Density Zoning, Neighborhood Type, and Exclusion by Income and Race

Authors

Noah J. Durst, Elise Breshears, Esther Sullivan, Danna Gutierrez Lanza

Corresponding author: Noah J. Durst (durstnoa@msu.edu)

Publication Note

This is the accepted version of the following paper:Noah J. Durst, Elise Breshears, Esther Sullivan, Danna Gutierrez Lanza, Density zoning, neighborhood type, and exclusion by income and race, Journal of Housing Economics, Volume 68, 2025, 102065, ISSN 1051-1377, https://doi.org/10.1016/j.jhe.2025.102065.

Abstract

In recent years, states and municipalities have taken steps to reform land use and zoning regulations. While prior research documents that density zoning contributes to residential segregation on the basis of income and race, the mechanisms remain largely unexplored. In this paper, we examine the relationship between density zoning, neighborhood type, and residential segregation. To do so, we use a national dataset of building footprints and machine learning to develop a neighborhood typology based on building characteristics. We then use land cover data to examine changes in building development in these neighborhoods between 2001 and 2019. Finally, we pair these data with demographics at the municipality level to examine changes in income and race between 2000 and 2020. In cross-sectional analyses, we find that density zoning is strongly associated with building characteristics and the presence of different neighborhood types. Although we find that density zoning is also associated with income and race, the effects are attenuated when accounting for neighborhood types. Our results provide new evidence into the "chain of exclusion" between density zoning and residential segregation, as we find that density zoning is primarily associated with reductions in the supply of single-family housing along the urban fringe. Lastly, we find that maximum density restrictions and changes in maximum density cannot explain the changes in demographics that we observe during this time period. We do, however, find some evidence of a relationship between changes in building development and changes in demographics. These results demonstrate the potential effects of upzoning policies.

Keywords

Building development, density zoning, neighborhood type, residential segregation, upzoning

1 Introduction

Over the past decade, states and municipalities have focused on reforms to land use and zoning regulations as potential policy solutions for the housing crisis in the United States. Following recommendations from the American Planning Association, municipalities have begun to update their land use and zoning regulations to allow for the development of a more diverse housing stock (Gabbe 2019). In addition to housing development, land use and zoning regulations address a range of issues, including affordable housing, climate mitigation, economic development, and residential segregation.

The relationship between land use regulations and residential segregation is well established in the academic literature. Strict density zoning is closely linked to high levels of income and racial segregation (Lens and Monkkonen 2016; Rothwell and Massey 2009), although the mechanisms by which it does so remain largely unexplored. One of the leading explanations in the literature is the "chain of exclusion" proposed by Pendall (2000). He purports that density zoning limits the supply of multifamily housing and rental options, which exacerbates income and racial segregation.

In this paper, we examine the potential impact of density zoning on income and racial segregation by analyzing variation in building characteristics and neighborhood types across municipalities in the 50 largest metropolitan areas in the country. We measure density zoning using estimates of the maximum number of dwelling units that are permitted per acre within a municipality, which are derived from the National Longitudinal Land Use Survey (Lo et al. 2019). To study neighborhood type, we use two-dimensional building footprints to measure the shape, size, and placement of buildings and how they vary within and between neighborhoods. We use machine learning (i.e., principal component analysis and unsupervised classification) to identify four distinct neighborhood types based on the physical and spatial characteristics of buildings: (1) Large-Lot Residential; (2) Small-Lot Residential; (3) Low-Intensity; and (4) High-Intensity neighborhoods. We then use land cover data to examine temporal changes in building development. To do this, we measure the share of land that was converted from undeveloped land to low-intensity, medium-intensity, or high-intensity development between 2001 and 2019, which we denote as new development. We then use this measure of new development to estimate the share of buildings that were developed between 2001 and 2019. Finally, we pair this with data from the U.S. Census Bureau to measure changes in demographics between 2000 and 2020. Specifically, we ask the following questions: (1) How does density zoning influence building characteristics and neighborhood types and at what density thresholds is zoning most impactful? (2) How does density zoning influence residential segregation and the income and racial composition of municipalities? (3) Do changes in building development influence changes in these demographics?

Our study contributes to the literature on density zoning and residential segregation in a number of ways. First, we use machine learning to develop a novel dataset of neighborhood types. We then examine the relationship between maximum density and the presence of each neighborhood type, illustrating the thresholds at which maximum density is most closely associated with changes in the housing stock. Second, we link this analysis of density zoning and neighborhood type with variation in the demographic composition of municipalities. Our results tell a different story than Pendall's (2000) chain of exclusion, as they indicate that the chain of exclusion functions primarily by limiting the supply of single-family housing units, particularly in suburban and exurban areas. Third, our findings provide insight into the potential effects of changes in zoning. While we use a coarse proxy for upzoning — changes in maximum density over time — we find some evidence that upzoning is associated with changes in neighborhood type but little to no evidence that it is associated with changes in demographics. Fourth, we contribute to the literature on the measures of density zoning by conducting a comparison of maximum density derived from the National Longitudinal Land Use Survey (Lo et al. 2019) and minimum lot size derived from the Wharton Residential Land Use Regulatory Index (Gyourko et al. 2008). Our estimates illustrate that maximum density is a more robust measure because it applies to both single-family and multifamily zones and captures greater variation in neighborhood type between municipalities. Lastly, we illustrate a novel method for observing the effect of density zoning via big data. To do so, we build on recent efforts that use building footprints to develop and study neighborhood types (Durst et al. 2024; Jochem and Tatem 2021). We find that neighborhood types are stronger predictors of demographic change than estimates of zoning change derived from municipal surveys. Our study accompanies recent efforts to refine the measurement of zoning across a broad geography, including the use of large language models to interpret municipal ordinances (Bartik et al. 2024) and crowd-sourcing to collect and standardize municipal zoning ordinances across the country via the National Zoning Atlas (Xu et al. 2023).

The remainder of the paper proceeds as follows. In section 2, we briefly review the existing literature on density zoning and residential segregation. We also discuss recent advancements in big data and machine learning that offer opportunities to overcome limitations in prior research. In section 3, we provide an overview of the data and conduct an exploratory analysis, highlighting stylized facts about density zoning, neighborhood type, and building development. In sections 4 and 5, we present the empirical approach and discuss the results. And, in section 6, we conclude and note potential policy implications of upzoning.

2 Literature Review

A substantial body of research examines the impact of land use and zoning regulations on housing markets. Prior work demonstrates that restrictive land use regulations limit housing supply in metropolitan areas (Been et al. 2019; Glaeser and Ward 2006; Kahn et al. 2010; Kok et al. 2014; Quigley and Raphael 2004), and in doing so, contribute to higher housing and land prices (Freemark 2023; Kok et al. 2014; Phillips et al. 2021; Quigley and Rosenthal 2005; Saiz 2010; Stacy et al. 2023; Wassmer and Williams 2021). As an example, Stacy et al. (2023) find that land use reforms that tighten restrictions and decrease densities are associated with increases in median rents. Meanwhile, an academic review on upzoning finds that increases in densities are associated with increases in housing supply, although there is mixed evidence on the impacts on housing costs, as the effects vary by neighborhood context and housing type (Freemark 2023).

Research also points to the link between land use and zoning regulations and residential segregation. Since the adoption of zoning in the early 20th century, density zoning has contributed to income and racial segregation (Goetz 2021; Lens and Monkkonen 2016; Manville et al. 2020; Pendall 2000; Rothstein 2017; Rothwell and Massey 2009; Trounstine 2020; Whittemore 2021). Lens and Monkkonen (2016), for example, find that strict density zoning is associated with the segregation of high-income and middle-income residents, while Rothwell and Massey (2009) find that it is associated with higher levels of Black segregation. The mechanisms by which density zoning impacts racial segregation are not fully understood. In Pendall's (2000) chain of exclusion, he argues that strict land use regulations create racial exclusion by structuring the distribution of the type of housing options, housing supply, and housing tenures. Specifically, he argues that density zoning reduces the share of multifamily housing, limiting the amount of rental units. This, in turn, reduces housing affordability and increases residential segregation.

One major limitation in the existing literature is that it is largely cross-sectional. Although the link between density zoning and residential segregation is well established, research documenting the impact of zoning change on neighborhood demographics is very limited (Freemark 2023). Understanding these potential impacts is particularly important given recent efforts by states and municipalities to eliminate single-family zoning. Research has found, for example, that the elimination of single-family zoning in Minneapolis, Minnesota, led to a 3–5% increase in housing prices (Kuhlmann 2021), but it is too early to identify impacts on changes in demographics (Furth and Webster 2023). In places like Oregon, states have effectively eliminated single-family zoning in large cities. Similar laws have been proposed in other states. These state laws have received public backlash because they often preempt local laws, calling into question their legal authority.[1] While the impacts of these recent state laws will not be apparent for several years (Garcia et al. 2020), there is interest among the academic community to explore their potential impacts through indirect methods.

A second important gap in the literature on density zoning and residential segregation is that, although both phenomena are spatial in nature, research on the relationship between the two is largely a-spatial. For example, zoning is explicitly spatial in at least two respects: 1) zoning enforces differential development restrictions across space (e.g., zones) and 2) zoning regulates the spatial characteristics of buildings, including building density, size, and setbacks. To the extent that there is an empirical link between density zoning and residential segregation, and to the extent that it functions through impacts on housing supply, these impacts should be observable in the characteristics of buildings and their variation between neighborhoods. Because each municipality implements zoning independently and each zone imposes distinct restrictions on development, prior research on the relationship between density zoning and residential segregation typically sacrifices either spatial precision or geographic coverage. Large, national analyses have relied on municipal-level estimates of minimum lot sizes (Lens and Monkkonen 2016) or maximum units per acre (Rothwell and Massey 2009) to examine the relationship between density zoning and income or racial segregation at the metropolitan level. These studies, however, often pay limited attention to variation in density zoning between municipalities. Similarly, analyses of the effects of zone-level variation in density restrictions are limited to single regions (Resseger 2022) due to the lack of national data. Although efforts are underway to develop a National Zoning Atlas (Xu et al. 2023), zone-level density restrictions are not yet available nationally.

Recent advancements in big data and machine learning have allowed for new approaches to measuring zoning and its potential impacts on development. For example, Lawrimore et al. (2024) develop machine learning models that predict zoning restrictions

[1] For example, some charter cities in California successfully petitioned that Senate Bill 9 (California Housing Opportunity and More Efficiency Act) infringed on their local zoning laws.

at the zone level as a function of population, community, built environment, and natural environment characteristics. Meanwhile, Mleczko and Desmond (2023) and Bartik et al. (2024) use natural language processing and large language models to create estimates of zoning restrictiveness directly from the text of municipal zoning ordinances. Additionally, recent advancements in software (Jochem and Tatem 2021; Fleischmann 2019) and the widespread availability of building footprints (Microsoft 2022) make it feasible to develop high-resolution estimates of the spatial characteristics of development and their potential impact on residential segregation. For example, Durst et al. (2023) use building footprints in five major metropolitan areas (Atlanta, Boston, Chicago, Houston, and Los Angeles) to demonstrate that building characteristics vary across the urban, suburban, and rural landscape and are associated with variation in the demographic and socioeconomic characteristics of neighborhoods. In this study, we pair land use and zoning surveys with building footprints and land cover data to examine the potential impacts of density zoning and changes in density zoning on building development and demographic change.

3 Data, Measurement, and Stylized Facts

3.1 Data

In this paper, we combine four types of data. First, we use building footprints to measure building characteristics and to develop a typology of neighborhoods. Second, we use land cover data to measure changes in building development. Third, we use responses from a national survey to measure land use and zoning regulations. Fourth, we use data from the U.S. Census Bureau to measure the income and racial composition of municipalities. We describe these data in more detail in the following sections.

3.1.1 Building Footprints

The main dataset that we use in this paper is building footprints, which we obtain from Microsoft (2022). Microsoft uses satellite imagery to identify building footprints and then makes these building footprints publicly available. Approximately half of the building footprints were created using 2019 and 2020 imagery; the remainder rely on imagery beforehand (for additional details, see Appendix A). This dataset includes about 130 million two-dimensional building footprints across the United States. From this dataset, we calculate a series of building characteristics to examine variation in the shape, size, and placement of buildings. We focus on three main building characteristics: area of buildings (square meters); density of buildings (buildings per square kilometer); and distance between buildings (meters). We then measure these building characteristics at the census block level. For area and distance, we calculate both the median and interquartile range, which represent the central tendency and variability of these building characteristics, respectively. Meanwhile, for density, we measure the number of buildings per square kilometer. As discussed in more detail below, we use these building characteristics to identify four distinct neighborhood types. In doing so, we consider an expanded set of building characteristics, including the angle of orientation, compactness index, length-equivalent-width ratio, perimeter, shape index, and total number of building footprints at the census block level.

3.1.2 Land Cover

We then use land cover data from the National Land Cover Database (NLCD) to measure changes in building development over time. The NLCD uses satellite imagery to identify different types of land cover at a 30×30 meter resolution. While this dataset is updated every few years, we use the 2001 and 2019 datasets to identify land cover changes over approximately 20 years. The NLCD contains many types of land cover, including developed land, forest, and water, among others. Developed land is further separated into four classes:

Figure 1: National Land Cover Data

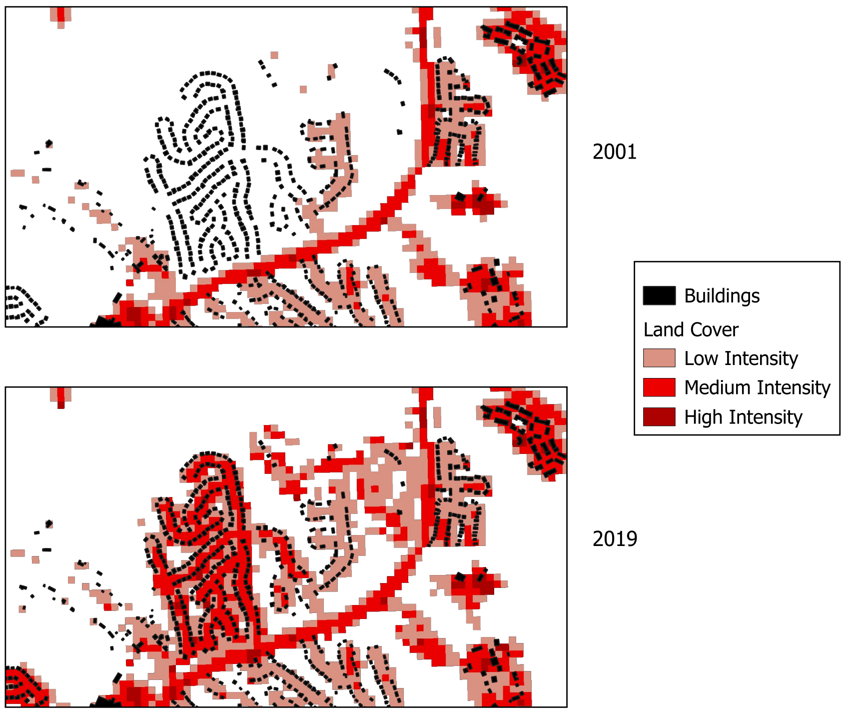

Notes: This figure illustrates the location of building footprints documented by Microsoft as well as changes in low-intensity, medium-intensity, and high-intensity land cover as measured via the NLCD. We use the data to estimate the proportion of buildings developed between 2001 and 2019.

open space; low-intensity; medium-intensity; and high-intensity. As seen in Figure 1, the low-intensity, medium-intensity, and high-intensity classes align closely with the location of building footprints, especially in the 2019 NLCD. We aggregate these three classes into a single class, denoted as LMH. We discuss how we use LMH to measure building development in more detail later.

3.1.3 Land Use and Zoning Regulations

We obtain data on land use and zoning regulations from the National Longitudinal Land Use Survey (NLLUS). This survey, which was conducted in 1994, 2003, and 2019, asks public officials a variety of questions about a municipality's land use and zoning regulations. The survey was distributed to municipalities in the 25 largest combined metropolitan statistical areas in 1994, and to municipalities in the 50 largest combined metropolitan statistical areas in 2003 and 2019. We focus on the 2003 survey because it is the most recent survey that was conducted prior to Microsoft collecting the building footprints.[2] The survey separates municipalities into three types: incorporated places (i.e., cities); unincorporated townships/towns; and counties. Our sample consists of 1,805 municipalities with 1,113 (62%) cities, 498 (28%) townships, and 194 (11%) counties. While the survey asks a variety of questions about land use and zoning regulations, we select a subset of questions that we suspect have the most influence on municipality-level variation in neighborhood type. These regulations include: adequate public facilities ordinance; affordable housing; growth limits; impact fees; maximum density restrictions; urban service area; voting referendum; and zoning prohibitions on mobile homes (for additional details, see Appendix B). Each of these regulations are binary (yes/no) variables with the exception of maximum density, which represents the maximum number of dwelling units permitted per acre. The survey measures maximum density with the following categories: < 4 units; 4–7 units; 8–15 units; 16–30 units; and > 30 units. For simplicity, we convert this categorical variable into a continuous variable by calculating the median number of units within each category. We cap the number of units at 31. In an alternative analysis, we create binary variables for these categories, comparing them to the base category of < 4 units.

In this paper, we focus on maximum density because, as we note earlier, the existing literature identifies density zoning as one of the primary land use and zoning regulations that drives residential segregation. In addition, density zoning shapes the spatial manifestation of development as observed via building characteristics (e.g., area, density, and distance). One thing to note about our measure of maximum density is that it need not apply to all residential zones in a municipality. Rather, it represents the maximum number of dwelling units permitted per acre in the highest residential density category in the municipality.

[2]As a robustness check, we perform similar analyses using the 2019 survey and find very similar results.

Thus, our measure of maximum density does not capture variation in density zoning within municipalities but rather variation between municipalities.

Another common application of density zoning is minimum lot size requirements. These requirements are analogous to maximum density restrictions in the sense that they also restrict residential development, just through a different mechanism. For example, a city may require all residential lots to be greater than 5,000 square feet. This, in turn, impacts the number of dwelling units that are permitted per acre. While the NLLUS does not collect information on minimum lot size requirements, the Wharton Residential Land Use Regulatory Index (WRLURI) survey collects this information. Similar to the NLLUS, the WRLURI asks public officials questions about different land use and zoning regulations. In Appendix C, we compare these two surveys and conduct a robustness check between the maximum density restrictions from the NLLUS and minimum lot size requirements from the WRLURI. While these two measures of density zoning have similar associations with variation in neighborhood type between municipalities, the maximum density restrictions from the NLLUS capture more of this variation. Because of this, we use the maximum density restrictions from the NLLUS in our main analysis.

3.1.4 Income and Racial Composition

To measure the income and racial composition of municipalities, we use a variety of datasets from the U.S. Census Bureau, compiled and published by Manson et al. (2022) via the National Historical Geographic Information System. While we focus our analysis on income and race, we also measure other demographic and socioeconomic characteristics. First, we use data from the 2000 decennial census and the 2018–2022 American Community Survey (ACS) to examine changes in income and race over time. We use the ACS for our later time period because income data are not available for the 2020 decennial census. We use 2018–2022 data because the midpoint is 2020, so our analyses of changes in demographics represent approximately two decades of change, between 2000 and 2020. Second, we use data from the 2016-2020 American Community Survey to evaluate contemporary differences in demographic and socioeconomic characteristics between neighborhood types. Specifically, we examine the building year of construction, housing value, homeownership rate, share of single-family and multifamily units, median household income, and share of residents who are non-Hispanic White. We use the 2016–2020 survey because it aligns mostly closely with the date of collection of the building footprints by Microsoft. Third, we use employment data from the Longitudinal Employer-Household Dynamics (LEHD) division. Specifically, we use the 2019 LEHD Origin-Destination Employment Statistics (LODES) dataset that measures the number of workers (i.e., jobs) in a census block. This employment data helps us determine whether our neighborhood types are residential or mixed-use areas.

3.2 Measurement

In this section, we describe how we use the building footprints and land cover data to measure neighborhood type and to examine patterns of new (post-2001) and old (pre-2001) development. We then discuss a series of stylized facts.

