Measuring Winners and Losers from Increasing Housing Supply

Vincent Rollet
MIT

Laura Weiwu
UC Berkeley

October 2025
Preliminary and Incomplete

Abstract

Local governments often restrict the construction of new housing through zoning and permitting hurdles. This resistance suggests that increasing the housing stock in an area may harm some of its existing residents. Understanding these losses is key to facilitating new housing construction. In this paper, we aim to measure who gains and who loses from a local increase in the housing supply, and why. We build a model of household location choices over time, allowing us to measure the effect of construction on amenities, prices, and local taxes. To estimate this model, we employ detailed data describing individual housing units, households and their migration patterns, local public finances, and zoning regulations.

JEL Classifications: R21, R23, R52, H71
Keywords: Zoning, housing, local public goods, housing demand

We thank David Atkin, David Autor, Arnaud Costinot, Dave Donaldson, and Tobias Salz, as well as seminar participants at MIT, Stanford, and UC Berkeley for helpful discussions and feedback. Teddy Maginn and Sara Shemali provided outstanding research assistance. This work benefited from the generous support of the George and Obie Shultz Fund, the C. Lowell Harriss Dissertation Fellowship, and the IHS. Author contacts: vrollet@mit.edu and lauraww@berkeley.edu.

1 Introduction

Many urban areas around the world face a mounting affordability crisis: housing production has stalled, and rents in most large cities have risen far faster than overall inflation. A growing body of empirical research points to binding land-use regulations—local rules that dictate what can be built where—as the chief culprit behind these supply shortfalls (Glaeser and Gyourko, 2018). The stringency of these regulations suggests that some incumbent residents expect to lose from an expanded housing stock, even when aggregate welfare would rise. In this view, the key constraint on construction is political, not technological or financial. This paper aims to measure who gains and loses from additional housing construction, and why. Such measurement is key to understanding the political economy of housing supply, and can help design politically feasible reforms of land use regulations.

New construction can affect households through several channels, described in Section 2. First, increases in housing supply will usually lead to a decrease in home prices and rents (Asquith, Mast and Reed, 2023). This is detrimental to homeowners but benefits renters. Second, new housing units can attract different types of households from those already living nearby. Existing residents may welcome or dislike these demographic shifts (Diamond and McQuade, 2019). Third, new construction can impose fiscal externalities on incumbent residents. Municipalities set a local property tax rate to finance local public goods, notably schools, while maintaining a balanced budget (Fischel, 2002). If new housing units are less valuable than existing ones, the revenue they generate may fall short of the cost of the additional public services they require. To close this gap, local governments may be forced to raise the property tax rate for all households. Fourth, new construction increases local density, which some households may dislike (Gyourko and McCulloch, 2024). These pecuniary and non-pecuniary externalities imposed by new construction explain the strict rules that municipalities enact to restrict new construction. Using detailed zoning maps from Zoneomics, covering 9,764 zoning jurisdictions across the United States and 75% of the country's population, we show that zoning favors the construction of large, single-family housing units on large plots of land—the types of housing units that impose the lowest costs on incumbent residents.

Section 3 introduces a simple framework for assessing the welfare effects of new construction. Adding housing in one location alters amenities (through demographics and density), prices, and taxes not only there but across all locations. An approximation of the impact of these shocks on a household's welfare can be decomposed into four components: an amenity effect, a price effect, a fiscal effect, and a relocation effect. Quantifying these channels requires information on (i) households' initial locations and real estate holdings; (ii) their migration probabilities across destinations over time; and (iii) how new construction shifts amenities, prices, taxes, and the likelihood of relocating in all locations.

While some of these variables can be directly observed in the data (such as households' initial locations and migration probabilities), measuring how increasing the housing stock in one location shifts prices, amenities, and tax burdens everywhere requires a model. We develop a model of household location choice in which, in each period, households select a housing unit to occupy. Housing units can be either rented or owned, and they differ in their location and characteristics. These attributes are valued heterogeneously by households. In particular, households vary in their preferences for neighborhood demographics and density, which can evolve over time as households relocate. Migration is costly, especially over longer distances and for older households. Municipalities fund local public goods through a property tax. Finally, the price of housing units is determined by market clearing.

We apply this framework to study the effects of housing construction in New Jersey, which offers an attractive case study. In this state, all land is fully incorporated into municipalities with stable boundaries, most local public goods are financed through property taxes, and school districts typically coincide with municipal borders. Our analysis requires detailed data (described in Section 4) on local housing supplies, household location choices, and municipalities' tax policies. Assembling data from commercial and publicly available sources (CoreLogic, Smarty, and the New Jersey Geographic Information Network), we build a dataset of all residential units in New Jersey, along with their characteristics. We associate these housing units with the households living in them, and track the movements of these households over time using data from Infutor. We measure the demographic characteristics of these households using data from the Census and ACS, and by matching their mortgages to data collected under the Home Mortgage Disclosure Act. We measure the value of housing units using listings and transactions data (again provided by CoreLogic). We finally gather data on local governments' finances from the Census of Governments and New Jersey's Department of the Treasury.

The model can be estimated using a revealed preferences approach similar to that of Bayer, Ferreira and McMillan (2007). Different types of households choose different locations to live in, revealing their valuation of different neighborhood and housing unit attributes. Some of these attributes (prices and local demographics) are endogenous, and measuring how households value them requires an instrumental variables strategy. For identification, we plan to draw on BLP instruments for the valuation of prices and a boundary design for the valuation of neighborhood demographics and density. Comparing the value of similar homes on opposing sides of a municipal border is informative on households' valuation of municipal characteristics, and comparing the value of homes located close to a municipal boundary to those located further away helps identify households' valuation of their immediate neighborhood. We find that households strongly value both the demographics of their immediate neighborhood and those of the broader municipality in which they live.

The estimated model will allow us to measure the effects of building an additional housing unit of a given type (e.g., a two-bedroom apartment) in a given location (e.g., on a specific street in Princeton) on the prices, neighborhood demographics, and property taxes of all housing units, and, in turn, on the welfare of any household. These welfare changes can be decomposed in a supply effect, an amenities effect, and a fiscal effect, allowing us to understand what drives the distributional effects from new construction.

To validate our model, we plan to leverage event study estimates of the effects of building a new multifamily building on the value of nearby preexisting buildings. In our data, we identify events in which a new multifamily building was completed and implement a staggered difference-in-differences design comparing areas that experienced construction earlier vs. later. This analysis shows that multifamily construction reduces the value of nearby single-family homes, but only in higher-income neighborhoods. We can simulate analogous construction events within our model and compare the predicted effects to these reduced-form estimates—consistency between the two would strengthen the credibility of our model's results.

Contribution to the literature. Studies evaluating the effects of housing supply expansions tend to measure the local price impacts of new construction, usually finding that additional housing lowers nearby property prices (e.g., Li, 2022; Asquith, Mast and Reed, 2023; Pennington, 2021). However, these effects vary widely depending on the type and location of construction (Diamond and McQuade, 2019; Blanco and Sportiche, 2024). Furthermore, due to migration, the price effects of new construction extend well beyond the immediate neighborhood where construction occurs (Mast, 2023). Finally, these price changes bundle the effects of changes in housing supply, amenities, and taxes, and reflect the preferences of the marginal household that chooses to live in the neighborhood. Our analysis confirms that new construction affects local prices heterogeneously, and builds a model of households' preferences and migration behavior. Our framework can rationalize the heterogeneous effects of construction through variation in household preferences. It can further evaluate the diffuse price effects of construction far from the construction site. Finally, it can disentangle the different channels through which construction affects prices and welfare, and evaluate the full distribution of the welfare effects of construction, including on the non-marginal residents of the neighborhood in which it takes place.

Disentangling the channels through which new construction affects welfare—and measuring the full distribution of these welfare effects—is central to understanding the political economy of land use regulation. A large literature articulates rationales for such regulation. In the homevoter view, zoning protects incumbent homeowners and their property values (Fischel, 2002). Research has also shown that households highly value living in higher-income municipalities (Boustan, 2013), and are especially sensitive to the demographic characteristics of their close neighbors (Bayer, Casey, McCartney, Orellana-Li and Zhang, 2022). These preferences give rise to an exclusionary motive for zoning, which largely contributed to the adoption of strict zoning rules in the postwar United States (Sahn, 2025; Cui, 2024). The concurrent fiscal motive for zoning was initially outlined by Hamilton (1975) and Oates and Mills (1975), and later refined by

Calabrese, Epple and Romano (2007, 2012), Barseghyan and Coate (2016), and Brueckner (2023). He, Nelson, Su, Zhang and Zhang (2022) show that fiscal motives are key to understanding zoning policy in China, and Krimmel (2021) shows that California's local governments adopted stricter zoning rules to preserve the quality of local public goods. Disentangling the rationales underlying land use regulation is challenging. For instance, building new multifamily housing may lower local property values through a supply effect, by increasing the share of lower-income and minority households in the neighborhood, or by increasing the property tax burden of incumbent residents. This paper provides a unified framework to jointly study and quantify these channels. By disentangling them, our contribution aims at illuminating which rationales for land use regulation may be the most potent. Understanding the rationales underlying land use regulation is key to understanding which policies are most likely to promote housing construction. For instance, if fiscal motives justify strong local opposition to development, redistribution between municipalities based on housing construction may prove efficient in promoting housing construction.

By measuring the interjurisdictional spillovers generated by new construction, our analysis informs the long-standing debate on the optimal level of policy decentralization (Tiebout, 1956; Olson, 1969; Oates, 1999; Besley and Coate, 2003).[1] This question is particularly important in the context of housing construction, where the households who lose from new development are few and geographically concentrated—making it easier for them to coordinate and lobby against it—while the gains from construction are widely dispersed. These correspond exactly to the conditions under which the harmful distributional coalitions, special interest groups who seek to protect benefits within the group at the expense of broader welfare, described by Olson (1982) are likely to emerge.

The framework we develop to measure the distributional effects of housing construction incorporates the key features required to accurately measure these effects. First, it represents locations with a very fine spatial granularity. Indeed, the effect of building new housing units is much stronger for households on the same street than on others living one half a mile away. Our model also features rich heterogeneity in both household characteristics and housing types. Indeed, whether households are homeowners or renters crucially affects whether they benefit or lose from changes in home prices. Beyond homeownership, households vary greatly in their valuation of different amenities, which leads them to sort into different types of neighborhoods and housing units. Accordingly, increasing a neighborhood's housing supply by 10,000 units will have very different effects on neighborhood demographics depending on whether the new construction consists of small apartments or large mansions, and local demographic shifts from construction matter much more to some households than to others. Finally, our model features

[1] This relates our work to Bordeu (2023), who measures interjurisdictional externalities in the context of infrastructure improvements, and a literature on municipal competition, summarized in Agrawal, Hoyt and Wilson (2022)

household mobility, which is essential to capture the diffuse effects of construction. Because of migration, increasing the housing supply in one neighborhood shifts prices and demographics everywhere, and hence indirectly affects all households. Furthermore, households initially far from the construction event are directly affected by the changes in the treated neighborhood to the extent that they may live there in the future.

To develop a model that incorporates these features, we build on two types of approaches in the literature: static models that accommodate many locations, housing types, and heterogeneous preferences (Bayer et al., 2007, e.g.,), and dynamic models that describe migration but typically restrict the number of locations for tractability (Almagro and Domínguez-Iino, 2025; Greaney, Parkhomenko and Van Nieuwerburgh, 2025). Our framework bridges these approaches: we extend static models to account for migration while maintaining tractability by assuming that households are myopic and that prices and amenities converge quickly to a new steady state after a shock.

The estimation of our model requires highly granular data on housing units and household location choices. While such datasets are now widely used in the literature, this paper introduces two data contributions. First, we provide new facts on zoning regulations across the United States using zoning maps from Zoneomics. Second, we build a dataset of all housing units in New Jersey, including those inside multifamily buildings. Such datasets have been challenging to build in the United States, where administrative data typically derives from tax assessment rolls, which enumerate parcels but rarely the units within them, limiting the coverage of apartments and pushing empirical work toward single-family homes. The procedure we follow to build a comprehensive dataset of housing units can easily be replicated in other states, and is valuable as multifamily buildings account for a large share of the housing stock, are key to densification and providing affordable housing, and are the building types that spur the strongest local opposition.

2 Rationales for Restricting Housing Supply

New housing construction affects home values, local amenities, and may trigger changes in property tax rates. We start our analysis by describing how increasing housing supply affects households through these channels, focusing on the potential losses from densification. New data from zoning maps allows us to describe the land use regulations that municipalities impose to protect incumbent residents.

2.1 How does new construction affect welfare?

Supply effects. Several studies have shown that new construction in an area typically lowers the value of preexisting buildings there (Li, 2022; Asquith et al., 2023; Pennington, 2021; Blanco and Sportiche, 2024), consistent with a higher supply of housing lowering its price. This benefits renters, but harms homeowners, who will realize a lower capital gain upon selling their home.

Demographics. Interestingly, a higher housing supply sometimes *increases* local home prices. Diamond and McQuade (2019), for instance, found that new construction in low-income neighborhoods attracted richer households and increased the value of nearby homes. However, building low-income housing units in more affluent neighborhoods lowered local home values. These results highlight that new construction can change neighborhood demographics, which in turn shift local amenities.

How new construction shifts local demographics depends crucially on the type of housing that is being built, and in particular whether it consists of single-family homes or multifamily buildings. Different types of homes attract different types of households (see Appendix Table B.1). In New Jersey, households in multifamily buildings have an average income 51% lower than those living in single-family homes in the same area. There are also 14 pp less likely to be white. Consistently, there is a strong correlation between the share of single-family homes in a municipality and that municipality's demographics—see Figure 1(a). Among single-family dwellings, larger homes built on larger plots of land will be more expensive and inhabited by richer households.

When households value living among wealthier neighbors, the construction of homes that are smaller and less expensive than existing ones will likely lower the area's average household income and lead to a decrease in local amenities. This creates an incentive for municipalities to resist densification, which typically requires building multifamily housing or homes on smaller lots.

Incumbent residents may resist the entry of lower-income residents because they directly reduce their utility, or because demographics affect other local public goods, such as schools. There is indeed a stark correlation between municipal demographics and school test scores, illustrated in Figure 1(b). Residents may value the quality of schools directly or indirectly, as higher test scores increase home values.

Externalities from density. Demographic change is not the only channel through which new construction can affect local amenities. Densification may directly decrease local amenities through increases in noise, pollution, and traffic (Kashner and Ross, 2025). Furthermore, construction can eliminate green spaces: Figure 1(c) shows that residents in municipalities with a higher proportion of single-family homes enjoy greater tree exposure, a feature households strongly value (Han, Heblich, Timmins and Zylberberg, 2024).

These negative impacts of construction may be mitigated by agglomeration externalities. Indeed, higher density in an area can trigger additional public investment there, and spur the growth of consumption amenities such as shops and restaurants.

