownership rate for Whites alone, was 74.5%, while that for Blacks and Hispanics was 44.5% and 49.1%, respectively.

[35] Finally, we note there is currently no consensus in the literature on how to disentangle these effects, nor is there formal modeling of the frictions required to generate these hyperlocal supply effects (see Furth (2024) for a review). Asquith et al. (2023) studies local effects of new apartment buildings and illustrates in a simple geometric exercise that if renters only look for housing in small geographic areas, this search friction can create differences in the probability a given housing unit is in competition with the new apartment building even across a small distance of space. However, in the model of Rosen (1974), it is the joint matchings of suppliers and consumers that generate the pricing function. In equilibrium, homeowners lie tangent to the pricing function based on their preferences for utility-bearing housing characteristics, and so the slope of the pricing function reflects those preferences according to Equation (1). In this way, supply conditions will determine what pricing function emerges in equilibrium and who chooses to live in that market, but the pricing function still identifies preferences for those homeowners who do choose to live there. This suggests that formal modeling of the frictions needed for local supply shocks to cause departures from the hedonic model's ability to capture individuals' preferences is needed.

Thus far, all estimates presented have relied on the parametric assumption of a log linear utility function following Bajari and Benkard (2005). In Appendix C, we explore how the average welfare loss estimates are affected by alternative restrictions on homeowners' utility. Appendix Table C1 first shows that assuming log linear utility yields a welfare loss estimate that is only modestly larger in magnitude than simply extrapolating linearly based on the observed MWTP. We then compare these estimates to an estimate retrieved following the maximum likelihood procedure detailed in Bishop and Timmins (2019). This approach yields a mean welfare loss estimate of -$33,984, or more than three times that from following Bajari and Benkard (2005). This pattern is very similar to that observed in Bishop and Timmins (2019)'s application, which suggests that traditional methods will tend to underestimate the willingness to pay to avoid increases in disamenities. However, because there is currently no consensus in the literature for which approach is the best practice for nonmarginal hedonic demand estimation (Bishop et al., 2020), we opt for the more conservative approach of Bajari and Benkard (2005), which yields estimates smaller in magnitude and stays closer to the price derivatives observed in the data, to generate our baseline results.

In sum, a host of robustness analyses yields results quite similar to our main findings which are: (a) average welfare losses are non-zero and economically meaningful; and (b) there is an especially strong distaste for density associated with rental units.

D *Preference heterogeneity*

While the typical individual in our sample dislikes housing density and loses welfare under the counterfactual higher density, there is significant heterogeneity in people's tastes according to our estimates. Figure 8 plots the empirical distributions of welfare changes based on the semiparametric specification estimated in Table 3.[36] The plot in the top panel for changes in total housing units shows that about two-thirds (65%) suffered a welfare loss with respect to a 1/2 unit per acre increase in total housing units, with one-third (35%) being indifferent or actually preferring more density. It is noteworthy that this distribution is skewed by those who have a particularly intense dislike for density. The median of -$5,164 is only 55% of the mean of -$9,462. The plot shows that this is driven by a longer and thicker left tail of those who dislike density. A few require nearly $150,000 in compensation for the hypothetical density increase.

[36] See Appendix Figures A8 and A9 for the analogous preference heterogeneity results for the linear and log specifications.