3.2.1 Neighborhood Types

To develop estimates of neighborhood type, we use principal component analysis and unsupervised classification (for additional details, see Appendix A). Unsupervised classification is a common machine learning technique that identifies patterns within a dataset and separates observations into different classes. In our case, an observation is a census block. Although we tested alternative classification algorithms and examined results from models with varying numbers of classes, we use the results from the K-means clustering model with four classes. In other words, we separate census blocks into four neighborhood types. Figure 2 provides an illustrative example of these neighborhood types. Meanwhile, Table 1 presents a variety of summary statistics by neighborhood type. Based on these illustrations and summary statistics, we denote the four neighborhood types as

Figure 2: Examples of the Neighborhood Types

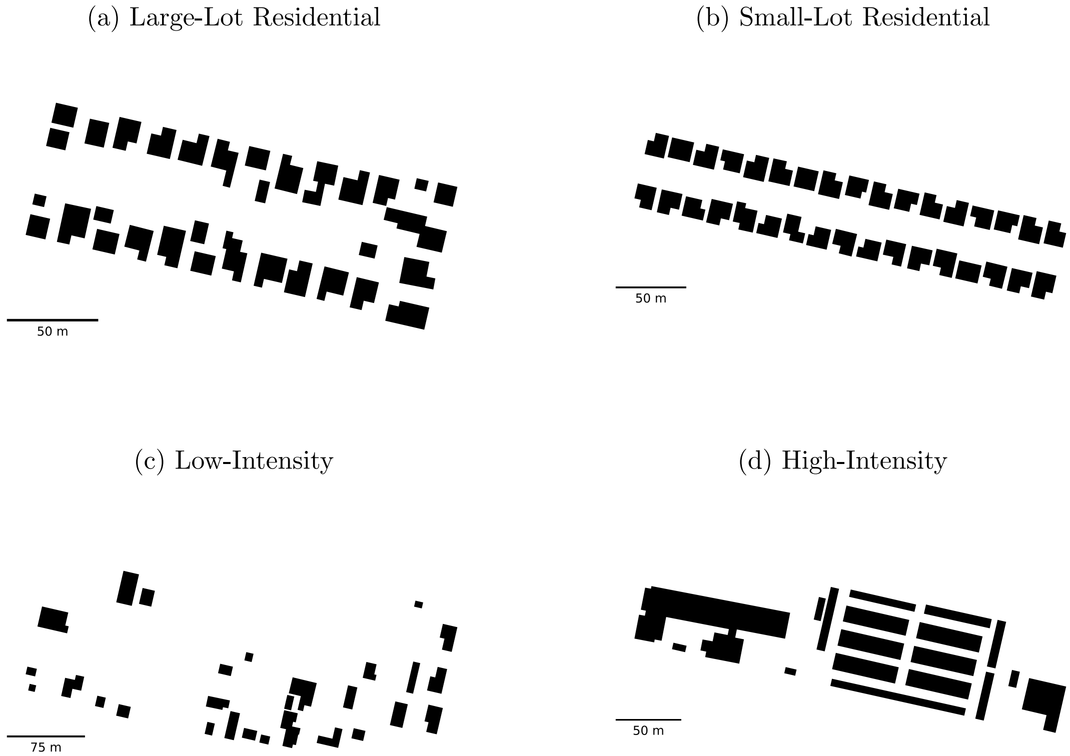

Notes: The figure presents examples of the neighborhood types that were identified from the unsupervised classification algorithm. Panels (a), (b), (c), and (d) represent Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively.

the following: (1) Large-Lot Residential; (2) Small-Lot Residential; (3) Low-Intensity; and (4) High-Intensity. While these four neighborhood types are an oversimplification of the complex and multidimensional nature of neighborhoods, we argue that they represent archetypal neighborhoods that provide insight into the varied patterns of development.

Table 1: Summary Statistics by Neighborhood Type

Characteristic	Neighborhood Type			
	Large-Lot Residential	Small-Lot Residential	Low-Intensity	High-Intensity
Panel A: Building Characteristics				
Area (Median)	232	169	202	852
Area (IQR)	73	51	159	1,011
Density	732	977	257	242
Distance (Median)	25	19	33	42
Distance (IQR)	5	4	17	15
Observations	409,916	621,143	343,194	174,116
(%)	(0.27)	(0.40)	(0.22)	(0.11)
Panel B: Block Socioeconomic Characteristics				
Residents	55.5	74.3	122.0	104.2
Workers	6.8	7.1	40.7	141.3
Homeownership Rate (%)	0.75	0.77	0.76	0.46
White (%)	0.64	0.59	0.75	0.54
Panel C: Block Group Socioeconomic Characteristics				
Year of Construction (Median)	1972	1969	1982	1976
Housing Value (Median)	441,731	309,360	340,471	471,487
Single-Family Detached (%)	0.59	0.85	0.63	0.11
Single-Family Attached (%)	0.07	0.04	0.05	0.13
Multifamily: 2-4 units (%)	0.10	0.06	0.06	0.10
Multifamily: 5-19 units (%)	0.06	0.02	0.09	0.16
Multifamily: 20+ units (%)	0.12	0.02	0.11	0.49
Household Income (Median)	87,805	89,160	82,939	81,890
Observations	553	4,501	989	2,294
(%)	(0.07)	(0.54)	(0.12)	(0.28)

Notes: The table presents summary statistics by neighborhood type. Panel (a) documents the building characteristics at the census block level. Panels (b) and (c) document the socioeconomic characteristics at the block and block group levels, respectively. Area is measured by square meters, density is measured by buildings per square kilometer, and distance is measured by meters to nearest building.

The first two neighborhood types are representative of residential neighborhoods. These neighborhoods have buildings that are relatively small (between 169 and 232 square meters) with limited variability in size. They are also more dense (732 and 977 buildings per square kilometer) and have closer spaced buildings. The demographic and socioeconomic characteristics provide further justification. These neighborhoods have a relatively high share of single-family housing units (66%–89%) and homeownership rate (75%–77%). Even more telling is the number of workers and the ratio of residents to workers. These neighborhoods have about 7 workers, on average, while the other neighborhood types have 41–141 workers. Similarly, these neighborhoods have a ratio of about 8–10 residents to workers, while the other neighborhood types have a ratio of about 0.75–3 residents to workers. Because of these facts, we comfortably argue that the first two neighborhood types are residential neighborhoods. Since the buildings are larger and spaced further apart from each other in the first neighborhood type, we refer to them as Large-Lot Residential neighborhoods. Alternatively, we refer to the second neighborhood type as Small-Lot Residential neighborhoods.

Next, we characterize the third neighborhood type. While buildings in these neighborhoods are comparable in size to Large-Lot Residential and Small-Lot Residential neighborhoods, there is greater variability in size. These neighborhoods are also less dense and have greater variability in distance. These neighborhoods do, however, share some characteristics with the residential neighborhoods discussed previously. For example, they too have a relatively high share of single-family housing units (68%) and homeownership rate (76%). Together, these facts suggest that these neighborhoods have a mix of both residential and non-residential low-intensity development. Because of this, we refer to them as Low-Intensity neighborhoods.

Finally, we characterize the fourth neighborhood type. These neighborhoods have the largest buildings and greatest variability in size. Like Low-Intensity neighborhoods, these neighborhoods are less dense and have greater variability in distance than Large-Lot

Residential and Small-Lot Residential neighborhoods. These neighborhoods, however, are vastly different than the other three neighborhood types with respect to the housing stock. While the other neighborhood types have a relatively high share of single-family housing units, the share of single-family units in these neighborhoods is only about 25%, with only 11% of units in detached single-family structures. The remainder are multifamily units, with about half of the housing units located in multifamily structures with 20+ units. Consequently, these neighborhoods are representative of areas with mixed-use and high-intensity development. Thus, we refer to them as High-Intensity neighborhoods.

There are other notable differences between these four neighborhood types in regard to housing values and the age of the housing stock. Housing values vary substantially, and are highest in High-Intensity ($471,487) and Large-Lot Residential ($441,731) neighborhoods and lowest in Small-Lot Residential ($309,360) and Low-Intensity ($340,471) neighborhoods. Meanwhile, housing units are older in Small-Lot Residential (1969) neighborhoods than Large-Lot Residential (1972), High-Intensity (1976), and Low-Intensity (1982) neighborhoods. Together, these findings suggest that Small-Lot Residential neighborhoods have older and more affordably priced housing, likely developed during the first wave of suburbanization. The substantial variation in the housing stock observed across the four neighborhood types indicates that building footprint-derived measures of neighborhood type successfully distinguish between neighborhoods with distinct housing characteristics.

To summarize, Large-Lot Residential and Small-Lot Residential neighborhoods are characteristic of urban and suburban areas, where housing units are predominantly single-family units. Meanwhile, Low-Intensity neighborhoods are characteristic of exurban areas with both residential and non-residential development. Lastly, High-Intensity neighborhoods are characteristic of mixed-use areas in the urban core with varied housing types, though they are predominantly multifamily. Later, we describe how we combine these neighborhood types with the land use and zoning regulations to examine how regulations, like maximum density restrictions, influence neighborhood type. In Appendix D, we present a small case study, looking at how these neighborhood types align with residential zones in Chapel Hill, North Carolina. Overall, these neighborhood types align closely with these residential zones and thus maximum density restrictions.

3.2.2 New and Old Development

In this section, we describe how we identify new (post-2001) and old (pre-2001) development. All of our development variables are evaluated at the municipality level. Our first development variable measures the presence of these different neighborhood types. We calculate this simply as the share of census blocks c in municipality m that are classified as neighborhood type n:

$$Neighborhood\ Share_{mn} = \frac{\sum_{c \in m} 1(c = n)}{\sum_{c \in m} 1} \tag{1}$$

Our other development variables measure the share of buildings that are observed in these neighborhood types.

As mentioned earlier, we use land cover data to examine changes in building development over time. Specifically, we estimate the number of buildings that were developed by 2001 and 2019. We then calculate the difference in these measures to obtain the number of buildings that were developed post-2001. To do this, we use our measure of low-intensity, medium-intensity, and high-intensity (LMH) land cover that was described earlier. For each census block, we count the number of pixels that are classified as LMH. We then use these pixels to measure new development, which is calculated as the following:

$$New\ Development_c = \max\left(\frac{LMH_{c,2019} - LMH_{c,2001}}{LMH_{c,2019}}, 0\right) \tag{2}$$

where $LMH_{c,2001}$ and $LMH_{c,2019}$ are the count of these pixels that are observed in 2001 and 2019, respectively. Thus, $New\ Development_c$ measures the share of development that occurred after 2001. Because our measure identifies new LMH development, it primarily captures greenfield development. This is illustrated in Figure 1 as the conversion of a white cell (i.e., undeveloped or developed open space) in 2001 to a red cell (i.e., low-intensity, medium-intensity, or high-intensity development) in 2019. We restrict this measure to be non-negative, replacing negative values with zero.[3] On average, 8% of pixels in a census block converted from undeveloped to LMH between 2001 and 2019. This was highest in Low-Intensity neighborhoods (14%), followed by High-Intensity (8%), Large-Lot Residential (6%), and Small-Lot Residential (5%) neighborhoods.

Next, we calculate the number of buildings. We begin with the number of buildings that were developed by 2019. For simplicity, we assume that the number of buildings that were developed by 2019 is the same as the number of building footprints that we observe in the Microsoft dataset.[4] Thus, the number of buildings in a census block that were developed by 2019 is calculated as the following:

$$Buildings_{c,2019} = \sum Building\ Footprints_c \tag{3}$$

We then use our measure of new development as a scaling factor to calculate the number of buildings that were developed by 2001 as well as the number of buildings that were developed post-2001. They are calculated as the following:

$$Buildings_{c,2001} = (1 - New\ Development_c) \times Buildings_{c,2019} \tag{4}$$

$$Buildings_{c,Post-2001} = (New\ Development_c) \times Buildings_{c,2019} \tag{5}$$

Finally, we aggregate the buildings b up to the municipality level to get the share of buildings by neighborhood type at each time period. For example, the share of buildings developed by 2001 is calculated as the following:

[3] Negative values represent a loss of developed land. This occurred for only about 10 census blocks.

[4] Note that some of the building footprints were collected in 2020. Because of this, our measure of the number of buildings developed by 2019 is not completely precise because some buildings may have been developed in 2020 after the 2019 NLCD dataset was constructed. Nonetheless, this measure is a close proxy.

Figure 3: Distribution of Neighborhoods and Buildings

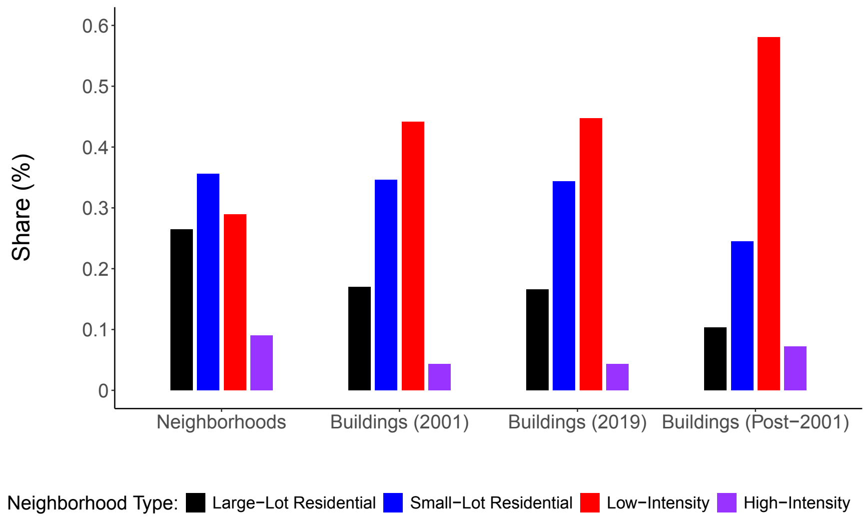

Notes: The figure plots the distribution of neighborhood types within the four development variables. The shares of the neighborhood types sum to one within each development variable.

$$Buildings\ Share_{mn,2001} = \frac{\sum_{b \in m} 1(b \in n)}{\sum_{b \in m} 1} \tag{6}$$

We calculate similar equations for the share of buildings developed by 2019 and the share of buildings developed post-2001. To summarize, our four development variables are $Neighborhood\ Share$, $Buildings\ Share_{2001}$, $Buildings\ Share_{2019}$, and $Buildings\ Share_{Post-2001}$.[5]

Figure 3 plots the distribution of development by neighborhood type. While Small-Lot Residential neighborhoods have the greatest neighborhood share, Low-Intensity neighborhoods have the greatest share of buildings that were developed by 2001 and 2019. Low-Intensity neighborhoods also have the greatest share of buildings that were developed post-2001 (58%). This is followed by Small-Lot Residential (24%), Large-Lot Residential (10%), and High-Intensity neighborhoods (7%). Thus, the majority of new development between 2001 and 2019 has occurred in the suburban and exurban fringe.

[5]For simplicity, we have dropped the mn subscripts.

3.2.3 Stylized Facts

In this section, we present a series of descriptive statistics and stylized facts about these data and measurements, aggregated at the municipality level. We aggregate up to the municipality level because that is the spatial unit of analysis that was used in the land use and zoning regulations survey. To obtain municipality-level measures of the neighborhood types, we aggregate the census blocks up to the municipality level based on the centroid of the census block. Because unincorporated and incorporated municipal boundaries often overlap, we only join census blocks to the unincorporated parts of townships and counties. In other words, if a census block is located in an incorporated municipality, it is not treated as being under the township or county's jurisdiction. We then measure the presence of the different neighborhood types by calculating the share of each neighborhood type within a municipality. Meanwhile, to obtain municipality-level measures of the building characteristics, we calculate the average building characteristics across census blocks within a municipality.

Table 2 presents summary statistics at the municipality level by municipality type. Panel (a) documents the land use and zoning regulations. While there is substantial variation in these regulations by municipality type, we focus on maximum density restrictions for reasons discussed previously. Recall that our measure of maximum density documents the maximum number of dwelling units permitted per acre in any residential zone in a municipality. While this is an imperfect measure, as it does not reflect variation within municipalities, it does provide insight about the variation between municipality types. The average number of units permitted per acre is 19 in cities, followed by 16 in counties and 9 in townships. This suggests that cities have the least restrictive density zoning, while townships have the most restrictive density zoning, given the caveat about within municipality variation. Panel (b) documents the share of neighborhood types. Cities have a substantially greater share of Small-Lot Residential neighborhoods (42%) than townships (26%) and counties (23%) but a lower share of Low-Intensity neighborhoods (19%) than townships (45%) and counties (48%). Municipalities also differ in the share of Large-Lot Residential and High-Intensity

Table 2: Summary Statistics by Municipality Type

	Municipality Type			
Characteristic	City	Township	County	Overall
Panel A: Land Use and Zoning				
Affordable Housing Funding Mechanism	0.27	0.15	0.22	0.23
Affordable Housing - Government	0.24	0.34	0.26	0.27
Affordable Housing - Private	0.36	0.24	0.27	0.32
APFO	0.45	0.60	0.58	0.51
Growth Limit Tool	0.09	0.01	0.22	0.08
Impact Fee	0.53	0.25	0.41	0.44
Mobile Home	0.48	0.48	0.85	0.52
Urban Service Area	0.07	0.02	0.17	0.07
Voting Referendum	0.28	0.44	0.16	0.31
Maximum Density (Continuous)	19	9	16	16
Maximum Density (Binary)				
$<$ 4 units	0.07	0.33	0.17	0.15
4–7 units	0.09	0.22	0.12	0.13
8–15 units	0.25	0.30	0.28	0.27
16–30 units	0.32	0.12	0.22	0.25
$>$ 30 units	0.27	0.03	0.21	0.20
Panel B: Neighborhood Type				
Large-Lot Residential	0.29	0.23	0.21	0.26
Small-Lot Residential	0.42	0.26	0.23	0.36
Low-Intensity	0.19	0.45	0.48	0.29
High-Intensity	0.11	0.06	0.08	0.09
Panel C: Building Characteristics				
Area (Median)	268	211	239	249
Area (IQR)	190	143	176	176
Density	736	358	299	585
Distance (Median)	25	33	34	28
Distance (IQR)	7	13	14	10
Observations	1,113	498	194	1,805
(%)	(0.62)	(0.28)	(0.11)	(1.00)

Notes: The table presents summary statistics by municipality type. Panels (a), (b), and (c) document land use and zoning regulations, neighborhood types, and building characteristics, respectively.

neighborhoods but to a lesser degree. Together, this suggests that most of the variation in neighborhood type between municipalities is driven by the share of Small-Lot Residential and Low-Intensity neighborhoods. While Small-Lot Residential neighborhoods are the defining neighborhood type in cities, it is Low-Intensity neighborhoods in townships and counties. Lastly, panel (c) documents the building characteristics. As expected, building density is higher in cities. Buildings are also larger and spaced closer together in cities.

Figure 4: Maximum Density - Building Characteristics

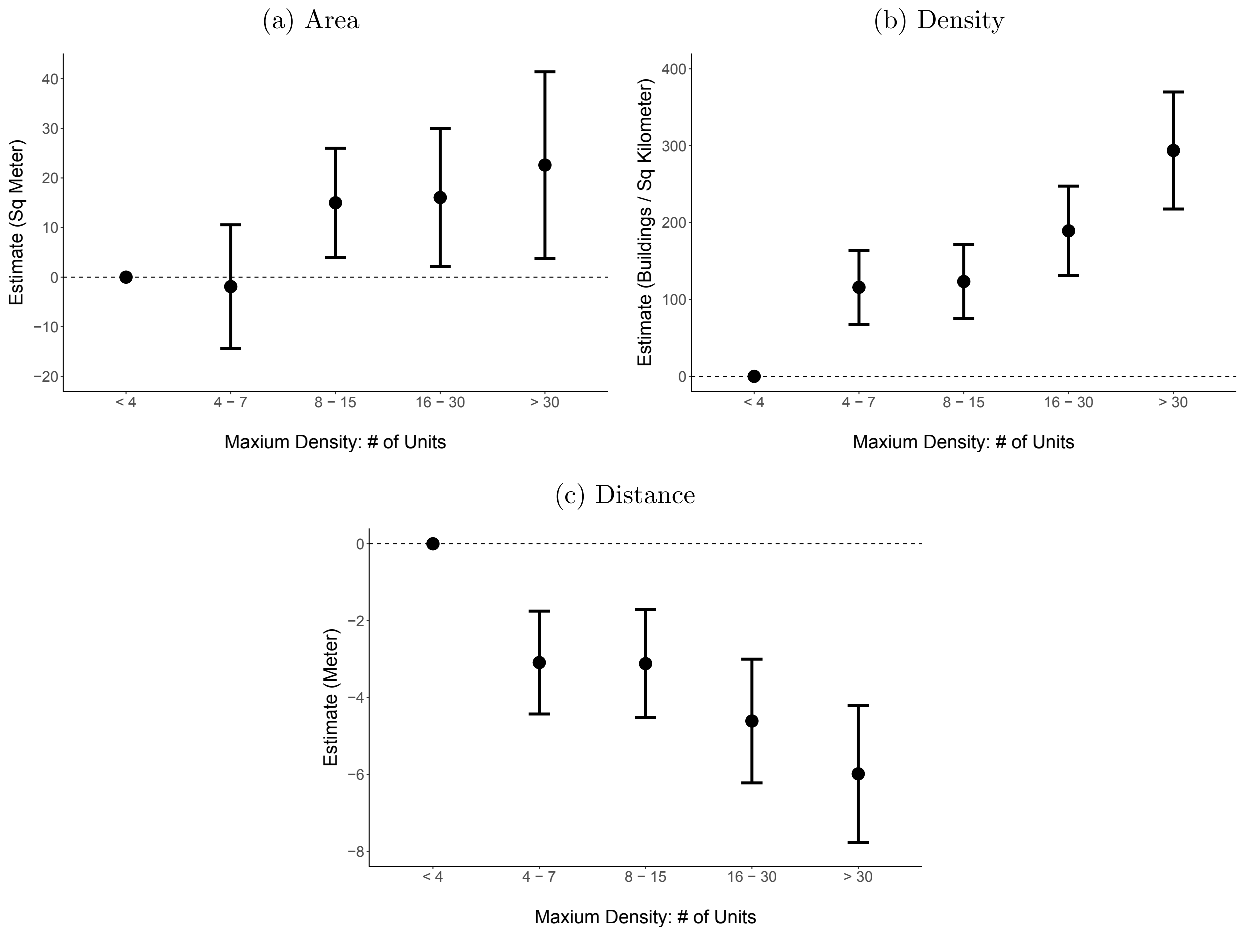

Notes: The figure plots the estimates of maximum density for the following building characteristics: area; density; and distance. The estimates are obtained from a regression similar to equation (7), where the building characteristics are the outcome variables.

