Housing Regulation and Neighborhood Sorting across the United States*

James Macek†

June 2026

Link to the most recent version.

Abstract

In this paper, I consider the effect of minimum lot size regulation on welfare and urban structure. I construct a general equilibrium model in which households of heterogeneous incomes choose cities and neighborhoods, value affluent neighbors, and are burdened differently by regulation. The model explains two salient facts about residential income sorting: household incomes decline with residential density within cities, and this gradient is steeper in more productive cities. I provide causal evidence in support. A counterfactual deregulation exercise shows significant and progressive welfare gains for households (exceeding $2000 per household per year) that offset the losses to landowners. The exercise also reveals two surprising results. First, any productivity gains that occur from the expansion of productive cities are largely nullified by the out-migration of affluent households who prefer regulated neighborhoods. Second, deregulation does not exacerbate the neighborhood choice externality arising from the demand for affluent neighbors. These results suggest that the most important consequence of deregulating housing markets is an increase in housing affordability. Other counterfactual exercises underscore the lack of incentives for cities to unilaterally deregulate and show a significant opportunity for improved spatial targeting.

*I am indebted to Nathaniel Baum-Snow, William Strange, Kevin Lim and Joseph Steinberg for their patience, guidance and support. I have benefited substantially from discussions with David Albouy, Boaz Abramson, Kristian Behrens, Sophie Calder-Wang, Ben Coulliard, Tom Davidoff, Devin Bissky-Dziadyk, Ruben Gaetani, Kris Gerardi, Ed Glaeser, Ingrid Gould Ellen, Brian Greaney, Nina Harari, Stephan Heblich, Guangbin Hong, Allan Hsiao, En Hua Hu, Hans Koster, Gabriel Kreindler, Jeffrey Lin, Florian Mayneris, Timothy McQuade, Peter Morrow, David Krisztian Nagy, Andrii Parkhomenko, Albert Saiz, Barry Scholnik, Tsur Somerville, Aradhya Sood, and Yue Yu. Special thanks to discussants Anais Fabre, John Mondragon, Fernando Ochoa, and to Tom Davidoff for very detailed comments on an earlier draft. I also thank all participants from the 2024 AREUEA Poster Session, 2024 UEA European Meeting, the UEA Summer School, 2025 North American UEA meetings, and the 2026 AREUEA National Conference for comments and feedback.

†Alberta School of Business, University of Alberta, Canada.

1 Introduction

In recent decades, the rapid rise of US housing prices has been ascribed to strict housing regulation (Molloy et al., 2022; Gyourko et al., 2013). However, these regulations have implications that extend beyond the issue of high housing prices; they have been found to slow aggregate growth by limiting density in large cities (Hsieh and Moretti, 2019; Duranton and Puga, 2023). A particular type of regulation – the minimum lot size – also causes differences in opportunity and affluence across cities and neighborhoods by excluding those who cannot afford large lots (Song, 2025; Kulka, 2019). In this paper, I ask how these minimum lot sizes shape housing affordability, welfare inequality, and income segregation within and across cities. Understanding minimum lot size restrictions in a way that accommodates migration both within and across labor markets is important because they are prevalent, vary substantially across the US, and are an actionable policy lever (Gyourko et al., 2021; Bartik et al., 2025). While other housing regulations have likely caused income sorting (e.g. building height restrictions), the evidence suggesting that minimal lots are particularly severe motivates the focus of this paper.

Previous work on the macroeconomics of housing regulation abstracts from the income sorting that these regulations cause. A convincing argument, emphasized by Hsieh and Moretti (2019) and Duranton and Puga (2023), is that regulation slows aggregate growth by preventing workers from accessing productive cities that are responsible for that growth. However, loosening regulations in productive cities in the presence of sorting causes high-skill, productive households to leave, attenuating productivity growth that would have been achieved in the absence of such sorting. Moreover, these migration responses may be reinforced by endogenous changes to residential amenity value, a point emphasized in recent literature (Diamond, 2016; Amalgro and Dominguez-Iino, 2021). These demand-side effects have also received little attention in computing the aggregate welfare consequences of housing regulations. In this study, I examine the extent to which large-scale deregulation affects aggregate productivity, particularly in relation to the accompanying increase in housing affordability.

Regulations also cause skill sorting across neighborhoods within cities. Accounting for this sorting is crucial for gauging the welfare impacts of these regulations because they alter neighborhood quality, conferring external costs or benefits on residents. The typically held view is that regulation is a tool to limit the negative externalities associated with lower income households free riding off amenities in rich neighborhoods (Calabrese et al., 2007; Hamilton, 1975). In this paper, I also ask by how much these externalities contribute to the costs of large-scale deregulation and whether this outweighs the accompanying increases in housing affordability.

To evaluate the welfare consequences of minimum lot size restrictions, I construct a general equilibrium model encompassing the metropolitan United States. In the model, households differ in skill, consume housing, and choose cities and neighborhoods subject to a varying intensity of regulation and non-homothetic housing demand. Minimum lot sizes impose a floor on housing consumption required to live in a neighborhood, as in Kulka (2019) and Calabrese et al. (2007). Tight regulation excludes the poor by constraining their choices over small and affordable housing, thereby increasing neighborhood affluence. The model incorporates rich heterogeneity across locations along two dimensions. First, cities differ in terms of labor productivity, so any changes in labor supply across cities affect aggregate productivity. Second, neighborhoods differ both exogenously and endogenously on amenity values by skill level. Increases in average neighborhood income cause increases in amenity value, as in Brueckner et al. (1999) or Guerrieri et al. (2013), with elasticities that vary by skill.

I show theoretically that the externality arising from these neighborhood preferences can justify the use of regulation. Regulation is desirable when targeted at neighborhoods that provide high exogenous amenity value to rich households. Regulating these neighborhoods can induce the movement of the poorest households in rich neighborhoods to become the richest households in poor neighborhoods, increasing the income and amenity value of the average neighborhood. This logic has been used to argue that fiscal centralization is typically more efficient than decentralization (Calabrese et al., 2011).

The model can explain two salient facts about residential income sorting. First, within cities, there is a negative gradient between household income and residential density across neighborhoods. Second, productive cities exhibit a steeper version of this gradient and host more affluent residents on average. The model suggests a new statistic to measure regulation – the *value of a minimal lot* – whose spatial variation can explain these facts. This suggests a new mechanism driving income sorting into lower-density neighborhoods that is distinct from access to public transportation (Glaeser et al., 2008), topographical and historical amenities (Brueckner et al., 1999), or filtering dynamics (Brueckner and Rosenthal, 2009). A shift-share empirical design that generates exogenous variation in regulatory stringency across cities supports a causal interpretation.

Using this model to study deregulation is challenging because of two methodological issues. First, minimum lot sizes are difficult to measure, especially with broad geographic coverage, because they vary by local jurisdiction and are, in most cases, not publicly available. I use a similar procedure to detect minimum lot sizes to that in Song (2025) and Cui (2023), leveraging CoreLogic's property assessment database. These minimum lot sizes enter directly into the calibration of the model, along with the estimates of housing supply elasticities from Baum-Snow and Han (2024).