Figure 8: Distribution of ΔWelfare estimates

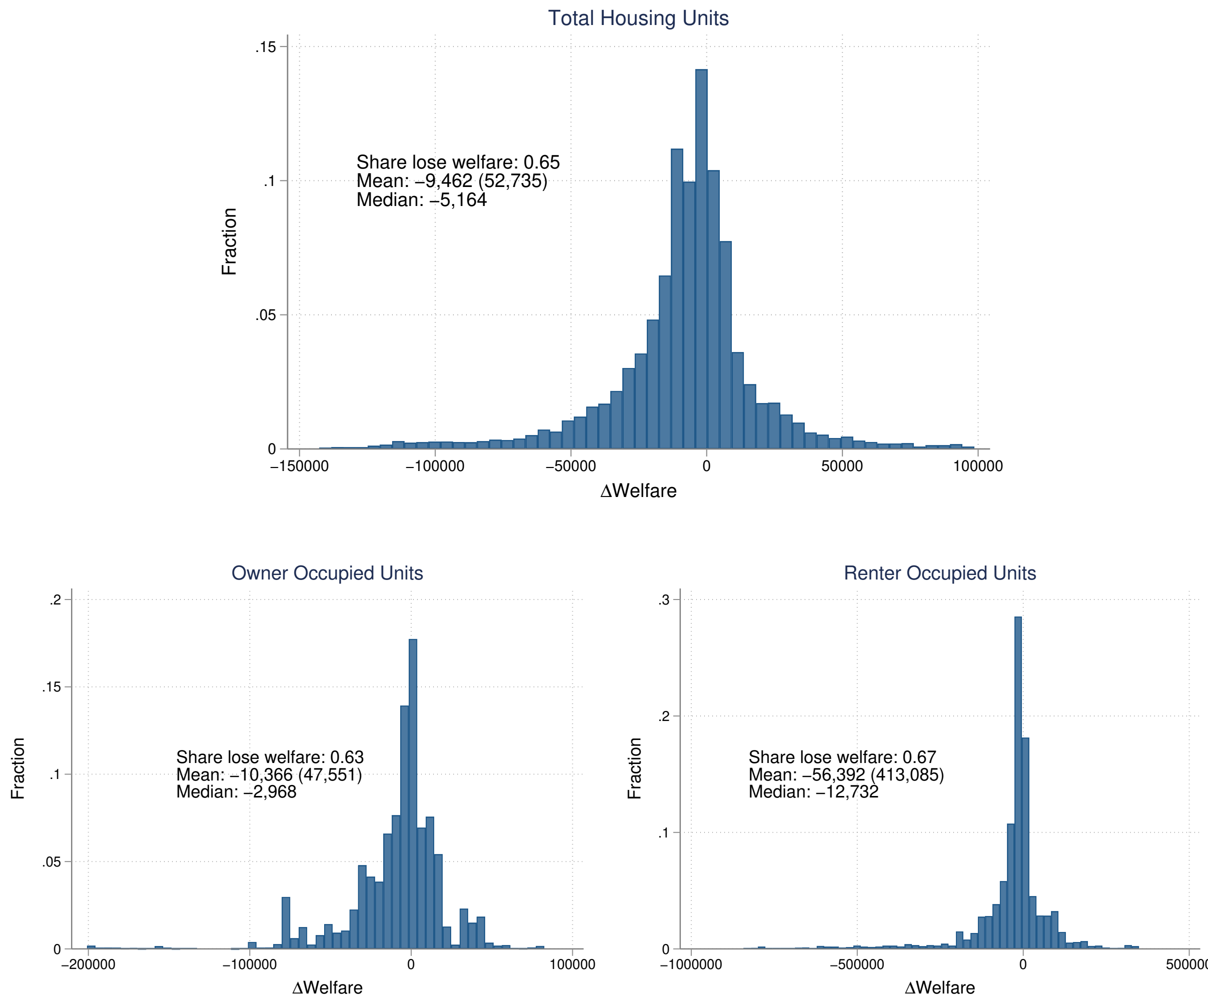

Note: This figure shows distributions of welfare change estimates under a counterfactual increase of 1/2 *unit per acre by housing type based on the baseline specification controlling for parcel characteristics, school test scores, and distance to local amenities (as in Table 3). The top panel plots the empirical distribution for the total housing unit density exposure from the 1st through 99th percentile. The bottom left panel plots the distribution of welfare changes from increasing owner occupied density exposure. The bottom right panel plots the distribution of welfare changes for increasing the renter occupied density exposure. Measures of central tendency for each distribution are reported in the accompanying text boxes.*

The bottom two plots in Figure 8 show the analogous distributions for increases in owner-occupied and renter-occupied density, respectively. They share important similarities with the top plot. The share of those experiencing welfare losses from the relevant density exposure increase are close to the 65% for total housing units. Hence, there are households who have a positive preference for renter-occupied density. However, both distributions are skewed by similarly longer and thicker tails of those who dislike density arising from any type of housing units. The coefficient of variation of the welfare losses associated with more renter-occupied density is far greater than that for the same amount of increased owner-occupied density. This reflects the fact that some people have an extremely strong dislike of rental density. The median distaste for density associated with more owner-occupied units drops to just under -\$3,000. The median distaste associated with more renter-occupied units is less than one-quarter of the mean value, but even so, it still is a healthy -\$12,700.

Overall, most homeowners have a moderate preference against density, some have a strong preference against density, and some moderately prefer density. In addition, the distaste for density associated with renter-occupied housing is much greater than that for owner-occupied housing on average.

Heterogeneity in the distaste for density clearly is important, and our empirical methodology allows us to compute estimates of average welfare change splitting the sample by differing neighborhood types and demographic breakdowns. Panel (a) of Table 4 reports summary statistics on density preference estimates by quartile of density exposure. Presuming that households will sort across neighborhoods based on their preferences, we would expect people living in low density areas to have stronger preferences against density. The results in the top panel are consistent with such sorting. While there are not stark differences between more or less dense neighborhoods in the share who have at least some dislike for density (i.e., it ranges from 62% to 70% in Panel (a)), there is a marked difference in the typical degree of dislike. The average welfare loss for our standard increase in the density of all housing units is nearly 5 times greater in the 25% least dense border areas compared to the 25% most dense border areas (i.e., -\$18,257 versus -\$3,696). That gap widens to 20x when comparing welfare losses associated with a density increase arising from more renter-occupied units. The results from this top panel of Table 4 shows that the -\$150,000 compensation for more rental-related density is coming from relatively low density areas.