The main objective of this paper is to examine the link between density zoning, neighborhood type, and residential segregation. We measure density zoning via maximum density restrictions. Thus, the first thing to examine is whether maximum density restrictions influence neighborhood type. We test this two ways. First, we examine the correlation between maximum density and the building characteristics. If maximum density is binding, then it should manifest in the physical and spatial characteristics of buildings. Figure 4 plots the relationship between maximum density and the following building characteristics: area, density, and distance. These results suggest that maximum density is positively correlated with area and density and negatively correlated with distance. In other words, buildings are larger and more dense in less restrictive municipalities that permit more dwelling units per acre. Meanwhile, buildings are spaced closer together in these less restrictive municipalities. Notably, increases in maximum density are binding on measures of building density and distance at the threshold between < 4 and 4–7 units per acre, suggesting even modest increases in maximum density are associated with changes in the housing stock. Conversely, increases in maximum density are binding on measures of building area only at 8–15 units per acre and above. This suggests that, at lower density thresholds, maximum density shapes the placement of dwellings but not their size.

Second, we examine the correlation between maximum density and the share of our different neighborhood types. Figure 5 presents a variety of plots, displaying the distribution of maximum density. Panel (a) plots the distribution of municipalities by maximum density category. Panel (b) plots the distribution of neighborhood types within each density category.[6] There are noticeable differences in these distributions. For example, the share of Low-Intensity neighborhoods decreases as maximum density increases, while the share of the other neighborhood types generally increase as maximum density increases. In our analysis, we explore heterogeneity by municipality type and census region. Panel (c) documents the variation in maximum density by municipality type. Cities predominantly have moderate to high measures of maximum density, whereas townships are the opposite. Counties vary, with a substantial share allowing moderate or high-density development, and a minority allowing only low-density development. Meanwhile, panel (d) documents the distribution of neighborhood types within each density category, disaggregated by municipality type. Cities, for example, are dominated by Large-Lot Residential neighborhoods, regardless of maximum density. Meanwhile, townships and counties tend to be dominated by Low-Intensity neighborhoods. Similarly,

[6]The neighborhood shares sum to one within each density category.

panel (e) plots maximum density by census region. A substantial share of municipalities in the Midwest and Northeast have maximum densities of less than 7 units per acre. This is less true for municipalities in the South and West. Lastly, panel (f) plots the distribution of neighborhood types within each density category, disaggregated by census region. While there are some differences across census regions, the distributions tend to follow a similar pattern as maximum density increases. Together, these results suggest that maximum density restrictions vary substantially across municipality types and census regions. This variation appears to shape both building characteristics and neighborhood types. We study the relationship between maximum density and neighborhood types in more detail in the following section using regression analysis. We also study the relationship between maximum density and the income and racial composition of municipalities using regression analysis, controlling for these neighborhood types.

In addition to maximum density, we also consider changes in maximum density. To do this, we select the sample of municipalities that responded to both the 2003 and 2019 NLLUS survey waves and then identify whether a municipality experienced a decrease, no change, or increase in maximum density between survey waves. Figure 6 presents a variety of plots, displaying the distribution of maximum density change. Of the 848 municipalities that responded to both surveys, 27%, 48%, and 25% experienced a decrease, no change, or increase in maximum density, respectively (see panel (a)). The distribution of neighborhood types is fairly consistent regardless of the change in maximum density (see panel (b)). There is much more variation by municipality type and census region, as seen in panels (c) and (e). For example, townships have the greatest share of municipalities that decreased maximum density (35%), counties have the greatest share of municipalities that experienced no change in maximum density (57%), and cities have the greatest share of municipalities that increased maximum density (28%). The share of municipalities that increased maximum density is fairly similar across regions, ranging from about 23% to 29%. This is not the case for municipalities that experienced a decrease or no change in maximum density. The share of municipalities that decreased maximum density ranges from 14% in the West to 36% in the Northeast. Meanwhile, the share of municipalities that experienced no change in maximum density ranges from 38% in the Northeast to 64% in the West. Similar to the full sample, there is little variation by neighborhood type within municipality types and census regions (see panels (d) and (f)).

Another objective of this paper is to examine the relationship between the change in building development and the change in the income and racial composition over time. We begin by looking at changes in building development. Similarly, Figure 7 presents a variety of plots, displaying the distribution of our measure of new development that was discussed previously. As seen in panel (a), the share of municipalities tends to decrease as new development increases. For example, 39% of municipalities experienced less than 5% new development, while 13% of municipalities experienced greater than 20% new development. While the shares of Small-Lot Residential and High-Intensity neighborhoods are fairly consistent by the degree of new development, this is not the case for Large-Lot Residential and Low-Intensity neighborhoods (see panel (b)). The share of Large-Lot Residential neighborhoods tends to decrease as new development increases, while the share of Low-Intensity neighborhoods tends to increase as new development increases. In other words, in municipalities that are developing rapidly, most new development is occurring in Low-Intensity neighborhoods, while in municipalities that are developing at a more modest pace, most new development is occurring in Large-Lot Residential neighborhoods. We similarly explore heterogeneity by municipality type and census region. As seen in panel (c), cities tend to develop at a slower pace, townships tend to develop at a moderate pace, and counties tend to develop at a faster pace. The increase in the share of Low-Intensity neighborhoods and the decrease in the share of Large-Lot Residential neighborhoods is strongest among townships and counties (see panel (d)). Together, these results are consistent with development in the suburban and exurban fringe. There is some variation by census regions. While the share of municipalities tends to decrease as new development

Figure 5: Maximum Density - Heterogeneity

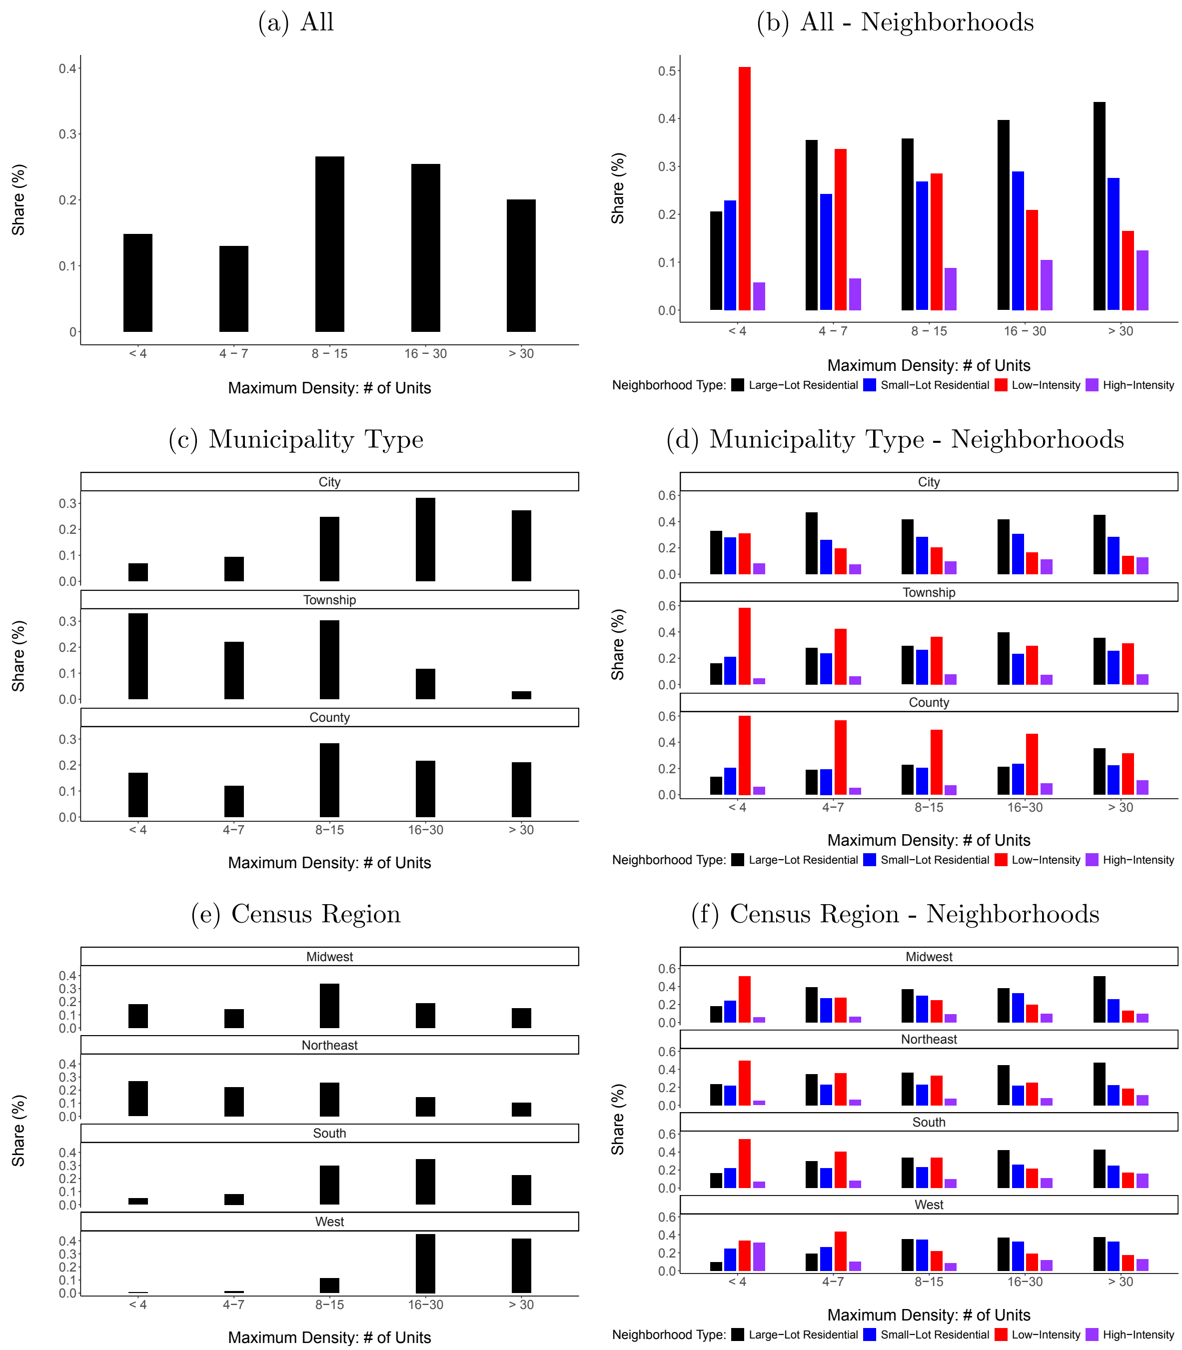

Notes: The figure presents multiple plots of the distribution of maximum density. Panel (a) plots the distribution for the full sample. Meanwhile, panel (b) plots the distribution of neighborhood types within each density category. Thus, the share of the neighborhood types sum to one within each density category. Panels (c) and (d) disaggregate the distributions by municipality type, while panels (e) and (f) disaggregate the distributions by census region.

Figure 6: Maximum Density Change - Heterogeneity

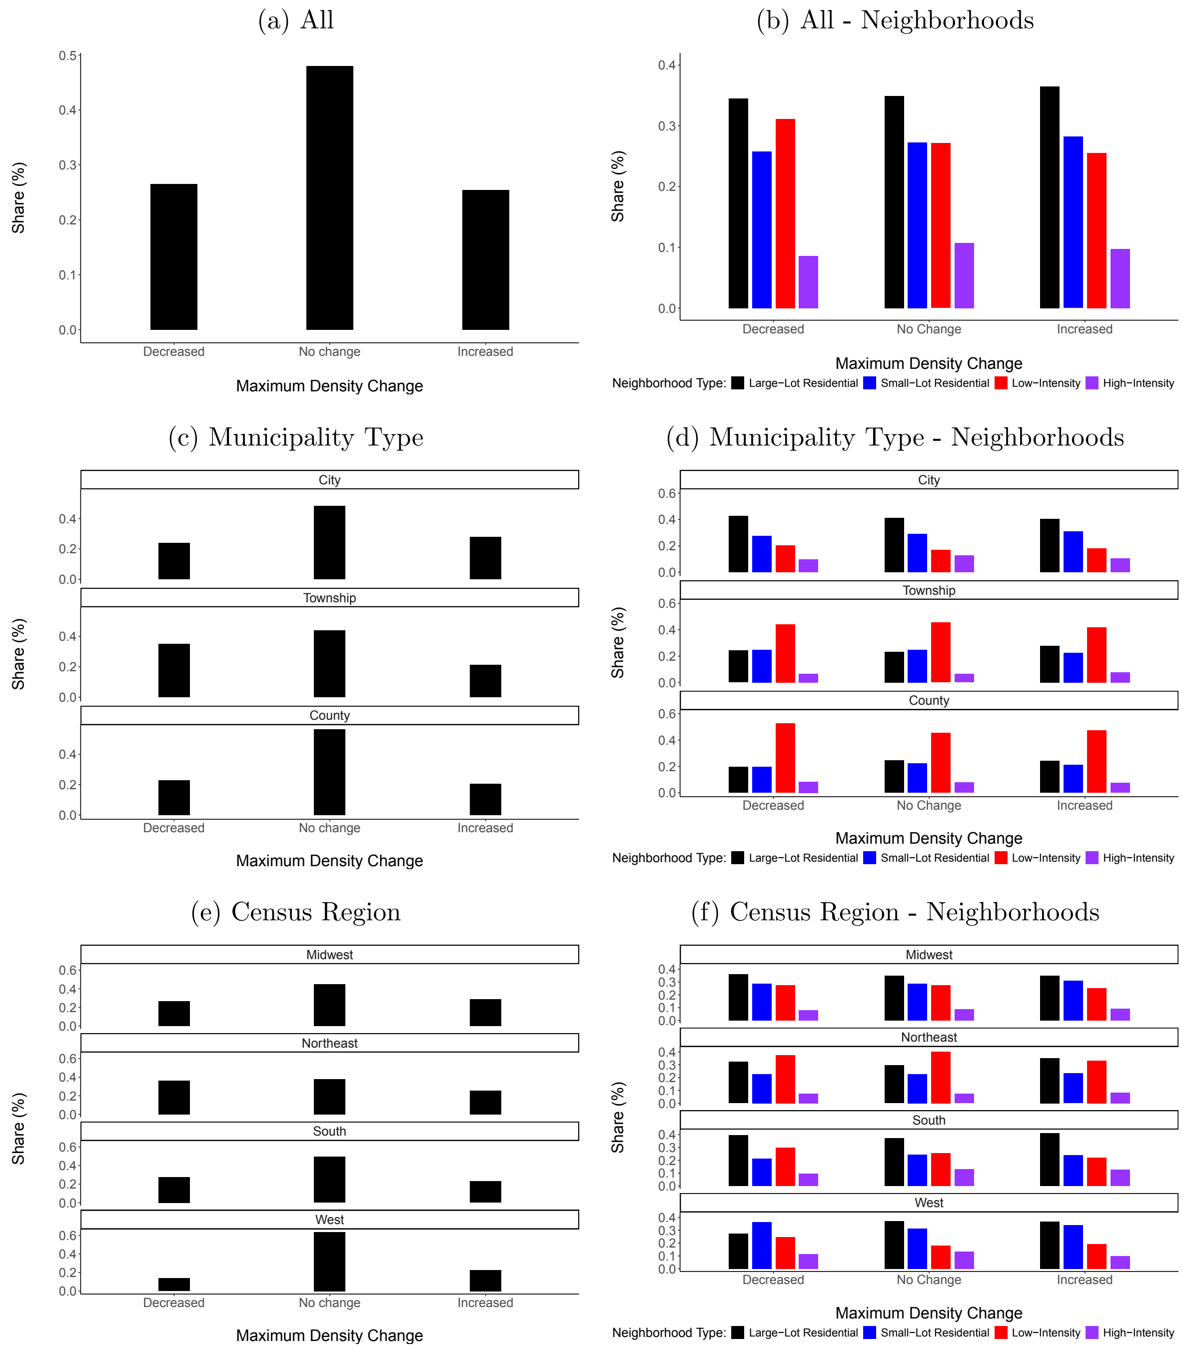

Notes: The figure presents multiple plots of the distribution of changes in maximum density (i.e., decreased, no change, and increased). Panel (a) plots the distribution for the full sample. Meanwhile, panel (b) plots the distribution of neighborhood types within each density category. Thus, the share of the neighborhood types sum to one within each density category. Panels (c) and (d) disaggregate the distributions by municipality type, while panels (e) and (f) disaggregate the distributions by census region.