Figure 1: Rationales for a municipality to restrict densification

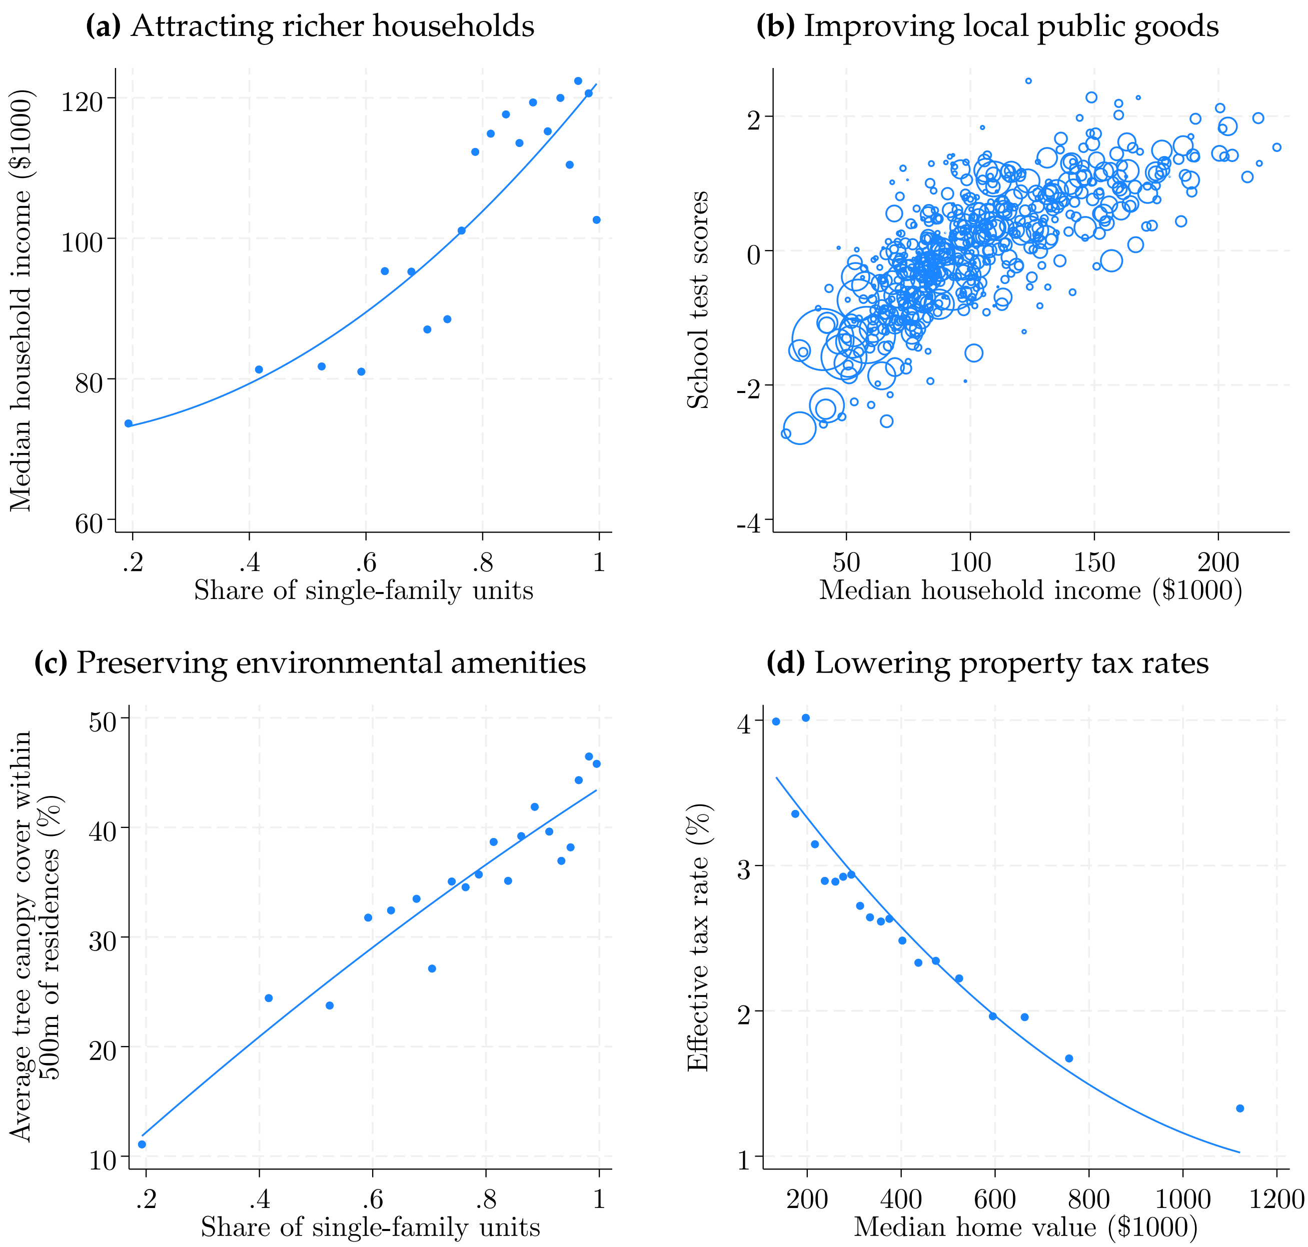

Notes: This Figure illustrates rationales for a municipality to resist densification. Panel (a) shows the correlation between the share of single-family units in a municipality and that municipality's median household income. Panel (b) shows the correlation between the median household income in a municipality and the test scores of its 4th-grade classes, weighting observations by the number of children in each municipality. Panel (c) shows the correlation between the share of single-family units in a municipality and the average tree canopy cover within 500 meters of its residential properties. Panel (d) shows the correlation between the median home value in a municipality and its effective property tax rate. The share of single-family units, household incomes, and home values in a municipality are extracted from the ACS. Other data sources are described in Appendix A.

Property taxes. Local governments rely heavily on property taxes to fund public services, with education being the most significant expenditure (see Appendix Figures B.1 and B.2). Consequently, areas with lower property values must implement higher tax rates to finance public goods, as illustrated in Figure 1(d).[2]

Because local governments rely on the property tax, new construction affects incumbent households through a fiscal channel. While new housing units expand the property tax base, they also generate new public expenditures (for example, by increasing enrollment in local schools). When the cost of these new services outweighs the additional revenue, the local government must increase the property tax rate on incumbent households to maintain budget balance.

Within a jurisdiction, the property tax will implicitly redistribute income from residents of large and expensive homes to residents in smaller, inexpensive dwellings. As shown by Huang (2025), these transfers are quantitatively large, and make fiscal externalities of new construction economically relevant. In New Jersey, the average homeowner pays \$9,600 in property taxes each year, corresponding to 7% of the state's average household income.

2.2 Protecting incumbent residents through zoning

Through land use regulations, residents of a municipality can influence the number and characteristics of new homes. These regulatory choices partly reveal the preferences of incumbent residents and are informative about the distributional effects of building different types of dwellings.

To understand how incumbent residents leverage regulation to shield themselves from the negative effects of construction outlined above, we rely on data from Zoneomics, which provides detailed zoning maps for 9,764 zoning jurisdictions across the United States, covering 75% of the country's population (see Appendix Figure B.4).[3] Zoning maps divide land into a set of zones, each being associated with rules on new construction.

The vast majority of jurisdictions (93%) reserve a fraction of their residential land for low-density structures (single-family or two-family dwellings only). In our sample, 79% of the residential land in the United States is exclusively zoned for low-density homes, and only 21% of the residential land allows the construction of higher-density, multifamily housing. On top of these restrictions privileging the construction of low-density housing, municipalities routinely impose policies that favor high consumption of land: minimum lot sizes, mandatory setbacks, and lot coverage limits (see Figure 2).

These zoning regulations are consistent with the expected effects of new construction described in Section 2.1. Indeed, the construction of homes on smaller areas of land, and the construction of multifamily housing in particular, is likelier to harm incumbent residents,

[2]Spending on local public goods weakly correlates with households' income. In New Jersey, municipalities with a median household income in the bottom quintile spent \$15,700 per household in 2017 on local public goods provision. Municipalities in the top quintile spent \$16,100 per household.

[3]While most zoning decisions are made by municipalities, other levels of government (counties, in particular) enact zoning rules on unincorporated land and also constitute zoning jurisdictions.

Figure 2: Prevalence of zoning instruments in the United States

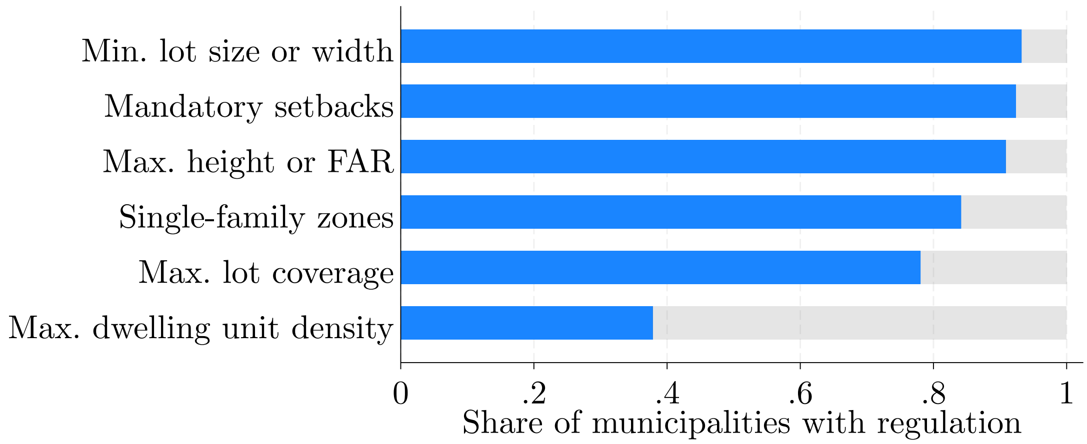

Notes: This Figure shows the share of municipalities in the Zoneomics dataset that have adopted different zoning instruments. A municipality is considered to have adopted an instrument if at least one zone in its zoning map imposes restrictions on residential construction through that instrument. Mandatory setbacks correspond to rules imposing a minimum front yard, rear yard, or side yard. The Floor Area Ratio (FAR) of a property is defined as the amount of floorspace on the property divided by its land area. A single-family zone is defined as a zone that allows single-family dwellings, but not two-family or multifamily dwellings. Maximum lot coverage rules impose a maximum share of a property that can be covered by buildings and other impervious surfaces. Finally, dwelling unit density limits impose a maximum on the number of allowed units per acre of land.

especially if they are homeowners. Housing units that use land less intensively tend to attract lower-income households, will raise less property tax revenue than large single-family homes, and impose larger externalities, for instance through reduced open space and higher local traffic. Accordingly, zoning regulations tend to restrict the construction of such housing while promoting that of single-family homes on large plots of land.

3 Theoretical Framework

The local effects of new construction on prices, amenities, and fiscal balance have been studied independently. The goal of this paper is to measure these effects jointly, with a high degree of granularity, and beyond the area immediately near construction.

3.1 The welfare effects of housing construction

We first develop a theoretical framework that can account for the different channels through which construction affects welfare. Building new housing influences housing prices, amenity levels, and property taxes, not only in the immediate neighborhood where construction occurs, but also in areas further away through spillover effects. Even in the absence of spillovers, households initially located far from the new units may still be affected over their life cycle. If they choose to move to the area where construction occurred, they will face different prices, amenities, and taxes than if new housing had not been built.

Exposure to amenity and rent shocks. Consider a simple model in which households live in periods $\{0, 1, ..., T\}$. At the beginning of each period, they choose a location to live in. Locations (which can represent, e.g., municipalities, housing units) are characterized by an amenity level B_k and a rent level R_k. Households have income level I, and their disposable income, used for consumption, is given by $C_k = I - R_k$. Moving between locations is costly, and households draw each period idiosyncratic preferences for each location. If a household starting a period in j decides to live in k, it will receive a flow utility

$$U_{jk} = B_k + \gamma C_k - \mu_{jk} + \varepsilon_k = u_{jk} + \varepsilon_k, \tag{1}$$

where γ captures the extent to which households value consumption relative to amenities, μ_{jk} is the cost of moving from j to k, and ε_k is an idiosyncratic preference shock for location k, distributed according to a Gumbel distribution with shape σ and mean zero. Rents are paid to an absentee landlord.

For simplicity, assume that households are myopic and choose each period the location that maximizes their flow utility. Then, the probability with which a household initially in j moves to k is $\mathcal{P}_{jk} = \frac{\exp(u_{jk}/\sigma)}{\sum_l \exp(u_{jl}/\sigma)}$, and the row vector $\boldsymbol{\pi}_t = \boldsymbol{\pi}_0 \boldsymbol{\mathcal{P}}^t$ measures the probability with which a household is in each location at any period (with $\boldsymbol{\pi}_0$ capturing the initial location of the household).

Households' expected lifetime welfare is given by $\mathbb{E}[W] = \sum_{t=0}^{T} \beta^t \left(\boldsymbol{\pi}_t \bar{\boldsymbol{U}}\right)$, where $\bar{U}_j = \sigma \log\left[\sum_k \exp(u_{jk}/\sigma)\right]$ the expected flow utility of a household starting a period in j. Consider a shock that permanently shifts amenities and rents just after households made their initial location choices. The approximation of the effect of this shock on expected lifetime utility is given by

$$\mathrm{d}\mathbb{E}[W] = \underbrace{\overbrace{\left(\sum_{t=0}^{T} \beta^t \boldsymbol{\pi}_t\right)}^{\text{Exposure vector}} \cdot \mathrm{d}\boldsymbol{B}}_{\text{Amenity effect}} - \underbrace{\left(\sum_{t=0}^{T} \beta^t \boldsymbol{\pi}_t\right) \cdot \gamma\, \mathrm{d}\boldsymbol{R}}_{\text{Price effect}} + \underbrace{\left(\sum_{t=0}^{T} \beta^t \mathrm{d}\boldsymbol{\pi}_t\right) \cdot \bar{\boldsymbol{U}}}_{\text{Relocation effect}}. \tag{2}$$

The first term captures households' direct exposure to amenity changes. Indeed, $E_i = \left(\sum_{t=0}^{T} \beta^t \boldsymbol{\pi}_t\right)_i$ captures the expected number of (time-discounted) periods the agent will spend in location i, and $E_i \mathrm{d}B_i$ captures the change in the agent's welfare due to their exposure to location i. Similarly, the second term captures households' direct exposure to rent changes: lowering the rent by a dollar in location i increases the agent's lifetime discounted consumption by E_i dollars, and increases their utility by that amount. The last term, equal to $\mathrm{d}\boldsymbol{E} \cdot \bar{\boldsymbol{U}}$, captures the fact that the shock shifts households' exposure to different locations. For instance, lowering rents in a high-amenity location (and hence a high-$\bar{U}$ location) will increase migration to that location. This increases households' exposure to high-amenity locations and increases welfare beyond the direct effect of rent decreases.

Accounting for homeownership. A major limitation of the model described above is its assumption that all households are renters. In reality, most households own their homes, and price changes in a location have opposite effects on the welfare of a household living there, depending on whether they are a homeowner or a renter. Indeed, increases in home values are immediately detrimental to renters, who have to dedicate a higher share of their income to rent payments and must reduce their consumption. Homeowners, on the other hand, see no immediate change in their housing costs, but will realize a higher capital gain upon migrating and selling their home. Therefore, higher home values increase their future liquid wealth and consumption, and therefore their welfare.

To see how equation (2) can be adjusted to account for homeownership, consider a variation of the model described above in which all households are homeowners. The simplest way to model homeownership is to assume that households hold liquid wealth A, and that liquid wealth decreases (resp., increases) when buying (resp., selling) a home by the transaction price of that home. A household's consumption is given by $C = I + rA$. Hence, a household initially starting with zero wealth buying a home at price P will have a consumption of $I - rP$, where $R = rP$ can be interpreted as a yearly mortgage payment for the home.

Again, we can measure the effect of a shock that shifts prices and amenities. This time, the approximate effect of the shock on households' welfare is given by

$$\mathrm{d}\mathbb{E}[W] = \underbrace{\left(\sum_{t=0}^{T} \beta^t \boldsymbol{\pi}_t\right) \cdot \mathrm{d}\boldsymbol{B}}_{\text{Amenity effect}} \quad + \quad \underbrace{\left(\sum_{t=0}^{T} \beta^t (\boldsymbol{\pi}_0 - \boldsymbol{\pi}_t)\right) \cdot \gamma \, \mathrm{d}\boldsymbol{R}}_{\text{Price effect}} \quad + \quad \underbrace{\left(\sum_{t=0}^{T} \beta^t \mathrm{d}\boldsymbol{\pi}_t\right) \cdot \bar{\boldsymbol{U}}}_{\text{Relocation effect}}. \tag{3}$$

The first and last term of this expression are identical to those of equation (2). The second term, however, now takes into account households' initial stake in the housing market. The ith element of $\sum \beta^t (\boldsymbol{\pi}_0 - \boldsymbol{\pi}_t)$ measures the household's net position as a seller versus a consumer of housing in location i, discounted over the household's lifetime.

Consider a shock that decreases the price of homes in location i by $\mathrm{d}P_i$. If a household starts in location i and never moves over the course of its lifetime, it never realizes any capital losses and its exposure to the price change is zero. Conversely, if that household moves to another location in the first period and never comes back, its liquid asset holdings will be lowered by $\mathrm{d}P_i$ for all periods $\{1, ..., T\}$ relative to a scenario where the shock did not take place. This lowers the household's consumption by $r \, \mathrm{d}P_i$ each period, for a total welfare loss of $\gamma \sum_{t=1}^{T} \beta^t r \, \mathrm{d}P_i$.