Second, inferring the causal effect of neighborhood affluence on amenities is difficult because unobserved amenities likely cause income sorting. Correctly identifying the strength of this relationship is important for welfare analysis because these preferences are the source of the neighborhood choice externality. I address this endogeneity issue by proposing an instrument based on terrain slopes. It has long been known that neighborhoods with steeper slopes have higher-income residents (Saiz, 2010), but these slopes are likely natural amenities and thus cannot be used as instruments alone (Davidoff, 2016). I argue that the amenity value of sloped terrain is derived from the sightlines generated within a home. This implies that the amenity value of sloped terrain decays with distance from the home, justifying the use of a neighborhood-level "donut" design: the income of a neighborhood is instrumented with the slopes of other neighborhoods that are within some distance band (Bayer et al., 2007; Anagol et al., 2021). Using the instrument, I find that a 1% increase in neighborhood income raises neighborhood value by $.2\%$ for a household with average income.[1] This elasticity increases with income, ranging from $.16\%$ for low-income households to $.31\%$ for high-income households. This implies that exogenous income changes cause self-reinforcing sorting, as in Diamond (2016) or Su (2022). The instrument corrects for a large downward bias, and the results are robust to a host of different controls, donut definitions, calibration strategies, and

[1] This is expressed in money metric units. That is, households are willing to take a 0.2% wage increase in lieu of living in a neighborhood with 1% greater income.

a placebo test that exploits time variation in neighborhood income changes.

I use the model to study the long-run implications of the nationwide elimination of lot size restrictions, paying special attention to the relative importance of its effect on housing affordability, the external costs of neighborhood choice, and aggregate labor productivity. The policy change delivers a large welfare benefit for the average household, exceeding $\$2000$ per year. On the other hand, there are large capital losses on land values of 7%.[2] Households of all income levels are made better off. In contrast with the literature, the counterfactual also reveals very little aggregate productivity gains associated with the expansion of productive cities, at 0.1%.[3] Instead, households primarily benefit from the opportunity to consume smaller and more affordable homes.

This policy change provides little evidence that the regulation corrects the neighborhood choice externality. The welfare benefits of deregulation remain essentially the same if neighborhood amenities are assumed to be exogenous. This means that minimum lot sizes are inefficient at correcting the externality: they do not appear to increase the amenity value of an average neighborhood by much relative to the accompanying distortions to housing consumption, as the theory otherwise suggests. This is in contrast to the quantitative findings in the local public finance literature (Calabrese et al., 2007). Instead, high-income households move to neighborhoods they would value in the absence of regulation, many of which I show are not strictly regulated in the data. Migration patterns within productive cities appear very similar to that of recent US gentrification: affluent households move toward central, high-density neighborhoods (Couture and Handbury, 2020; Baum-Snow and Hartley, 2020). This result is motivated by the empirical observations, which suggest that the large income-density gradient in expensive cities is caused by minimum lot size regulation.[4] Overall, these results suggest that housing affordability is the most important consequence of large-scale deregulation.

Motivated by recent policy changes in California, I also use the model to study the unilateral halving of minimum lot sizes in San Francisco. The welfare benefits of this policy change are substantial. Households of all skill levels benefit from this policy change, gaining welfare equivalent to 0.1% of their income. Aggregate land value in the city also increases modestly, at 1.5%. However, there are large distributional consequences across neighborhoods within the city. These distributional consequences are driven by endogenous changes in neighborhood amenity values. Initially stringent, low-density neighborhoods bear concentrated drops in land value as their areas become less exclusive to high-skill households. Meanwhile, high-skill households move into high-density neighborhoods, as suggested by the motivating empirical evidence. Moreover, falling land values in low-density neighborhoods attenuate aggregate land value growth in the city. Had amenity values been exogenous, land values in the city would have risen by 17% instead. This counterfactual explains why local political resistance to deregulation can arise even when the city as a whole appears to benefit. Strong preferences for affluent neighbors exacerbate the incentives for political resistance.

[2] To weigh landowner losses against household welfare, I model the disutility associated with capital losses on a land portfolio for each skill level, with more details in Section 6. The welfare benefit with this approach is smaller but still sizable, exceeding \$1500 per household per year.

[3] Skill sorting responses to deregulation drive this low value. Assuming cities changed at their predicted levels and holding city income fixed, aggregate productivity growth would instead be 1.6%. These results are also robust to considering agglomeration economies at typical values (Combes and Gobillon, 2015), production complementarities between low and high skill labor (Card, 2009), and skill-augmenting agglomeration economies (Diamond, 2016; Baum-Snow et al., 2018).

[4] In particular, I find that incomes in the highest density neighborhoods of productive cities increase by up to 30%, which is enough to invert the observed negative relationship between income and density within them.

Finally, I use the model to study the consequences of permuting observed regulation across space to target neighborhoods that are fundamentally valued by rich households, as the theory suggests to do. This policy change delivers welfare gains exceeding $\$2,200$ for the average household, benefits households of all skill levels, and – in contrast to complete deregulation – generates capital gains on national land values of 4.1%. Welfare benefits arise for significantly different reasons than that of complete deregulation. Low-income households benefit from access to affordable neighborhoods that they value. In contrast, high-income households benefit from endogenously higher amenities in neighborhoods that they value, which become protected by regulation. This policy change induces low-skill migration toward productive cities and the gentrification of high-density neighborhoods within them: a set of predictions similar to that of complete deregulation. Taken together, these results suggest that minimum lot sizes are desirable for policymakers, but only if they target the right neighborhoods. Low-income households value productive cities and low-density neighborhoods sufficiently that their exclusion from them is unjustified.

This paper builds upon several strands of literature within macroeconomics, housing regulation, urban, and public economics. First, this paper challenges the idea that housing deregulation must lead to the growth of productive cities. This is recognized by virtually all work in the macroeconomics of housing regulation as a large benefit of deregulation (Hsieh and Moretti, 2019; Duranton and Puga, 2023; Parkhomenko, 2023; Herkenhoff et al., 2018; Bunten, 2017), with the notable exception of Martellini (2022).[5] This model yields an opposing conclusion because of household skill heterogeneity and endogenous neighborhood amenity value. Stringent regulation in productive cities affect both their size and skill composition in opposing directions.

This paper also complements theory and evidence of the effects of housing regulation on housing prices, city structure and income segregation (Molloy, 2020; Gyourko and Molloy, 2015; Turner et al., 2014; Glaeser and Gyourko, 2018; Brueckner and Singh, 2020; Anagol et al., 2021; Bertaud and Brueckner, 2005; Mills, 2005; Hilber and Robert-Nicoud, 2013; Ortalo-Magné and Prat, 2014), and particularly the minimum lot size (Zabel and Dalton, 2011; Song, 2025; Kulka, 2019; Cui, 2023; Davidoff et al., 2022; Molloy et al., 2022; Kulka et al., 2026; Grieson and White, 1981; White, 1975).[6] Using a discontinuity design, Song (2025) shows extreme income and racial sorting on minimum lot sizes. I use these estimates to discipline key model parameters. Mei (2022) uses a synthetic control method to show that minimum lot size deregulation in Houston decreased prices for small homes where regulation was more likely to bind. I complement this empirical work with a model that can fit the observed regulation and quantify the severity of a salient neighborhood choice externality for the entire US economy. Generally, this literature abstracts from externalities that justify regulation in the first place.

This paper builds upon relatively recent work studying spatial sorting on income and other demographics (Diamond, 2016; Baum-Snow and Hartley, 2020; Couture and Handbury, 2020; Couture et al., 2023; Gyourko et al., 2013; Su, 2022; Baum-Snow and Pavan, 2011; Brueckner and Rosenthal, 2009; Fogli and Guerrieri, 2019; Glaeser et al., 2008; Brueckner et al., 1999;

[5] Martellini (2022) shows that deregulation in productive cities attracts less skilled workers because of selection in the labour market, and this attenuates productivity growth through knowledge diffusion. We arrive at similar conclusions but differ entirely on the model mechanisms underpinning them. I consider how regulations directly cause income sorting through preferences for housing consumption and how they affect residential amenity value. Our work jointly stresses the importance of household skill heterogeneity for understanding how deregulation affects productivity through separate and complementary channels.

[6] A small theoretical literature also examines the incentives for housing developers to impose growth controls absent local governments (Helsley and Strange, 1995; Henderson and Thisse, 2001).