Figure 7: New Development - Heterogeneity

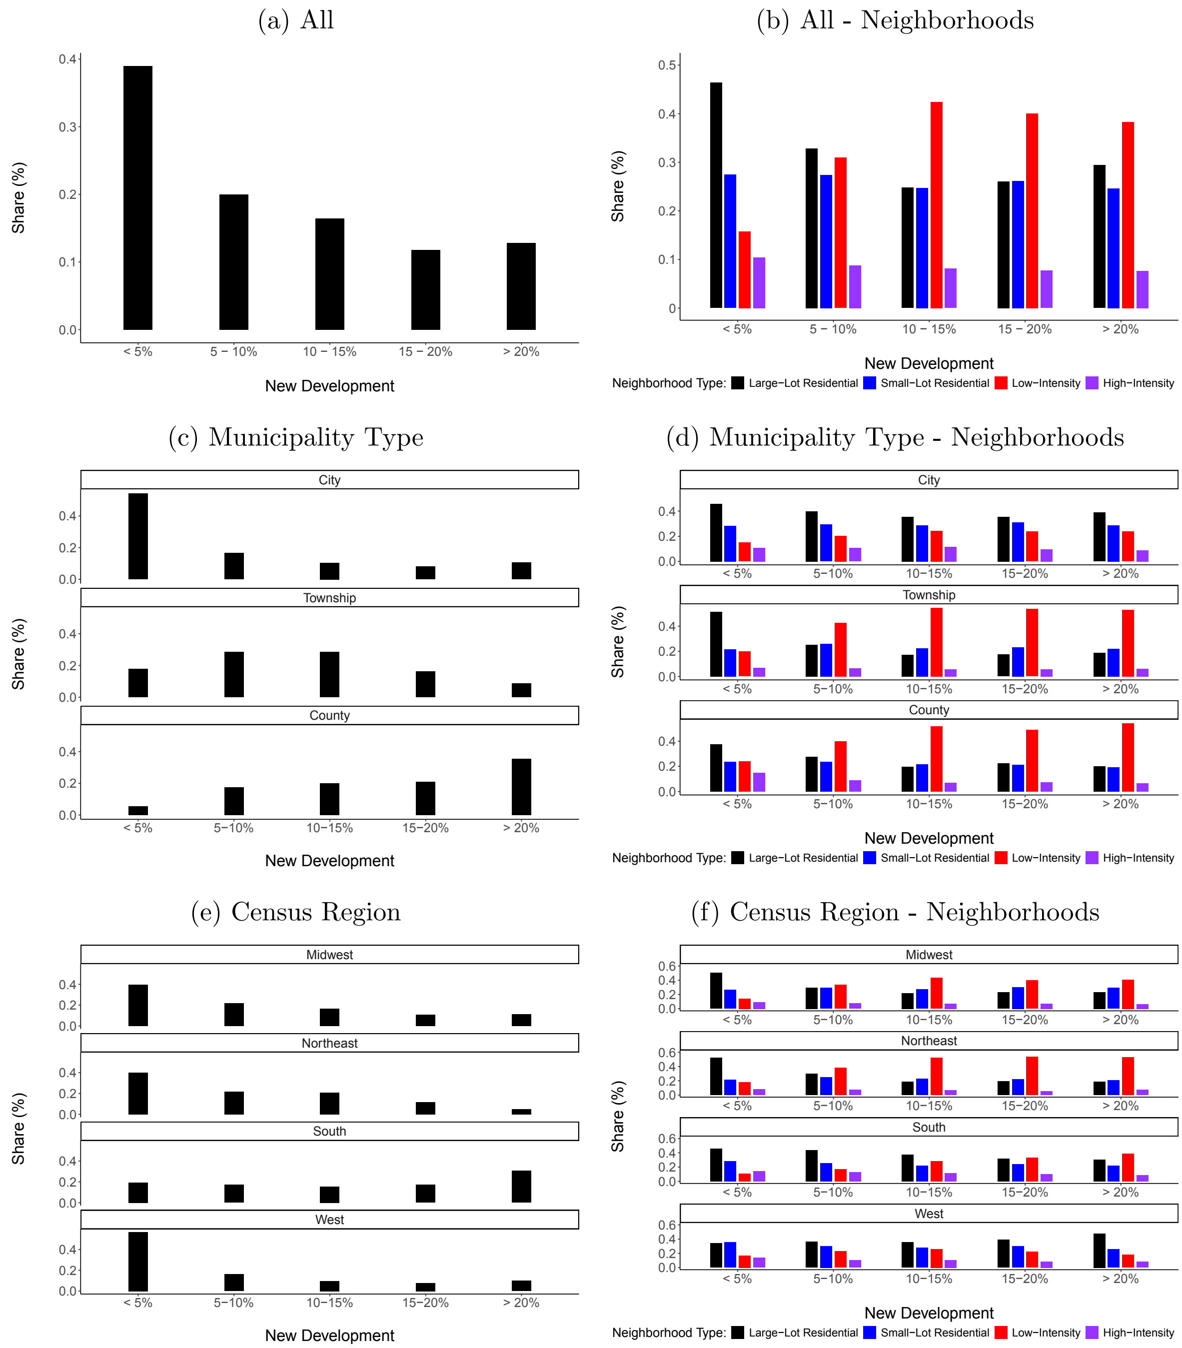

Notes: The figure presents multiple plots of the distribution of new development. Panel (a) plots the distribution for the full sample. Meanwhile, panel (b) plots the distribution of neighborhood types within each density category. Thus, the share of the neighborhood types sum to one within each density category. Panels (c) and (d) disaggregate the distributions by municipality type, while panels (e) and (f) disaggregate the distributions by census region.

Figure 8: New Development - Demographics

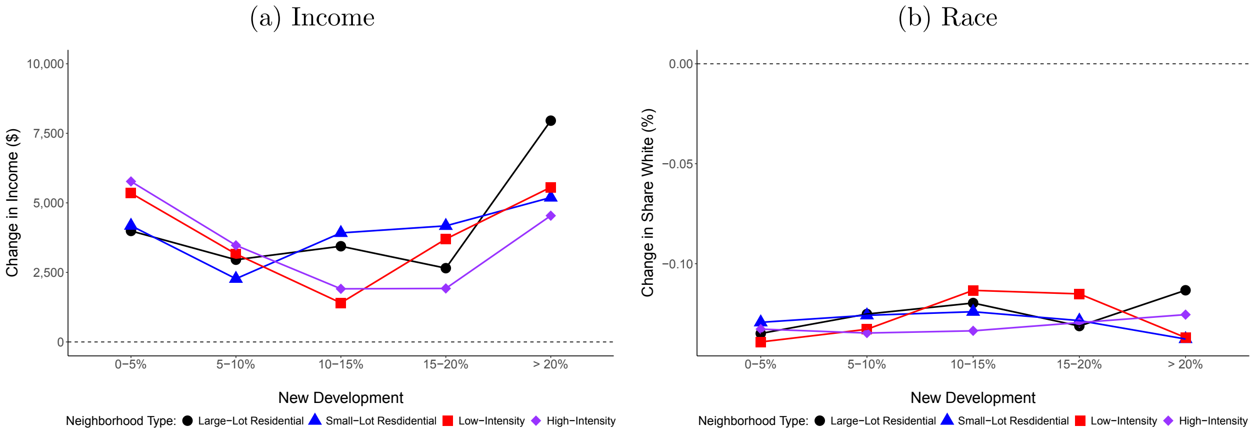

Notes: Panels (a) and (b) plot the mean change in income and share of residents who are non-Hispanic White from 2000 to 2020, respectively, disaggregated by neighborhood type and the share of new development.

increases, this is not the case for municipalities in the South (see panel (e)). Thus, municipalities in the South have experienced the most development between 2001 and 2019. While these other municipalities have experienced less development, they tend to be developing in Low-Intensity neighborhoods at a greater rate, with the exception of municipalities in the West (see panel (f)).

Finally, we pair the change in building development with the change in the income and racial composition of municipalities. Panel (a) of Figure 8 plots the change in income by the share of new development. On average, income increased by about \$4,000 between 2000 and 2020. This increase in income varies both between neighborhood types and the share of new development. For example, income increased between about \$3,000 to \$8,000 for Large-Lot Residential neighborhoods, with the greatest increase occurring in municipalities that had $>$ 20% new development. Meanwhile, income increased between \$2,000 to \$6,000 for High-Intensity neighborhoods, with the greatest increase occurring in municipalities that had 0–5% new development. Similarly, panel (b) plots the change in the share of residents who are non-Hispanic White by the share of new development. On average, the share of White residents decreased by about 13% between 2000 and 2020. Now, there is little variation between neighborhood types and the share of new development, as the share of

White residents decreased by about 11% to 14%. We study the change in these demographics more carefully in the next section, looking at the impact of maximum density as well as changes in maximum density.

From these figures and descriptive statistics, we highlight the following stylized facts:

1. Cities have the least restrictive maximum density restrictions, as they permit the greatest number of dwelling units per acre. Meanwhile, townships have the most restrictive. Cities were also the most likely to have relaxed these restrictions and increased maximum density between 2001 and 2019. Townships, in contrast, were the most likely to have tightened restrictions and decreased maximum density over this time period.

2. Municipalities in the South and West have the least restrictive maximum density restrictions, while municipalities in the Midwest and Northeast have the most restrictive. Municipalities in the Northeast have experienced the most change in maximum density, though they were more likely to tighten restrictions than loosen.

3. Counties have experienced the most new development, followed by townships and cities.

4. Municipalities in the South have experienced the most new development.

5. Most new development has occurred in suburban and exurban areas in Small-Lot Residential and Low-Intensity neighborhoods, respectively.

4 Empirical Approach

In the remainder of the paper, we examine the relationship between density zoning, neighborhood type, and residential segregation. We begin by conducting a series of cross-sectional analyses of the relationship between density zoning and neighborhood type in 2001 and 2019. We also conduct cross-sectional analyses of the relationship between density zoning and residential segregation in 2000 and 2020. Next, we examine the relationship between maximum density and changes in building development between 2001 and 2019. Lastly, we examine changes in maximum density and changes in income and racial composition between 2000 and 2020.

We measure density zoning via maximum density restrictions.[7] We begin by estimating the following equation:

$$y_{amt} = \beta Maximum\ Density_m + \gamma Z_m + \alpha_a + \alpha_t + \varepsilon_{amt} \tag{7}$$

where y_{amt} is the outcome variable for municipality m of type t in metropolitan area a. As discussed in the previous section, we have four main outcome variables that we use to examine neighborhood type:

- $Neighborhood\ Share$: Share of a given neighborhood type.
- $Buildings\ Share_{2001}$: Share of buildings that were developed by 2001 in a given neighborhood type.
- $Buildings\ Share_{2019}$: Share of buildings that were developed by 2019 in a given neighborhood type.
- $Buildings\ Share_{Post-2001}$: Share of buildings that were developed post-2001 (i.e., between 2001 and 2019) in a given neighborhood type.

The coefficient of interest is β, which represents the association between maximum density and these outcome variables. Although our primary focus is maximum density, we control for additional land use and zoning regulations, represented by the vector Z_m. We include all land use and zoning regulations listed in Table 2. We also include fixed effects, controlling for variation within metropolitan area α_a and municipality type α_t. When estimating the equation for $Buildings\ Share_{Post-2001}$, we also control for building density (i.e., buildings per square kilometer) in 2001 and new development, which was defined previously. We

[7] See Appendix C for a comparison with minimum lot size.

do this to account for the land available for development as of 2001 and the rate of new development since then. We cluster the standard errors by metropolitan area. We estimate these equations individually for each neighborhood type.

We then estimate similar equations to examine the relationship between density zoning and residential segregation. In doing so, we control for the share of neighborhood types. Since we can only include three out of the four neighborhood types in an equation without issues of multicollinearity, we drop Large-Lot Residential neighborhoods. Thus, we estimate the following equations:

$$\begin{aligned} y_{amt} = &\beta Maximum\ Density_m + \delta_1 Small\text{-}Lot\ Residential_m + \delta_2 Low\text{-}Intensity_m + \\ &\delta_3 High\text{-}Intensity_m + \gamma Z_m + \alpha_a + \alpha_t + \varepsilon_{amt} \end{aligned} \tag{8}$$

where y_{amt} is median household income or the share of residents who are non-Hispanic White in a municipality. We estimate this equation separately for these demographics in 2000 and 2020. In these equations, the neighborhood types represent the share of buildings that were developed by 2001 and 2019 in a given neighborhood type, respectively. We also estimate this equation for the change in demographics between 2000 and 2020. In this equation, the neighborhood types represent the share of buildings that were developed post-2001 in a given neighborhood type. The δs represent the association between these shares and the demographics. When estimating the change in demographics, we also control for building density in 2001, new development, and demographics in 2000. For example, when estimating the change in income, we control for income in 2000. We do so to account for variation in the demographic composition of municipalities prior to our estimates of new development.

In our main specifications, we use a continuous measure of maximum density (i.e., number of dwelling units permitted per acre). This yields a linear relationship between maximum density and the outcome variables. To examine non-linear relationships, we perform an alternative specification, where we replace the continuous measure of maximum density with a set of binary variables corresponding to the categories of maximum density that were reported in the survey (e.g., 4–7 units). In this specification, we compare the estimates to the base category of < 4 units. This allows us to examine the particular density thresholds at which density zoning is binding.

Finally, we conduct a series of additional specifications to study heterogeneity. As seen in the previous section, there is substantial variation in maximum density and new development by municipality type and census region. We explore this variation more carefully here by subsetting the sample by municipality type and estimating the equations for these subsamples individually. We do the same for census regions. Next, we conduct a separate heterogeneity analysis to study the potential impacts of upzoning. As described previously, we identify whether a municipality experienced a decrease, no change, or increase in maximum density between the 2003 and 2019 NLLUS survey waves. In this specification, we include binary variables that indicate these changes in maximum density. Lastly, as a robustness check, we test alternative measures of density zoning, replacing maximum density with minimum lot size from the WRLURI. These results are in Appendix C.

5 Results

5.1 Neighborhood Type

We begin by examining the relationship between density zoning and neighborhood type. Again, we measure density zoning via maximum density restrictions. The results are presented in Table 3. The panels display the results for a given neighborhood type, each estimated by an individual regression. Meanwhile, the columns display the results for our outcome variables: $Neighborhood\ Share$, $Buildings\ Share_{2001}$, $Buildings\ Share_{2019}$, and $Buildings\ Share_{Post-2001}$. Thus, column (1) measures the share of neighborhood types, while the other columns measure the share of buildings within neighborhood types. While the estimates are qualitatively similar between columns, the estimates tend to be greater in magnitude and more precise when measuring the share of buildings within neighborhood types. Since we are ultimately interested in changes in building development, we focus our discussion on the share of buildings within neighborhood types, comparing results for buildings developed by 2001 and buildings developed post-2001. Generally, the estimates for buildings developed by 2001 are greater in magnitude than buildings developed post-2001.[8]

The results indicate that maximum density is associated with relatively small increases in the share of buildings in Large-Lot Residential. A one unit increase in maximum density is associated with about a 0.07 to 0.11 percentage point (0.4% to 1.1%) increase in the share of buildings in Large-Lot Residential neighborhoods. This range of estimates is based on buildings developed by 2001 and buildings developed post-2001. The relationship between maximum density and the share of buildings in Small-Lot Residential and Low-Intensity neighborhoods is stronger. A one unit increase in maximum density is associated with about a 0.25 to 0.54 percentage point (1.0% to 1.5%) increase in the share of buildings in Small-Lot Residential neighborhoods. In contrast, a one unit increase in maximum density is associated with about a –0.71 to –0.43 percentage point (–1.6% to –0.7%) decrease in the share of buildings in Low-Intensity neighborhoods. Meanwhile, a one unit increase in maximum density is associated with about a 0.05 to 0.10 percentage point (1.3% to 1.4%) increase in the share of buildings in High-Intensity neighborhoods. These results suggest that maximum density restrictions have the largest absolute (though opposing) impact on the share of Small-Lot Residential and Low-Intensity neighborhoods. That said, although the impact of maximum density on the presence of High-Intensity neighborhoods is smaller, because High-Intensity neighborhoods represent less than 10% of all neighborhoods, increases in maximum density are associated with relative effects (i.e., percent changes) on the share of these neighborhoods.

[8]This may be attributable to the large number of municipalities with low levels of new development and the variability in neighborhood type in municipalities with different amounts of new development (see panels (a) and (b) in Figure 7).

Table 3: Estimates - Neighborhood Type

		Buildings			
	Neighborhoods	2001	2019	Post-2001	Post-2001
	(1)	(2)	(3)	(4)	(5)
Panel A: Large-Lot Residential					
Maximum Density	0.0001	0.0011***	0.0012***	0.0007*	0.0007
	(0.0003)	(0.0004)	(0.0003)	(0.0004)	(0.0005)
Change: No					-0.0069
					(0.0088)
Change: Increase					-0.0014
					(0.0087)
Mean	0.26	0.17	0.17	0.10	0.10
Adjusted R^2	0.20	0.33	0.35	0.25	0.28
Panel B: Small-Lot Residential					
Maximum Density	0.0038***	0.0054***	0.0054***	0.0025***	0.0033***
	(0.0006)	(0.0007)	(0.0007)	(0.0008)	(0.0011)
Change: No					0.0407***
					(0.0149)
Change: Increase					0.0385**
					(0.0181)
Mean	0.36	0.35	0.34	0.24	0.24
Adjusted R^2	0.37	0.40	0.40	0.33	0.37
Panel C: Low-Intensity					
Maximum Density	-0.0049***	-0.0071***	-0.0070***	-0.0043***	-0.0051***
	(0.0005)	(0.0006)	(0.0007)	(0.0009)	(0.0010)
Change: No					-0.0388**
					(0.0146)
Change: Increase					-0.0472**
					(0.0181)
Mean	0.29	0.44	0.45	0.58	0.58
Adjusted R^2	0.61	0.60	0.61	0.46	0.50
Panel D: High-Intensity					
Maximum Density	0.0011***	0.0005***	0.0005***	0.0010***	0.0012**
	(0.0002)	(0.0001)	(0.0001)	(0.0003)	(0.0005)
Change: No					0.0051
					(0.0080)
Change: Increase					0.0101
					(0.0106)
Mean	0.09	0.04	0.04	0.07	0.07
Adjusted R^2	0.20	0.28	0.23	0.17	0.17
Controls: Land Use and Zoning	✓	✓	✓	✓	✓
Fixed Effect: Municipality Type	✓	✓	✓	✓	✓
Fixed Effect: Metropolitan Area	✓	✓	✓	✓	✓
Observations	1,620	1,620	1,620	1,614	1,614

*p<0.1; **p<0.05; ***p<0.01

Notes: The table presents the estimates of the continuous measure of maximum density as well as the change in maximum density for the neighborhood types. Standard errors are clustered by metropolitan area. Panels (a), (b), (c), and (d) document Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively. Each panel is estimated individually.

Figure 9: Estimates - Neighborhood Type (Binary)

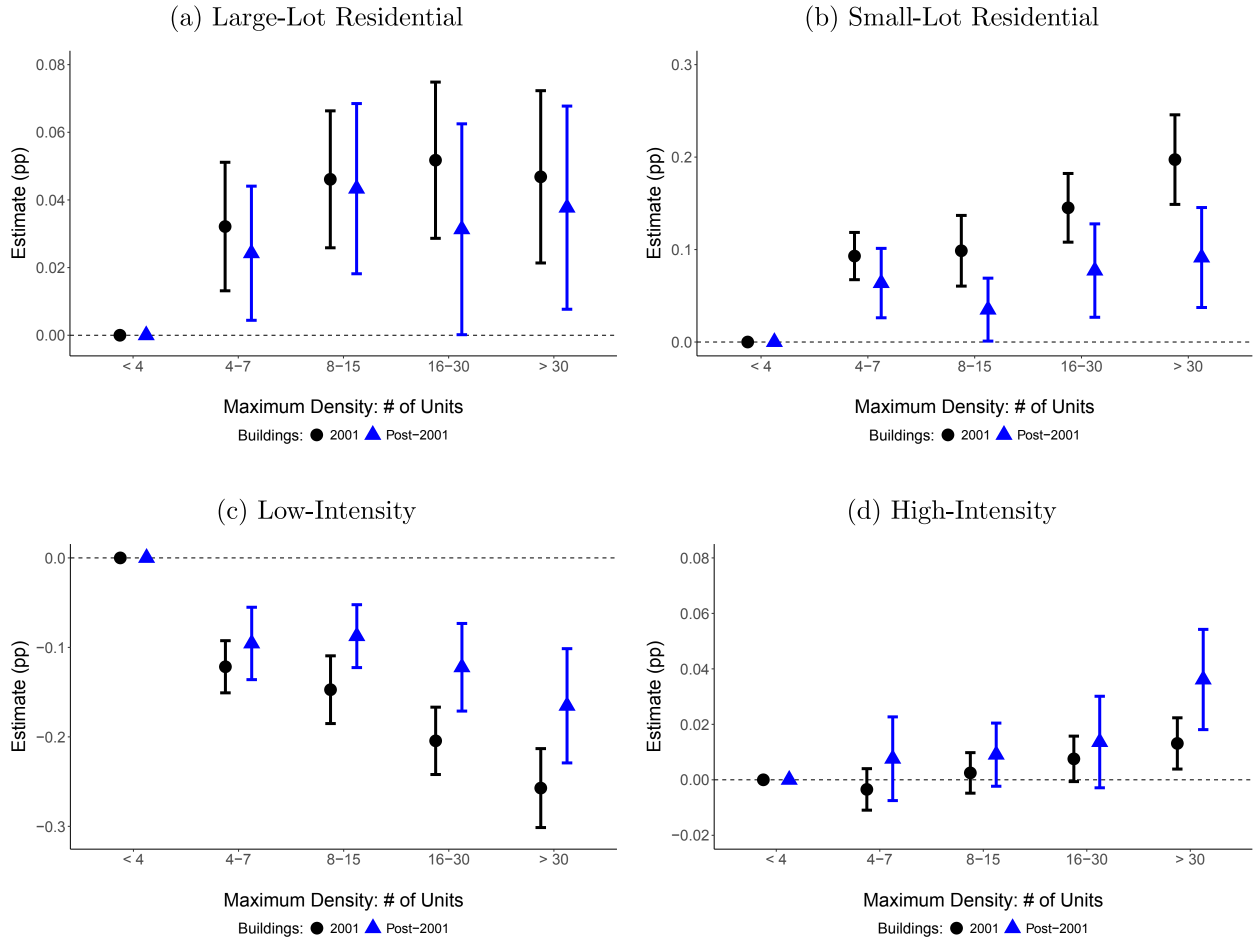

Notes: The figure presents the estimates of the binary measure of maximum density for the neighborhood types. The base category is < 4 units per acre. Standard errors are clustered by metropolitan area. Panels (a), (b), (c), and (d) document Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively.