When prices drop in a location, this leads to a concentrated welfare loss for incumbent homeowners in that location, and a diffused welfare gain among households in other locations, who will face lower housing costs if they move to the affected location at one point over their life cycle. The welfare gains fully offset welfare losses. Indeed, summing $(\boldsymbol{\pi}_0)_i$ and $(\boldsymbol{\pi}_t)_i$ across households yields the number of housing units in location i in both cases, and the second term of equation (3) therefore averages zero across households.

General framework. The model developed above can be generalized to allow households to choose between renting and owning their home, to account for the payment of property taxes, and to relax the assumption of an absentee landlord. To this end, suppose that the locations of the model are individual housing units. Units are either always rented out or always owner-occupied. Unit owners must pay each year a property tax and maintenance costs, totaling T and possibly varying from one unit to the next. The yearly cost of housing is T for a homeowner and $R + T$ for a renter. Homes are transacted at a price $P = R/r$. The approximation of the effect of a shock on the household's expected welfare is given by

$$\mathrm{d}\mathbb{E}[W] = \underbrace{\left(\sum_{t=0}^{T} \beta^t \boldsymbol{\pi}_t\right) \cdot \mathrm{d}\boldsymbol{B}}_{\text{Amenity effect}} + \underbrace{\left(\sum_{t=0}^{T} \beta^t (\boldsymbol{\pi}_0^O - \boldsymbol{\pi}_t)\right) \cdot \gamma\, \mathrm{d}\boldsymbol{R}}_{\text{Price effect}} - \underbrace{\left(\sum_{t=0}^{T} \beta^t \boldsymbol{\pi}_t\right) \cdot \mathrm{d}\boldsymbol{T}}_{\text{Fiscal effect}} + \underbrace{\left(\sum_{t=0}^{T} \beta^t \mathrm{d}\boldsymbol{\pi}_t\right) \cdot \bar{\boldsymbol{U}}}_{\text{Relocation effect}}, \tag{4}$$

where $\boldsymbol{\pi}_0^O$ denotes the asset holdings of the household at time 0. $(\boldsymbol{\pi}_0^O)_i$ is equal to one if the household owns home i and zero otherwise. The introduction of taxes and maintenance adds a final term to the approximate effect of a shock: an increase in T in one location decreases the household's expected welfare proportionally to the its exposure to that location.

New housing construction. In this framework, the construction of a new housing unit can be modeled as the introduction of a new location and the expansion of households' migration choice set. In this location, π_t can be positive following construction, instead of equal to zero in all periods.

New construction affects welfare through several channels. First, by changing local density and demographics (depending on who migrates to the new housing unit), it may affect local amenities. Second, construction may prompt municipalities to adjust their property tax rates. Third, the increase in local housing supply, as well as accompanying changes in amenities and taxes, will change market clearing rents in all locations. On top of these channels, described in Section 2, equation (4) points to a final effect of new construction: the relocation effect. New construction will increase overall utility if it shifts households towards high-$\bar{U}$ locations.[4]

To quantify the welfare effects of building new housing, we need estimates of γ, $\boldsymbol{\pi_t}$, and $\bar{\boldsymbol{U}}$, as well as measures of the local shocks caused by new construction ($\Delta\boldsymbol{B}$, $\Delta\boldsymbol{R}$, $\Delta\boldsymbol{T}$, and $\Delta\boldsymbol{\pi_t}$). Apart from $\Delta\boldsymbol{R}$ and $\Delta\boldsymbol{T}$ (which are common shocks across households) all of these variables will typically vary from one household to the next.

[4]If migration were costless and agents had no idiosyncratic preferences, spatial equilibrium would imply a uniform $\bar{U}$ across locations, making the relocation effect vanish.

3.2 Household location choices

To evaluate how new construction shifts amenities, rents, taxes, migration behaviors, and ultimately welfare gains and losses, we extend the framework and specify a model of household location choice that determines how location characteristics evolve.

Model setup. We consider the housing choices of a finite set of households $i \in \mathcal{I}$. At the beginning of each period t, they choose whether to remain in their current housing unit or move to another. A household starting in unit j and deciding to spend the period in location $k \in \mathcal{H}$ (where k can be j) receives a flow utility

$$u_{it}(j \to k) = B_{it}(k) + \gamma_{it}\, C_{it}(k) - \mu_{it}(j \to k) + \varepsilon_{itk}, \tag{5}$$

where $B_{it}(k)$ is the amenity level of unit k, $C_{it}(k)$ is the consumption level of household i if it decides to live in k, $\mu_{it}(j \to k)$ is a moving cost associated with moving from j to k, and ε_{itk} is an idiosyncratic preference shock. Each housing unit has a fixed tenure type, denoted by the binary variable $O_k \in \{0,1\}$ where $O_k = 1$ if the unit is owner-occupied instead of rented. Households are myopic: they ignore future periods when making their decisions, and assume that future conditions will be similar to current conditions. $\boldsymbol{z}_{it}$ are demographic characteristics of household i, such as income level, race, and age of the household head.

Amenities. The amenity level associated with unit k is

$$B_{it}(k) = \boldsymbol{\alpha}_{it}^{\top} \boldsymbol{X}_{kt} + \xi_k, \tag{6}$$

where $\boldsymbol{X}_{kt}$ are observable characteristics of units and ξ_k is an amenity shock that reflects characteristics of the unit that are unobservable to the econometrician, including unobservable characteristics of the municipality in which it is located. $\boldsymbol{X}_{kt} \equiv \left[\boldsymbol{X}_k^{\text{unit}}; \boldsymbol{X}_{kt}^{\text{location}}\right]$ includes both unit characteristics $\boldsymbol{X}_k^{\text{unit}}$ (such as the unit's square footage and number of bedrooms and bathrooms), and location characteristics $\boldsymbol{X}_{kt}^{\text{location}}$ (such as local demographics and density). The weights households place on observable amenities are a linear function of household demographics, with

$$\boldsymbol{\alpha}_{it} = \boldsymbol{\alpha}_0 + \boldsymbol{A}\, \boldsymbol{z}_{it}. \tag{7}$$

Where the (l, m) coefficient in matrix $\boldsymbol{A}$ measures the extent to which the mth household demographic (e.g., income) affects its valuation of the lth location characteristic (e.g., square footage).

Cost of housing. A household's consumption is its income I_i minus the cost of housing, $H_{it}(k)$. Housing costs sum over rent R, property taxes and maintenance costs T, or interest paid on real estate debt A.

$$C_{it}(k) = I_i - H_{it}(k) = I_i - ((1 - O_k)R_{kt} + T_{kt} + r\,A_{it}). \tag{8}$$

Rents are only paid for rented units (with $O_k = 0$). We assume property taxes and maintenance costs T_{kt} are paid by the units' occupants, whether they are owners or renters. Accordingly, the rent R_{kt} excludes T_{kt}. Finally, real estate debt A evolves upon buying or selling a home. Buying (resp., selling) a home at price P increases (resp., decreases) A by that amount. The price of a housing unit is given by a no-arbitrage condition: it is the discounted sum of the rental income stream a (myopic) investor would expect from the unit: $P_{kt} = \mathcal{O}_t R_{kt}/r$, where $\mathcal{O}_t$ is the expected occupancy rate of the unit.

A_{it} corresponds to real estate debt holdings during period t (measured after buying and selling decisions have been made), and $r\,A_{it}$ can be interpreted as annual mortgage payments. A_{it} can be negative (for instance, if a households decides to sell their home to rent another)—in that scenario, $r\,A_{it}$ can be interpreted as interest received on liquid assets.

The relative weight γ_{it} households place on consumption relative to amenities is a linear function of demographics, with

$$\gamma_{it} = \gamma_0 + \boldsymbol{\gamma}^\top \boldsymbol{z}_{it}. \tag{9}$$

Migration costs. When a household changes housing units, it must pay the migration cost $\mu_{it}(j \to k)$, with

$$\begin{aligned} \mu_{it}(j \to k) = &\mu_0 + \mu_d\, d(j,k) + \mu_{\text{age}}\, \text{age}_{it} + \mu_{\text{own}}\, O_j + \\ &\mu_{\text{mun}} \mathbb{1}\{m(k) \neq m(j)\} + \mu_{\text{state}} \mathbb{1}\{s(k) \neq s(j)\}. \end{aligned}$$

Migration costs increase with the geographical distance between the two units $d(j,k)$, the head of household's age, tenure at a location, and if the original unit is owned instead of rented. Moving across municipalities ($m(k) \neq m(j)$) or across states ($s(k) \neq s(j)$) is also associated with higher migration frictions. Households who decide to stay in their housing unit pay no migration costs, so $\mu_{it}(j \to j) = 0$.

Property taxes and maintenance costs. Each period, municipalities (indexed by m) must recover $N_{mt}T_m$ of tax revenue, where N_{mt} is the number of housing units in m, and T_m is the average property tax paid across housing units. Municipalities impose a property tax rate τ_{mt} such that the property tax contribution of housing unit k is $T_{kt}^{\text{tax}} = \tau_{mt} P_{kt}$. To recover an average tax revenue T_m across units, the municipality sets $\tau_{mt} = T_m / \bar{P}_{mt}$, where $\bar{P}_{mt}$ is the average price of housing units in m. This expression makes clear that adding to a municipality housing units that are cheaper than the average preexisting housing unit increases the municipality's equilibrium tax rate.[5]

Moving probabilities and equilibrium rents. When idiosyncratic shocks ε_{itk} are i.i.d. according to a standard type-I extreme value distribution, moving probabilities are given by

$$P_{it}(j \to k) \;=\; \frac{\exp\{B_{it}(k) + \gamma_{it}\, C_{it}(k) - \mu_{it}(j \to k)\}}{\sum_{\ell \in \mathcal{H}} \exp\{B_{it}(\ell) + \gamma_{it}\, C_{it}(\ell) - \mu_{it}(j \to \ell)\}}\,. \tag{10}$$

As income is not tied to the location choice, these moving probabilities can also be expressed as a function of solely amenities, housing costs, and migration costs.

$$P_{it}(j \to k) \;=\; \frac{\exp\{B_{it}(k) - \gamma_{it}\, H_{it}(k) - \mu_{it}(j \to k)\}}{\sum_{\ell \in \mathcal{H}} \exp\{B_{it}(\ell) - \gamma_{it}\, H_{it}(\ell) - \mu_{it}(j \to \ell)\}}\,. \tag{11}$$

Equilibrium rents are determined by market clearing. Because there are more housing units than households, all housing units are not systematically occupied. However, we assume that rents adjust such that each period, the expected number of households choosing each unit equals the overall occupancy rate $\mathcal{O}_t$: $\sum_{i \in \mathcal{I}} P_{it}(j_{i,t-1} \to k) = \mathcal{O}_t$.

4 Data and Context

New Jersey offers an attractive case study for quantifying the distributional effects of new construction and exploring the channels through which it influences welfare. Indeed, the state is fully incorporated into municipalities with stable boundaries over recent decades, most local public goods are financed through municipal property taxes, and school districts generally coincide with municipal borders. This Section summarizes the rich data we gathered to measure households' location choices in New Jersey, as well as municipalities' tax and zoning policies. Appendix A provides details about the data we collected, and Appendix B provides additional descriptive statistics.

Housing units and home values. To measure the supply of housing in different locations and its price, we primarily rely on proprietary datasets provided by CoreLogic. CoreLogic

[5] Brosy and Langley (2025) show that municipalities in New Jersey adjust property tax rates in response to changes in property values in order to keep the property tax levy approximately constant, consistent with our model. By assuming that new construction does not alter the average tax revenue municipalities must raise per housing unit, we are further assuming constant returns to scale in the provision of local public goods. This assumption aligns with the literature on local public service production, which generally finds limited economies of scale, particularly once jurisdictions reach moderate size (Gómez-Reino et al., 2023). Consistent with this evidence, in New Jersey we observe that school districts' per-student expenditures remain flat beyond an enrollment of roughly 400 students (see Appendix Figure B.3).

aggregates data from property tax assessments, sales records, and Multiple Listing Services. Using their proprietary datasets, as well as datasets listing all addresses in New Jersey from the open-source NJGIN and the commercial vendor Smarty, we build a dataset of all residential units in the state. We associate them with characteristics such as whether the unit is a single-family home, a duplex, or an apartment, the unit's number of bedrooms and bathrooms, the unit's size, and the size of the lot on which the unit is located. Accurately capturing the total number of residential units is challenging but required for estimating the model, which features households that choose from all available units. Existing studies do not quantify the full supply of residential housing. For example, in Diamond and McQuade (2019), unit counts of LIHTC structures are unobserved while in Asquith et al. (2023), unit information is only measured for large apartment buildings.

CoreLogic further provides data on real estate listings and sales, which allows us to measure the price at which these housing units are traded.

Migration, ownership, and housing demand. To measure households' demand for housing, we rely on revealed preferences and study households' location choices over time. In particular, we leverage data from Infutor that documents, for most households in the United States, the list of their successive addresses as well as when they moved from one location to the next. This address history additionally enables us to infer whether households are homeowners or renters, as we link their address to the unit from CoreLogic and the housing type (single-family, apartment, condo, etc.) of the unit.

We associate households with demographics using several sources. First, we link households with the mortgages they obtained using CoreLogic's mortgage data. We then link these mortgages with the Loan Application Register (LAR) database. This data, collected under the Home Mortgage Disclosure Act (HMDA), lists mortgage applications submitted to financial institutions, associating each with applicants' race and income. While this data is anonymized, it indicates each mortgage's census tract, transaction date, loaned amount, loan type, and the lender's name. We use this data to match the LAR with CoreLogic mortgages, and hence to households, using a procedure similar to that in Bayer et al. (2022). For households not covered by the HMDA dataset, we recover demographics through a Bayesian procedure. Data from the ACS allows us to infer the income and race of the household in a given housing unit (Cook, 2025, uses a similar approach). Following Diamond et al. (2019), we use the names of household members, available in the Infutor data, to infer individuals' race.

Local government finances. To account for the fiscal effects of new construction, we gather data on local governments' finances. The Census of Governments tracks the revenue and spending of states, counties, municipalities, school districts and special districts, as well as the intergovernmental transfers between these entities. Most special districts and school districts serve a single municipality. We map these specialized local governments to the municipalities they cover, and aggregate local governments' revenue and expenditure to the municipal level.

5 Model Estimation

The granular datasets we have assembled on migration history, housing units, and neighborhood and municipal characteristics allow us to estimate a rich model of household location choice. We describe how our empirical strategy isolates quasi-random variation for key parameters, and then we present the estimation steps for matching moments in the model to the data.

5.1 Evidence from a border design

We first illustrate the sources of variation that contribute to estimating some of the parameters in the utility function. This variation leverages differences in housing regulation across municipal borders, which in turn affects various neighborhood characteristics that households have preferences over. These preferences appear through shifts in home prices at the borders that empirically indicate how neighborhood features shape housing demand. We exploit our rich dataset on land use regulations to extend the findings of previous studies that also employ the border design in their empirical approach (Song, 2021; Gyourko and McCulloch, 2024).

Border discontinuities. For this design, we focus on areas within a 500-meter radius of the boundaries that separate municipalities after removing sections that are near more than two municipal jurisdictions (see Appendix Figure C.1 for an illustration). Required for identification of the effects of differences by the border, we assume that unobserved location characteristics such as natural amenities and market access are smooth across the border. However, other neighborhood characteristics, which we aim to study, such as regulation and endogenous amenities of demographics or density, discontinuously change.

We start by measuring the impact of regulation on some of the channels that rationalize restricting housing supply in the border discontinuity specification below:

$$Y_i = \underbrace{\delta_{b(i)}}_{\text{Boundary FE}} + \beta D_i + \gamma_1 \text{dist}_i + \gamma_2 (D_i \times \text{dist}_i) + \varepsilon_i$$

When the treatment is regulatory restrictiveness, the treated side ($D_i = 1$) of the border is the municipality with more stringent regulations, and the other side is the one with less stringent regulations. Stringency is classified by aggregating over several dimensions of regulatory restrictiveness (share SF and the municipal median value for: minimum lot area, front yard setbacks, maximum height, and maximum lot coverage). We include boundary fixed effects to account for differences across municipal comparisons that are correlated with zoning.