Lee and Lin, 2017; Gyourko and McCulloch, 2024). I add to this literature by providing theory and evidence that housing regulation causes income sorting into the low-density neighborhoods of productive cities. A major lesson from this literature is that exogenous demographic changes can be positively reinforced by the endogenous supply of amenities. I also highlight the role of endogenous amenities in both sorting patterns and welfare gains after deregulation and carefully identify this relationship.

Finally, this paper builds on the local public finance literature, particularly the idea of housing regulation as an efficient substitute for head taxation (Hamilton, 1976; Calabrese et al., 2007; Fernandez and Rogerson, 1996, 1997; Epple and Platt, 1998; Calabrese et al., 2011; Barseghyan and Coate, 2016). In this model, the relationship between a neighborhood's amenities and its skill composition can be interpreted in the context of local public goods, creating a similar neighborhood choice externality. However, I do not endogenize the choice of regulation and instead consider it a lever that can be freely changed by a social planner. This paper makes two contributions. First, my model is calibrated with direct measures of minimum lot sizes, rather than being inferred indirectly in a political economy equilibrium. I find that regulation corrects the neighborhood choice externality by considerably less than these papers suggest, but at a similar cost of distorting housing consumption. Second, my model can match observed neighborhood skill distributions through variation in neighborhood amenities. I show both theoretically and with the calibrated model that this feature matters for understanding how well regulation corrects the neighborhood choice externality. In particular, I show that regulation is misallocated across space: large welfare benefits could be achieved by accurately targeting regulation to rich neighborhoods.

The remainder of this paper is organized as follows. Section 2 introduces the data sources. Section 3 introduces and validates the model. Section 4 calibrates the model using US cross-sectional data. Section 5 estimates the relationship between neighborhood amenity values and income. Section 6 evaluates counterfactual reforms. Section 7 concludes the paper.

2 Data

This paper draws on two main sources of data used for empirical work and to calibrate the model. Full details of the data construction are provided in Appendix A.

Geography and Demographics The primary unit of analysis is the 2020-definition census block group, which I refer to as a *neighborhood* in the model.[7] Each block group is assigned to one *city*, defined as a 2013-definition Metropolitan Statistical Area (MSA), which I treat as a self-contained labor market.[8] I take household income distributions and demographic information from the pooled 2016–2020 American Community Survey (ACS),[9] housing unit counts from the 2020 Census, and various neighborhood-level controls from the 2016–2017 National Neighborhood Data Archive (NaNDA). I also draw on the 2008–2012 ACS and the 2010 Census, both for model validation and basic empirical facts.

[7] I think of Census block groups as representing the smallest geographical unit by which there is meaningful variation in location characteristics that factor into housing demand. In addition, block groups often adhere to political boundaries and are likely to have little variation in both regulations and choices of residential structures within. Finally, block groups are small enough to reasonably capture demand spillovers that may arise from the presence of affluent neighbors.

[8] There are approximately 196,000 block groups allocated to 377 cities.

[9] This includes all sources of income reported in the ACS, including capital gains and rental income.

Property Assessments and Transactions Local jurisdictions collect detailed data on residential structures—including lot size, construction material, and building systems—to calculate property taxes. These records are digitized and harmonized by CoreLogic. I take the most recent assessment of each residential property as of December 2022, matched to 2020-definition block groups using CoreLogic's parcel coordinates. I supplement these with CoreLogic's universe of arms-length transactions from 2016 to 2022 (and 2008–2012 for robustness), which together provide lot sizes, housing values, and other raw inputs for measuring minimum lot sizes.

Minimum Lot Sizes To assess the macroeconomic consequences of minimum lot size regulation, a national database is required. Housing regulations, including the minimum lot size, are typically set by regional planning authorities and vary widely within and across planners' jurisdictions. These jurisdictions are usually incorporated municipalities. Comprehensive data on minimum lot sizes are not publicly available at the scale required for macroeconomic analysis. Therefore, I infer them from the observed *bunching* in the lot size distribution, following the well-validated approaches of Song (2025) and Cui (2023). All additional details are provided in Appendix A. The identifying logic is that if building below the legal minimum is costly, developers will partition lots that bunch at that minimum; therefore, the mode of the lot size distribution reveals the binding constraint. Figure 1 provides an example from Hayward, California, where this bunching around the mode is visible and accurately suggests a 5000 square foot minimum lot size. Throughout the paper, I assume minimum lot size regulation applies only to *regulated structures*—single-family homes, duplexes, triplexes, and quadriplexes (1–4 units per lot). This assumption is validated by a textual analysis of zoning ordinance documents from 13 municipalities. [10]

Figure 1: Bunching of the lot size distribution in Hayward, California

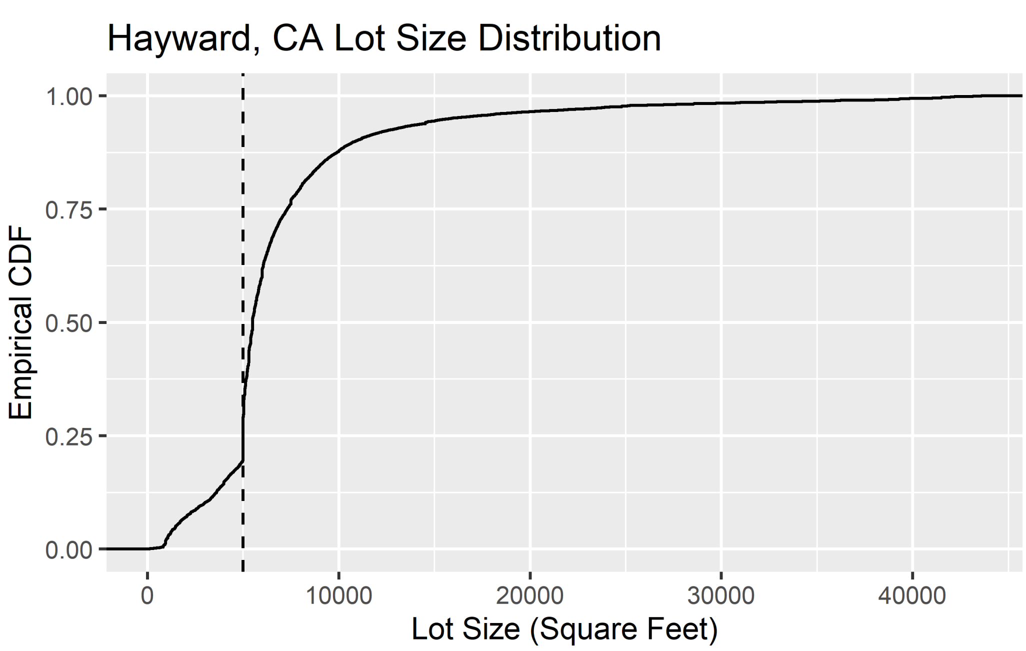

The distribution includes all structures with 1-4 housing units within the Hayward municipality. The data are censored at 1 acre.

[10]Feeding hand-collected zoning ordinance documents into a LLM reveals that more than 80% percent of zoning districts do not mention housing unit density restrictions greater than 4 units per lot. Of the 20% of zoning districts that do, restrictions on the maximum number of housing units per acre are 43 units on average, or 1 unit per 1000 square feet of lot size. This is an order of magnitude larger than the average density restriction of 4 units per acre derived from this methodology. These municipalities are Albany, GA; Atlanta, GA; Berkeley, CA; Cleveland, OH; Columbus, OH; Hayward, CA; Mesa, AZ; Miami, FL; Minneapolis, MN; New Orleans, LA; Oakland, CA; and Scottsdale, AZ.