Next, we examine the non-linear effects of maximum density. Figure 9 displays the estimates from the equation with the binary measure of maximum density (see also Appendix Table 1). These figures plot the estimates for both buildings developed by 2001 and buildings developed post-2001. All estimates are relative to the base category of < 4 units per acre. As seen in panel (a), municipalities with higher maximum density tend to have a greater share of buildings in Large-Lot Residential neighborhoods. The magnitude of these estimates are smaller than the other neighborhood types, ranging from about 3 to 5 percentage points (18.9% to 30.4%) for buildings developed by 2001 and about 2 to 4 percentage points (24.3% to 43.3%) for buildings developed post-2001. The marginal effects of these estimates attenuate substantially, beginning with 8–15 units per acre. Again, the relationship between maximum density and Small-Lot Residential neighborhoods is stronger, and the estimates continue to increase as maximum density increases (see panel (b)). The estimates for buildings developed by 2001 range from about 9 to 20 percentage points (26.6% to 56.4%), while the estimates for buildings developed post-2001 range from about 4 to 9 percentage points (14.6% to 38.1%). These effects are large and immediate, beginning with 4–7 units per acre. As an example, the estimate of 4–7 units is about half the magnitude of the estimate of > 30 units for buildings developed by 2001. While the marginal effects attenuate somewhat, the estimates tend to increase as maximum density increases. The relationship between maximum density and Low-Intensity neighborhoods is also strong but in the opposite direction (see panel (c)). Now, municipalities with higher maximum density tend to have a lower share of buildings in Low-Intensity neighborhoods. The estimates for buildings developed by 2001 range from about –26 to –12 percentage points (–58.0% to –27.7%), while the estimates for buildings developed post-2001 range from about –17 to –10 percentage points (–37.6% to –16.5%). Again, the marginal effects are immediate and attenuate with further increases in maximum density. Lastly, panel (d) displays the estimates for High-Intensity neighborhoods. While the share of buildings in High-Intensity neighborhoods tends to increase as maximum density increases, the effect is delayed until > 30 units per acre. In other words, a modest increase maximum density from < 4 units has no impact on the share of buildings in these neighborhoods. The estimates range from about 0 to 1 percentage points (–8.8% to 32.8%) for buildings developed by 2001 and about 0 to 4 percentage points (10.1% to 51.6%) for buildings developed post-2001.

Taken together, these findings provide insight into the non-linear ways in which maximum density shapes the share of buildings in a given neighborhood type and the thresholds at which maximum density is most impactful. Maximum density shapes neighborhood type primarily through the distribution of Small-Lot Residential and Low-Intensity neighborhoods. These distributions are particularly influenced by differences in the base category (< 4 units) and the second-lowest density category (4–7 units). Thus, it is the tradeoff between low and low-moderate maximum densities that influences the share of these two types of neighborhoods. Maximum density also influences the share of High-Intensity neighborhoods, though it only occurs at high densities. The results of our analysis of minimum lot size from the WRLURI lead to a similar conclusion (see Appendix C).

For our heterogeneity analyses, we estimate similar equations with the binary measure of maximum density. First, we examine heterogeneity by municipality type. Appendix Figures 1 and 2 plot the estimates for buildings developed by 2001 and buildings developed post-2001, respectfully. While the relationship between maximum density and the share of buildings in a given neighborhood type is similar between municipality types, the magnitude of the effects tend to be greater in cities and townships. The estimates for counties are attenuated, especially for Large-Lot Residential neighborhoods. This relationship is true for both buildings developed by 2001 and buildings developed post-2001. These results suggest that the majority of the effect that we observed in the full sample of municipalities is driven by cities and townships. Second, we examine heterogeneity by census region. Similarly, Appendix Figures 3 and 4 plot the estimates for buildings developed by 2001 and buildings developed post-2001, respectively.[9] The estimates do not vary much by census region for buildings developed by 2001. There is more variation for buildings developed post-2001. For example, the magnitude of the effect for Large-Lot Residential and Small-Lot Residential neighborhoods is the greatest in the Midwest and South, respectively. The latter result coincides with our stylized fact that most new development has been occurring in suburban areas, particularly the South. Meanwhile, the magnitude of the effect for Low-Intensity

[9] We do not display the estimates for municipalities in the West because there are a limited number of municipalities in the base category of maximum density with < 4 units (see panel (e) of Figure 5). These estimates tend to have large standard errors, which severely skew the y-axis of the figures.

neighborhoods is the smallest in the Northeast. There is little variation across regions for High-Intensity neighborhoods.

Lastly, we examine changes in maximum density. Recall that we measure whether a municipality experienced a decrease, no change, or increase in maximum density between the 2003 and 2019 NLLUS survey waves. Column (5) of Table 3 reports the estimates for buildings developed post-2001 with indicators for no change and an increase in maximum density. Thus, the base category is a decrease in maximum density. The first thing to note about these regressions is that the estimates are robust to the regressions without these indicators, as seen in column (4). The estimates of these indicators are small and statistically insignificant for Large-Lot Residential and High-Intensity neighborhoods. Meanwhile, the estimates of these indicators are about 4 percentage points (16.0% to 17.0%) for Small-Lot Residential neighborhoods, though they are not statistically different from each other. A similar result occurs for Low-Intensity neighborhoods, just in the opposite direction. Now, the estimates of these indicators are about –5 to –4 percentage points (–8.1% to –6.7%). Together, these results suggest that the potential impact of upzoning may be limited, as we observed no statistical difference between municipalities that reported no change in maximum density and those that reported an increase in maximum density. Rather, the results suggest that downzoning (i.e., decreasing maximum density) may have a larger impact on patterns of development. This may be attributable, in part, to our measure of new development, which primarily captures development in greenfield sites rather than redevelopment or the intensification of development.

5.2 Income and Racial Composition

We now examine the relationship between density zoning and residential segregation, as measured by the income and racial composition of municipalities. We also explore the role of neighborhood type in this relationship. Table 4 presents the estimates for median household income. In panel (a), we do not include the neighborhood types. Here, a one unit increase in maximum density is associated with about a –\$374 (–0.4%) and –\$248 (–0.3%) decrease in income in 2000 and 2020, respectively. In panel (b), we include the neighborhood types as control variables. For income in 2000, we measure this by the share of buildings developed by 2001 in a given neighborhood type. Similarly, for income in 2020, we measure this by the share of buildings developed by 2019 in a given neighborhood type. Now, the effects of maximum density are attenuated to –\$170 (–0.2%) and –\$34 (< –0.1%), respectively, and are no longer significant. These results suggest that neighborhood type partially mediates the relationship between maximum density and the income composition of municipalities. As seen in panel (a) of Figure 10 and Appendix Table 1, most of this effect for income in 2000 is driven by municipalities that allow > 30 units per acre. These municipalities tend to have about \$9,920 less income than municipalities that allow < 4 units. Heterogeneity analyses suggest that this is driven by cities and townships as well as municipalities in the Northeast (see Appendix Figures 5 and 6), where density zoning is particularly strict (see Figure 6). We additionally examine the change in income between 2000 and 2020. The results indicate that a one unit increase in maximum density is associated with about a \$102 (2.6%) increase in income over this time period. While the estimate of maximum density is not attenuated when including the neighborhood types, the neighborhood types remain largely significant. For example, a one percentage point increase in the share of buildings developed post-2001 in Small-Lot Residential neighborhoods is associated with about a –\$116 (–3.0%) decrease in the change of income over this time period. As seen in panel (a) of Figure 10, there is little variation in the estimates between maximum density categories, as they are all centered around zero. Additionally, there is little variation by municipality type and census region (see Appendix Figures 5 and 6). Lastly, we examine the change in maximum density. In column (4), we include indicators for no change and an increase in maximum density. We find no evidence that a change in maximum density over this time period influences the change in income. Together, these results suggest that changes in building development are more correlated with changes in income than are measures of change in maximum density.

This further suggests that changes in neighborhood type, as measured by building footprints, are more successful at capturing changes in income than are measures of maximum density derived from the NLLUS.

Table 5 presents the estimates for the share of residents who are non-Hispanic White. Here, a one unit increase in maximum density is associated with about a –0.3 percentage point (–0.5% to –0.4%) decrease in the share of White residents in both 2000 and 2020. Again, the effects of maximum density are attenuated after controlling for the neighborhood types. Similar to income, the effect of maximum density on race is driven by municipalities that allow $>$ 30 units per acre (see panel (b) of Figure 10 and Appendix Table 2). These municipalities tend to have about 9 percentage points (11.2%) less White residents than municipalities that allow $<$ 4 units. Heterogeneity analyses suggest that this relationship is generally true for all municipality types and census regions, but it is particularly strong in counties and municipalities in the Northeast (see Appendix Figures 5 and 6), where density zoning is particularly strict. We find that there is no observable relationship between maximum density and changes in the share of White residents between 2000 and 2020. We do, however, observe a small negative effect (–3 to –2 percentage points) for municipalities that have maximum density between 4 and 30 units per acre. This effect is not observed for municipalities with $>$ 30 units per acre, suggesting that the impact of maximum density on the change in race may be offset in municipalities that allow very high densities. This may point toward potential gentrification of these municipalities rather than their racial diversification. Similar to income, we find no evidence that changes in maximum density are associated with changes in the share of White residents, but do observe a relationship between changes in the share of White residents and our measures of new development. Again, this suggests that changes in neighborhood type are better proxies for changes in development than measures of maximum density derived from the NLLUS.

Table 4: Estimates - Income

	Year			
	2000 (1)	2020 (2)	2000–2020 (3)	2000–2020 (4)
Panel A: Without Neighborhood Types				
Maximum Density	-373.6***	-248.2**	101.54**	191.5**
	(102.1)	(93.6)	(47.8)	(89.2)
Change: No				273.6
				(1,056.3)
Change: Increase				1,701.3
				(1,626.7)
Mean	93,872	96,958	3,859	3,859
Adjusted R^2	0.28	0.32	0.33	0.32
Panel B: With Neighborhood Types				
Maximum Density	-170.2	-34.2	110.5**	207.8**
	(102.1)	(114.4)	(51.2)	(98.1)
Change: No				646.1
				(1,169.8)
Change: Increase				2,049.8
				(1,728.2)
Small-Lot Residential (2001)	-208.6**			
	(93.86)			
Low-Intensity (2001)	182.7			
	(136.1)			
High-Intensity (2001)	721.6**			
	(349.8)			
Small-Lot Residential (2019)		-267.3**		
		(114.3)		
Low-Intensity (2019)		133.9		
		(154.7)		
High-Intensity (2019)		250.2		
		(331.8)		
Small-Lot Residential (Post-2001)			-115.7**	-197.7**
			(56.4)	(81.3)
Low-Intensity (Post-2001)			-91.5	-148.8
			(59.6)	(90.9)
High-Intensity (Post-2001)			-147.8**	-225.8**
			(62.8)	(95.5)
Mean	93,872	96,958	3,859	3,859
Adjusted R^2	0.33	0.36	0.34	0.33
Controls: Land Use and Zoning	✓	✓	✓	✓
Fixed Effect: Municipality Type	✓	✓	✓	✓
Fixed Effect: Metropolitan Area	✓	✓	✓	✓
Observations	1,638	1,597	1,566	807

*p<0.1; **p<0.05; ***p<0.01

Notes: The table presents the estimates of the continuous measure of maximum density as well as the change in maximum density for income. Standard errors are clustered by metropolitan area. Panel (a) does not control for the neighborhood types but panel (b) does. For income in 2000 and 2020, we control for the share of buildings developed in each neighborhood type by 2001 and 2019, respectively. Meanwhile, for the change in income between 2000 and 2020, we control for the share of buildings developed post-2001.

Table 5: Estimates - White

	Year			
	2000 (1)	2020 (2)	2000–2020 (3)	2000–2020 (4)
Panel A: Without Neighborhood Types				
Maximum Density	-0.0032***	-0.0031***	0.0002	0.0004
	(0.0007)	(0.0008)	(0.0003)	(0.0005)
Change: No				-0.0009
				(0.0061)
Change: Increase				-0.0157
				(0.0104)
Mean	0.80	0.67	-0.13	-0.13
Adjusted R^2	0.46	0.50	0.18	0.18
Panel B: With Neighborhood Types				
Maximum Density	-0.0015*	-0.0006	0.0004	0.0006
	(0.0008)	(0.0008)	(0.0003)	(0.0005)
Change: No				0.0028
				(0.0059)
Change: Increase				-0.0120
				(0.0093)
Small-Lot Residential (2001)	-0.0008*			
	(0.0004)			
Low-Intensity (2001)	0.0016***			
	(0.0005)			
High-Intensity (2001)	-0.0031			
	(0.0022)			
Small-Lot Residential (2019)		-0.0012**		
		(0.0005)		
Low-Intensity (2019)		0.0025***		
		(0.0005)		
High-Intensity (2019)		-0.0034		
		(0.0022)		
Small-Lot Residential (Post-2001 Development)			-0.0006**	-0.0008**
			(0.0002)	(0.0003)
Low-Intensity (Post-2001 Development)			0.0003	0.0002
			(0.0002)	(0.0003)
High-Intensity (Post-2001 Development)			-0.0000	-0.0001
			(0.0003)	(0.0003)
Mean	0.80	0.67	-0.13	-0.13
Adjusted R^2	0.50	0.57	0.22	0.21
Controls: Land Use and Zoning	✓	✓	✓	✓
Fixed Effect: Municipality Type	✓	✓	✓	✓
Fixed Effect: Metropolitan Area	✓	✓	✓	✓
Observations	1,638	1,644	1,613	838

*p<0.1; **p<0.05; ***p<0.01

Notes: The table presents the estimates of the continuous measure of maximum density as well as the change in maximum density for the share of residents who are non-Hispanic White. Standard errors are clustered by metropolitan area. Panel (a) does not control for the neighborhood types but panel (b) does. For the share of residents who are non-Hispanic White in 2000 and 2020, we control for the share of buildings developed in each neighborhood type by 2001 and 2019, respectively. Meanwhile, for the change in the share of residents who are non-Hispanic White between 2000 and 2020, we control for the share of buildings developed post-2001.

Figure 10: Estimates - Demographics (Binary)

(a) Income (b) White

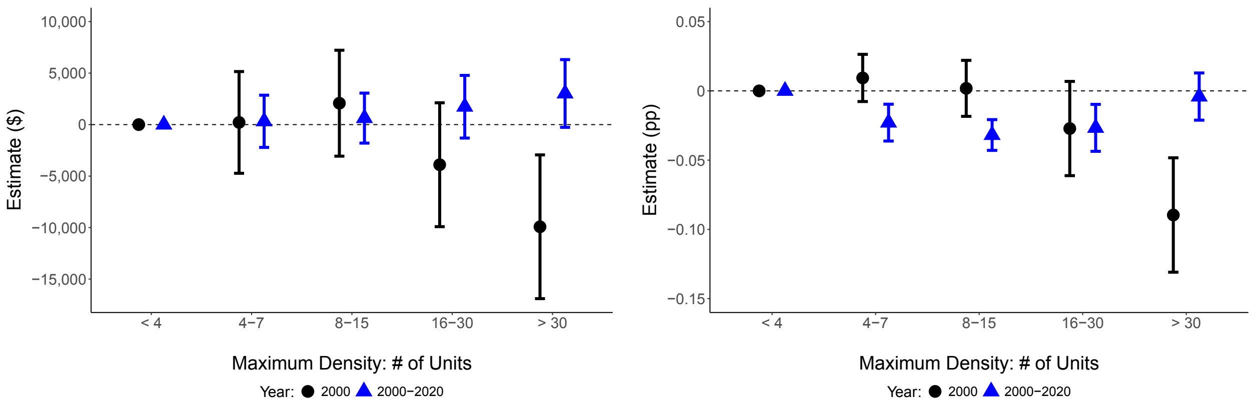

Notes: The figure presents the estimates of the binary measure of maximum density for the demographics. The base category is <4 units per acre. Standard errors are clustered by metropolitan area. Panels (a) and (b) document income and the share of residents who are non-Hispanic White, respectively.

6 Conclusion

In this paper, we examine the relationship between density zoning, neighborhood type, and residential segregation. We measure density zoning via maximum density restrictions reported in a national land use and zoning survey. Meanwhile, we use building characteristics derived from two-dimensional building footprints and machine learning to develop the following neighborhood types: (1) Large-Lot Residential; (2) Small-Lot Residential; (3) Low-Intensity; and (4) High-Intensity neighborhoods. We then use land cover data to estimate the share of buildings that were developed by 2001 and 2019 in each of these neighborhood types. Finally, we pair this with data from the U.S. Census Bureau to study the change in the income and racial composition of municipalities between 2000 and 2020.

Broadly, we find that maximum density is correlated with both building characteristics and neighborhood types. We also find that there is a tradeoff between neighborhood types, particularly Small-Lot Residential and Low-Intensity neighborhoods, which are representative of suburban and exurban areas, respectively. Specifically, an increase in maximum density is associated with an increase in the share of buildings located in Small-Lot Residential neighborhoods but a decrease in the share of buildings located in Low-Intensity neighborhoods. The largest marginal effect occurs at maximum densities of approximately 4–7 units per acre. We also find that maximum density is correlated with cross-sectional measures of income and race, and that the effect is diminished when accounting for neighborhood types. This suggests that neighborhood type mediates some of the relationship between density zoning and residential segregation. Additionally, it suggests that density zoning contributes to a chain of exclusion by restricting the development of Small-Lot Residential neighborhoods along the urban fringe. We find no evidence of a link between changes in maximum density and demographics over time, although we do observe a relationship between changes in building development and demographics between 2000 and 2020. Lastly, our analysis provides insight into the measurement of density zoning via municipal surveys. We find that the National Longitudinal Land Use Survey's (Lo et al. 2019) estimates of maximum units per acre are a more robust municipal-level measure of density zoning than are estimates of minimum lot size from the Wharton Residential Land Use Regulatory Index (Gyourko et al. 2008) because the former measure density zoning in both single-family and multifamily zones and capture more variation in neighborhood types.

The results from this paper point toward two potential avenues for zoning reform. First, counties and townships along the urban fringe could pursue modest increases in allowable density to accommodate new single-family residential development on greenfield sites. This could be achieved, for example, by increasing the maximum density to 4 or more units per acre or by reducing the minimum lot size to less than half acre. Second, cities could pursue more substantial increases in allowable density to promote multifamily and mixed-use development, either in selected areas or citywide. This might occur, for example, by upzoning to allow accessory dwelling units or duplexes on individual lots or by increasing the allowable units per acre or eliminating minimum lot size restrictions to allow redevelopment at higher densities than is typically observed in single-family neighborhoods (e.g., greater than 15 units per acre).

Both of these zoning reforms have the potential to increase the housing stock and to diversify housing options, potentially alleviating concerns about housing affordability and residential segregation, but the first type of zoning reform is likely to have the greatest impact. Our findings suggest that even modest increases in maximum density on greenfield sites along the urban fringe would be accompanied by substantial changes in new development. Both our exploratory analyses and regression analyses indicate that modest changes in density zoning along the urban fringe — for example, increasing maximum density from < 4 units to 4–7 units per acre — would substantially increase the housing stock in these areas by facilitating the development of new Small-Lot Residential development.[10] Given the lower housing values observed in these neighborhoods, this would likely increase the supply of affordable homeownership options within the municipality. Moreover, given the association we observe between changes in the share of Small-Lot Residential neighborhoods and changes in the income and racial composition of municipalities, facilitating new Small-Lot Residential development would likely contribute to decreases in income and racial segregation. Notably, these changes could be achieved with only minimal increases in maximum density and thus without necessitating the construction of multifamily buildings and the opposition that often accompanies it (Scally 2013).

The impact of upzoning in other areas is less clear. For example, we find little evidence to suggest that modifying zoning ordinances to allow for multifamily housing in traditional single-family neighborhoods would have widespread impact. Although our cross-sectional analyses do point to a relationship between maximum density and new development in High-Intensity neighborhoods, the effect is only observed at 15 or 30 units per acre. It is unclear

[10]This could be, in part, an artifact of our measure of new development, which primarily captures greenfield development along the urban fringe. However, as our analysis of the relationship between maximum density and building characteristics shows, the largest marginal changes in building density and distance also occur between < 4 and 4–7 units per acre.

whether allowing such densities in single-family neighborhoods would be politically feasible. Moreover, we find no evidence that changes in maximum density over time are associated with new development in High-Intensity neighborhoods. Thus, at least between 2001 and 2019, upzoning does not appear to have contributed to more High-Intensity development in municipalities across the 50 largest metropolitan areas. This may be attributable, in part, to how we measure new development and zoning change. We measure new development as the conversion from undeveloped land to developed land. Since we do not measure variation in the intensity of development (e.g., low, medium, and high), we likely primarily capture greenfield development rather than redevelopment of existing neighborhoods, such as via the addition of accessory dwelling units or the development of duplexes. In addition, we measure these changes in land cover at the census block level rather than individual buildings. Future research may develop more refined estimates of new development by identifying land cover changes at the building-pixel level.