Panel (b) of Table 4 shows similar summary statistics, but now splitting the data into subsamples based on the median household income of a house's census tract. Individuals in higher income neighborhoods tend to have much larger welfare losses under the counterfactual higher density exposure. Moving from the first quartile (low income) to the fourth quartile (high income), average welfare losses increase sharply in magnitude from -\$4,494 to -\$19,003 for the all housing units case. As with density, the range widens considerably for changes associated with increased rental density.[37]

Panel (c) splits the sample by quartile of the share of the population that is white in census blocks within 500 meters of the focal house. Here, we find less evidence of a clear link between preferences for density and race arising from the total or owner occupied housing units. The top and bottom quartiles by share white have almost identical mean welfare losses from our standard increase in density via 'all housing units'. However, there is a big difference for the loss

[37]In Appendix Table A4, we split the sample by the share of the 25 year old+ population possessing a bachelor's degree or higher. As was the case with high income areas, we find large shares of people with preferences against density in neighborhoods where most of the population has completed higher education. This is unsurprising if the high income neighborhoods are also the highly educated neighborhoods, which tends to be the case.

associated with higher density driven by more renter-occupied units, with predominantly white neighborhoods disliking such density much more.[38]

Table 4: Preference estimates by neighborhood type

		Q1	Q2	Q3	Q4
Panel (a)	**Avg. density exposure (total units per acre)**	0.99	2.07	3.07	4.91
	Share who dislike density	0.65	0.70	0.65	0.62
	Avg. ΔWelfare, All	-\$18,257	-\$8,873	-\$6,967	-\$3,696
	Avg. ΔWelfare, Owner	-\$18,850	-\$9,698	-\$7,947	-\$4,920
	Avg. ΔWelfare, Renter	-\$151,842	-\$44,347	-\$21,419	-\$7,409
		Q1	Q2	Q3	Q4
Panel (b)	**Avg. household income (\$ 2021)**	\$50,972	\$78,998	\$105,680	\$157,261
	Share who dislike density	0.65	0.65	0.64	0.68
	Avg. ΔWelfare, All	-\$4,494	-\$6,936	-\$7,423	-\$19,003
	Avg. ΔWelfare, Owner	-\$6,092	-\$8,377	-\$9,759	-\$17,246
	Avg. ΔWelfare, Renter	-\$10,038	-\$27,849	-\$44,834	-\$142,957
		Q1	Q2	Q3	Q4
Panel (c)	**Avg. share white ($<$ 500 meters)**	0.55	0.80	0.89	0.95
	Share who dislike density	0.65	0.68	0.68	0.61
	Avg. ΔWelfare, All	-\$6,762	-\$12,991	-\$11,266	-\$6,830
	Avg. ΔWelfare, Owner	-\$9,415	-\$17,454	-\$9,846	-\$4,748
	Avg. ΔWelfare, Renter	-\$24,478	-\$52,476	-\$77,724	-\$70,896

Note: This table shows summary statistics on individual preference estimates by quartile subgroup using the baseline specification of Table 3. Panel (a) shows the share of individuals who dislike density (MWTP < 0) and the average welfare changes from increasing density exposure 1/2 unit per acre splitting the sample by quartile of density exposure. Panel (b) shows the same summary statistics but splitting the sample into quartiles by the median household income of a house's census block group. Panel (c) splits the sample by quartile of the share of people identifying as white in nearby census blocks.

Given the strong correlation between the the preference for density with density itself and household income, Figure 9 delves more deeply into the intersection of that heterogeneity. It presents a breakdown by both density exposure and household income quartile. The y-axis measures the income quartile and the x-axis the density exposure quartile such that household income decreases as one moves down the y-axis and density increases moving right along the x-axis. Regardless of whether one is measuring welfare losses from more units in total, only owner-occupied units or only renter-occupied units, the largest welfare losses always are concentrated in high income, low density areas.

[38] In Appendix Table A4, we further split the sample by the share of households with children. As with race, we do not observe much of a pattern between family structure and preference for density for total or owner occupied but a pattern emerges looking at rental unit density.