Detecting Minimum Lot Sizes Applying the bunching estimator requires a geographic boundary within which regulation is approximately uniform – referred to as a *zoning district* (Song, 2025). For approximately two-thirds of block groups, populated zoning codes in the CoreLogic assessments directly identify these boundaries.[11] For the remainder, I cluster block groups within municipalities into zoning districts using the clustering algorithm of Chavent et al. (2018). The algorithm weighs geographic proximity and the mode of the local lot size distribution to define the clusters. These weights are the hyperparameters of the algorithm. Appendix A.2 details how hyperparameters are chosen and tests alternative jurisdiction definitions. Both empirical facts and model output are robust across the entire range of these hyperparameters.[12]

For each structure type (1, 2, 3, and 4 units) and zoning district, I estimate the minimum lot size as the smallest mode of the lot size distribution. I divide this mode by the number of implied units per lot to obtain a *housing unit density restriction*. For example, the minimum land per housing unit for a duplex is half the lot size, since two units occupy the lot.[13] The neighborhood-level unit density restriction is taken to be the minimum across all structure types.

Validating Minimum Lot Sizes To select among clustering hyperparameters and to test the accuracy of the algorithm, I use two data sources. The first is the Terner California Land Use Survey, which covers practically all municipalities in California. The algorithm predicts this data with a 7% and 16% median absolute error for single-family and multi-family lot size regulation, respectively. However, minimum lot sizes in this data are aggregated at the municipal level, while in most cases, they vary within municipalities. To test how well the algorithm works in capturing sub-municipality variation, I overlay predicted lot sizes with the MAPC Zoning Atlas. The Atlas covers all zoning districts in Massachusetts, including 504 districts in the Boston metropolitan area. Across these zoning districts, the median absolute error of the algorithm is larger but very accurate at 18.4%.[14]

Despite its precision, there are two objections to the link between observed bunching and official regulatory constraints in the literature. First, lots observed below the inferred minimum may reflect grandfathered structures built before widespread regulation was adopted, which could bias the true level of regulatory stringency (Cui, 2023). The existence of small lots may also reflect exemptions to stated regulations granted by planning authorities on a parcel-by-parcel basis – often referred to as zoning *variance*. Second, one may worry that the bunching estimator captures demand patterns or engineering constraints rather than binding legal minimums; for instance, developers may uniformly subdivide lots to take advantage of scale economies. Each can rationalize the existence of lots below the bunching point (Figure 1). I address these issues in two ways. First, I re-estimate minimum lot sizes using only structures built after 1940 and 1970; the empirical facts and model results are unchanged. Second, I re-classify all lots observed below the inferred minimum as unregulated structures, which is a conservative treatment that leaves all results qualitatively intact.

[11] I use the modal zoning code in a block group to aggregate from the parcel level to the block group level. Missing codes are omitted from the calculation of the mode.

[12] For properties with missing municipality information, I assume the county is responsible for setting regulation. This is typical for many unincorporated locations.

[13] This yields a measure of minimum land per housing unit that is comparable across structure types.

[14] In Appendix A.4, I detail how the estimated minimum lot sizes are merged to each of these data sources. Figure 8 in Appendix A.2 shows the optimal set of zoning districts for Hayward, California. The algorithm places high weight on geographic proximity in defining clusters and has an average cluster size of roughly 3 block groups, with a standard deviation of approximately 10.

In the following section, I introduce a structural framework used to assess the welfare consequences of deregulation. The model will be disciplined by this national regulatory data.

3 Theoretical framework

In this section, I present a theoretical framework that encompasses the metropolitan United States. The central mechanism of the framework is that minimum lot sizes make housing effectively indivisible: to live in a regulated neighborhood, a household must purchase at least as much housing as a minimal lot provides, regardless of what they would otherwise choose. This distinguishes minimum lot sizes from other housing regulations studied in the literature, such as floor area ratios or supply elasticity restrictions, which limit the density of floorspace but do not impose a floor on housing consumption per household. I embed this mechanism in a quantitative spatial model with heterogeneous cities and neighborhoods, validate it against two empirical facts about residential income sorting on density, and use it to study the implications of regulation for aggregate productivity, neighborhood quality, and the efficient spatial allocation of regulation. The framework is designed to transparently fit the spatial distribution of minimum lot size regulation and neighborhood income distributions in the data.

3.1 A Quantitative Spatial Model of Density Regulation

I consider a finite set of cities $\mathbb{C}$ indexed by c, which map to MSAs in the data. These cities are self-contained labor markets. Each city c has an exogenous and finite set of neighborhoods $N(c)$. I use the index i to denote a typical neighborhood from any city, $i \in \cup_{c \in \mathbb{C}} N(c)$, and define the map $\mathbb{C}(i)$ to be the city associated with i.

Neighborhoods correspond to Census block groups. Each of these neighborhoods has an exogenous amount of land T_{ir} zoned for *regulated structures* and land T_{iu} zoned for any structure.[15] I use the notation $\boldsymbol{o} \in \{r, u\}$ to index regulated and unregulated zones, respectively. Structures built in regulated zones are subject to a maximum density of housing units per unit of land.

Developers Each neighborhood and zone consists of a continuum of *sites* on the unit interval $[0, 1]$, each site containing a mass $T_{i\boldsymbol{o}}$ of land. Developers produce housing services on these sites using a constant-returns Cobb-Douglas technology over land and (numeraire) capital. This choice is well motivated by virtually all estimates of the housing production function (Epple et al., 2010; Combes et al., 2021). A consequence of the constant returns technology is that the density of housing services per unit of land is independent of the land used in production. Conditional on developing a given site, the density of housing services per unit of land is

$$A_{i\boldsymbol{o}} = \lambda_i P_{i\boldsymbol{o}}^{\epsilon_i} \tag{1}$$

where $P_{i\boldsymbol{o}}$ is the price of a unit of housing services in zone $\boldsymbol{o}$, ϵ_i is the supply elasticity of housing services conditional on site development, and λ_i is an exogenous supply shifter. With Cobb-Douglas housing production, land values are a fraction $\frac{1}{1+\epsilon_i}$ of the value of housing services $P_{i\boldsymbol{o}} A_{i\boldsymbol{o}}$. The remaining share is the cost of capital in construction.

[15] For the empirical implementation of the model, I assume regulated structures are single family homes, duplexes, triplexes, and fourplexes (structures with 1-4 housing units per lot). See the discussion in Section 2.

In a world without minimum lot sizes, developers' construction costs are independent of the allocation of housing services across housing units. This means that the profit-maximizing developer may build many small houses or few large ones, provided that the density of housing services on developed sites is given by (1). In contrast, developers in zone r respect the minimum amount of land that can be allocated to a housing unit, denoted by l_i. This is precisely what was measured in Section 2. Combined with Equation (1), this implies the minimum amount of housing services per housing unit is

$$A_{ir}\, l_i = \lambda_i\, P_{ir}^{\epsilon_i}\, l_i \tag{2}$$

Define the quantity

$$R_i = P_{ir}\, A_{ir}\, l_i \tag{3}$$

which is the cost of housing services on a minimal lot in a regulated zone. In the unregulated zone U, I assume that there are no restrictions on housing unit density. This is equivalent to imposing $l_i = 0$.

Equation (2) captures the essential difference between lot size regulation and other housing regulations studied in recent quantitative models. Contrast the equation with that of a Floor Area Ratio restriction studied in Brueckner and Singh (2020), which puts limits on the density of floorspace (or housing services) on each parcel. In this framework, there are only restrictions on the number of housing units (or households) that can occupy a given unit of land. This distinction is forcefully argued in Grieson and White (1981). Most work studying the aggregate implications of housing regulation assume that regulation affects the floorspace supply elasticity ϵ_i or construction productivity λ_i (Hsieh and Moretti, 2019; Parkhomenko, 2023; Herkenhoff et al., 2018). In this paper, minimum lot sizes block the development of low-quality housing units that would otherwise appear on small lots. The interpretation is that minimum lot size regulation only affects the supply elasticity of low-quality units.