Similarly, our measure of density zoning has a number of limitations that may limit our ability to capture the potential impacts of upzoning on existing residential neighborhoods. We measure density zoning via maximum density restrictions at the municipality-level rather than at the zone-level. This limitation is not unique to our analysis, however, as it is a challenge experienced by many researchers who have examined the relationship between zoning and segregation across the country (Lens and Monkkonen 2016; Rothwell and Massey 2009). While measures of land use and zoning regulations are not currently available nationwide at the zone-level, recent advancements in data collection, and in particular the compiling of zone-level data via the National Zoning Atlas (Xu et al. 2023), will allow for analyses of the impacts of zoning at the neighborhood or property level. A second limitation regarding our measure of density zoning is that we rely on a coarse measure of zoning change — the reported change in maximum density between the two waves of the NLLUS — which may not capture all increases in allowable density across the jurisdiction.

These findings are accompanied by a couple of additional caveats. First, the neighborhood types that we identified through machine learning represent an oversimplification of the diversity in neighborhoods and housing options that are observed across metropolitan areas in the United States. Although we find clear and intuitive links between density zoning, building characteristics, neighborhood types, and residential segregation, a more detailed set of building characteristics or a more nuanced typology of neighborhoods may reveal additional ways in which density zoning impacts new development and thus residential segregation. Second, we cannot rule out the possibility of confounding factors, as our controls do not capture all relevant measures of the regulatory environment. We do, however, find similar results using minimum lot size and after controlling for a variety of additional land use and zoning regulations (e.g., delays in the permitting process) from the WRLURI (Gyourko et al. 2021). Lastly, we also cannot rule out the possibility of reverse causality, especially in our cross-sectional analyses. For example, municipalities may be more or less likely to adopt strict land use and zoning regulations depending on their preexisting income and racial compositions (Locke and Rissman 2015). This is less of a concern in our analyses estimating the relationship between new development and demographics. In the future, advances in data collection, such as the National Zoning Atlas (Xu et al. 2023), may facilitate more robust empirical approaches (e.g., spatial regression discontinuity designs) that can better address the issue of confounding.

Data Availability

The census block-level and municipal-level data used in this study are available on the Inter-university Consortium for Political and Social Research (ICPSR) at https://www.openicpsr.org/openicpsr/project/208149/version/V1/view.

Acknowledgments

This study was funded by the National Science Foundation, Award # 2048562.

References

Bartik, A., Gupta, A. and Milo, D., The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement (September 14, 2024). Available at SSRN: https://ssrn.com/abstract=4627587 or http://dx.doi.org/10.2139/ssrn.4627587

Been, V., Ellen, I. G., and O'Regan, K. 2019. Supply Skepticism: Housing Supply and Affordability. Housing Policy Debate, 29(1), 25–40. https://doi.org/10.1080/10511482.2018.1476899

Durst, N.J. Safarova, B., Gutierrez Lanza, D., and Xu, C., 2024. Neighborhood Morphology and Residential Segregation in Metropolitan America. Available at SSRN: https://ssrn.com/abstract=4905718

Durst, N.J., Sullivan, E., and Jochem, W.C., 2024. The spatial and social correlates of neighborhood morphology: Evidence from building footprints in five U.S. metropolitan areas. PLoS ONE 19(4): e0299713.

Freemark, Y. 2023. Zoning Change: Upzonings, Downzonings, and Their Impacts on Residential Construction, Housing Costs, and Neighborhood Demographics. Journal of Planning Literature 38(4), 548-570.

Furth, S., Webster, M. (2023). Single-Family Zoning and Race: Evidence From the Twin Cities. Housing Policy Debate, 33(4), 821–843. https://doi.org/10.1080/10511482.2023.2186750

Gabbe, C.J., 2019. Changing Residential Land Use Regulations to Address High Housing Prices: Evidence From Los Angeles. Journal of the American Planning Association 85(2):152–68. doi: 10.1080/01944363.2018.1559078.

Garcia, D., Tucker, J., Schmidt, I. 2020. Single-Family Zoning Reform: An Analysis of SB 1120. Terner Center for Housing Innovation. University of California Berkeley. hhttps://www.habitatca.org/wp-content/uploads/2020/08/Single-Family_Zoning_Reform_An_Analysis_of_SB_1120.pdf

Glaeser, E.L., Ward, B.A., 2009. The causes and consequences of land use regulation: Evidence from Greater Boston. Journal of Urban Economics, 65 (3): 265-278.

Goetz, E., 2021. Democracy, Exclusion, and White Supremacy. Urban Affairs Review 57 (1): 269–83.

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, Vol. 45, Iss. 3, pp. 693-729.

Gyourko, J., Hartley, J.S., 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 (July) 103337.

Jochem W.C., Tatem, A.J., 2021 Tools for mapping multi-scale settlement patterns of building footprints: An introduction to the R package foot. PLoS ONE 16(2): e0247535. doi.org/10.1371/journal.pone.0247535

Kahn, M.E., Vaughn, R., and Zasloff, J., 2010. The housing market effects of discrete land use regulations: Evidence from the California coastal boundary zone. Journal of Housing Economics,19 (4): 269-279.

Kok N., Monkkonen P., Quigley J.M., 2014. Land use regulations and the value of land and housing: An intra-metropolitan analysis. Journal of Urban Economics, 81, 136–148.

Kuhlmann, D. (2021). Upzoning and Single-Family Housing Prices: A (Very) Early Analysis of the Minneapolis 2040 Plan. Journal of the American Planning Association, 87(3), 383–395. https://doi.org/10.1080/01944363.2020.1852101

Lens, M.C., and Monkkonen, P., 2016. Do Strict Land Use Regulations Make Metropolitan Areas More Segregated by Income? Journal of the American Planning Association 82(1):6–21. doi: 10.1080/01944363.2015.1111163.

Locke, C.M., Rissman, A.R., 2015. Factors influencing zoning ordinance adoption in rural and exurban townships. Landscape and Urban Planning 134 (February), 167-176. doi.org/10.1016/j.landurbplan.2014.10.002

Lo, L., Gallagher, M., Pendall, R., Hariharan, A., Davis, C. 2019. National Longitudinal Land Use Survey. The Urban Institute: Washington, DC. Accessible from: https://datacatalog.urban.org/dataset/national-longitudinal-land-use-survey-nllus. Data developed at the Urban Institute, and made available under the ODC-BY 1.0 Attribution License.

Manson, S., Schroeder, J. Van Riper, D., Knowles, K., Kugler, T., Roberts, F., and Ruggles, S., 2022. IPUMS National Historical Geographic Information System: Version 18.0. Minneapolis, MN: IPUMS.. http://doi.org/10.18128/D050.V18.0

Manville, M., Monkkonen, P., and Lens, M., 2020. It's Time to End Single-Family Zoning. Journal of the American Planning Association 86(1):106–12. doi: 10.1080/01944363.2019.1651216.

Lawrimore M. A. Sanchez G.M., Cothron C., Tulbure M.G., BenDor T.K., Meentemeyer R.K. 2024. Creating spatially complete zoning maps using machine learning, Computers, Environment and Urban Systems 11. doi: 10.1016/j.compenvurbsys.2024.102157

Mleczko M. and Desmond M. 2023. Using natural language processing to construct a National Zoning and Land Use Database. Urban Studies, 60(13), 2564-2584. doi:

10.1177/00420980231156352

Pendall, R., 2000. Local Land Use Regulation and the Chain of Exclusion. Journal of the American Planning Association 66(2):125–42. doi: 10.1080/01944360008976094.

Phillips, S., Manville, M., and Lens, M., 2021. Research Roundup: The Effect of Market-Rate Development on Neighborhood Rents. Los Angeles, CA: UCLA Lewis Center. February.

Quigley J.M., Raphael S., 2004. Is housing unaffordable? Why isn't it more affordable? Journal of Economic Perspectives, 18(1), 191–214.

Quigley J.M., Rosenthal L.A., 2005. The effects of land use regulation on the price of housing: What do we know? What can we learn? Cityscape (Washington, DC), 69–137.

Resseger M., 2022. The Impact of Land Use Regulation on Racial Segregation: Evidence from Massachusetts Zoning Borders. Mercatus Research Paper, Available at SSRN: https://ssrn.com/abstract=4244120

Rothstein, R., 2017. The Color of Law: A Forgotten History of how our Government Segregated America. First ed. London;New York;: Liveright Publishing Corporation.

Rothwell, J., and Massey, D.S., 2009. The Effect of Density Zoning on Racial Segregation in U.S. Urban Areas. Urban Affairs Review 44(6):779–806. doi: 10.1177/1078087409334163.

Saiz A., 2010. The geographic determinants of housing supply. The Quarterly Journal of Economics, 125(3), 1253–1296.

Scally, C.P., 2013. The Nuances of NIMBY: Context and Perceptions of Affordable Rental Housing Development. Urban Affairs Review, 49(5), 718-747. https://doi.org/10.1177/1078087412469341

Stacy, C., Davis, C., Freemark, Y. S., Lo, L., MacDonald, G., Zhen, V., and Pendall, R., 2023. Land-Use Reforms and Housing Costs: Does Allowing for Increased Density Lead to Greater Affordability? Urban Studies 004209802311595. doi: 10.1177/00420980231159500.

Timberlake, J.M., and Ignatov, M.D., 2014. Residential Segregation. obo in Sociology. doi: 10.1093/obo/9780199756384-0116

Trounstine, J., 2020. The Geography of Inequality. American Political Science Review 114 (2): 443–55.

Wassmer R., Williams J., 2021. The influence of regulations of residential land prices in United States metropolitan area, Cityscape: A Journal of Policy Development and Research, 23(1), 9-36.

Whittemore, A., 2021. Exclusionary Zoning: Origins, Open Suburbs, and Contemporary Debates. Journal of the American Planning Association 87 (2): 167–80.

Xu, W., Markley, S., Bronin, S.C., and Drogaris, D., 2023. A National Zoning Atlas to Inform Housing Research, Policy, and Public Participation. Cityscape 25, no. 3 (2023): 55-72

Appendix Tables

Appendix Figure 1: Estimates - Neighborhood Type (Binary)

	Neighborhoods	Buildings		
		2001	2019	Post-2001
	(1)	(2)	(3)	(4)
Panel A: Large-Lot Residential				
Maximum Density: 4 - 7	0.0299***	0.0321***	0.0300***	0.0243**
	(0.0106)	(0.0095)	(0.0090)	(0.0099)
Maximum Density: 8 - 15	0.0410***	0.0461***	0.0452***	0.0433***
	(0.0110)	(0.0101)	(0.0097)	(0.0125)
Maximum Density: 16 - 30	0.0347***	0.0517***	0.0502***	0.0313**
	(0.0121)	(0.0115)	(0.0113)	(0.0155)
Maximum Density: > 30	0.0161	0.0468***	0.0478***	0.0377**
	(0.0132)	(0.0127)	(0.0126)	(0.0149)
Panel B: Small-Lot Residential				
Maximum Density: 4 - 7	0.0837***	0.0930***	0.0948***	0.0637***
	(0.0101)	(0.0127)	(0.0131)	(0.0187)
Maximum Density: 8 - 15	0.0828***	0.0987***	0.0959***	0.0350**
	(0.0132)	(0.0190)	(0.0188)	(0.0169)
Maximum Density: 16 - 30	0.1139***	0.1451***	0.1442***	0.0773***
	(0.0149)	(0.0185)	(0.0185)	(0.0251)
Maximum Density: > 30	0.1447***	0.1973***	0.1945***	0.0914***
	(0.0203)	(0.0241)	(0.0241)	(0.0269)
Panel C: Low-Intensity				
Maximum Density: 4 - 7	-0.1099***	-0.1217***	-0.1215***	-0.0956***
	(0.0127)	(0.0145)	(0.0146)	(0.0201)
Maximum Density: 8 - 15	-0.1350***	-0.1473***	-0.1441***	-0.0874***
	(0.0140)	(0.0188)	(0.0188)	(0.0175)
Maximum Density: 16 - 30	-0.1669***	-0.2044***	-0.2022***	-0.1222***
	(0.0142)	(0.0187)	(0.0189)	(0.0244)
Maximum Density: > 30	-0.1920***	-0.2572***	-0.2554***	-0.1653***
	(0.0162)	(0.0220)	(0.0223)	(0.0318)
Panel D: High-Intensity				
Maximum Density: 4 - 7	-0.0037	-0.0035	-0.0033	0.0076
	(0.0057)	(0.0037)	(0.0038)	(0.0075)
Maximum Density: 8 - 15	0.0112**	0.0025	0.0031	0.0091
	(0.0055)	(0.0036)	(0.0036)	(0.0057)
Maximum Density: 16 - 30	0.0183***	0.0076*	0.0078*	0.0136
	(0.0060)	(0.0041)	(0.0041)	(0.0082)
Maximum Density: > 30	0.0313***	0.0131***	0.0131***	0.0361***
	(0.0078)	(0.0046)	(0.0046)	(0.0090)
Controls: Land Use and Zoning	✓	✓	✓	✓
Fixed Effect: Municipality Type	✓	✓	✓	✓
Fixed Effect: Metropolitan Area	✓	✓	✓	✓
Observations	1,620	1,620	1,620	1,614

*p<0.1; **p<0.05; ***p<0.01

Appendix Figure 2: Estimates - Demographics (Binary)

	Year		
	2000	2020	2000-2020
Panel A: Income			
Maximum Density: 4 - 7	212.7	643.1	319.2
	(2,457.4)	(2,855.2)	(1,261.5)
Maximum Density: 8 - 15	2,077.4	2,890.4	629.8
	(2,558.9)	(2,784.8)	(1,208.0)
Maximum Density: 16 - 30	-3,895.7	-1,697.9	1,735.5
	(2,992.4)	(3,127.4)	(1,513.8)
Maximum Density: > 30	-9,920.2***	-6,088.8*	3,020.3
	(3,472.4)	(3,326.9)	(1,636.8)
Panel B: White			
Maximum Density: 4 - 7	0.0093	-0.0174*	-0.0229***
	(0.0085)	(0.0087)	(0.0066)
Maximum Density: 8 - 15	0.0018	-0.0322**	-0.0319***
	(0.0101)	(0.0126)	(0.0055)
Maximum Density: 16 - 30	-0.0272	-0.0582***	-0.0267***
	(0.0169)	(0.0197)	(0.0084)
Maximum Density: > 30	-0.0896***	-0.0988***	-0.0041
	(0.0205)	(0.0222)	(0.0085)
Controls: Land Use and Zoning	✓	✓	✓
Fixed Effect: Municipality Type	✓	✓	✓
Fixed Effect: Metropolitan Area	✓	✓	✓

*p<0.1; **p<0.05; ***p<0.01

Appendix Figures

Appendix Figure 1: Estimates - Neighborhood Type by Municipality Type (2001)

(a) Large-Lot Residential (2001)

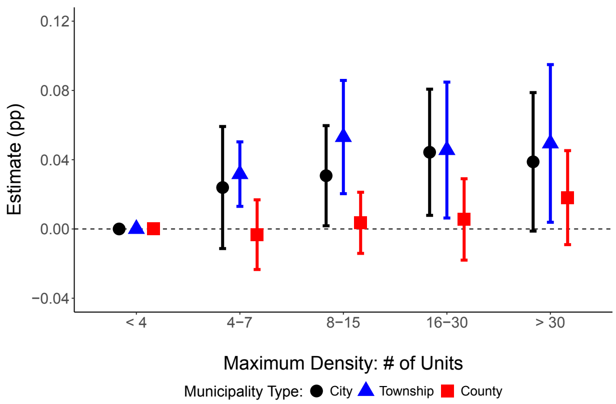

(b) Small-Lot Residential (2001)

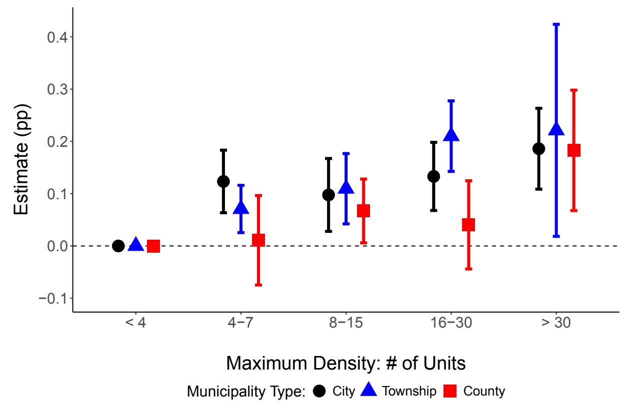

(c) Low-Intensity (2001)

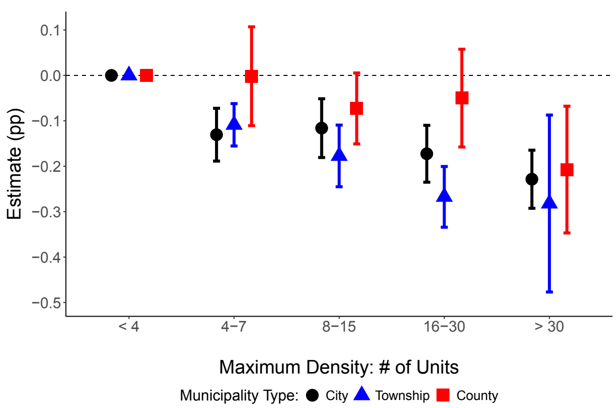

(d) High-Intensity (2001)

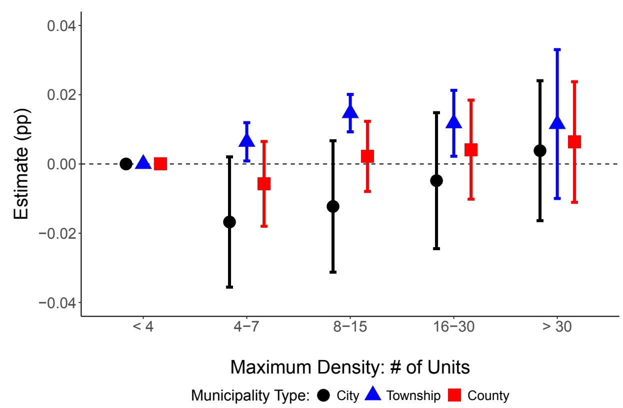

Notes: The figure presents the estimates of the binary measure of maximum density for the neighborhood types, disaggregated by municipality type. The estimates are based on buildings developed by 2001. The base category is < 4 units per acre. Standard errors are clustered by metropolitan area. Panels (a), (b), (c), and (d) document Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively.

Appendix Figure 2: Estimates - Neighborhood Type by Municipality Type (Post-2001)

(a) Large-Lot Residential (Post-2001)

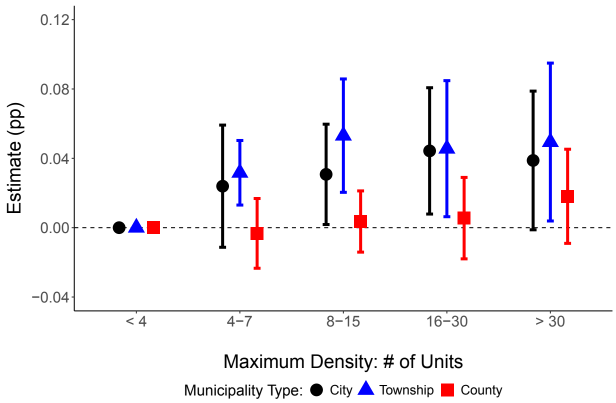

(b) Small-Lot Residential (Post-2001)

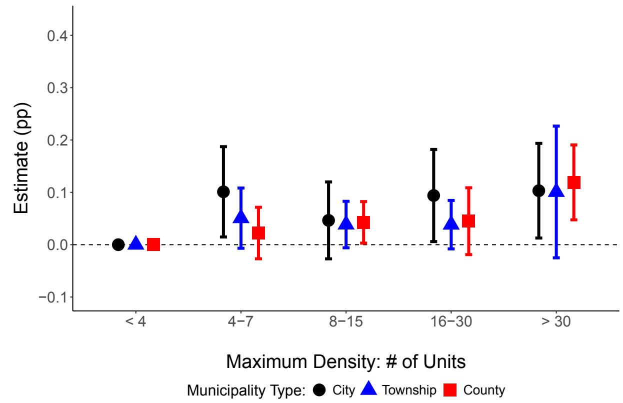

(c) Low-Intensity (Post-2001)

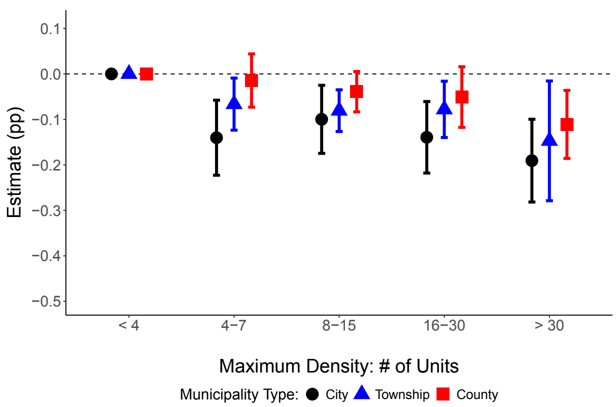

(d) High-Intensity (Post-2001)

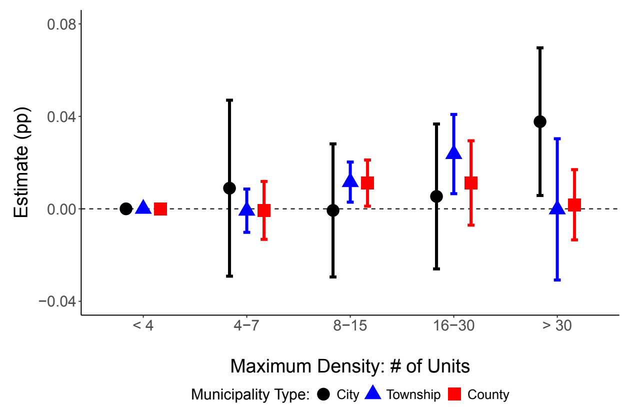

Notes: The figure presents the estimates of the binary measure of maximum density for the neighborhood types, disaggregated by municipality type. The estimates are based on buildings developed post-2001. The base category is < 4 units per acre. Standard errors are clustered by metropolitan area. Panels (a), (b), (c), and (d) document Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively.