Figure 9: ΔWelfare estimates by density and income quartile

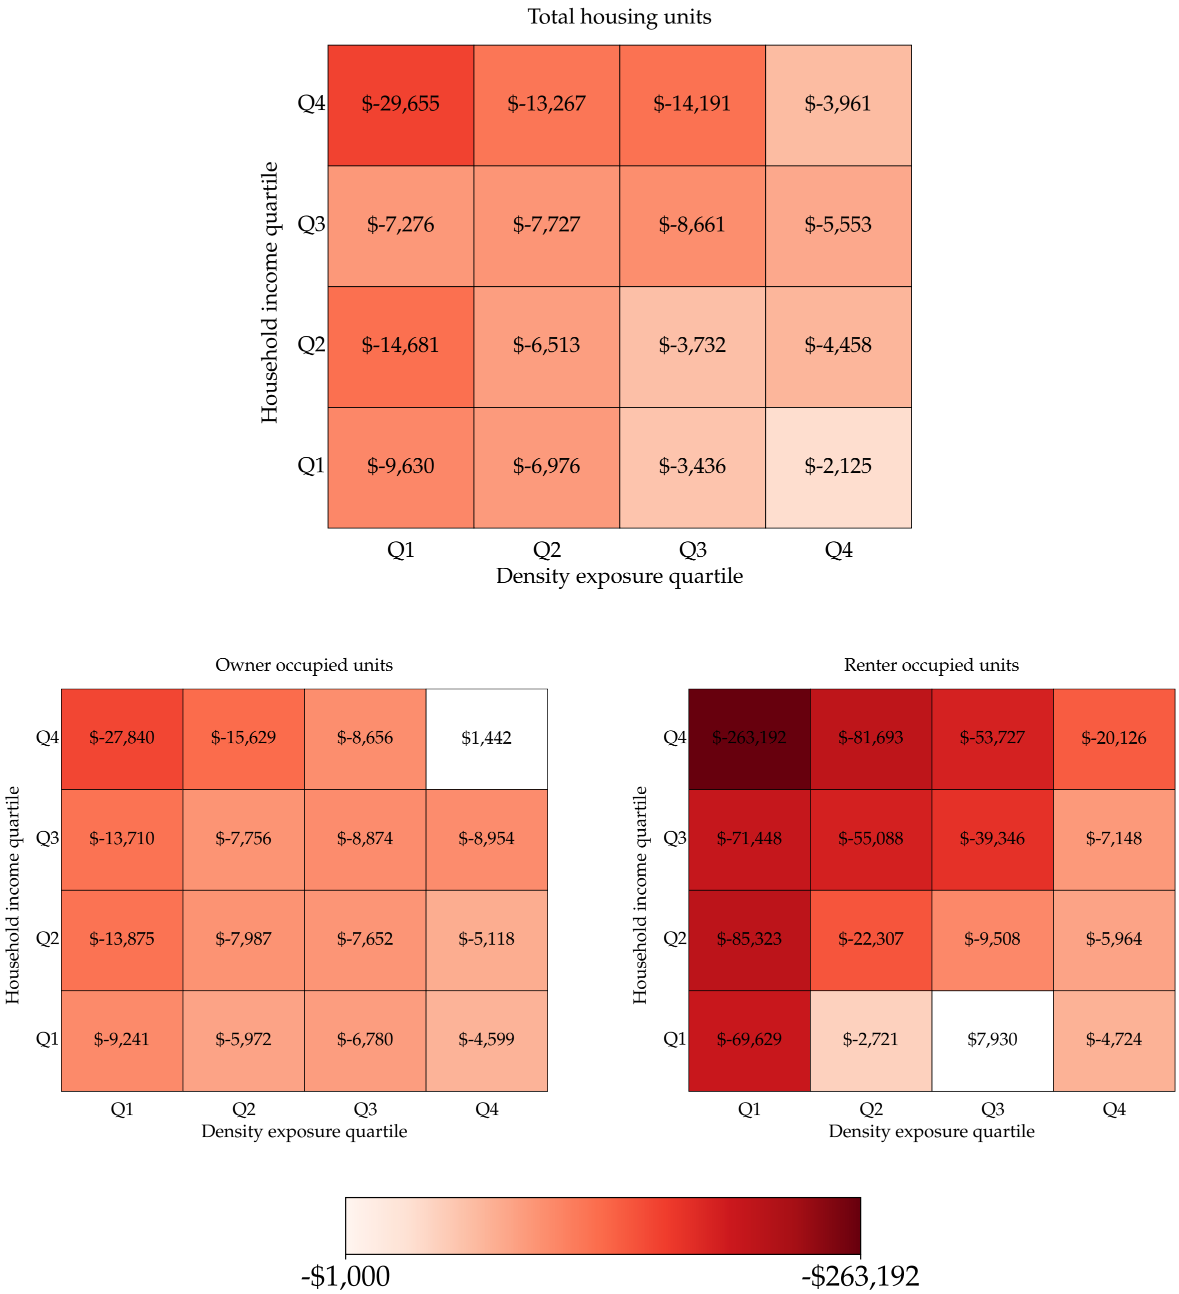

Note: This figure shows the average welfare change from a 1/2 unit per acre increase in density exposure by both census block group household income quartile and a house's density exposure quartile. The y-axis is the income quartile and the x-axis is the density exposure quartile. Household income decreases moving down the y-axis and density increases moving right along the x-axis, so the top left quadrant represents houses in high income and low density areas, while the bottom right quadrant represents houses in low income and high density areas. Each cell represents an average welfare change estimate for a given income quartile and density quartile. Cells are colored by the magnitude of the welfare change estimate with darker red cells indicating a larger welfare loss and whiter cells indicating a smaller or positive welfare change. The first grid shows the welfare changes from a 1/2 unit per acre increase in total housing unit density exposure. In the second row, the left and right grids tabulates the analogous averages for owner occupied and renter occupied density increases respectively. All grids are plotted on the same color scale shown at the bottom.

With respect to increased density due to more owner-occupied units, individuals in places that are in both the top income quartile (Q4 in the figure) and the most dense quartile (Q1 in the figure) have average welfare losses of nearly -$28,000, while those in places that are in the lowest income quartile and least dense quartile have average welfare losses below -$5,000. The aversion by the highest income people living in the least dense areas is particularly extreme with respect to density arising from more rental units. Their welfare loss from increased apartment density is over one-quarter million dollars, while those in the lowest income and highest density places have welfare losses of just under -$5,000.

E ***Moving to avoid density***

Thus far we have assumed all households are forced to live under the counterfactual 1/2 acre increases in housing unit density. In reality, local homeowners could move to avoid neighborhood change, but changing houses can impose significant costs due to a variety of factors well documented in the literature (Quigley, 2002). Hence, we augment our counterfactual to allow homeowners to move at a cost that is proportional to their house value. If homeowners choose to move, they get to consume their original bundle but must pay a moving cost. They no longer bear the utility change from the increase in density and instead only experience the price change and a moving cost.