Household consumption Households have Stone-Geary preferences over a freely traded numeraire good g and housing services A. Households differ in *skill level* indexed by $z \in Z$ and hold no land wealth. Deferring location choice for a moment, suppose a household of skill z has chosen neighborhood i and zone $\boldsymbol{o}$. Given the city $\mathbb{C}(i)$, households receive a wage $w(\mathbb{C}(i)) := w_i$ per unit of skill and enjoy a neighborhood-skill-specific *amenity value* $b_i(z)$. Given the wage and amenity, the household of skill z chooses numeraire good g and housing services A to maximize

$$V_{i\boldsymbol{o}}(z) := \underbrace{\max_{A,g} z^{-1}\beta^{-\beta}(1-\beta)^{-(1-\beta)}(A-\bar{A})^{\beta}g^{1-\beta}}_{\text{Consumption value}} \quad + \quad \underbrace{\log b_i(z)}_{\text{Amenity value}} \tag{4}$$

subject to the standard budget constraint and the *minimum lot size constraint*.

$$P_{i\boldsymbol{o}}A + g \leq w_i z \quad \text{and}$$
$$P_{i\boldsymbol{o}}A \geq R_i, \quad \text{if } \boldsymbol{o} = R.$$

The minimum lot size constraint requires that expenditures on housing be no less than the cost to consume housing that appears on a minimal lot. Preference parameters β and $\bar{A}$ are also crucial determinants of housing consumption. β governs the average expenditure share on housing services desired by households of all skill levels. $\bar{A} > 0$ governs the rate at which the desired expenditure share on housing declines in household wealth. This is a well-documented empirical relationship (Finlay and Williams, 2022; Couture et al., 2023).

In what follows, denote $C_{io}(z) := V_{io}(z) - \log b_i(z)$ as the consumption value component of the utility V_{io}. To understand how regulation distorts neighborhoods' consumption value, $C_{iR}(z)$ can be multiplicatively decomposed into two components when $\bar{A} = 0$:

$$C_{iR}(z) = \underbrace{\frac{w_i}{P_{iR}^{\beta}}}_{\text{Undistorted utility}} \times \underbrace{\left[\frac{\frac{R_i}{w_i z}}{\beta}\right]^{\beta}\left[\frac{1-\frac{R_i}{w_i z}}{1-\beta}\right]^{1-\beta}}_{\text{Distortion factor}} \tag{5}$$

whenever $\beta w_i z < R_i$, so that the desired spending on housing is smaller than the cost to consume housing on a minimal lot. The first component is the utility of consumption for a household if regulation did not bind. The second component represents by how much regulation binds: it is always less than one and decreases in the distance between the desired spending share on housing β and the spending share on housing $\frac{R_i}{w_i z}$ induced by regulation.[16]

Regulation reduces residential density by distorting housing consumption, but less so for affluent households. To see this, consider how much a household of skill z needs to be compensated, in terms of a lower per-unit price of housing P_{iR}, to be indifferent to living in a neighborhood with more expensive minimal lots R_i. In Appendix B.2, I show that

$$\frac{\partial}{\partial z}\left[\frac{\partial C_{iR}(z)}{\partial R_i} \,/\, \frac{\partial C_{iR}(z)}{\partial P_{iR}}\right] < 0 \tag{6}$$

whenever regulation is binding at z and $w_i z > R_i$. Equation (6) states that a higher-skilled household requires a lower reduction in housing prices to be indifferent to a given increase in the price of a minimal lot.

Neighborhood and zone choice Households choose neighborhoods and zones that offer high consumption or amenity value. Beyond these characteristics, households also draw idiosyncratic and additive preference shocks over neighborhoods and over zones. These shocks follow a Gumbel distribution, and are correlated across neighborhoods and zones in the same city, but not in different cities. This means that the mass of skill z households who choose neighborhood i and zone $\mathbf{o}$ is

$$L_{i\mathbf{o}}(z) = \left[\frac{W(\mathbb{C}(i), z)}{\boldsymbol{W}(z)}\right]^{\theta}\left[\frac{\exp[V_{i\mathbf{o}}(z)]}{W(\mathbb{C}(i), z)}\right]^{\rho}\bar{L}(z) \tag{7}$$

where

$$W(\mathbb{C}(i), z) = \left[\sum_{i' \in N(\mathbb{C}[i]),\, \mathbf{o}' \in \{R,U\}} \left(\exp[V_{i'\mathbf{o}'}(z)]\right)^{\rho}\right]^{\frac{1}{\rho}}$$

and

$$\boldsymbol{W}(z) = \left[\sum_{c \in \mathbb{C}} W(c, z)^{\theta}\right]^{\frac{1}{\theta}}. \tag{8}$$

$\log \boldsymbol{W}(z)$ is the expected utility of a skill-z household before the preference shock is realized, and $\bar{L}(z)$ is the mass of households of skill z nationally. $\log \boldsymbol{W}(z)$ is the standard measure of household welfare used throughout the paper. θ governs the responsiveness of migration

[16] The same logic applies in the empirically-relevant case where $\bar{A} > 0$. However, the formula is more difficult to interpret because the desired spending share on housing is not a simple transformation of β and $\bar{A}$. I derive this general formula in Appendix B.1.

across cities to changes in city value. ρ governs the responsiveness of migration across neighborhoods and zones within any given city. In what follows, I let $L_i(z) = \sum_{\mathbf{o}\in\{R,U\}} L_{i\mathbf{o}}(z)$ be the total population of z-households in neighborhood i. Moreover, let

$$C_i(z) = \frac{1}{\rho}\log\left[\sum_{\mathbf{o}'\in\{R,U\}}\left(\exp[V_{i\mathbf{o}'}(z)]\right)^{\rho}\right] - \log b_i(z)$$

be the aggregated consumption value of a neighborhood i, which summarizes the consumption value offered by each zone.

When choosing zones within a neighborhood, there is an important trade-off. In spatial equilibrium, the disutility of an expensive minimal lot must be compensated by a lower price per unit of housing services in the regulated zone relative to the unregulated zone. High-skill households are more willing to trade off expensive minimal lots for lower per unit prices. However, this logic does not hold when making comparisons of regulation *across* neighborhoods. This is because neighborhood amenities $b_i(z)$ respond endogenously to the level of regulation, thus serving as another compensating differential.

Endogenous Amenities Holding amenity values fixed, regulation must suppress utility because it constrains the housing consumption possibilities of some households. This is not true when neighborhood quality responds endogenously to regulation. I assume that amenity values $b_i(z)$ depend on the average income of the neighborhood:

$$\log b_i(z) \quad = \quad \Omega(z)\log \text{Inc}_i \; + \; \log \nu_i(z) \tag{9}$$

where $\text{Inc}_i := \frac{\sum_{z'\in Z} w_i z' L_i(z')}{\sum_{z'\in Z} L_i(z')}$ is the average income of neighborhood i. I refer to the $\nu_i(z)$ as *fundamental amenities*, which contain all other observed or unobserved neighborhood demand factors that can be reasonably taken as exogenous with respect to this model, including commuting time and natural amenities. $\Omega(z)$ governs the elasticity of amenity values to income and will be a key input to predict the counterfactual effects of deregulation. When $\Omega(z) > 0$, Equation (9) implies an externality because residents are not compensated for the deterioration of amenities caused by the location choice of low-skill households. At the end of this section, I show that this externality is more severe in neighborhoods that provide high fundamental amenity value to high-skilled households. Lower-skill households free-ride on the amenities that appear endogenously in such neighborhoods. Regulation could be a tool to stop free-riding.