Appendix Figure 3: Estimates - Neighborhood Type by Census Region (2001)

(a) Large-Lot Residential (2001)

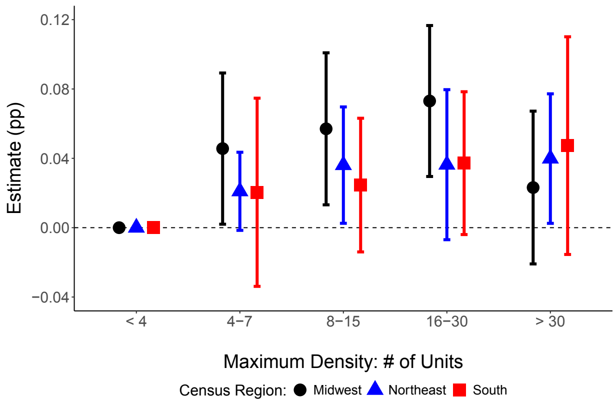

(b) Small-Lot Residential (2001)

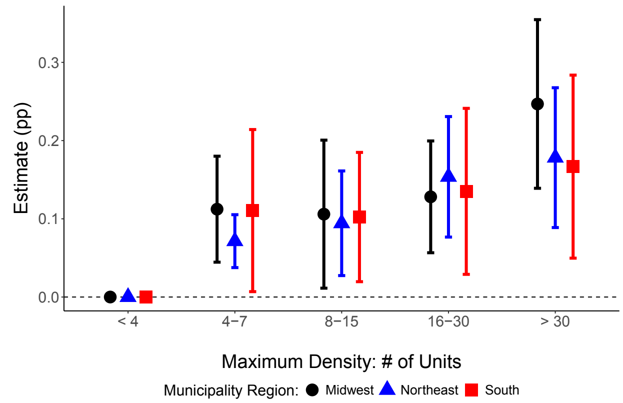

(c) Low-Intensity (2001)

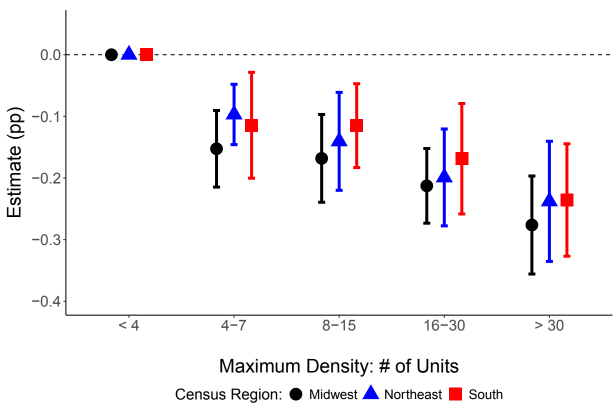

(d) High-Intensity (2001)

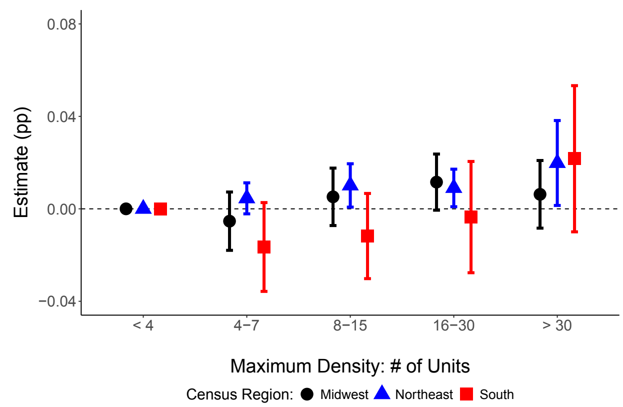

Notes: The figure presents the estimates of the binary measure of maximum density for the neighborhood types, disaggregated by census region. The estimates are based on buildings developed by 2001. The base category is < 4 units per acre. Standard errors are clustered by metropolitan area. Panels (a), (b), (c), and (d) document Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively.

Appendix Figure 4: Estimates - Neighborhood Type by Census Region (Post-2001)

(a) Large-Lot Residential (Post-2001)

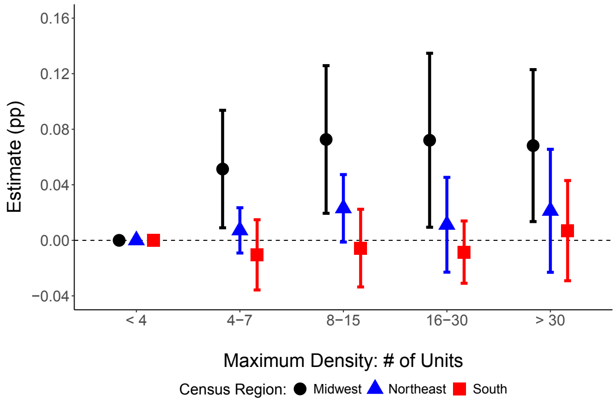

(b) Small-Lot Residential (Post-2001)

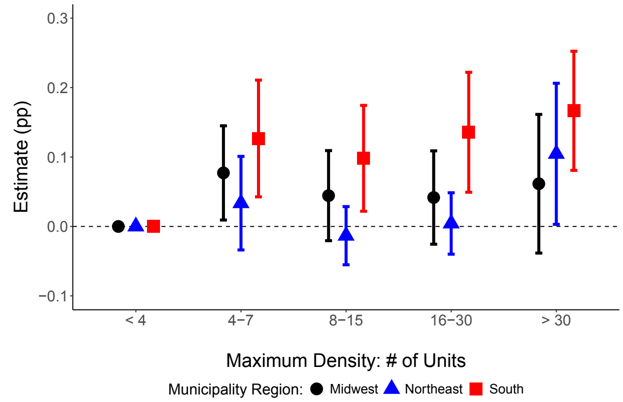

(c) Low-Intensity (Post-2001)

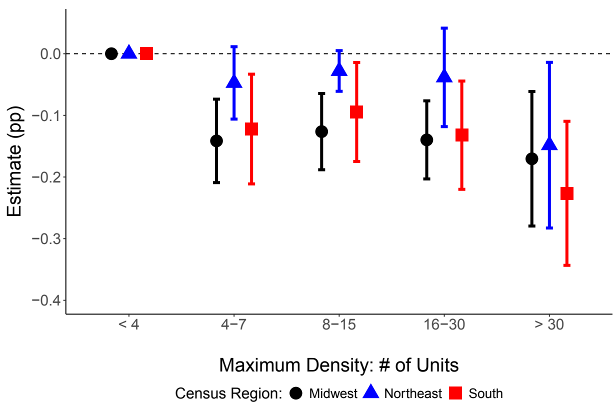

(d) High-Intensity (Post-2001)

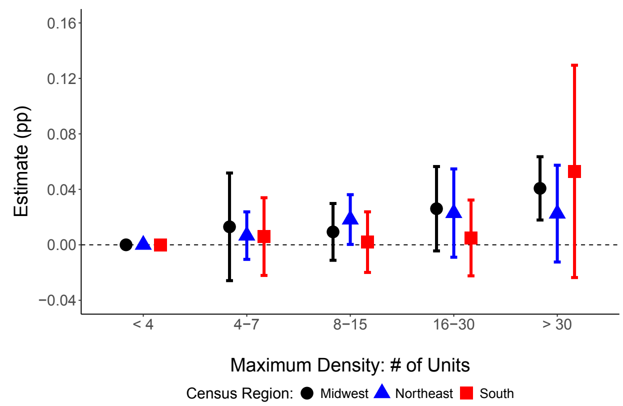

Notes: The figure presents the estimates of the binary measure of maximum density for the neighborhood types, disaggregated by census region. The estimates are based on buildings developed post-2001. The base category is < 4 units per acre. Standard errors are clustered by metropolitan area. Panels (a), (b), (c), and (d) document Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively.

Appendix Figure 5: Estimates - Demographics by Municipality Type

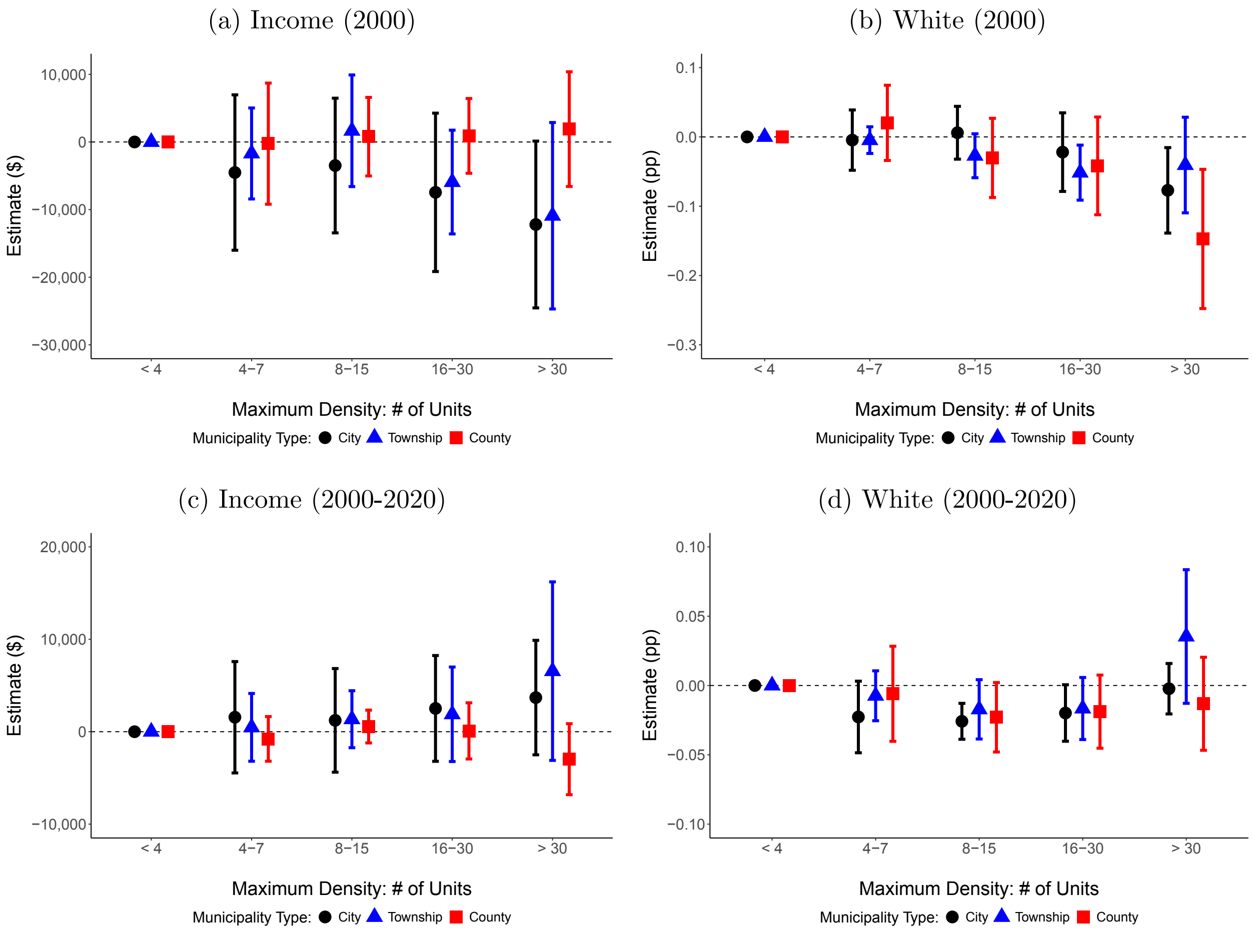

Notes: The figure presents the estimates of the binary measure of maximum density for the demographics, disaggregated by municipality type. The base category is < 4 units per acre. Standard errors are clustered by metropolitan area. Panels (a) and (b) document income and the share of residents who are non-Hispanic White in 2000, respectively. Meanwhile, panels (c) and (d) document the change in income and the share of residents who are non-Hispanic White between 2000 and 2020, respectively.

Appendix Figure 6: Estimates - Demographics by Census Region

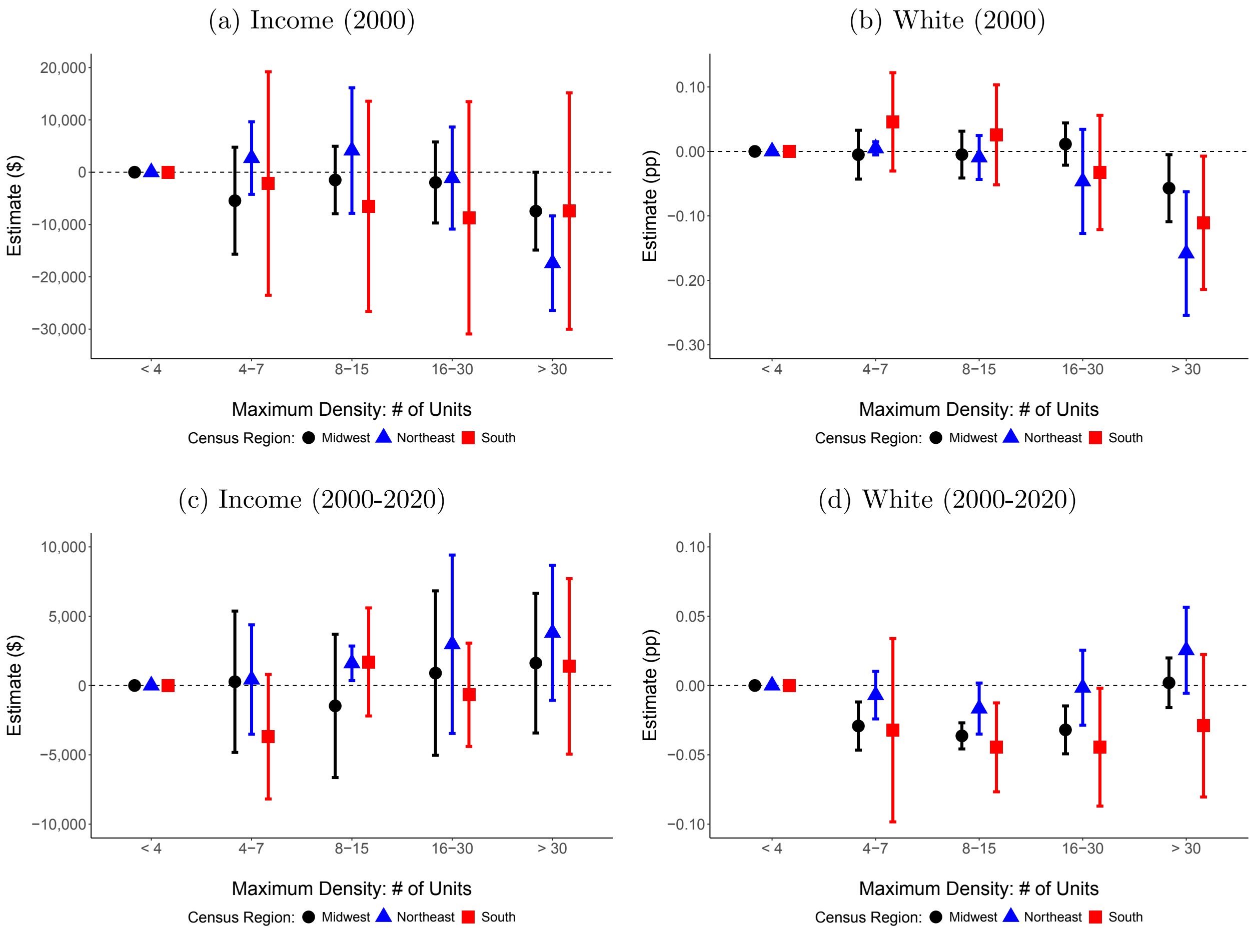

Notes: The figure presents the estimates of the binary measure of maximum density for the demographics, disaggregated by census region. We do not present the estimates for municipalities in the West because they skew the x-axis. The base category is < 4 units per acre. Standard errors are clustered by metropolitan area. Panels (a) and (b) document income and the share of residents who are non-Hispanic White in 2000, respectively. Meanwhile, panels (c) and (d) document the change in income and the share of residents who are non-Hispanic White between 2000 and 2020, respectively.

Appendix A

Measurement of Neighborhood Type

Microsoft's Building Footprints dataset is a publicly available source of building footprints, containing approximately 130 million computer-generated building footprints from across the United States. The dataset, derived using computer vision algorithms applied to satellite imagery, is available individually for each state and the District of Columbia. Microsoft leveraged a two-stage process involving Semantic Segmentation and Polygonization for building extraction. Approximately half of the building footprints were created using 2019 and 2020 imagery; the remainder rely on imagery from 2012, on average. The precision of the data is measured at 98.5% with a recall of 92.4%. The dataset is licensed under the Open Data Commons Open Database License (ODbL) and is free to download and use. See https://github.com/microsoft/USBuildingFootprints for additional details. We acquired the data used for this study in March of 2022.

We use measures of neighborhood type developed by Durst et al. (2024) across the largest 100 metropolitan areas. To do so, we use principal components analysis (PCA). First, we calculate a series of building characteristics at the census block level using the "foot" package in R (Jochem and Tatem, 2021). We remove buildings from the sample if their area is less than 25 square meters, eliminating sheds and other small buildings. Before using PCA, we normalize the distribution of each building characteristic, then standardize and mean center them. To reduce noise in and the dimensionality of the dataset, we select the first eight factors, representing approximately 95% of the variation in the building characteristics, for use in the unsupervised classification. We test both K-Means and Gaussian Mixture Models (GMM) with K=2 through K=10 classes. Since the K-Means and GMM methods produce similar results, we present the results from K-Means given the ubiquity of the method in the literature. Across both methods, the model with two classes produce the highest silhouette scores (average of approximately .25), with silhouette scores leveling off around 5 clusters.

After reviewing descriptive statistics for the building characteristics for clusters with differing numbers of classes, we select the results with four classes because they yield a conceptually distinct set of classes. Appendix Table A.1 displays the mean for each building characteristic by class (i.e., neighborhood type).

Appendix Table A.1: Building Characteristics by Neighborhood Type

	Neighborhood Type			
Building Characteristic	Large-Lot Residential	Small-Lot Residential	Low-Intensity	High-Intensity
Angle of Orientation (Entropy)	0.31	-0.19	-0.51	0.57
Area (Maximum)	-0.22	-0.47	0.60	1.61
Area (Median)	0.11	-0.39	-0.02	2.18
Area (IQR)	-0.29	-0.63	0.50	1.97
Compactness Index (Median)	-0.65	0.79	-0.08	-1.52
Compactness Index (IQR)	0.00	-0.68	0.59	1.07
Distance (Median)	-0.26	-0.71	0.45	0.96
Distance (IQR)	-0.41	-0.66	0.76	0.43
Footprints (Count per block)	-0.23	0.26	0.61	-0.77
Footprints (Density per square kilometer)	0.40	0.79	-0.63	-0.58
Length-equivalent-Width Ratio (Median)	0.43	-0.66	-0.02	1.60
Length-equivalent-Width Ratio (IQR)	-0.02	-0.67	0.47	1.16
Perimeter (Median)	0.16	-0.46	-0.02	2.27
Perimeter (IQR)	-0.31	-0.64	0.57	1.72
Shape Index (Median)	-0.67	0.82	-0.11	-1.34
Shape Index (IQR)	-0.06	-0.63	0.60	0.80
Observations	409,916	621,143	343,194	174,116
(%)	(0.27)	(0.40)	(0.22)	(0.11)

Notes: The table presents the mean of the standardized building characteristics that are used in the principal components analysis and unsupervised classification. Angle of orientation is measured by entropy, area is measured by square meters, distance is measured by meters to nearest building, and perimeter is measured by meters.