Table 5 presents welfare change estimates analogous to Table 3's, but homeowners can move at a cost proportional to their house's value. Panel (a) repeats the results from the baseline specification while Panel (b) reports the average welfare change if homeowners can move at 15% of house value.[39] The biggest impact of mobility is associated with a density increase due to more renter-occupied units. With moving costs equal to 15% of house value, 17% of homeowners in our sample choose to move at a typical cost of roughly -$56,000. Allowing these moves reduces the average welfare loss by about 85%, from -$56,392 to -$8,808. Obviously, the outliers with the most extreme aversion to more apartment units avoid that cost by moving in this counterfactual, and the largest outliers involve changes in rental density exposure. In contrast, there is not much reduction in the average welfare loss due to increases in total or owner occupied housing density because only 3%-4% of homeowners sustain a utility loss from the density increase that exceeds the cost to move.

[39]In Appendix Table A5, we also report results for higher and lower moving costs equal to 10% and 20% of house value.

Table 5: Average ΔWelfare estimates with household mobility

Tenure	Avg. ΔWelfare	% move	Avg. moving cost
Panel (a): Baseline, no moving allowed			
All housing units	-$9,462	0%	.
Owner occupied	-$10,366	0%	.
Renter occupied	-$56,392	0%	.
Panel (b): Moving costs 15% of house value			
All housing units	-$8,151	3%	-$40,995
Owner occupied	-$9,147	4%	-$35,202
Renter occupied	-$8,808	17%	-$55,582

Note: This table presents estimates of the average welfare change for individual homeowners from increasing their density exposure by 1/2 unit per acre depending on if homeowners can avoid utility loss by moving. The first panel repeats the mean welfare change using the baseline specification with no moving allowed (corresponding to the estimates in Table 3). Panel (b) display the average welfare losses when homeowners can move by paying a price at 15% of their home value. The first column displays the average welfare change estimates. The second column shows the percent of homeowners that move under each counterfactual increase. The third column shows the average moving cost among the homeowners who move.

Finally, Figure 10 tabulates the average welfare change from a 1/2 unit per acre increase in rental unit density allowing moving at 15% of house value by both household income and density quartile (analogous to Figure 9). Although the same qualitative patterns hold as before, namely that high income, low density areas sustain the largest welfare losses, we again observe large reductions in the size of these welfare losses by allowing homeowners to move. Saliently, the average welfare loss in the highest income quartile and lowest density quartile falls by an order of magnitude from -$260,000 to -$29,000, with similarly large reductions in the size of the average welfare loss across quartile groups. This large reduction in the average welfare loss is due to nearly 1/3 of homeowners in that quartile group choosing to move.

Figure 10: Renter exposure ΔWelfare estimates by density and income quartile, moving costs 15% of house value