Figure 3: Regulation stringency affects characteristics by municipal border

(a) Housing types

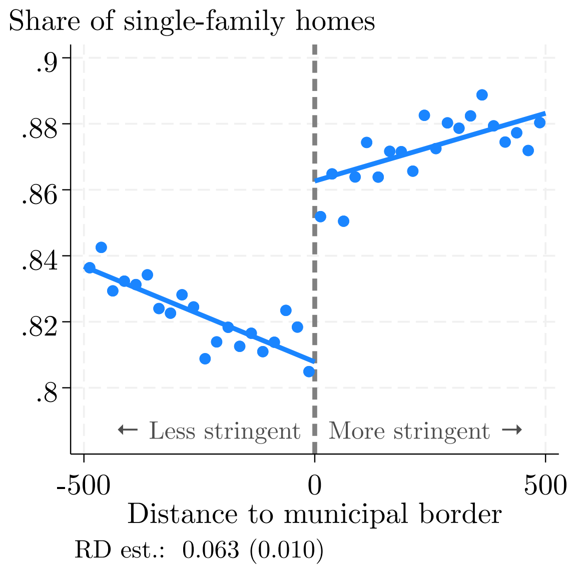

(b) Income sorting

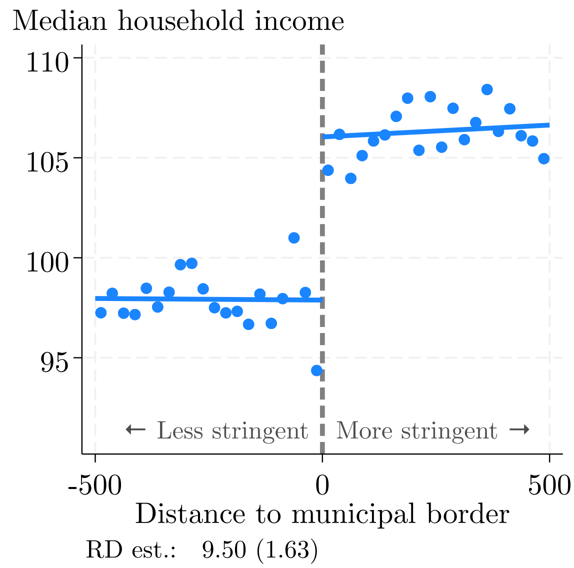

(c) Local taxation

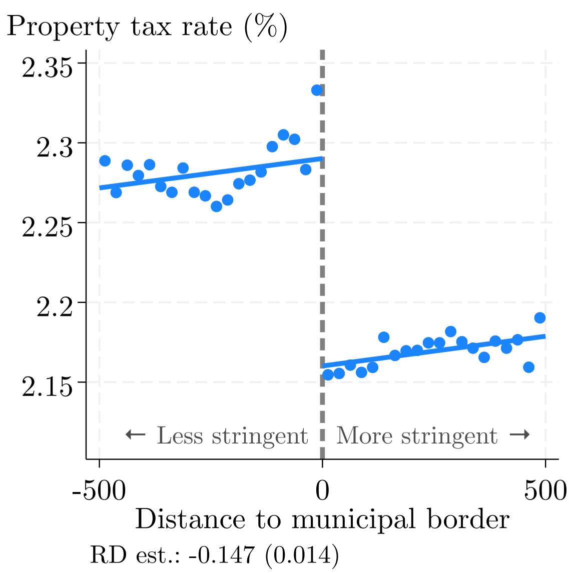

(d) Single-family home values

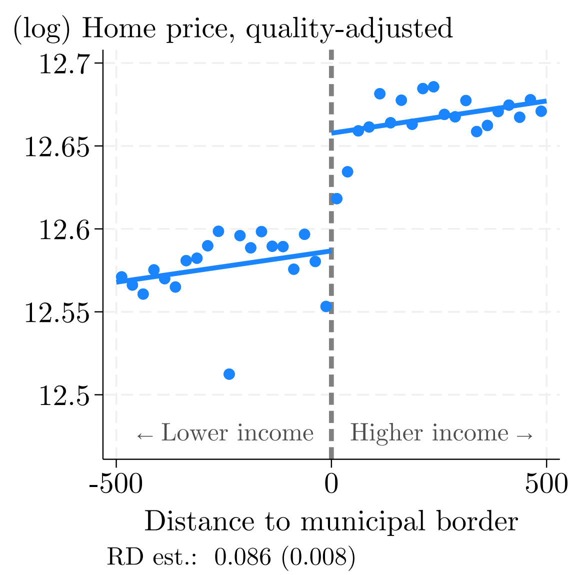

Notes: Observations are housing units from CoreLogic. A single-family zone is defined as a zone that allows single-family dwellings, but not two-family or multifamily dwellings. The share of single-family homes is measured by census block, median household income is measured by census block group, and property tax rate is measured at the municipal level. Home prices in panel (d) are measured by housing units (limited to only SF) and quality-adjusted for the housing characteristics: number of bedrooms, number of bathrooms, age of home, log of floorspace, and log of lot area.

In Figure 3(a), we show how the more stringent municipality requires a greater share of homes to be zoned for single-family housing. As a consequence of these regulations, Figures 3(b) and 3(c) depict how the median household income (a demographic characteristic correlated with other amenities such as public schools) is higher, and property tax rates are lower. Mechanically, the tax rate should be reduced since zoning forces homes to be larger in size and priced higher per unit. Additionally, home values rise because demand for these locations increases from the improved amenities/lower taxation.

This greater demand is reflected in the *quality-adjusted* home values, shown in Figure 3(d), where the running variable is household income, rather than regulatory stringency, to measure how preferences for higher income neighbors appear in prices. Adjusting for characteristics such as number of bedrooms/bathrooms and floorspace eliminates the portion of prices that is mechanically higher due to how zoning affects the size of the home.

In addition to the increase in home prices when entering the more stringent municipality, home values also slope upwards on both sides, and the discontinuity at the border is fuzzy rather than sharp. The gradual change in home prices may reflect how exposure to high-income neighbors is also gradual, since immediately on the higher-income side, homes are still exposed to neighbors on the lower-income side. Likewise, homes on the lower-income side near the border are exposed to neighbors on the higher-income side. In our next specification, we leverage this variation from areas adjacent to the border.

Valuation of local and municipal characteristics. While the border discontinuity design provides evidence on the valuation of *municipal* characteristics, it is also informative for understanding preferences for hyper-*local* characteristics. Demand for homes immediately by the boundaries is affected by their exposure to the other side within a narrow radius. We capture this local variation in the following hedonic regression of prices on municipal characteristics and local exposure, calculated as the average of neighborhood characteristics within a 500-meter radius. The sample is limited to housing units within 500 meters of each border, and border fixed effects absorb differences in levels across municipal comparisons.

$$\log\left(\text{price}_i\right) = \underbrace{\delta_{b(i)}}_{\text{Border FE}} + \underbrace{\boldsymbol{\beta}^{\text{local}}\boldsymbol{X}_i^{\text{local}}}_{\text{Local exposure}} + \underbrace{\boldsymbol{\beta}^{\text{muni}}\boldsymbol{X}_{m(i)}^{\text{muni}}}_{\text{Municipal chars}} + \underbrace{\gamma\boldsymbol{X}_i^{\text{unit}}}_{\text{Controls}} + \varepsilon_i$$

In this equation, $\boldsymbol{X}$ contains the variables of average household income, share white, and density (units per acre) at both the local and municipal level. These characteristics are standardized to allow for comparison of the magnitudes of the coefficients. We include quality controls for the housing unit in $\boldsymbol{X}^{\text{unit}}$, so the outcome variable of price represents how demand changes as a result of neighborhood and municipality features rather than the characteristics of the house, which may be changing due to regulations correlated with local demographics and density.

Using the discontinuity in demographics and density at the border, which further shifts local exposure, we find in Table 1 that home values respond more to racial composition, i.e., share white compared to average income for hyper-local exposure. Yet, for municipal characteristics, average income is valued more than racial composition since income level may more strongly determine public goods and services provided by the jurisdiction. In this hedonic regression, we do not find that housing demand is affected by density, i.e., units per acre. This last finding is somewhat similar to that of Gyourko and McCulloch (2024) as they show how, after controlling for demographics, the distaste for density is lowered.

Table 1: Preferences for demographics & density using border exposure design

	(log) Home price	
Local exposure (within 500m)		
Average household income	0.040***	(0.009)
Share white	0.101***	(0.011)
Units per acre	-0.013	(0.010)
Municipal characteristics		
Average household income	0.059***	(0.010)
Share white	0.030**	(0.009)
Units per acre	-0.013	(0.010)
Boundary FE	✓	
Observations	342,637	

Notes: For local exposure, average household income is at the census block group level, share white is at the block level, and units per acre is at the block level. Home prices are quality-adjusted for the features: number of bedrooms, number of bathrooms, age of home, log of floorspace, and log of lot area.

Now that we have presented where the empirical variation is derived from, we return to the model set-up laid out in Section 3.2 and describe the steps to infer model parameters.

5.2 Estimation procedure

We estimate the model using a two-step procedure similar to that of Bayer et al. (2007). The first step uses individual location choices to measure heterogeneity in household preferences. The second step recovers unobservable amenities associated with each housing unit, using an instrumental variables strategy to account for the correlation between these unobservable characteristics and home prices.

Step 1: Taste heterogeneity and mean utilities. For estimation, it is convenient to collect all components of $u_{it}(j \to k)$ that are common across households for a given unit–time pair (k, t) into a mean utility

$$\delta_{kt} \equiv \alpha_0^\top X_{kt} - \gamma_0 \left(H_t(k)\right) + \xi_k, \tag{12}$$

where $H_t(k) = (1 - O_k)R_{kt} + T_{kt}$ is the component of housing costs that is common across households. ξ_k captures unobserved amenities. We then write the individual component of utility that is related to location choice,

$$\tilde{u}_{itk} \equiv (\boldsymbol{A}\boldsymbol{z}_{it})^\top X_{kt} - (\gamma^\top \boldsymbol{z}_{it}) H_{it}(k) - \gamma_0(rA_{it}) - \mu_{it}(j \to k), \tag{13}$$

Utility can then be formulated as the sum of common and individual-specific components

$$u_{itk} = \delta_{kt} + \gamma_{it} I_i + \tilde{u}_{itk} + \varepsilon_{itk} \tag{14}$$

and migration probabilities are given by

$$P_{it}(j \to k) = \frac{\exp(\delta_{kt} + \tilde{u}_{itk})}{\sum_{\ell \in \mathcal{H}} \exp(\delta_{\ell t} + \tilde{u}_{it\ell})}. \tag{15}$$

Let θ_1 denote the parameters in $\boldsymbol{A}$, γ, and in the moving-cost specification $\mu_{it}(\cdot)$. The first-step estimator maximizes the log-likelihood

$$\mathcal{L}(\theta_1, \{\delta_{kt}\}) = \sum_{i,t} \log P_{it}\Big(j_{i,t-1} \to k_{it}^{\text{obs}}\Big), \tag{16}$$

jointly over θ_1 and the vector $\{\delta_{kt}\}$.

Step 2: Mean Preferences. In the second step we decompose the estimated mean utilities into observables and unobservables:

$$\widehat{\delta}_{kt} = \beta^\top X_{kt} - \lambda H_t(k) + \eta_{b(k)} + \zeta_{kt}, \tag{17}$$

where β recovers mean preferences for amenities, λ is the mean marginal utility of money, $\eta_{b(k)}$ are boundary fixed effects, and ζ_{kt} absorbs ξ_k and any remaining shocks. The boundary fixed effects implement a municipal boundary discontinuity design.

Because the endogenous variables of amenities X_{kt} and the cost of housing $H_t(k)$ are equilibrium outcomes that generally covary with the unobserved component ζ_{kt} in (17), we estimate (17) using instrumental variables. We discuss the instruments for X_{kt} and $H_t(k)$ in sequence and how they address bias due to correlation with the error term.

Instruments for location characteristics. Recall that X_{kt} contains both unit and location characteristics $[\boldsymbol{X}_k^{\text{unit}}; \boldsymbol{X}_{kt}^{\text{location}}]$. We leave the time-invariant unit characteristics $\boldsymbol{X}_k^{\text{unit}}$ as controls, and focus on the time-varying, endogenous location features in $\boldsymbol{X}_{kt}^{\text{location}}$. Following the findings from the boundary exposure design, these features can be divided into local exposure $\boldsymbol{X}_{kt}^{\text{local}}$ and municipal characteristics $\boldsymbol{X}_{mt}^{\text{muni}}$. Variation for estimation derives from local exposure to discontinuities across municipal borders and municipal-level shocks in characteristics.

In the border design, we focus on areas within a narrow band where we assume unobserved natural amenities that are neighborhood-specific (such as availability of lakes) are the same and thus do not bias the estimation of β. However, an identification concern with employing border discontinuities is that unobserved *municipal* amenities may also change sharply at the border. For example, higher-income municipalities often exhibit better governance because high-income households choose well-managed jurisdictions, which importantly implies the correlation between income and governance need not reflect a causal effect of income on governance. Without an instrument, our estimated coefficient for the preferences for municipal average income may also contain the preference for good governance.

To address this omitted variables bias, we construct instruments from a combination of sources. For local exposure, the full set of zoning regulations can be used as shifters of municipal traits that affect the bordering neighborhoods outside of the municipality. Because zoning is endogenous to unobserved municipal characteristics, we include border fixed effects to absorb the level differences across municipalities. The instruments operate through how the relative difference in regulation across municipalities changes the relative difference in population characteristics.

Importantly, we need more than a single instrument as we have several characteristics in $\boldsymbol{X}_{kt}^{\text{location}}$ such as income, racial composition, and density. Let r denote the type of regulation, such as minimum lot size, share zoned SF, floorspace, etc. For each unit k, define $D_k^r \in \{0,1\}$ where $D_k^r = 1$ if k is on the side that is more stringent. Define the difference in stringency between one side and the other side as $S_{b(k)}^r$. We instrument the local exposure measure $\boldsymbol{X}_{kt}^{\text{local}}$ using the vector $Z_k^{\text{local}} = \left\{\text{exposure}_k\left(D_k^r S_{b(k)}^r\right)\right\}_r$ which is calculated as the 500 meter radius of exposure to the vector of municipal regulatory jumps at the border. The various dimensions of regulation provide variation to separately shift the multiple characteristics of income, race, and density if, for example, height restrictions target density, minimum lot size targets income, and share zoned SF targets race (through the wealth channel).

However, municipal differences in housing regulations cannot be used as instruments for $\boldsymbol{X}_{kt}^{\text{muni}}$ since the regulations may arise from unobserved municipal features such as local government capacity. Instead we construct Bartik-style municipal-level shocks Z_k^{muni} from the interaction of national (leave-out) changes in population characteristics and the initial shares in each municipality.

Our approach differs from the canonical border design study of Bayer et al. (2007), which does not employ instruments in its estimation of preferences for demographics and instead relies on the assumption that, after controlling for school quality, no differences in confounders across municipalities remain. However, as noted earlier, public services and administrative capacity are unobservable variables that can be correlated with demographics and may not be fully absorbed by school quality. Additionally, without separate instruments for each characteristic of income and race, the relative size of the coefficients can be affected by which variable has more measurement error.

Instruments for housing prices. For the costs of housing $C_t(k)$, we construct two instruments following Bayer et al. (2007) to address how home prices are affected by unobserved characteristics. The first is "far-ring" characteristics (shares of land use and characteristics of the housing stock in concentric rings 3–10 miles around k) Z_{kt}^{far} that shift prices through market competition. The estimating equation requires the inclusion of a rich set of "near–ring" controls Z_{kt}^{near} defined analogously to the *far-ring* instruments at 1–3 miles from k, the municipal–boundary fixed effects $\eta_{b(k)}$, and the other observables X_{kt}. The *near-ring* controls address unobservable common shocks that covary with both unit k and the *far-ring* instruments. Conditional on these variables, the exclusion restriction is that the *far-ring* instruments only affect utility through the price channel.

The second instrument for prices is a *model–predicted market–clearing* user–cost, $\widetilde{H}_{kt}^{mc}$, obtained by solving the model for the price vector that clears markets when the unobserved choice component is set to zero ($\xi_k \equiv 0$), holding fixed the first–step preference estimates and the exogenous characteristics of locations. $\widetilde{H}_{kt}^{mc}$ also satisfies the exclusion restriction once we condition on X_{kt}, near–ring controls, and $\eta_{b(k)}$ because it is a function only of exogenous spatial features (in particular, far–ring features) and the first–step parameters.