I emphasize one microfoundation for Equation (9). Local governments provide congested public goods financed through property or income taxes (Calabrese et al., 2007, 2011). In this case, income per capita is replaced with property tax revenue per capita. In a model with Cobb-Douglas preferences over housing, no minimum lot sizes, and random heterogeneity in property tax rates, this is, on average, identical to income per capita. With Stone-Geary preferences and minimum lot sizes, the relationship between property tax revenue and neighborhood income is typically not linear, but they are still highly correlated. In Appendix B.3, I provide microfoundations for this mechanism, but I do not limit the interpretation of (9) to these. Instead, I consider $\Omega(z)$ to reflect all factors that could be caused by the compositional effects of affluence. Apart from what was mentioned, these may include reduced crime, peer effects (Chetty and Hendren, 2018), or a general taste for affluent neighbors (Guerrieri et al., 2013; Brueckner et al., 1999).[17]

[17] Alternatively, I could have specified that amenities depend only on the average skill (and not income, which

How households trade off consumption and amenities When making location decisions, households trade off high-amenity, high-income neighborhoods with high per-unit housing prices and/or regulation. Households differ when making this tradeoff, and this difference is governed by the elasticity $\Omega(z)$. To see this, consider how large a change in consumption value must be to make a z-household indifferent to a neighborhood with marginally higher average skill. This is

$$\frac{\partial C_i(z)}{\partial \log \text{Inc}_i} = -\Omega(z) \tag{10}$$

which decreases in $\Omega(z)$.[18] In Section 5, I causally estimate that $\Omega(z)$ is strictly increasing in skill relative to the marginal utility of income. Taken together with Equation (10), the model predicts that high-income households sort into expensive, regulated, and high-amenity neighborhoods, as in the Tiebout sorting literature (Epple and Platt, 1998).

Production In each city c, production of the numeraire good g takes place competitively with a constant-returns technology

$$g_c = \iota_c \left[\sum_{i \in N(c)} \sum_{z \in Z} z L_i(z) \right] \tag{11}$$

where ι_c is the exogenous labor productivity in city c. In equilibrium, it must be that $\iota_c = w_c$ and so I refer to both interchangeably. Following Hsieh and Moretti (2019), I define aggregate labor productivity as the productivity of the numeraire good sector:

$$\tilde{g} = \frac{\sum_{c \in C} g(c)}{\bar{L}} \tag{12}$$

where $\bar{L}$ is the total mass of households nationally. Aggregate productivity is high when productive cities host more households (Duranton and Puga, 2023; Hsieh and Moretti, 2019; Herkenhoff et al., 2018). This is achieved when deregulation expands the housing supply in productive cities. Additionally, I argue that differences in the average skill level between cities matter when assessing the impacts of deregulation. This is because deregulation also causes low-skill households to sort into productive cities and high-skill households to sort out, highlighting a new and important margin.

With the housing and labor markets defined, I turn to the definition of an equilibrium with exogenous productivity.

Definition 1. *An equilibrium is defined as a set of housing prices $P_{i\boldsymbol{o}}$, neighborhood population allocations $L_{i\boldsymbol{o}}(z)$, amenities $b_i(z)$ such that*

1. *Labour Markets clear: Given indirect utility $V_{i\mathbf{o}}(z)$ solving (4), and amenities $b_i(z)$ solving (9), labour supply per household skill $L_{i\mathbf{o}}(z)$ solves (7) at wages $w_c = \iota_c$.*

2. *Housing Markets clear: Given population $L_{i\boldsymbol{o}}(z)$, the neighborhood demand for housing services derived from (4) equals the total supply of housing services in every neighborhood i and zone $\boldsymbol{o}$. The total supply of housing services is given by the density of housing services (1) multiplied by the mass of land $T_{i\boldsymbol{o}}$.*

includes the effect of city productivity). In the baseline model, city wages are exogenous. Consequently, this alternative model makes numerically identical predictions of the effects of deregulation. I confirm that a model deviating from the assumption of exogenous wages (e.g., agglomeration economies) makes virtually identical predictions. This distinction is immaterial for the conclusions of the paper.

[18] The formula assumes a given household draws the same idiosyncratic preference shocks for all neighborhoods. Alternatively, this corresponds to a thought experiment faced by the household with average draws of these preference shocks.

In Section 6, I use the equilibrium model to assess the effects of deregulation on welfare, welfare inequality, and the spatial distribution of income. To motivate this assessment, I show that this model can explain the observed income sorting on residential density, and provide evidence that this relationship is plausibly causal.

3.2 The model can explain nationwide skill sorting on density

Through Equation 5, the model makes a stark prediction: the extent to which regulation inflates land consumption in regulated zones, reduces residential density, and causes skill sorting is entirely dependent on the value of a minimal lot R_i. The housing unit density restriction l_i is itself not a sufficient statistic. This is because a large minimum lot can only bind housing consumption when land values are high. A one-acre minimal lot in San Francisco binds more than a one-acre minimal lot in Tulsa, Oklahoma.

This reasoning raises a question. How well does the empirical counterpart of R_i explain the variation in household income and residential density, both within and across cities? Is this relationship plausibly causal? To answer these questions, I propose the following statistic $\tilde{R}_{ic}$ that can be constructed for each MSA c and block group i. The statistic is the product of three components:

$$\tilde{R}_{ic} = \text{HousingValueDensity}_{ic} \times \text{UnitDensityRestriction}_{ic} \times \text{FractionUnitsRegulated}_{ic} \quad (13)$$

where $\text{HousingValueDensity}_{ic}$ is the housing value per acre constructed from the CoreLogic data.[19] $\text{UnitDensityRestriction}_{ic}$ measures the minimum land per housing unit and is the product of the structural break detection algorithm from Section 2. $\text{FractionUnitsRegulated}_{ic}$ is the ACS fraction of households living in *regulated structures*, which I define as those between 1 and 4 housing units per lot. I assume that density regulation does not apply to lots containing more than 4 units. Through the lens of the model, $\tilde{R}_{ic}$ can be interpreted as the expected price of a minimal lot faced by a randomly selected household in i whom may choose between structures in regulated and unregulated zones.

Empirically, $\tilde{R}_{ic}$ is a strong predictor of neighborhood affluence. A regression of average income on $\tilde{R}_{ic}$ using block group level variation yields an R^2 of 40%. Crucially, the prediction capabilities of $\tilde{R}_{ic}$ are not driven by any one component in Equation 13, but by the interaction of all three. Adding each combination of components and their pairwise interactions in this regression increases the R^2 by only 2 percentage points. Additionally, I show that the variation in $\tilde{R}_{ic}$ can explain two features of the relationship between income and residential density, as summarized below.

Fact. *(The geography of residential income sorting)*

1. *Within cities, there is a negative gradient between household income and residential density across neighborhoods.*

2. *Productive cities: 1) have higher income residents and 2) exhibit steeper income-density gradients.*

[19] I measure housing values per acre using a interpolation procedure detailed in Appendix A.3. To be consistent with the model, $\text{HousingValueDensity}_{ic}$ is not adjusted for quality variation across neighborhoods (i.e., through a hedonic regression). This is because large minimal lots beget higher quality, expensive housing, and housing quality is one of the margins by which regulation binds. There is extensive causal evidence of the effect of density regulation on housing quality (Song, 2025; Kulka et al., 2026).