Appendix B

Land Use and Zoning Regulations

The land use and zoning regulations data comes from the 2003 National Longitudinal Land Use Survey, available here: https://datacatalog.urban.org/dataset/national-longitudinal-land-use-survey-nllus. The survey was administered to municipalities from the top 50 Combined Metropolitan Statistical Areas. While the survey was often reserved for municipalities with a population greater than 10,000 residents, it was also administered to smaller municipalities on a case-by-case basis. The following survey questions are used to construct the variables used in the analysis:

- Affordable Public Facilities Ordinance: "Does your county have an adequate public facilities ordinance that makes development permission contingent on the levels of offsite public services?"
- Affordable Housing Funding Mechanism: "Does your county have a local affordable housing funding mechanism (e.g., housing trust fund)?"
- Affordable Housing - Government: "Does your county have government assisted affordable housing?"
- Affordable Housing - Private: "Does your county use any incentives or requirements to encourage private-sector builders to develop affordable housing?"
- Growth Limit Tool: "Urban growth boundary in place since ____ (year)"
    - Converted to binary variable: "Jurisdiction has urban growth boundary in place."
- Impact Fee: "Does your jurisdiction charge impact fees?"
- Maximum Density: "What is the theoretical maximum number of dwelling units that may be constructed per net acre in areas to which your county zoning ordinance applies, in areas zoned in the highest residential density category?"

- Responses are in the following ranges:
  - < 4 units
  - 4–7 units
  - 8–15 units
  - 16–30 units
  - > 30 units
- Mobile Home: "Does your county permit the placement of new mobile homes, either on a single lot or in a mobile home park?"
- Urban Service: "Urban service area / urban service boundary in place since ____ (year)"
  - Converted to binary variable: "Jurisdiction has urban service area / urban service boundary."
- Voting Referendum: "Does your county require a popular vote of the county's residents as a precondition to rezoning?"

Appendix C

NNLUS and WRLURI Comparison

We use a measure of maximum density restrictions from the National Longitudinal Land Use Survey (NLLUS) to document density zoning in municipalities across the United States. This differs from other measures of density zoning, like minimum lot size requirements, which are examined in other national surveys such as the Wharton Residential Land Use Regulatory Index (WRLURI) (Gyourko et al., 2008). In this section, we compare the measures of density zoning from these two surveys.

We conduct a similar process to develop municipal-level estimates for the set of municipalities that responded to the WRLURI. We focus on muncipalities in the 100 largest metropolitan areas, where we developed estimates of neighborhood type at the census block level. As described earlier, we calculate the share of buildings developed by 2001 and the share of buildings developed post-2001 in each neighborhood type at the municipal level. We then present descriptive statistics on variation in neighborhood type and replicate our regression analyses using the WRLURI measure of minimum lot size. The findings suggest that the NLLUS measure of maximum density captures more of the variation in density zoning and its impact on neighborhood type.

Both the NLLUS and WRLURI measure density zoning using categorical responses from municipal officials. The NLLUS relies on surveys distributed to public officials in cities, counties, and townships across the 50 largest combined metropolitan statistical areas in 2003, while the WRLURI was distributed to public officials in a representative sample of cities and townships (but not counties) nationwide in 2006. The NLLUS asks respondents about the maximum number of dwelling units that are permitted per acre anywhere within a municipality, ranging from < 4 units per acre to > 30 units per acre. Meanwhile, the WRLURI asks respondents about the minimum lot size with categories ranging from > 2 acres to "no restriction" (see Appendix Table C.1). The NLLUS is a more comprehensive measure of density zoning because restrictions on units per acre are applicable to both single-family and multifamily housing, whereas lot size is only applicable to single-family housing. In addition, the two measures capture different ends of the distribution of relative allowable density. For example, the lowest density category in the NLLUS (< 4 units per acre) is equivalent to a minimum lot size of .33 acres or more. There are four categories on the WRLURI that cover some or all of this density range: < 0.5 acres through > 2 acres. The inverse is also true, as the two highest density categories in the WRLURI (< 0.5 acres and no restriction) are equivalent to all five categories in the NLLUS.

To examine whether the NLLUS measure is more successful at capturing variation in neighborhood type, we estimate the share of buildings developed by 2001 and post-2001 in each neighborhood type (see Appendix Table C.1 and C.2, respectively). The results suggest that both measures of density zoning successfully align with variation in the distribution of buildings across neighborhood types. For example, increasing measures of relative density either on the WRLURI (i.e., decreasing the minimum lot size) or the NLLUS (i.e., increasing the maximum number of dwelling units permitted per acre) is associated with increases in the share of buildings in Large-Lot Residential, Small-Lot Residential, and High-Intensity neighborhoods and decreases in the share of buildings in Low-Intensity neighborhoods. This is true for both buildings developed by 2001 and buildings developed post-2001. The measure of maximum density provided in the NLLUS, however, captures more variation in density zoning, as illustrated in the max-min difference. This is particularly the case for Small-Lot Residential and Low-Intensity neighborhoods.

As a robustness check, we replicate our main regression analysis, replacing maximum density restrictions with minimum lot size requirements. In these regressions, we control for other indices that are reported in the WRLURI: Local Political Pressure Index (LPPI), State Political Involvement Index (SPII), State Court Involvement Index (SCII), Local Zoning Approval Index (LZAI), Local Project Approval Index (LPAI), Local Assembly Index (LAI), Supply Restrictions Index (SRI), Open Space Index (OSI), and Approval

Delay Index (ADI). These estimates are presented in Appendix Figure C.1. The base category is the lowest density category ($>$ 2 acres). The relationship between minimum lot size and the neighborhood types is similar to that observed with maximum density. Here, an increase in density (i.e. decrease in minimum lot size) is associated with a greater share of development in Large-Lot Residential, Small-Lot Residential, and High-Intensity neighborhoods and a lower share of development in Low-Intensity neighborhoods. This effect, however, is observable primarily at the highest density categories ($<$ 0.5 acres and no restriction). There are little to no observable differences in neighborhood type between municipalities with the three lowest density categories. These results suggest that maximum density restrictions from the NLLUS capture more of the variation in neighborhood types than minimum lot size requirements from the WRLURI. Because of this, we argue that maximum density restrictions are a better measure of density zoning than minimum lot size requirements for our analysis.

Appendix Table C.1: Share by Neighborhood Type (2001)

Survey	Measurement	Large-Lot Residential	Small-Lot Residential	Low-Intensity	High-Intensity
WRLURI	> 2 Acres	0.144	0.282	0.529	0.045
	1 - 2 Acres	0.125	0.226	0.605	0.044
	0.5 - 1 Acres	0.188	0.303	0.457	0.052
	< 0.5 Acres	0.196	0.388	0.367	0.049
	No Restriction	0.208	0.402	0.325	0.065
	Max-Min Difference	0.064	0.120	-0.204	0.020
NLLUS	< 4 Units	0.130	0.193	0.649	0.029
	4 - 7 Units	0.147	0.344	0.480	0.029
	8 - 15 Units	0.182	0.375	0.407	0.037
	16 - 30 Units	0.226	0.423	0.305	0.047
	> 30 Units	0.216	0.484	0.245	0.055
	Max-Min Difference	0.086	0.291	-0.404	0.026

Appendix Table C.2: Share by Neighborhood Type (Post-2001)

Survey	Measurement	Large-Lot Residential	Small-Lot Residential	Low-Intensity	High-Intensity
WRLURI	> 2 Acres	0.085	0.194	0.655	0.066
	1 - 2 Acres	0.071	0.156	0.709	0.063
	0.5 - 1 Acres	0.102	0.200	0.620	0.077
	< 0.5 Acres	0.121	0.278	0.528	0.073
	No Restriction	0.110	0.262	0.514	0.115
	Max-Min Difference	0.025	0.068	-0.141	0.049
NLLUS	< 4 Units	0.081	0.134	0.746	0.039
	4 - 7 Units	0.102	0.266	0.576	0.056
	8 - 15 Units	0.140	0.276	0.523	0.061
	16 - 30 Units	0.143	0.332	0.449	0.076
	> 30 Units	0.133	0.356	0.404	0.107
	Max-Min Difference	0.052	0.222	-0.342	0.068

Notes: These figures present the share of neighborhood types by minimum lot size from the WRLURI and maximum density from the NLLUS. Appendix Table C.1 is based on buildings developed by 2001 and Appendix Table C.2 is based on building developed post-2001.

Appendix Figure C.1: Estimates - Neighborhood Type by Minimum Lot Size

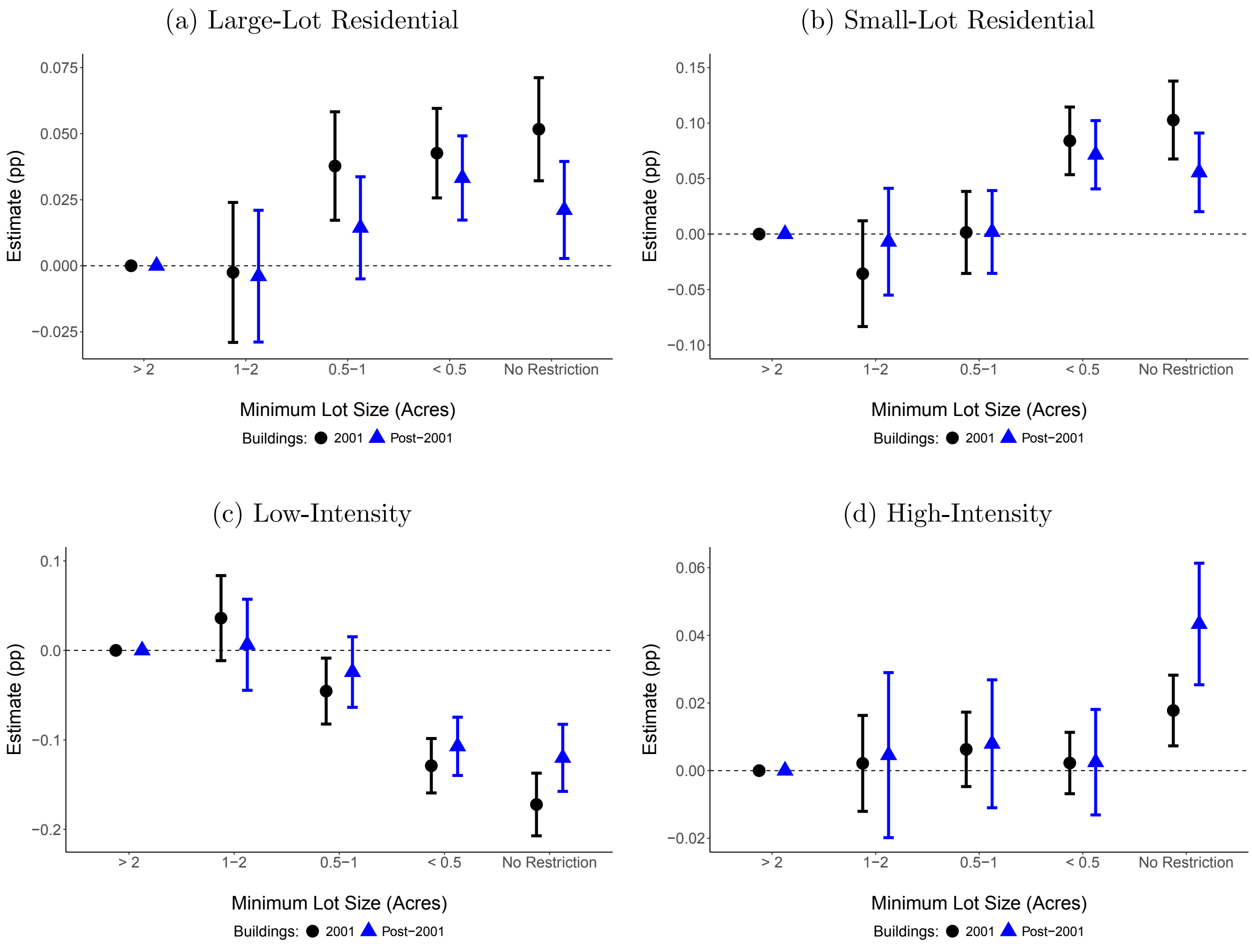

Notes: The figure presents the estimates of the binary measure of minimum lot size from the WRLURI for the neighborhood types. The base category is > 2 acres. Standard errors are clustered by metropolitan area. Panels (a), (b), (c), and (d) document Large-Lot Residential, Small-Lot Residential, Low-Intensity, and High-Intensity neighborhoods, respectively.

Appendix D

Case Study: Chapel Hill, North Carolina

To illustrate the alignment between our neighborhood types and density zoning at the level of individual zones, we present a brief case study, using Chapel Hill, North Carolina as an example. Here, we combine neighborhood types with data on density restrictions from local zoning maps. We select Chapel Hill for multiple reasons. First, Chapel Hill is a relatively small city. This makes it easy to infer the relationship between zoning and neighborhood types through visual comparison. This would be more difficult to do in a larger city with a broader geography. Second, the population in Chapel Hill has grown considerably since 1950, as illustrated by data from the U.S. Census Bureau and the North Carolina Office of State Planning, increasing from about 10,000 to 62,000 residents in 2020. This reduces the likelihood that extant patterns of neighborhood type are attributable to older land use and zoning regulations that may no longer be enforced. Third, we suspect that Chapel Hill has higher rates of High-Intensity neighborhoods than comparably sized cities because of development around the University of North Carolina, a large public university. For these reasons, Chapel Hill provides a good illustrative case of the relationship between zoning and neighborhood type.

Appendix Figure D.1 displays two maps of Chapel Hill. Panels (a) and (b) document the neighborhood types and zoning, respectively. In Chapel Hill, there are 6 residential zones (R-1 to R-6) as well as non-residential zones. In general, there is broad alignment between the distribution of distinct residential zones identified within the Chapel Hill zoning ordinance and the distribution of our four types of neighborhood. To illustrate, Appendix Table D.1 shows a cross-tabulation by zone and neighborhood type. Along the top, we show the maximum density allowed in each zone, which ranges from 3 dwelling units per acre in R-1 zones to 15 units per acre in R-6 zones (non-residential zones, labeled "Other", do not have a specified maximum number of units). Large-Lot Residential (52%) and Low-Intensity neighborhoods (55%) are most frequently located in the R-1 zone, where a maximum of 3 units per acre are permitted. This aligns with our measures of density derived from building footprints (approximately 1 and 3 buildings per acre in Low-Intensity and Large-Lot Residential neighborhood, respectively; see Table 1). Meanwhile, substantial shares of Small-Lot Residential neighborhoods are located in R-2 and R-3 zones (47% and 16%, respectively), where a maximum of 4 and 7 units are permitted per acre. Lastly, High-Intensity neighborhoods are primarily sited in R-3 to R-6 zones (29%) that permit 7–15 units per acre and in non-residential zones (62%). These results suggest that our neighborhood typology maps intuitively onto zone-level density restrictions.

Appendix Figure D.1: Case Study - Neighborhood Types and Zoning

(a) Neighborhood Types

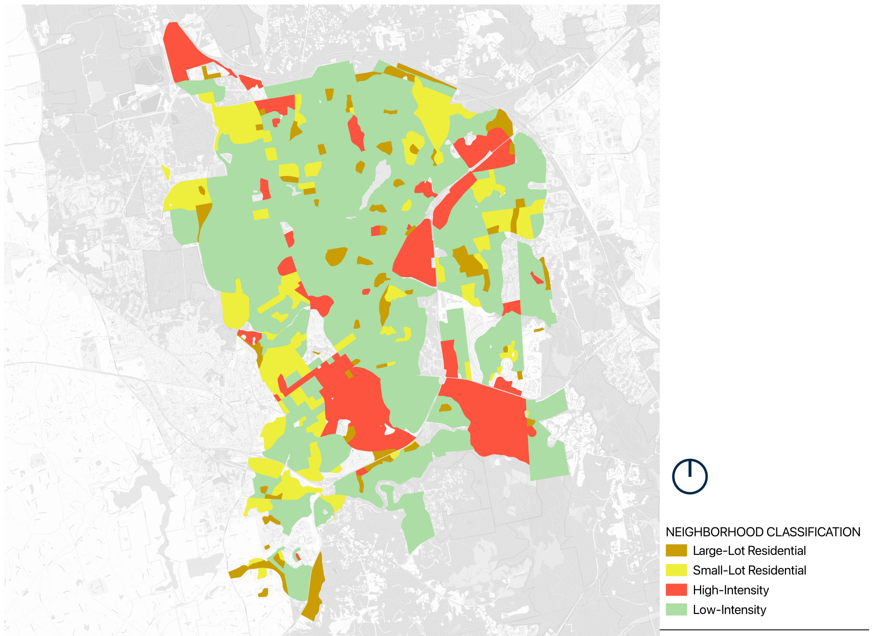

(b) Zoning

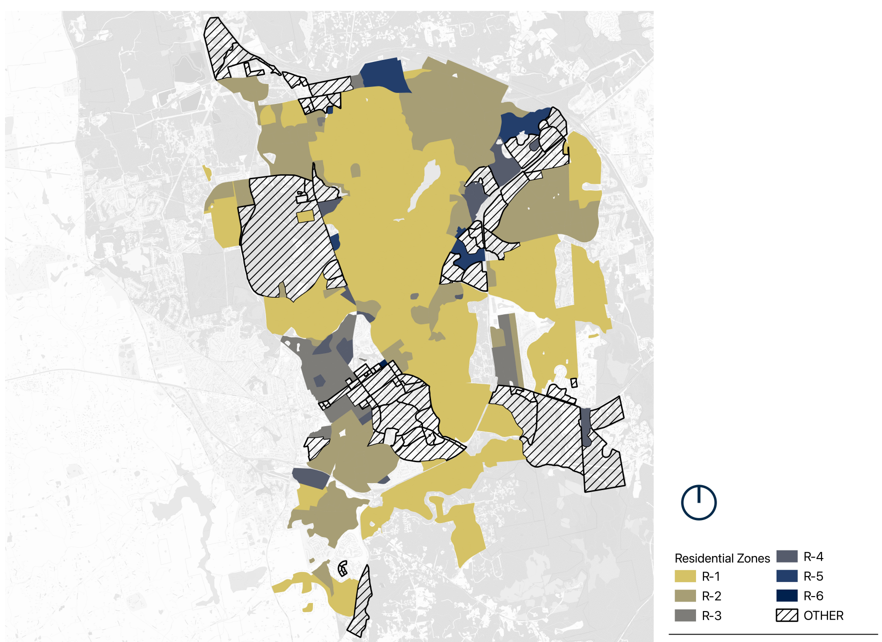

Notes: The figure presents maps for the case study of Chapel Hill. Panel (a) maps the neighborhood types while panel (b) maps the residential zones. The "Other" zone represents non-residential zones.

Appendix Table D.1: Case Study - Neighborhood Types and Zoning

	Zone Type							
	R-1	R-2	R-3	R-4	R-5	R-6	Other	Total
Maximum Units	3	4	7	10	15	15	–	–
Neighborhood Type								
Large-Lot Residential	44	26	3	2	1	0	9	85
	(0.52)	(0.31)	(0.04)	(0.02)	(0.01)	(0.00)	(0.10)	(1.00)
Small-Lot Residential	42	70	24	8	0	0	4	148
	(0.28)	(0.47)	(0.16)	(0.05)	(0.00)	(0.00)	(0.03)	(1.00)
Low-Intensity	83	35	8	7	2	0	15	150
	(0.55)	(0.23)	(0.05)	(0.05)	(0.01)	(0.00)	(0.01)	(1.00)
High-Intensity	5	1	9	6	5	1	44	71
	(0.07)	(0.01)	(0.12)	(0.08)	(0.07)	(0.01)	(0.62)	(1.00)
Total	174	132	44	23	8	1	72	454
	(0.38)	(0.29)	(0.10)	(0.05)	(0.02)	(0.00)	(0.16)	(1.00)

Notes: The table presents the marginal distribution of residential zones in Chapel Hill, including the count and percentage. The row-wise percentage is listed in parentheses. The “Other” zone represents non-residential zones.