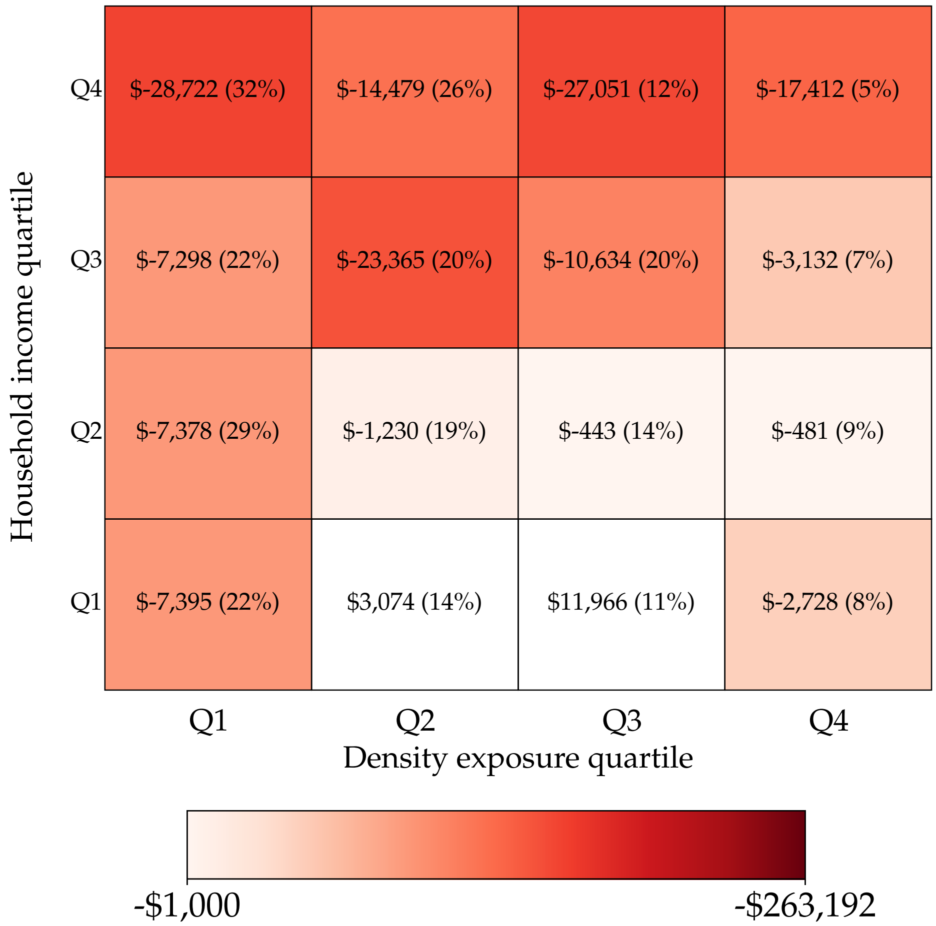

Note: This figure shows the average welfare change from a 1/2 unit per acre increase in rental unit exposure when households are allowed to move at 15% of their house value by both census block group household income quartile and a house's density exposure quartile. The y-axis is the income quartile and the x-axis is the density exposure quartile. Household income decreases moving down the y-axis and density increases moving right along the x-axis, so the top left quadrant represents houses in high income and low density areas, while the bottom right quadrant represents houses in low income and high density areas. Each cell represents an average welfare change estimate for a given income quartile and density quartile. The percentage of homeowners that move is shown in parantheses. Cells are colored by the magnitude of the welfare change estimate with darker red cells indicating a larger welfare loss and whiter cells indicating a smaller or positive welfare change. The color scale is shown at the bottom.

F *Discussion: Implications for Regulatory Change*

That increased density decreases local homeowners' welfare considerably is driven by a majority of homeowners at least moderately disliking density, with a small proportion having a very large willingness to pay to avoid density. Homeowners with these extreme preferences are disproportionately represented in affluent neighborhoods.

This has a number of potential implications for policy makers as they consider ways to persuade citizens to allow more building. We have shown in principle that household mobility can limit the more extreme preferences observed against rental density, but it well could be that existing owners will decide the minimum cost strategy is to fight prospective densification via the political process. Although one should not literally generalize our results to the broader community, as we rely on comparisons within a small area to justify the reliability of our density preference estimates, the distribution of preferences suggests that as a likely outcome. The combination of economically meaningful average losses combined with heterogeneity resulting in more negative outliers than positive ones seems ripe to create an organized opposition. Under the current political system which tends to locate control over zoning and land use regulation at the individual community level, these results raise the possibility that there may be no reason we should expect the typical suburb to densify of its own accord.

Furthermore, under purely local control the option for some type of Coasian bargain looks unlikely. Based on our results, a government transfer policy seeking to compensate incumbent homeowners to allow higher density would be both expensive and highly regressive. Recent policy actions have sought to increase housing supply by arresting control from the local level and moving it to the state level. Our results suggest this may be one of the few feasible policy recourses toward solving issues of housing supply in built-up areas.[40]

While we reiterate that our results only capture the changes to incumbent homeowners' welfare, we have shown their losses are not insignificant and there is heterogeneity across neighborhoods. From the perspective of place-based policy, neighborhoods with lower welfare losses are arguably better candidates for densification. However, this tradeoff is complicated by the regressivity of such a policy and the potential gains to lower income people being able to move into the new units in high opportunity areas. Deriving a rule for optimal allocation of densification that balances enhancing affordability, protecting amenities, and social marginal utility in a unified way constitutes yet another direction for future research.

G *What Mediates the Distaste for Density? A First Look at Neighbor Preferences*

Another obvious next step for research is to identify how the preference for density is mediated. There are a number of potential factors, including crime, traffic, noise, pollution, neighbor preferences, etc., and a deeper understanding of them is important economically and socially. In this subsection, we study one of these potential factors–neighbor preference. We explore this potential mechanism by augmenting our hedonic specification with demographic controls for race, income, education, and family structure.

[40]Reassigning control to the state level could create the opportunity for more Coasian exchange. As an example, when California's state housing mandate demanded that Atherton, CA, an affluent suburb of San Francisco, build town houses, many residents advocated paying a $9 million fine for violation of the mandate rather than build denser housing. `https://www.sfchronicle.com/bayarea/article/Atherton-to-consider-building-town-homes-to-avoid-17169242.php`

Table 6 presents results analogous to those in Table 3 that also include the neighborhood demographic variables in the hedonic regressions. Column (1) repeats the baseline set of results already presented in Table 3. Column (2) adds controls for neighborhood race. Specifically, we control for the share of people that are Black and Hispanic in nearby census blocks using exposure measures constructed analogously to the density exposure. Adding these controls moderately reduces the magnitude of the estimate for total housing unit density from -$9,462 to -$7,771 and rental unit density from -$56,392 to -$48,593 and moderately increases the owner occupied density estimate. Column (3) includes additional controls for median household income in a parcel's block group. Again, inclusion of more demographic variables moderately reduces the magnitude of total and rental occupied density estimates while increasing the magnitude of the owner occupied estimate. Finally in Column (4), we add controls for the share of people who are married, with child, and middle aged in nearby census blocks and the share of adult individuals with a college degree or more in a parcel's block group. Collectively, this portfolio of added controls is associated with large changes in our estimated average welfare effects. The mean change associated with our hypothetical 1/2 unit increase in total housing units falls by just over one-half from our baseline estimate of -$9,462 to -$4,648. The required compensation for a pure rental density increase falls by even more—by 63% to $20,998. On the other hand, the implied welfare change for a pure owner-occupied density increase is 23% greater than our baseline result in column 1.[41]