Identification requires: (i) smoothness of housing and neighborhood characteristics within municipal–boundary bands (so that $\eta_{b(k)}$ absorb slow–moving unobservables at the boundary); (ii) that near–ring characteristics capture the local neighborhood attributes that enter utility; and (iii) that far–ring characteristics shift equilibrium prices but, conditional on (i)–(ii), do not affect utility directly. The construction of $\widetilde{H}_{kt}^{mc}$ concentrates the information in the far–ring variables into a single instrument.

Estimating equations. Estimation of parameters (β, γ) happens simultaneously, so for the two endogenous variables, the first stage equations are then

$$\begin{aligned} X_{kt}^{location} &= \Psi_1 Z_m^{\text{muni}} + \Psi_2 Z_k^{\text{local}} + \omega\, \widetilde{H}_{kt}^{mc} + \Omega_1\, Z_{kt}^{\text{far}} + \Omega_2\, Z_{kt}^{\text{near}} + \eta_{b(k)} + \omega_{\text{unit}} X_k^{\text{unit}} + \zeta_{kt} \\ H_t(k) &= \Phi_1 Z_m^{\text{muni}} + \Phi_2 Z_k^{\text{local}} + \pi\, \widetilde{H}_{kt}^{mc} + \Pi_1\, Z_{kt}^{\text{far}} + \Pi_2\, Z_{kt}^{\text{near}} + \eta_{b(k)} + \pi_{\text{unit}} X_k^{\text{unit}} + v_{kt} \end{aligned} \tag{18}$$

where the excluded instruments are $\left(Z_m^{\text{muni}}, Z_k^{\text{local}}, \widetilde{H}_{kt}^{mc}, Z_{kt}^{\text{far}}\right)$ and the included controls are $\left(Z_{kt}^{\text{near}}, \eta_{b(k)}, X_k^{\text{unit}}\right)$.

We then estimate (17) by 2SLS using the fitted values $\widehat{X}_{kt}^{location}, \widehat{H}_t(k)$ from (18) and the included controls. These values deliver mean preference parameters $(\widehat{\beta}, \widehat{\lambda})$. Given the first-step estimates of $\boldsymbol{A}$ and of the parameters inside $\gamma(\cdot)$, the implied distribution of marginal willingness to pay (MWTP) for component m of X_{kt} across households with characteristics z is

$$\text{MWTP}_m(z) \;=\; \frac{\beta_m + (\boldsymbol{A}_m z)}{\gamma(z)}\,. \tag{19}$$

5.3 Building counterfactuals

With the estimated model in hand, we first consider the initial (t_0) allocation of households across housing units, measured in 2020. Assuming that amenity levels $\boldsymbol{B}(z)$, rents $\boldsymbol{R}$, and taxes $\boldsymbol{T}$ are at their steady-state levels, we compute the migration probability matrix $\boldsymbol{\mathcal{P}}(z)$. This matrix is a function of household demographics to the extent that location amenities and migration costs vary with demographics. This matrix, along with households' initial location, allows us to compute each household's probability $\boldsymbol{\pi}_{it}^0$ of being in any housing unit in subsequent years.

We then compute in the model the effect of building an additional housing unit with given characteristics (location, type, size, and tenure) on $\boldsymbol{\mathcal{P}}$, and on amenities, prices, and taxes in the first post-treatment period, t_1. We use these effects as an approximation of the permanent shocks $\Delta\boldsymbol{B}$, $\Delta\boldsymbol{R}$, $\Delta\boldsymbol{T}$, and $\Delta\boldsymbol{\pi}_t$ included in the computation of welfare effects. To validate this approximation, we show in simulations that the effects of construction on amenities, prices, and taxes in the longer run are very close to the short-run effects.

5.4 Model validation against event study findings

With the model's predictions for the impacts of building additional units on surrounding areas, we then turn to event study specifications to measure how the predictions compare to the effects of multifamily construction that took place in New Jersey. These comparisons validate whether the parameterization of the model is consistent with un-targeted moments in the data.

We estimate dynamic treatment effects of nearby multifamily construction on single-family home values using a staggered difference-in-differences (DiD) design following de Chaisemartin and D'Haultfœuille (2024). This event-study specification captures the evolution of outcomes before and after the onset of nearby construction activity. Appendix Figure C.2 shows an example of a construction event we exploit. The sample is restricted to single-family homes that are eventually treated, i.e., those located within 500 meters of a multifamily construction event at some point during the study period. Identification relies on variation in treatment timing rather than cross-sectional differences between treated and untreated locations. We estimate the following empirical specification:

$$\log\left(\text{price}_{it}\right) = \alpha_i + \lambda_t + \sum_{k\neq -1} \rho_k \, \mathbf{1}\{\text{EventTime}_{it} = k\} + \varepsilon_{it},$$

where $\log\left(\text{price}_{it}\right)$ is the transacted sales price for home i in year t, α_i and λ_t are home and year fixed effects, and $\mathbf{1}\{\text{EventTime}_{it} = k\}$ is an indicator for being k years relative to the start of nearby multifamily construction (with $k = -1$ omitted as the reference period). The coefficients ρ_k trace the dynamic effects of treatment.

We further examine heterogeneity in treatment effects by estimating the same event-study specification separately for neighborhoods of differing income levels. Specifically, we split the sample into low-income and high-income neighborhoods, defined as those in the bottom and top halves of the distribution of neighborhood median household income. The results across types of neighborhoods are shown in Figure 4.[6]

In our model framework, heterogeneity in households' preference parameters within the location-choice model implies that multifamily construction should generate differential changes in housing demand across neighborhoods. The magnitude of the empirical heterogeneity we observe in the event-study estimates should align with the magnitude of demand heterogeneity implied by the model-based predictions.

Figure 4: Multifamily construction affects home values in high-income neighborhoods

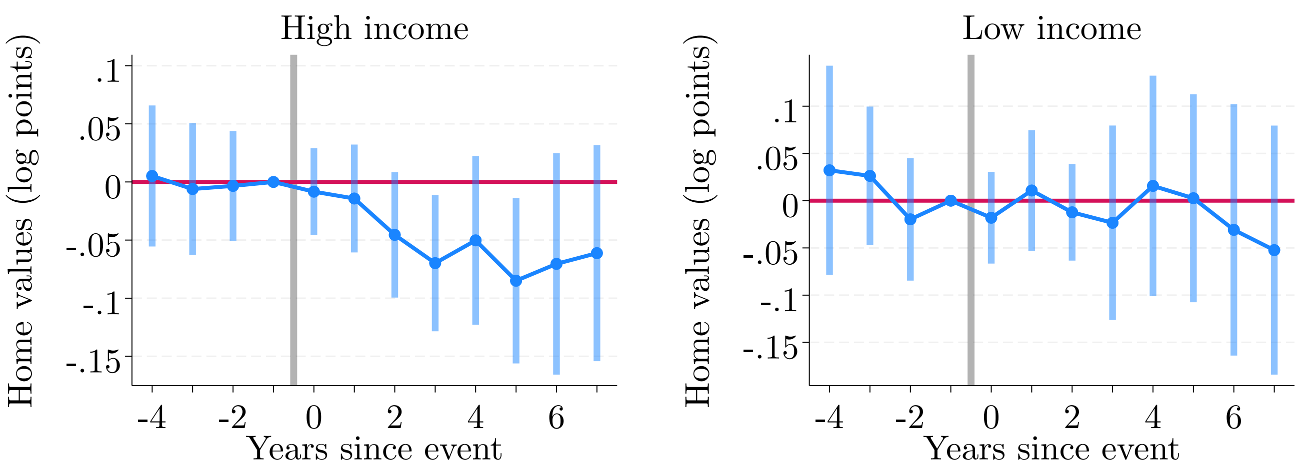

Notes: Home prices are quality-adjusted for the features: number of bedrooms, number of bathrooms, age of home, log of floorspace, and log of lot area. Treatment is defined as a multifamily construction event where at least 10 units are built within 500 meters of the single-family home. Income of the neighborhood is defined at the block group level, and high/low-income categories are designated as the top/bottom half of the neighborhood income distribution.

6 Conclusion

The rate at which new dwellings are being added to the housing stock has been declining precipitously in the United States since the Second World War (Glaeser and Gyourko, 2025). This slowdown is widely attributed to stringent land use regulations that restrict increases in housing supply and reflect political opposition from incumbent residents, who may lose from new construction through adverse shifts in prices, amenities, or local taxes.

We develop a unified framework to measure the distributional effects from new housing construction, both in the immediate vicinity of new development and in more distant locations. This

[6]Reassuringly, the adverse effects of multifamily construction we find are driven by large developments. Indeed, the construction of multifamily structures with fewer than 10 housing units do not have a discernable effect on the value of nearby single-family homes (see Appendix Figure C.3).

framework exploits increasingly available microdata on household characteristics, migration, housing units, and local governments. It is designed to clarify the reasons for which incumbent residents of an area tend to oppose construction, and measure how new development affects welfare across jurisdictions and household types.

The losses from new construction are likely concentrated in the neighborhood in which it is taking place, while gains are much more diffusely distributed. This leads housing markets to be particularly vulnerable to the formation of special interest groups who lobby against development and institutional sclerosis (in the language of Olson, 1982). We hope to shed light on the political economy forces that sustain restrictive land-use regulation, and better understand which policies are likely to make new construction more politically palatable. More broadly, our framework offers a template for evaluating the interjurisdictional consequences of local policy, contributing to ongoing debates on the optimal scale of governance in urban policy.

References

Agrawal, David R, William H Hoyt, and John D Wilson, "Local Policy Choice: Theory and Empirics," *Journal of Economic Literature*, 2022, *60* (4), 1378–1455.

Almagro, Milena and Tomás Domínguez-Iino, "Location Sorting and Endogenous Amenities: Evidence From Amsterdam," *Econometrica*, 2025, *93* (3), 1031–1071.

Asquith, Brian J, Evan Mast, and Davin Reed, "Local Effects of Large New Apartment Buildings in Low-Income Areas," *Review of Economics and Statistics*, 2023, *105* (2), 359–375.

Barseghyan, Levon and Stephen Coate, "Property Taxation, Zoning, and Efficiency in a Dynamic Tiebout Model," *American Economic Journal: Economic Policy*, 2016, *8* (3), 1–38.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 2007, *115* (4), 588–638.

_ , Marcus D Casey, W Ben McCartney, John Orellana-Li, and Calvin S Zhang, "Distinguishing Causes of Neighborhood Racial Change: A Nearest Neighbor Design," 2022.

Besley, Timothy and Stephen Coate, "Centralized versus Decentralized Provision of Local Public Goods: A Political Economy Approach," *Journal of Public Economics*, 2003, *87* (12), 2611–2637.

Blanco, Hector and Noémie Sportiche, "Local Effects of Bypassing Zoning Regulations," 2024.

Bordeu, Olivia, "Commuting Infrastructure in Fragmented Cities," 2023.

Boustan, Leah Platt, "Local Public Goods and the Demand for High-Income Municipalities," *Journal of Urban Economics*, 2013, *76*, 71–82.

Brosy, Thomas and Adam H. Langley, "When Property Values Rise, Do Property Taxes Rise Too?," Technical Report, Lincoln Institute of Land Policy 2025.

Brueckner, Jan K, "JUE Insight: Zoning and Property Taxation Revisited—Was Hamilton right?," *Journal of Urban Economics*, 2023, *133*, 103393.

Calabrese, Stephen, Dennis Epple, and Richard Romano, "On the Political Economy of Zoning," *Journal of Public Economics*, 2007, *91* (1-2), 25–49.

Calabrese, Stephen M, Dennis N Epple, and Richard E Romano, "Inefficiencies from Metropolitan Political and Fiscal Decentralization: Failures of Tiebout Competition," *The Review of Economic Studies*, 2012, *79* (3), 1081–1111.

Cook, Cody, "Heterogeneous Preferences for Neighborhood Amenities: Evidence from GPS Data," *Review of Economics and Statistics*, 2025, pp. 1–43.

Cui, Tianfang, "Did Race Fence Off The American City? The Great Migration and the Evolution of Exclusionary Zoning," 2024.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Difference-in-Differences Estimators of Intertemporal Treatment Effects," *The Review of Economics and Statistics*, 02 2024, pp. 1–45.

Diamond, Rebecca and Tim McQuade, "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development," *Journal of Political Economy*, 2019, *127* (3), 1063–1117.

_ , _ , and Franklin Qian, "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, 2019, *109* (9), 3365–3394.

Fischel, William A, *The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies*, Harvard University Press Cambridge, MA, 2002.

Glaeser, Edward and Joseph Gyourko, "The Economic Implications of Housing Supply," *Journal of Economic Perspectives*, 2018, *32* (1), 3–30.

Glaeser, Edward L and Joseph Gyourko, "America's Housing Supply Problem: The Closing of the Suburban Frontier?," 2025.

Gómez-Reino, Juan Luis, Santiago Lago-Peñas, and Jorge Martinez-Vazquez, "Evidence on Economies of Scale in Local Public Service Provision: A Meta-Analysis," *Journal of Regional Science*, 2023, *63* (4), 793–819.

Greaney, Brian, Andrii Parkhomenko, and Stijn Van Nieuwerburgh, "Dynamic Urban Economics," 2025.

Gyourko, Joseph and Sean E McCulloch, "The Distaste for Housing Density," 2024.

Hamilton, Bruce W, "Zoning and Property Taxation in a System of Local Governments," *Urban studies*, 1975, *12* (2), 205–211.

Han, Lu, Stephan Heblich, Christopher Timmins, and Yanos Zylberberg, "Cool Cities: The Value of Urban Trees," 2024.

He, Zhiguo, Scott T Nelson, Yang Su, Anthony Lee Zhang, and Fudong Zhang, "Zoning for Profits: How Public Finance Shapes Land Supply in China," 2022.

Huang, Zong, "The Welfare Effects of Local Property Taxation," 2025.

Kashner, Zane and Brad Ross, "Measuring and Mitigating Traffic Externalities," 2025.

Krimmel, Jacob, "Reclaiming Local Control: School Finance Reforms and Housing Supply Restrictions," 2021.

Li, Xiaodi, "Do New Housing Units in your Backyard Raise your Rents?," *Journal of Economic Geography*, 2022, *22* (6), 1309–1352.

Mast, Evan, "JUE Insight: The Effect of New Market-Rate Housing Construction on the Low-Income Housing Market," *Journal of Urban Economics*, 2023, *133*, 103383.

Oates, Wallace and Edwin S. Mills, *Fiscal Zoning and Land Use Controls: The Economic Issues*, Saxon House, 1975.

Oates, Wallace E, "An Essay on Fiscal Federalism," *Journal of Economic Literature*, 1999, *37* (3), 1120–1149.

Olson, Mancur, "The Principle of "Fiscal Equivalence": The Division of Responsibilities Among Different Levels of Government," *The American Economic Review*, 1969, *59* (2), 479–487.

—, *The Rise and Decline of Nations*, Yale University Press, 1982.

Pennington, Kate, "Building New Housing Cause Displacement? The Supply and Demand Effects of Construction in San Francisco," 2021.

Sahn, Alexander, "Racial Diversity and Exclusionary Zoning: Evidence From the Great Migration," *The Journal of Politics*, 2025, *87* (4).

Song, Jaehee, "The Effects of Residential Zoning in US Housing Markets," *Available at SSRN 3996483*, 2021.

Tiebout, Charles M, "A Pure Theory of Local Expenditures," *Journal of Political Economy*, 1956, *64* (5), 416–424.

Appendix

A Data

A.1 CoreLogic-HMDA matching

This Section describes how we infer the race and income of homeowners in CoreLogic by matching their mortgage deed information to mortgage loans recorded by the Home Mortgage Disclosure Act (HMDA), which documents the demographics of the borrower. As the public version of HMDA does not disclose the exact address of the purchased home, we match on Census tract. Since the borrowers are anonymized, we fuzzy match on other identifiers such as transaction date, loaned amount, loan type, and lender name. Below, we describe the data cleaning process for the two datasets that are merged and the algorithm for matching. The match results are presented at the end.