To demonstrate these facts, I proceed in two steps. First, I group each MSA into a "superstar" sample containing cities that are in the top quartile of average residualized wages, and a "non-superstar" sample containing all other cities.[20] Second, using the 2016-2020 ACS sample and the 2020 Census, I estimate the following partially linear model separately for each city group:

$$\log \text{Income}_{ic} = \mathbf{S}(\text{HousingDensityQuantile}_{ic}) + \sum_{x \in \mathbf{Controls}} \beta_x x_{ic} + \text{Error}_{ic} \tag{14}$$

where $\text{HousingDensityQuantile}_{ic}$ is the quantile of the density distribution of housing units within the city, and **S** is a flexible spline function.[21] For the baseline specification, **Controls** include only CBD distance rankings (all results are robust to their exclusion). Each variable, including controls, are demeaned at the MSA level. The objective is to document the shape of the function **S** (the income-density gradient) across samples, which is shown in Figure 2, Panel A.

Figure 2: The geography of residential income sorting

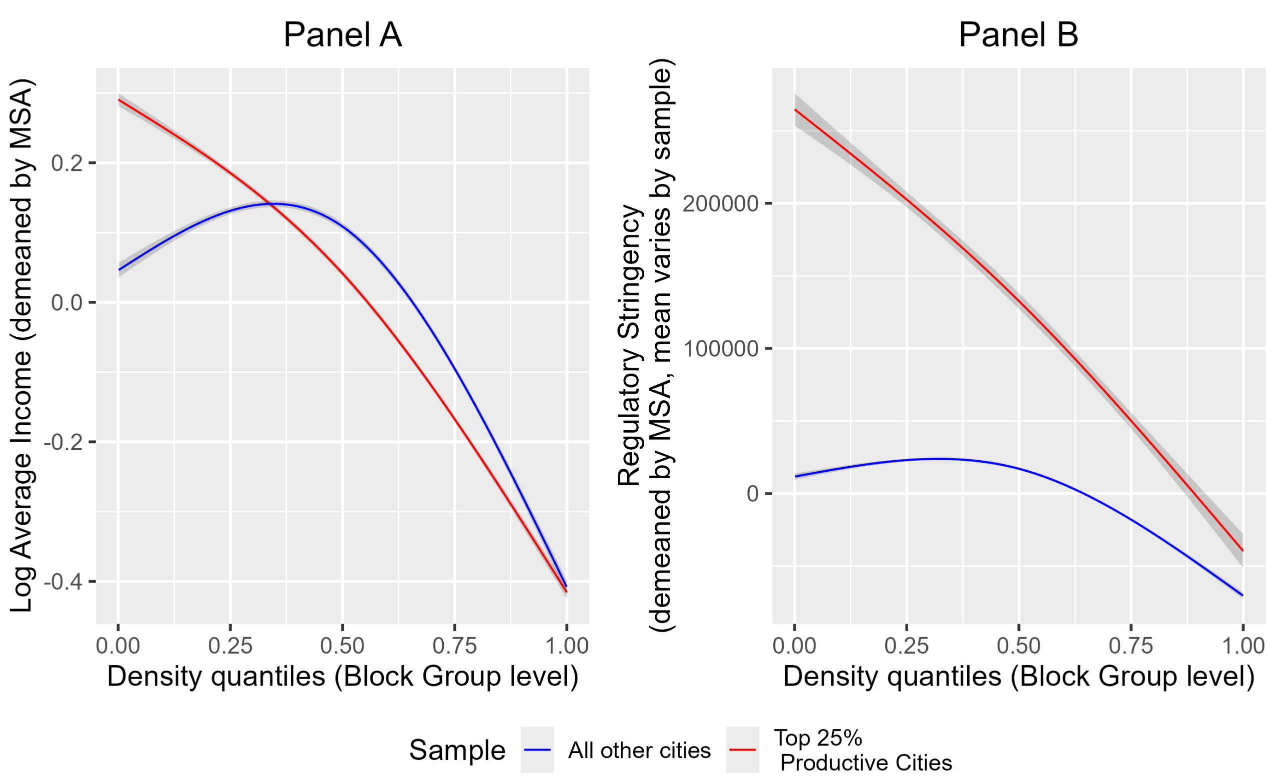

Plot of cublic spline **S** from Equation (14) with 5 knots, estimated with the *mgcv* package in R (Wood, 2012). 95% confidence sets are reported. Panels A and B report models that only control for CBD distance. Results are quantitatively similar when including a full set of controls for the share of white households, average building age, household size, share of households using public transport to commute, share of households using cars to commute, and average commuting time from the ACS; density of performing arts, spectator sports, casinos, recreational activities, restaurants, fast food, bars, and coffee shop establishments from NaNDA; and the density of EPA toxic releases and the density of public transport stops from NaNDA. The results are also robust to controlling for the variance in log wages at the city level, demonstrating that the steeper income-density gradient is not driven by greater wage inequality in productive cities. All variables from the ACS vary at the block group level and at the tract level for NaNDA. Panel B regressions are defined by taking the estimates using MSA demeaned variables and adding back the difference in raw means of $\tilde{R}_{ic}$ between each sample to the superstar cities estimate.

Panel A of Figure 2 reports **S** for superstar cities (in red) and non-superstar cities (in blue). I emphasize two observations. First, there is a sizable negative income-density gradient for the majority of neighborhoods in all cities. Differences in income between the lowest and highest density neighborhoods exceed 0.4 log points, or 50%. Second, there are differences in the income-density gradient across samples: low-density neighborhoods in productive cities

[20] Residualized wages are derived from a standard Mincer regression that parses worker characteristics from city wage fixed effects. They measure differences in productivity across cities net of sorting on household skill. I provide additional details of this regression in Section 4.

[21] Density is defined as the total number of housing units divided by the total land mass of the block group.

are relatively more affluent than those in unproductive cities, and vice versa. Visually, this corresponds to the blue curve being below the red curve for neighborhoods at the bottom of the density distribution. The differences in these gradients are economically meaningful. In "superstar" cities, the lowest density neighborhoods are 25% richer relative to their cities' average income. For non-superstar cities, this is less than 5%.

To show that the variation in $\tilde{R}_{ic}$ across neighborhoods can explain these facts, I estimate (14) after replacing $\log \text{Income}_{ic}$ with $\tilde{R}_{ic}$. To emphasize across-city variation, I additionally allow for the mean of $\tilde{R}_{ic}$ to differ across city samples.[22] The results are reported in Panel B of Figure 2. Qualitatively, Panel B exhibits patterns similar to those of A. $\tilde{R}_{ic}$ decreases in residential density for most neighborhoods, and the gradient is steeper in productive cities. $\tilde{R}_{ic}$ is also higher on average in productive cities, which explains some sorting on household skill into productive cities (Diamond, 2016; Baum-Snow and Pavan, 2011; Combes et al., 2008). Taken together, Panel B shows that the gap in $\tilde{R}_{ic}$ between these samples is driven by the lowest-density neighborhoods. In the highest-density neighborhoods, this gap virtually disappears.[23]

These facts bear an important implication for the model. In Section 6, I show that in an equilibrium with no regulation, differences in income-density gradients across city samples disappear entirely. The model predicts that, in the absence of regulation, these facts are no longer facts.

Causal Evidence In the model, regulation causes skill sorting by pricing out households who cannot afford large lots. I argued this can explain observed skill sorting into productive cities, and steeper income-density gradients in productive cities. However, interpreting this causally is difficult when regulation is set by local voting – an issue of reverse causality (Parkhomenko, 2023; Calabrese et al., 2007; Hilber and Robert-Nicoud, 2013). Residents of high-value neighborhoods have an incentive to impose regulation to increase land values (Fischel, 2001).

To provide plausibly causal evidence, I test a unique prediction made by the model. Minimum lot sizes have remained relatively stable over time, if not increased (Gyourko et al., 2021). As a result, the model predicts that city-wide growth in housing prices should make neighborhoods with large minimal lots more expensive and thus more stringent (a consequence of Equation 13). The Facts establish that these tend to be low-density neighborhoods.