This admittedly simple mediation analysis still highlights that some—but not all—of the preference against housing density (particularly from rental units) is explained by neighbor preference. It further suggests that controls for family structure, age and education are similarly influential drivers of the distaste for density as race and income. That said, there certainly could be other externalities such as those noted above mediating the preference for low housing density. Evaluating the contribution of these other potential factors in a more formal mediation analysis (e.g., Gelbach (2016)), and how they interact, is an important direction for future research.

[41] The increase in the owner occupied welfare loss estimate may be due to differential ownership propensities by demographic groups. For example, data from the St. Louis Federal Reserve Bank's FRED files shows the overall homeownership rate in the fourth quarter of 2020 to be 65.6%. For Whites alone, it was 74.5%, while that for Blacks and Hispanics was 44.5% and 49.1%, respectively.

Table 6: Average ΔWelfare estimates by housing tenure, Mediation analysis

Tenure	(1)	(2)	(3)	(4)
All housing units	-$9,462	-$7,771	-$6,820	-$4,648
	($1,457)	($1,515)	($1,547)	($1,664)
Owner occupied	-$10,366	-$11,777	-$11,821	-$12,748
	($3,053)	($2,958)	($2,986)	($2,494)
Renter occupied	-$56,392	-$48,953	-$43,201	-$20,998
	($12,552)	($11,979)	($11,506)	($10,848)
Housing traits	Y	Y	Y	Y
Elevation/slope	Y	Y	Y	Y
Distance to amenities	Y	Y	Y	Y
School test score	Y	Y	Y	Y
Race		Y	Y	Y
Income			Y	Y
Other demographics				Y

Note: This table presents estimates of the average change for individual homeowners from increasing their density exposure by 1/2 *unit per acre by housing tenure. Each column shows average welfare change estimates under a different set of controls included in the hedonic regressions. All regression estimates were retrieved separately by border area using 263,340 houses spread throughout the 217 border areas. Standard errors are calculated by cluster resampled bootstrap. The 217 border areas were resampled with replacement for 1,000 bootstrap replications.*

8. Conclusions and Next Steps

Using data on density restrictions in the form of minimum lot sizes combined with micro data on house prices, we provide the first estimates of the value homeowners place on exposure to density. We do so by creating a novel measure of density exposure that varies across the border of contiguous communities with different minimum lot size restrictions. We find that house prices are about $40,000 more costly and lots are 3,000ft^2 larger on the side of the border with a more restrictive minimum lot size regulation. Price effects of exposure are estimated in narrowly defined border areas using the variation from the discontinuities in lot size. These price effects are connected to a structural hedonic model of housing choice suggested by Bajari and Benkard (2005).

This empirical strategy is used to establish a set of stylized facts about the preference for density. First, the distaste for density is widespread. Nearly two-thirds of existing owners in our border areas suffer a loss from densification; only one-third are indifferent or benefit from higher density. The average owner suffers around a $9,500 welfare loss if density exposure increases by 1/2 unit per acre, but that average masks interesting and important heterogeneity. The median loss is roughly half that amount, but there is a longer and thicker tail of those suffering much larger losses than there is for those experiencing welfare gains from higher density. Across individuals, some have a severe distaste for density and require much higher compensation for the same increase in exposure. However, more have only a modest distaste, while another group has a slight preference for density.

We also find a much larger distaste for density associated with more multifamiliy housing. The average distaste for the same 1/2 unit per acre increase arising from more renter-occupied units is 5-6 times greater than if it arises from more owner-occupied homes. The impacts of sorting are evident in the heterogeneity across types of neighborhoods, too. There is a far higher distaste for density in places that are both in the top quartile of the income distribution and bottom quartile of the density distribution. People who live in such areas have systematically higher aversion to densification.

An obvious implication of our findings is that the distaste for density provides a clear underpinning of the demand for regulation. There certainly can be other other motivations, but the combination of an economically meaningful average welfare loss plus a long left tail of those with an extreme aversion to density indicates this is an important one. Second, and related, the compensation per household needed to ameliorate those who dislike density is quite large compared to the costs of other housing policies such as the Low Income Housing Tax Credit program or the voucher program. One way to try to encourage higher density is for government to provide compensation to those communities allowing more building. Our estimates suggest any such intervention will have to be large to be effective. It is also likely to be highly regressive, as indicated by the largest welfare losses reported for increasing density in high income and low density neighborhoods. Third and finally, the mean loss and the nature of the distribution about that mean seem ripe for generating local coalitions to fight densification. Most people do not like it, and some are extremely harmed by it. This may go a long to helping us understand why it so difficult to change building conditions on the ground in American communities. That is an important local political economy issue that we hope our research can inform going forward.

References

Anenberg, E. and Kung, E. (2014). Estimates of the size and source of price declines due to nearby foreclosures. *American Economic Review*, 104(8):2527–2551.