Pre-processing of HMDA data. We construct the HMDA dataset used for matching from public HMDA data on loan originations covering 2007 to 2023. We keep only loans that were originated and drop loans with missing Census tract information. We also restricted income values below $10,000 or above $15 million to remove outliers. We applied lender name standardization to remove corporate suffixes and standardized common abbreviations. For example, all mentions and variations of the terms ″LLC″, ″corporation″, ″NA″, and ″LTD″ were removed. Abbreviations such as ″BK″, ″SVGS″, and ″MTG″ were expanded to ″bank″, ″savings″, and ″mortgage″, respectively. After all of these steps, we only keep observations uniquely identified by the combination of year, lender name, census tract, loan amount, loan purpose, loan type, and co-applicant sex. This resulted in 3,777,663 HMDA loans eligible for matching.

Pre-processing of CoreLogic data. We construct the CoreLogic dataset used for matching to match the HMDA dataset as closely as possible. To this end, we apply lender name standardization to remove all variations of corporate suffixes and unabbreviated terms. The CoreLogic data had many more uncommon abbreviations (e.g. ″INVS″ for ″investment″, ″SVCNG″ for ″servicing″, ″HSNG″ for ″housing″, etc.) which we unabbreviated. As a next step, we dropped mortgages that belonged to corporate borrowers, and removed non-residential properties (e.g. agriculture, commercial, etc.) in line with reporting standards in HMDA. We also dropped HUD-originated loans as they are not included in HMDA. This resulted in 4,658,294 loans from CoreLogic eligible for matching. For dating the mortgage, we used the mortgage date if available, and if not, then the mortgage recording date, and finally the transaction batch date if other date variables were missing. Addresses were matched to Census tracts based on the HMDA schedule of tract adoptions. This means for mortgages from 2007-2011, the 2000 Census tract matching was used. For 2012 - 2021, mortgages were matched using the 2010 Census mapping, and finally for 2022 - 2023, the 2020 Census tract mapping was used. We created a co-applicant indicator that equals 1 if the second borrower name field was not empty or if the borrower et al code indicated multiple borrowers. We also standardized CoreLogic loan amounts to match public HMDA reporting conventions. This entailed rounding mortgage amounts to the nearest \$1,000 for loans before 2018. For 2018 onwards, mortgages under \$10,000 were rounded to the midpoint of the nearest thousand dollars. Mortgages over \$10,000 were rounded to the nearest \$5000 (the midpoint of the \$10,000 bins).

Matching algorithm. The matching algorithm closely follows Bayer et al. (2017). The variables used for matching are year, Census tract, loan amount, lender name, loan type, loan purpose, and presence of a co-applicant. There are seven rounds of matching which progressively relax matching criteria. Year and Census tract are required to match exactly in each round. In rounds 1 through 4, lender name is matched with a fuzzy matching command from the `fuzzywuzzy` package in Python. In the first round, an exact match is required on each variable. In the second round, we relax the co-applicant requirement. In the third round, the co-applicant requirement is reinstated and loan type requirement is relaxed. In the fourth round, co-applicant, loan type, and loan purpose requirements are relaxed. Round 5 requires an exact match on all variables except it employs a Soundex matching function from the `jellyfish` package in Python rather than fuzzy matching for the lender name. Round 6 utilizes Soundex matching and relaxes the loan type, purpose, and co-applicant criteria. The final round uses fuzzy matching on lender name and requires matching on all variables except co-applicant. In this round, the loan amount criteria is relaxed to allow matches for loans in a band of +/- \$2,000 for loans originated before 2018 and +/- \$5,000 for loans after 2018. Additionally, in the rare case of one to many matches in which one CoreLogic loan matches to multiple HMDA candidates, this is counted as a match if the race in HMDA is the same and the incomes in HMDA fall within a \$10,000 band. In that case, the match is made and the mean of the incomes is used. In sum, the iterative matching process results in a match rate of 62.56%. The bulk of these matches are completed in rounds 1 and 2, which is consistent with prior literature.

Matching results. Out of our sample of CoreLogic loans which contains 4,658,294 mortgage originations, the algorithm was successful in matching 2,914,011 loans (62.56%) to HMDA records. Of these matches, 1,383,337 (29.70%) were found in Round 1 with exact matches on all variables. Round 2 yielded a further 971,389 (20.85%). Round 3 and 4 found 23,645 (0.51%) and 271,376 matches (5.83%), respectively. Round 5 yielded a match rate of 102,142 matches (2.19%) using Soundex matching, while Round 6 identified 113,559 matches (2.44%) with Soundex and relaxed requirements. Finally, Round 7 found 48,563 matches (1.04%).

A.2 Local government finances

Overview of the Census of Governments. The Census Bureau conducts the Census of Governments every five years. We extract data on local governments' finances from the 2017 census, which documents the revenues and expenditures of the state's 21 counties and 565 municipalities (which can be cities, boroughs, townships, towns, or villages).[7] It also describes the finances of 529 school districts and 223 special districts.

School districts. Education is the most important local public expenditure in New Jersey (see Appendix Figure B.1). In 2024, there were 687 public school districts in New Jersey. Not all are included in the Census of Governments, as school districts do not all act as independent governments. Some school districts' finances are fully controlled by another local government (typically, a municipality). These school districts' revenues and expenditures are consolidated with those of the parent government in the Census of Governments.

We have associated the 529 school districts in the Census of Governments with the municipalities they serve. Some school districts (often providing secondary education) cover several municipalities. For instance, the Buena Regional School District covers both Buena Borough and Buena Vista Township. The vast majority (82%) of school districts, however, are associated with a single municipality, and the average school district serves 1.4 municipalities.

Special districts. Special districts are independent, local units of government created for a specific, limited purpose. Fire districts are the most common special districts in New Jersey (71% of the total), followed by transportation (9%), sewerage (8%), soil conservation (7%), utilities (4%), and water districts (1%). Although there are many special districts, they only account for a residual fraction of local governments' total spending (see Appendix Figure B.1).

These districts vary widely in the number of municipalities they cover. Fire districts overwhelmingly cover a single municipality, while soil conservation districts cover one or more counties.

Consolidation of municipal finances. In our analysis, we aggregate local governments' revenues and spending to the municipal level, as most public goods are provided by municipalities or other local governments that only serve a single municipality. We distribute the revenue and expenditures of counties and special districts to the municipalities they cover, proportionally to population. Similarly, we distribute the revenue and expenditures of school districts proportionally to the population below the age of 18 in the municipalities they serve.

[7] The number of municipalities in New Jersey decreased by one between 2017 and today due to a merger. In 2021, the Borough of Pine Valley (population 13) was absorbed by the adjacent Borough of Pine Hill (population approx. 10,700).

Additional local public finance data. We obtained municipality-level estimates of the average effective property tax rate (the total annual property tax as a fraction of market value) from the New Jersey Division of Taxation.

A.3 Additional data sources

School test scores. To measure local school quality, we used 4th-grade English Language Arts (ELA) and math scores from the 2019 New Jersey Student Learning Assessments (NJSLA), obtained from the state's Department of Education. We first calculated an overall score for each school district by averaging its standardized ELA and math results. We then aggregate these scores at the municipal level by calculating an average of the scores of its school districts, weighted by enrollment.

Tree canopy cover. We gather data on tree canopy cover in 2019 from the Forest Service-Field Services & Innovation Center. Its CONUS dataset, built using satellite imagery, documents the tree canopy cover at a 30-meter resolution.

B Descriptive Statistics

B.1 Local public finances

Figure B.1: Local government expenditures

(a) By government level **(b)** By expenditure category

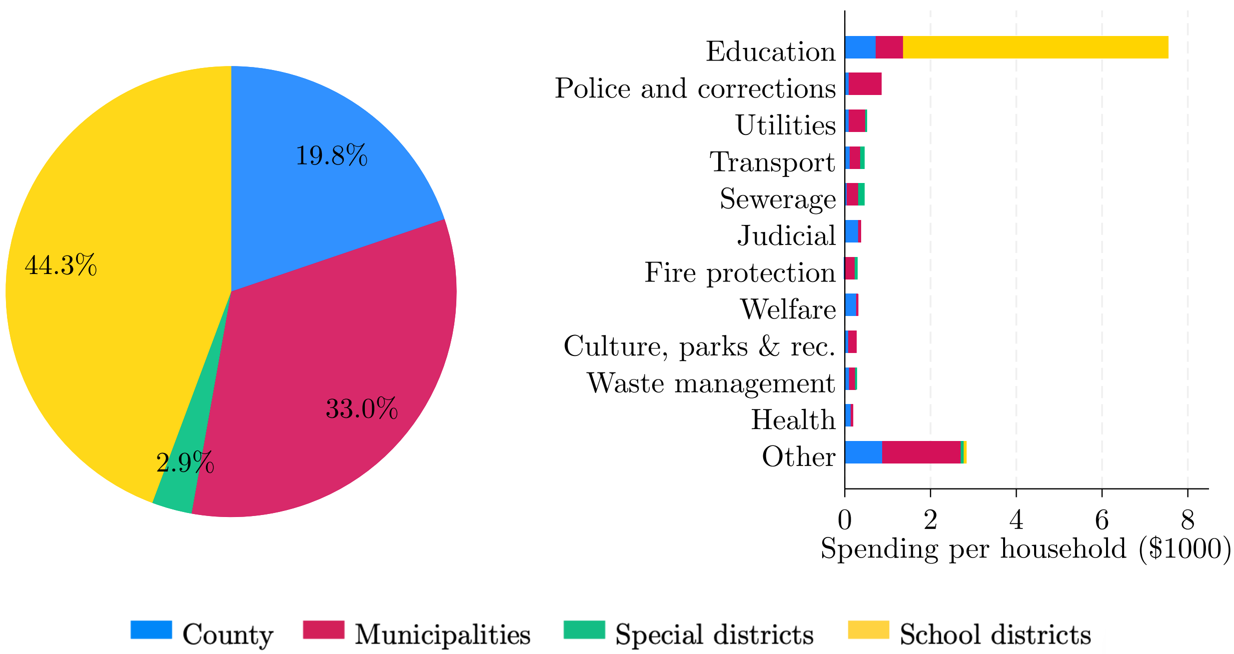

Notes: This Figure describes local governments' direct expenditures. Panel (a) shows the share of local governments' expenditures associated with each government level (county, municipality, special district, and school district). Panel (b) shows how expenditures are distributed across spending categories.

Figure B.2: Local government revenue

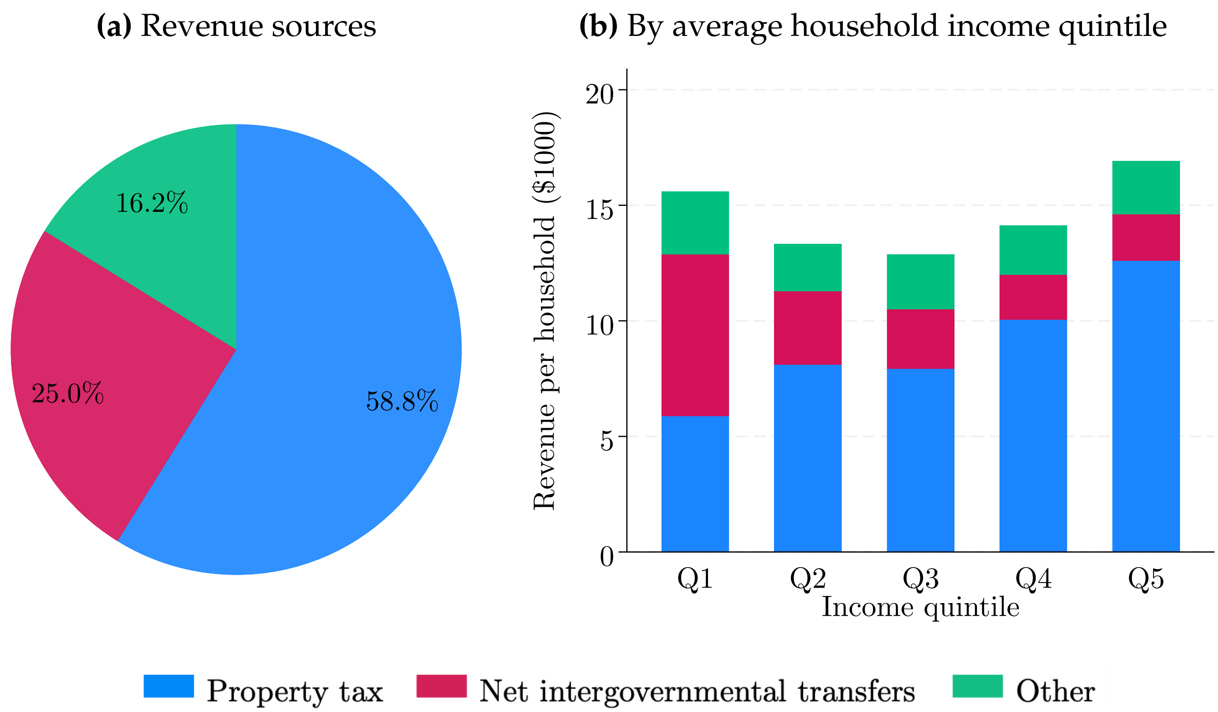

Notes: This figure describes the origin of local governments' revenue, both in aggregate for the entire state of New Jersey (panel a), and by income quintile (panel b). To build this figure, we aggregate the total revenue from property taxes, net government transfers, and other sources for counties, municipalities, special districts and school districts. In panel (b), we consolidate municipal finances by allocating counties', special districts', and school districts' revenue to the municipalities they cover, proportionally to population. We then split municipalities in quintiles of their average household income.

Figure B.3: Limited economies of scale in the provision of schooling

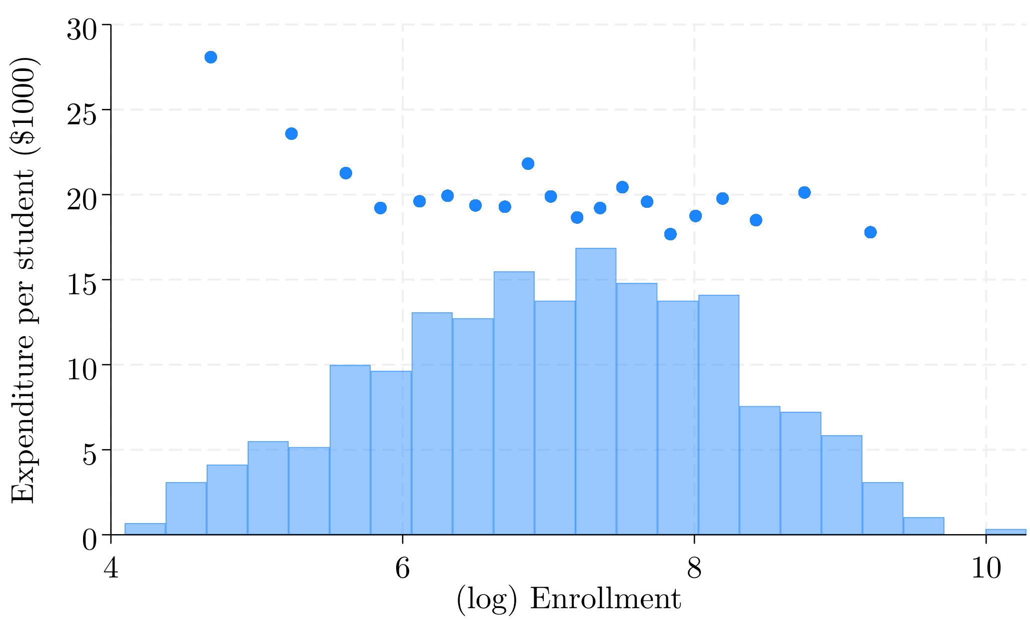

Notes: This figure shows, for the school districts included in the Census of Governments, the average yearly expenditure per enrolled student as a function of the district's total enrollment. The overlayed histogram shows the distribution of enrollment across school districts. The average school district has about 2,100 students.