Asquith, B., Mast, E., and Reed, D. (2023). Local effects of large new apartment building in low-income areas. *Review of Economics and Statistics*, 105(2):359–375.

Bajari, P. and Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of Political Economy*, 113(6):1239–1276.

Bajari, P. and Kahn, M. (2005). Estimating housing demand with an application to explaining racial segregation in cities. *Journal of Business & Economic Statistics*, 23(1):20–23.

Bartik, T. J. (1987). The estimation of demand parameters in hedonic price models. *Journal of Political Economy*, 95(1):81–88.

Baum-Snow, N. (2023). Constraints on city and neighborhood growth: The central role of housing supply. *Journal of Economic Perspectives*, 37(2):53–74.

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

Bishop, K. and Timmins, C. (2019). Estimating the marginal willingness to pay function without instrumental variables. *Journal of Urban Economics*, 109:66–83.

Bishop, K. C., Kuminoff, N. V., Banzhaf, H. S., Boyle, K. J., von Gravenitz, K., Pope, J. C., Smith, V. K., and Timmins, C. D. (2020). Best practices for using hedonic property value models to measure willingness to pay for environmental quality. *Review of Environmental Economics and Policy*, 14(2):260–281.

Blanco, H. and Sportiche, N. (2024). Local effects of bypassing zoning regulations. *Manuscript*.

Calonico, S., Cattaneo, M., and Farrell, M. (2019). nprobust: Nonparametric kernel-based estimation and robust bias-corrected inference. *Journal of Statistical Software*, 91(8):1–33.

Cattaneo, M., Crump, R., Farrell, M., and Feng, Y. (2024a). Binscatter regressions. *The Stata Journal*, Forthcoming.

Cattaneo, M., Crump, R., Farrell, M., and Feng, Y. (2024b). On binscatter. *American Economic Review*, 114(5):1488–1514.

Davidoff, T., Pavlov, A., and Somerville, T. (2022). Not in my neighbour's back yard? Laneway homes and neighbours' property values. *Journal of Urban Economics*, 128.

Diamond, R. and McQuade, T. (2019). Who wants affordable housing in their backyard? An equilibrium analysis of low-income property development. *Journal of Political Economy*, 127:1063–1117.

Duranton, G. and Puga, D. (2023). Urban growth and its aggregate implications. *Econometrica*, 91(6):2219–2259.

Epple, D. (1987). Hedonic prices and implicit markets: estimating demand and supply functions for differentiated products. *Journal of Political Economy*, 95(1):59–80.

Furth, S. (2024). Substitutability in the demand for housing over small distances. *Manuscript*.

Gelbach, J. (2016). When do covariates matter? And which ones, and how much? *Journal of Labor Economics*, 34(2):509–543.

Glaeser, E. and Gyourko, J. (2018). The economic implications of housing supply. *Journal of Economic Perspectives*, 32(1):3–30.

Greenstone, M. (2017). The continuing impact of Sherwin Rosen's "Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55.

Gyourko, J., Hartley, J., and Krimmel, J. (2021). The local residential land use regulatory environment across U.S. housing markets: Evidence from a new Wharton index. *Journal of Urban Economics*, 124.

Gyourko, J. and Krimmel, J. (2022). The impact of local residential land use restrictions on land values across and within single family housing markets. *Journal of Urban Economics*, 126.

Gyourko, J. and McCulloch, S. (2023). Minimum lot size restrictions: Impacts on urban form and house price at the border. *NBER Working Paper*.

Gyourko, J., Saiz, A., and Summers, A. (2008). A new measure of the local regulatory environment for housing markets: The Wharton residential land use regulatory index. *Urban Studies*, 45(3):693–729.

Hartley, D. (2014). The effect of foreclosures on nearby housing prices: Supply or dis-amenity? *Regional Science and Urban Economics*, 49:108–117.

Hsieh, C.-T. and Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39.

Kulka, A., Sood, A., and Chiumenti, N. (2024). Under the (neighbor)hood: Understanding interactions among zoning regulations. *Review of Economics and Statistics*, Forthcoming.

Li, X. (2022). Do new housing units in your backyard raise your rents? *Journal of Economic Geography*, 22(6):1309–1352.

McConnell, V. and Walls, M. (2005). The value of open space: Evidence from studies of nonmarket benefits. *Resources for the Future Report*.

Pennington, K. (2021). Does building new housing cause displacement? The supply and demand effects of construction in San Francisco. *Manuscript*.

Quigley, J. (2002). Transactions costs and housing markets. In O'Sullivan, T. and Gibb, K., editors, *Housing Economics and Public Policy*, chapter 4. Blackwell Science Ltd.

Robinson, P. (1988). Root-n-consistent semiparametric regression. *Econometrica*, 56(4):931–954.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55.

Song, J. (2024). The effects of residential zoning in U.S. housing markets. *Manuscript*.

Turner, M., Haughwout, A., and van der Klaauw, W. (2014). Land use regulation and welfare. *Econometrica*, 82(4):1341–1403.

Uribe, J. P. (2022). Equilibrium effects of housing subsidies: Evidence from a policy notch in colombia. *Manuscript*.