B.2 Zoning

Figure B.4: Zoneomics coverage

(a) Contiguous United States **(b)** New Jersey

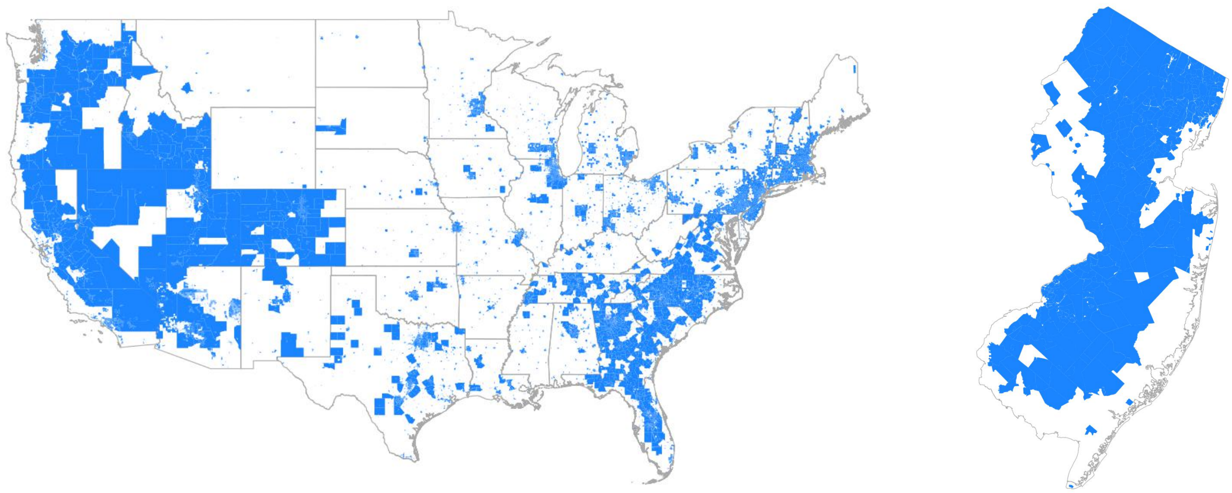

Notes: This figure maps the jurisdictions for which Zoneomics zoning data is available, both for the contiguous United States (panel a) and New Jersey (panel b).

Figure B.5: Zoned land uses

(a) Nationwide **(b)** New Jersey

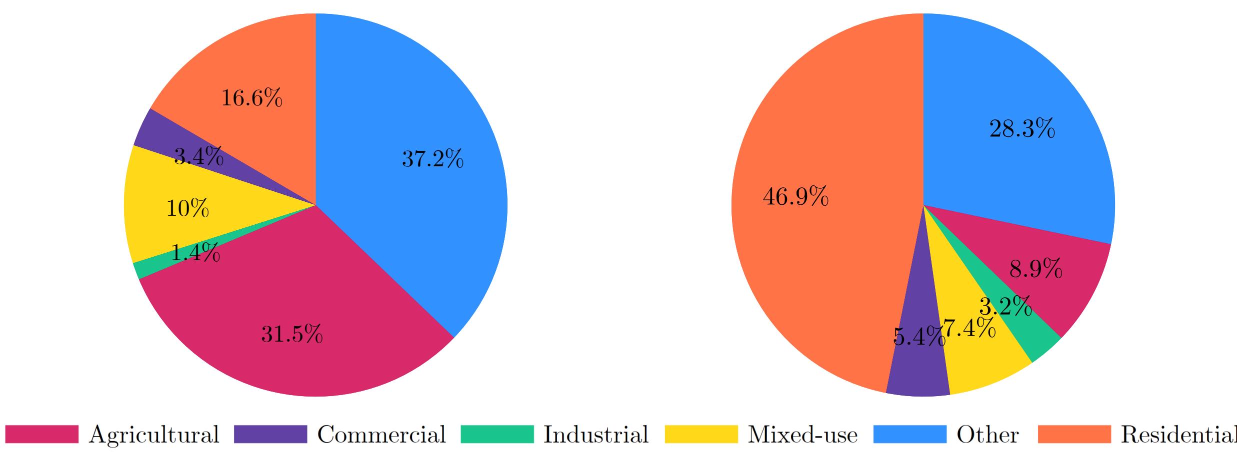

Notes: This figure shows the share of land zoned for different uses across the United States (panel a) and in New Jersey (panel b). "Other" includes protected or conservation areas, parks, government buildings, and other special-purpose zones, such as cemeteries and recreational facilities.

Figure B.6: Share of land zoned for low-density use across municipalities

(a) Nationwide

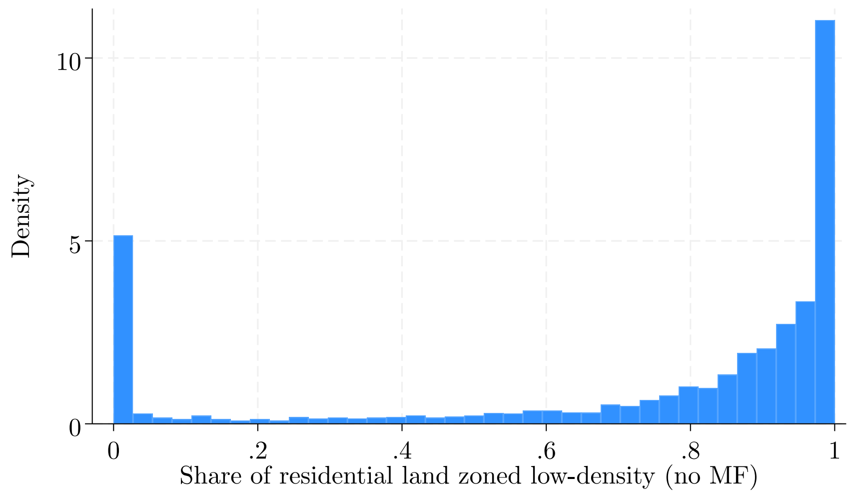

(b) New Jersey

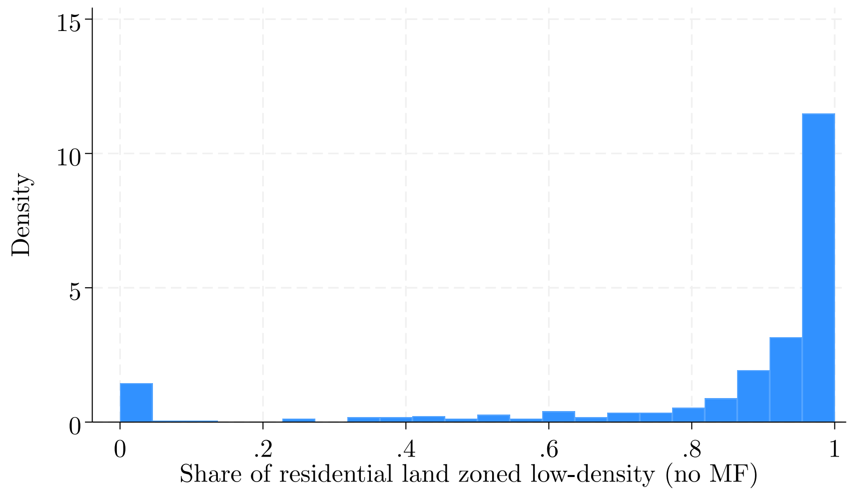

(c) New Jersey vs. United States

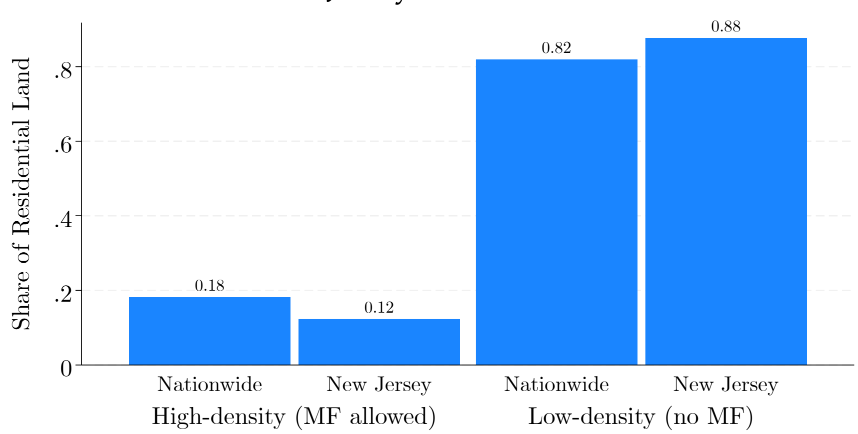

Notes: This figure shows the distribution of the share of residential land zoned for low-density use across municipalities in the United States (panel a) and New Jersey (panel b). Panel c compares the share of residential land zoned for low- versus high-density use in the United States and New Jersey. “Low-density” refers to structures accommodating up to two families, including single-family homes, ADU’s, and duplexes. “High-density” refers to structures accommodating more than two families.

Figure B.7: Stringency of zoning regulations

(a) Minimum lot size

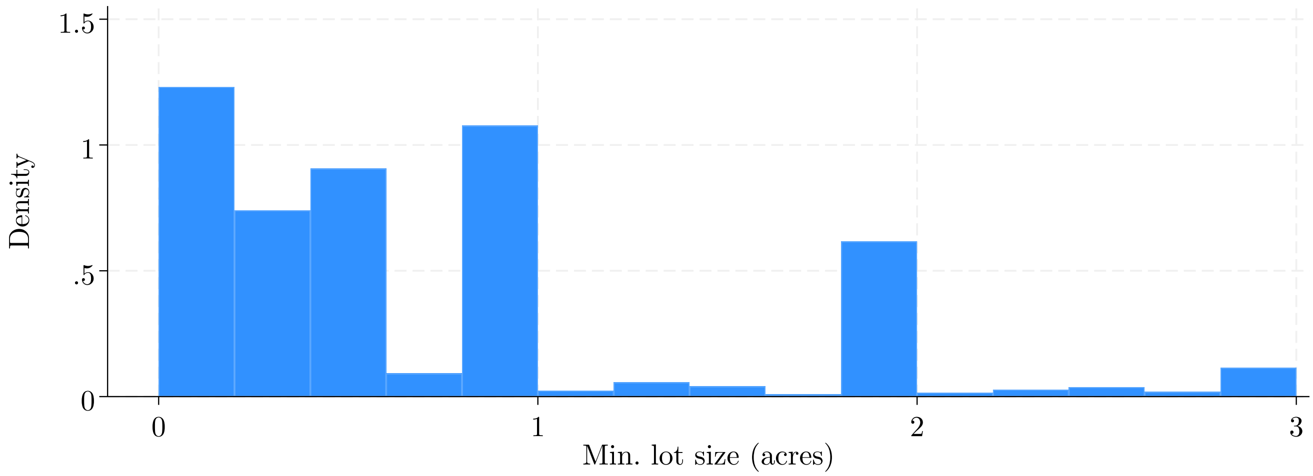

Regulation covers 53% of municipalities and 35% of land
Mean: 0.8, Median: 0.5, SD: 0.7
Variance: 40.8% between metros, 38.5% between municipalities within metros, 20.6% within municipalities

(b) Maximum building height

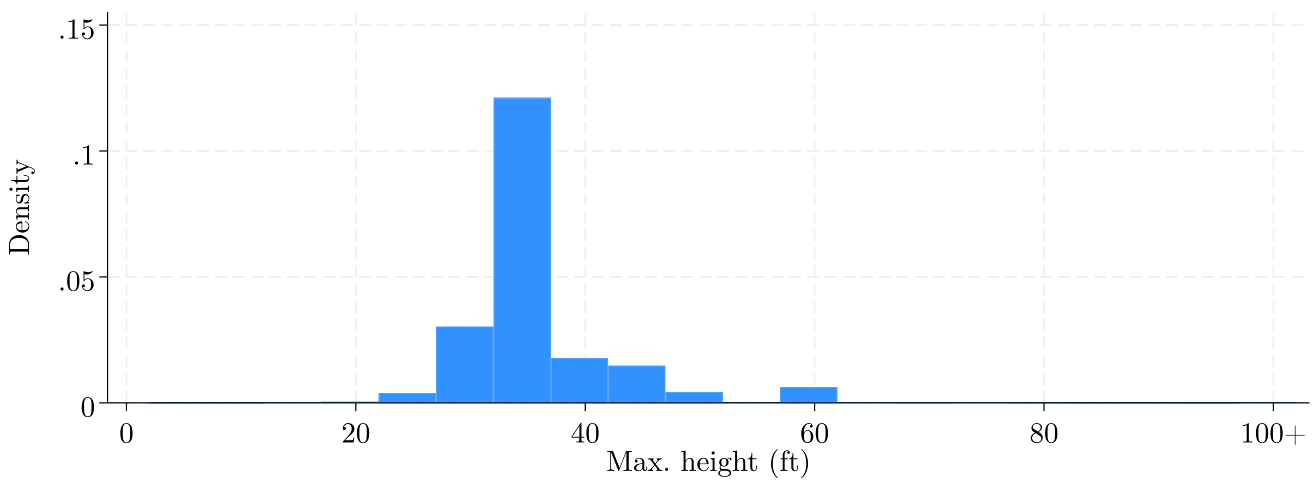

Regulation covers 57% of municipalities and 52% of land
Mean: 36.4, Median: 35.0, SD: 7.3
Variance: 33.2% between metros, 33.9% between municipalities within metros, 32.9% within municipalities

(continued)

(c) Minimum setback

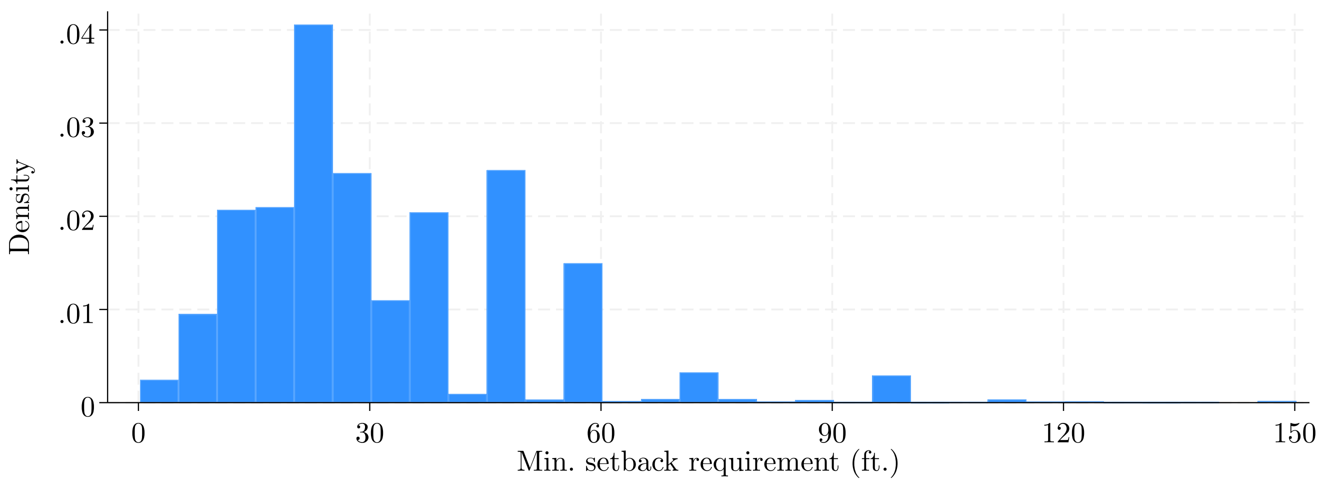

Regulation covers 60% of municipalities and 53% of land
Mean: 33.8, Median: 30.0, SD: 18.6
Variance: 42.8% between metros, 43.1% between municipalities within metros, 14.1% within municipalities

(d) Maximum density

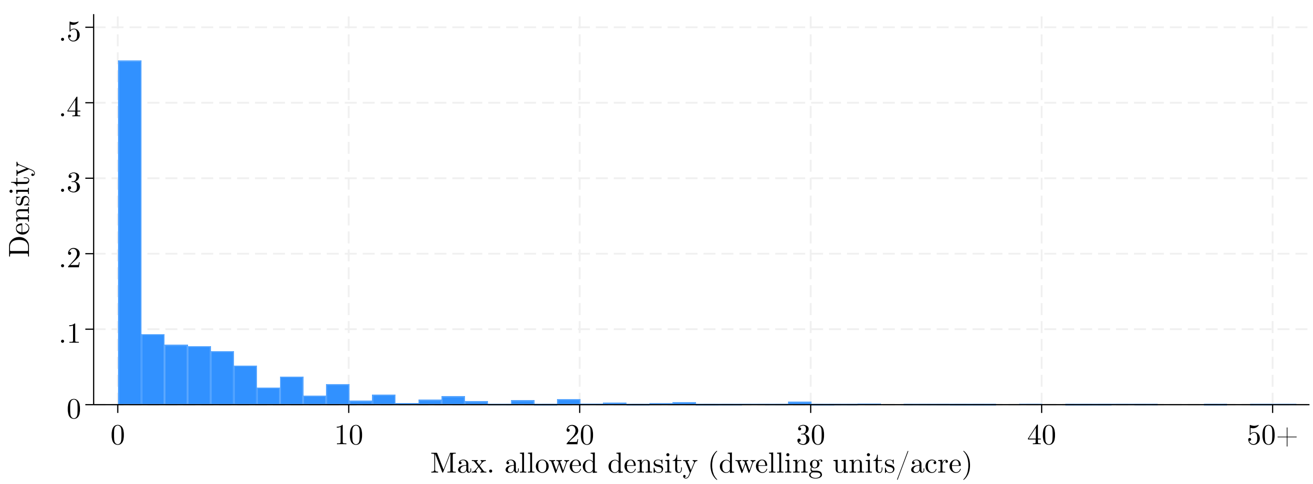

Regulation covers 15% of municipalities and 8% of land
Mean: 3.9, Median: 1.7, SD: 5.6
Variance: 31.5% between metros, 25.8% between municipalities within metros, 42.6% within municipalities

Notes: This figure shows the area-weighted distribution of the stringency of various zoning regulations across zones in the Zoneomics data. Panel (a) shows the distribution of minimum lot sizes, panel (b) shows the distribution of maximum building heights, panel (c) show the distribution of minimum setbacks (computed as the maximum of the required front yard, side yard, and back yard setbacks), and panel (d) shows the distribution of maximum densities (caps on the number of dwelling units per acre). For each regulation, we indicate the share of municipalities using it in its zoning code, the share of land subject to it, the area-weighted mean, median, and standard deviation of the regulation stringency. We also show the share of total variance attributable to differences beween CBSAs, differences between municipalities within the same CBSA, and differences between municipalities.

Figure B.8: Relationship between zoning and demographics

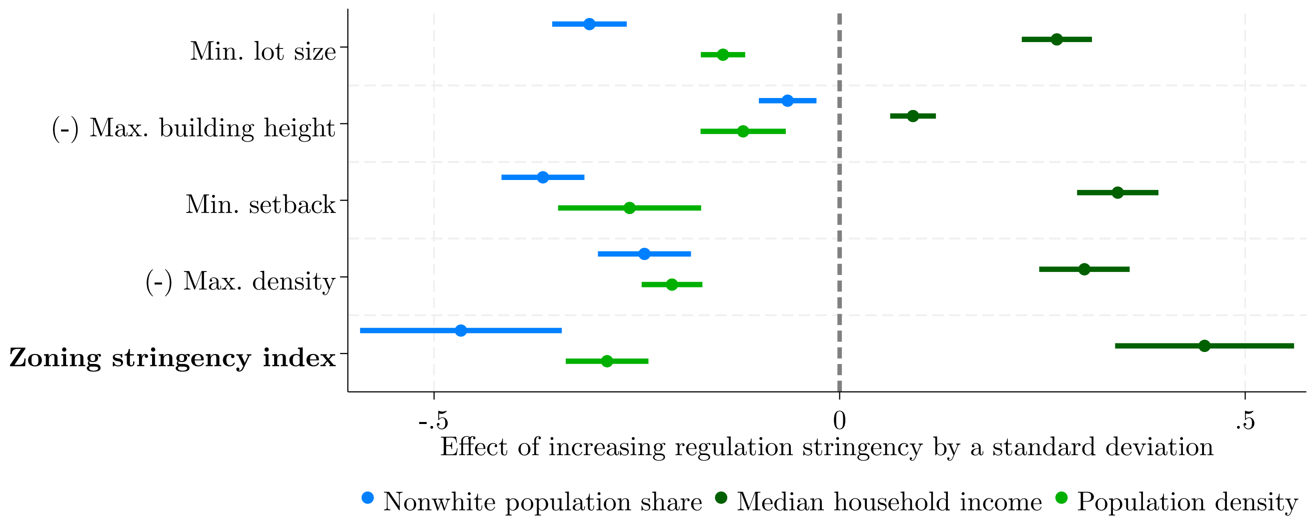

Notes: This figure reports OLS estimates and 90% confidence intervals of the effect of a one standard deviation increase in zoning stringency on demographic outcomes (also expressed in standard deviations). Each estimate corresponds to a different regression where we control for CBSA fixed effects and cluster standard errors at the municipality level. We define the minimum setback requirement as the maximum of a zone's front, rear, and side yard minimum setbacks. Maximum building height and allowable density are multiplied by –1 so that greater values indicate greater regulatory stringency. The zoning stringency index is a composite of the other four regulations, calculated for each zone as the standardized mean of their individual standardized values.

Figure B.9: Zoning stringency across CBSAs

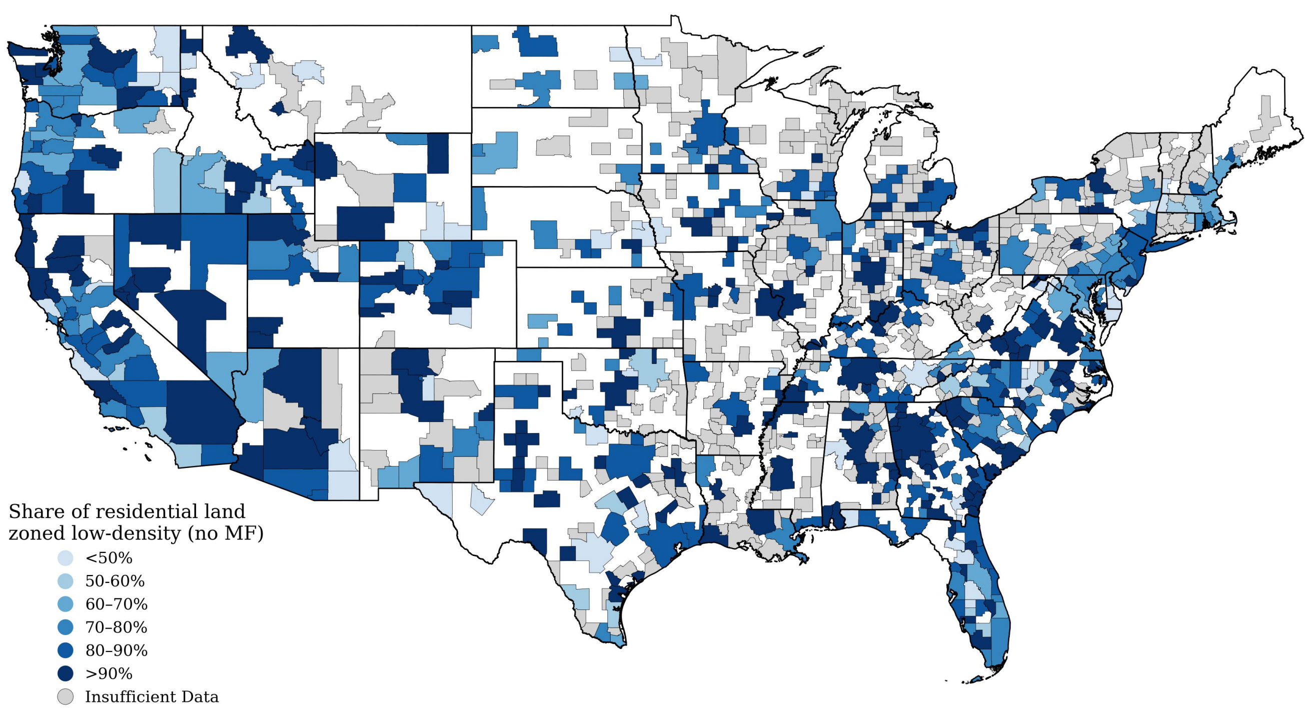

Notes: This figure maps the share of residential land zoned for non-multifamily uses across CBSAs. CBSAs for which the zoning data covers less than 50% of the CBSA's population are labeled "Insufficient Data".

B.3 Household sorting

Table B.1: Demographics of households in single-family vs. multifamily housing

	(1) (log) HH income	(2) White	(3) College	(4) (log) Home value	(5) (log) Rent	(6) (log) Prop. taxes
Panel A: Nationwide						
Multifamily	-0.690***	-0.142***	-0.069***	-0.203***	-0.263***	-0.158***
	(0.001)	(0.000)	(0.001)	(0.002)	(0.001)	(0.002)
Observations	5,605,208	5,671,001	5,671,001	3,994,343	1,579,496	3,860,480
PUMA FE	✓	✓	✓	✓	✓	✓
Panel B: New Jersey						
	(1) (log) HH income	(2) White	(3) College	(4) (log) Home value	(5) (log) Rent	(6) (log) Prop. taxes
Multifamily	-0.716***	-0.143***	-0.115***	-0.448***	-0.349***	-0.295***
	(0.006)	(0.003)	(0.003)	(0.008)	(0.007)	(0.006)
Observations	155,904	157,360	157,360	111,355	44,315	109,615
PUMA FE	✓	✓	✓	✓	✓	✓

Notes: This table shows regressions of household characteristics on a dummy indicating residence in a multifamily dwelling and Public Use Microdata Area (PUMA) fixed effects. This regression leverages microdata from the 2020 5-year ACS for the entire United States (top panel) and for New Jersey (bottom panel).

C Additional Figures

Figure C.1: Border areas in Middlesex County

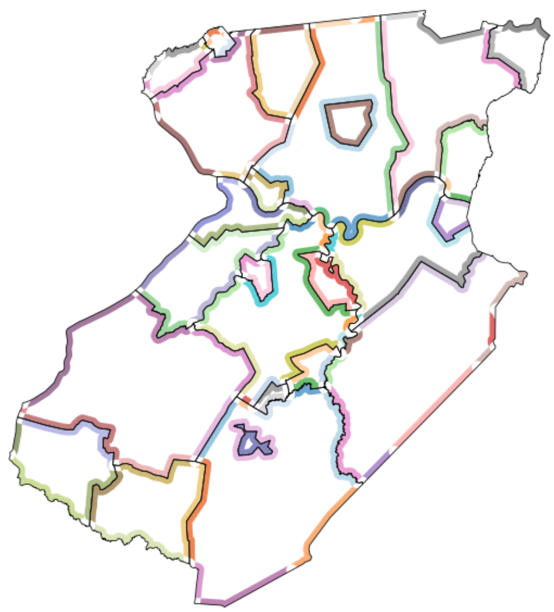

Notes: This figure shows the areas within 500 meters of a municipal boundary in Middlesex County, NJ.

Figure C.2: Example of a multifamily construction event

(a) 2007

(b) 2015

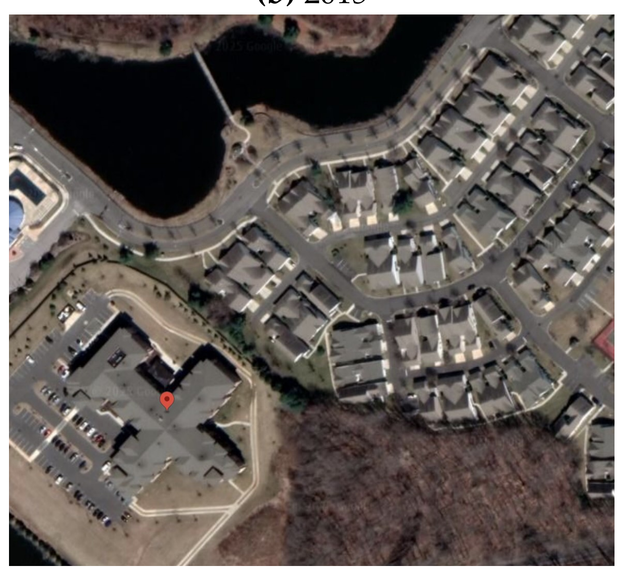

Notes: This figure illustrates one of the multifamily housing construction events we consider in the empirical design of Section 5.4. Panel (a) (resp., b) shows a 2007 (resp., 2015) satellite view of 1 Overlook Drive in Monroe Township, where a large multifamily structure was built in 2009.

Figure C.3: Effects of multifamily construction on nearby home values: Additional results

(a) Full sample **(b)** Small developments

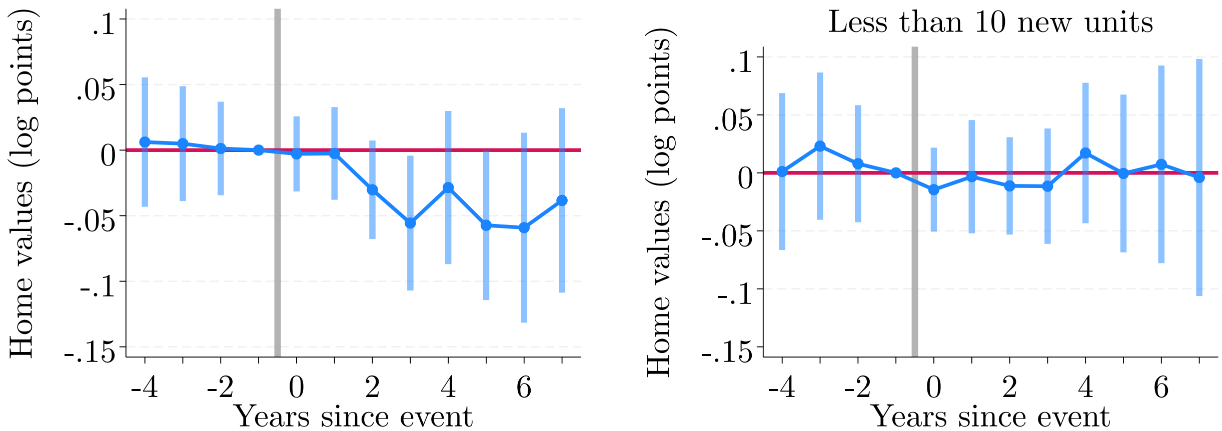

Notes: This figure shows additional results from the event study design we describe in Section 5.4, measuring the effect of new multifamily construction on the value of nearby single-family homes. Panel (a) shows the effects measured using the full sample of construction events, and panel (b) shows effects when restricting the sample to small developments, defined as those with fewer than 10 housing units.

D Estimation

Sampled alternatives

This appendix documents an acceleration of the computation of likelihoods in equation (15) based on sampled alternatives. The baseline estimation described in the main text uses the full set of housing units as the choice set for every household and period. When computational constraints make the full choice set infeasible, it is standard to replace the (very large) set of nonchosen alternatives by a random sample of size C for each observation, always retaining the chosen alternative. Under random sampling independent of utilities, the MLE for the taste parameters remains consistent as $N \to \infty$ even with fixed C. Asymptotic arguments and implementation details are provided in the Technical Appendix of Bayer et al. (2007), which relies on McFadden's (1978) "sampled choice set" results.

Construction. For each (i, t):

1. Include the chosen unit k_{it}^{obs}.
2. Draw C units without replacement from $\mathcal{H} \setminus \{k_{it}^{\text{obs}}\}$ according to a fixed sampling scheme (e.g., uniform within market t, or stratified by distance bands).
3. Form the working choice set $\mathcal{H}_{it}^{\text{wrk}} = \{k_{it}^{\text{obs}}\} \cup \{\text{the } C \text{ draws}\}$.

Likelihood with sampling correction. Let π_{kt} denote the sampling probability for unit k in period t. The logit likelihood with sampled alternatives uses the weighted utilities $u_{itk}^{\star} \equiv u_{itk} - \log \pi_{kt}$; equivalently, add $-\log \pi_{kt}$ to δ_{kt} in the sampled likelihood. The resulting weighted exogenous sampling MLE preserves the probability ratios that identify the taste parameters, ensuring consistency. In practice:

$$P_{it}^{\star}(j \to k) = \frac{\exp\left(\delta_{kt} - \log \pi_{kt} + \tilde{u}_{itk}\right)}{\sum_{\ell \in \mathcal{H}_{it}^{\text{wrk}}} \exp\left(\delta_{\ell t} - \log \pi_{\ell t} + \tilde{u}_{it\ell}\right)}, \tag{D.1}$$

$$\mathcal{L}^{\star}(\theta_1, \{\delta_{kt}\}) = \sum_{i,t} \log P_{it}^{\star}\Big(j_{i,t-1} \to k_{it}^{\text{obs}}\Big). \tag{D.2}$$

Updating $\{\delta_{kt}\}$ by contraction within the sampled sets reproduces the same fixed point conditions as in the full set, up to the known sampling offsets $-\log \pi_{kt}$.