[22] I provide summary statistics for this measure and its components in Appendix A.5; they are additionally broken down by superstar city sample, reported for select cities, and plotted against measures of city productivity.

[23] Conditioning each regression on an additional set of controls listed in Figure 2 yields qualitatively identical results, so the associated plots are omitted. Quantitatively similar results are also obtained in the absence of controls for the CBD distance ranking. These facts are also robust to alternative weights, clustering schemes, definitions of superstar cities using housing prices, density, or productivity (wages) alone, and various combinations of control variables, time periods, and the omission of neighborhoods with a density of less than 500 households per square mile. They are also robust to: 1) using demeaned residential density instead of quantiles as the regressor and 2) controlling for city-level variance in log wages. This allays concerns that these observations are driven mechanically by greater income inequality in productive cities (Parkhomenko, 2025; Eeckhout et al., 2014). However, these facts do not hold when density rankings are replaced with distances to the central business district. These robustness checks are discussed in more detail in Appendix A.6.

Common objections to the structural break method in Section 2 also apply to these facts. Reassuringly, these facts are robust to 1) treating housing units on lots with sizes below the detected minimal lot size as "unregulated structures" in the computation of FractionRegulated, and 2) using properties built after thresholds of 1940 and 1970 in the measurement of minimum lot sizes. The latter robustness check allays concerns that these facts are driven only by developers' incentives to subdivide lots uniformly (e.g., to take advantage of economies of scale in home design).

Finally, I note the non-superstar sample estimate in Panel A of Figure 2 exhibits a non-monotonic relationship. Interestingly, the same non-monotonic relationship occurs with $\tilde{R}_{ic}$ in Panel B.

Consequently, city-wide price growth should make the income-density gradient steeper.

To test this, I consider an MSA-level regression of the form

$$\Delta\text{IncomeDensityGradient}_c = \beta \times \Delta \log P_c + \text{Error}_c \tag{15}$$

where $\Delta\text{IncomeDensityGradient}_c$ is generated from a linear regression of city-demeaned log income on density quantiles (after controlling for CBD distance), and taking long differences between the 2020 and 2010 ACS samples. $\Delta \log P_c$ is the long difference in a repeat sales housing index from the CoreLogic data over the same period. The model predicts that $\beta < 0$: price growth makes the negative gradient steeper. However, OLS is likely biased because cities that experienced high price growth saw rising incomes in central, high-density neighborhoods for reasons independent of regulation (Baum-Snow and Hartley, 2020; Couture and Handbury, 2020; Couture et al., 2023).

Identification requires an instrumental variable that 1) shocks housing prices and 2) is unrelated to unobserved factors causing changes to the income-density gradient. I rely on housing demand shocks generated from a striking national trend: the employment of individuals above the age of 55 increased by over 50% between 2010-2020. This reflects the aging population and increasingly delayed retirement. The instrument takes a shift-share form:

$$\tilde{l}_c = \sum_{k \in \boldsymbol{K}} \omega_{kc} \hat{L}_{kc} \tag{16}$$

where $\hat{L}_{kc}$ is the gross national employment growth of workers in age bracket k between 2010-2020 leaving out city c's contribution, ω_{kc} is the initial 2010 share of employed workers of city c who are in age bracket k, and $\boldsymbol{K}$ denotes the set of age brackets.[24] These data are derived from the LODES database. Identification requires that the initial employment by age shares ω_{kc} are uncorrelated with unobserved factors shifting the income-density gradient. This assumption alone is likely to be violated because households of different age groups have different incomes and sort differently on density. For example, there has been a net migration of young people toward downtown neighborhoods (Couture and Handbury, 2020). Moreover, housing price growth causes greater wage inequality, which may mechanically cause stronger income sorting across neighborhoods (Parkhomenko, 2025). Therefore, identification relies on controlling for changes in the age-density gradient and changes in the within-city variance of wages across households. The estimates of this regression are reported in Panel A of Table 1.

Columns (1) and (2) begin with OLS estimates. They confirm that the OLS relationship between price growth and the income-density gradient is small, consistent with reasons for upward bias that were described above. The IV estimates are reported in Columns (3) and (4), with the independent variable standardized. The estimates are negative, substantially larger in magnitude relative to OLS, and statistically significant. The point estimate of $-.10$ is economically significant. A one-standard-deviation increase in housing price growth implies that the gap in income between high- and low-density neighborhoods increases by approximately 10%. Adding controls in Column (4) reduces the magnitude of the point estimate, but only slightly. The stability of the IV coefficient allays concerns that the instrument, which is correlated with age demographics, is correlated with the non-regulatory determinants of income sorting on density.[25]

[24] I use three age brackets to define the shift-share that are available in the LODES dataset. These are: 1) 29 or younger, 2) 30-54 and 3) 55+.

[25] Panel B of Table 1 reports the estimates of the first stage regression. Large Kleibergen-Paap F statistics

Table 1: Housing Price Growth and the Income-Density Gradient

	OLS		IV	
Panel A: IV	(1)	(2)	(3)	(4)
Dep. Var: Δ Inc. Density Gradient				
$\Delta \log P_c$ (standardized)	-0.010	-0.019**	-0.108***	-0.096**
	(0.011)	(0.009)	(0.038)	(0.040)
Δ Age-density gradient		-0.013**		-0.011*
		(0.005)		(0.006)
Δ Var(log wages)		-1.124***		-0.868*
		(0.380)		(0.517)
N	361	361	361	361
R^2	0.005	0.141	-	-
Panel B: First Stage				
Dep. Var: $\Delta \log P_c$ (standardized)				
Shift-share age shock			-0.362***	-0.337***
			(0.104)	(0.103)
Δ Age-density gradient				-0.010
				(0.041)
Δ Var(log wages)				2.417
				(3.701)
Kleibergen-Paap First Stage F-statistic			30.8	32.49

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Panel A, Columns (1) and (2) report OLS regressions of specification (15), where the independent variable of interest $\Delta \log P_c$ is standardized to have a standard deviation of one. Columns (3) and (4) report IV regressions where the employment shift share (16) is used as an instrument. Columns (2) and (4) additionally control for housing prices and populations in 2020 (omitted). Panel B reports the first-stage regressions for IV models (3) and (4). All regressions are weighted by the 2020 city population. First stage power is reliant on weighting by city population. Results are virtually identical when defining the income density gradient to be unconditional on CBD distance rankings.

Replicating the Facts with the model I have established that expensive cities host higher skilled households and exhibit stronger negative skill sorting on residential density, and that this can be explained by spatial variation in the prices of minimal lots. With little structure on the model, I construct an equilibrium by which a chosen set of these values reproduces both facts.

To this end, I consider the following stylized geography. Suppose there are C cities indexed by $c \in \{1 \dots C\}$, with two skill levels z_0 and z_1 such that $z_0 < z_1$. Cities differ in productivity, which is ordered as $\iota_c < \iota_{c'}$ whenever $c < c'$. Each city has two neighborhoods indexed by $i \in \{0, 1\}$, which do not differ in any characteristic except for regulation. I also abstract from the distinction between regulated and unregulated zones. Cities and neighbor-

confirm instrument relevancy. Interestingly, the shift-share instrument $\tilde{l}_c$ is strongly and *negatively* correlated with employment and price growth. This deviates from the typical logic of the industry employment shift share: national employment growth in an industry (e.g., software services) would manifest in the disproportionate expansion of employment in cities that initially specialize in it (for example, San Francisco or San Jose). This means that cities with initially young workers experienced the greatest housing price growth. I argue that this negative correlation can be rationalized by the greater spatial mobility of young working adults relative to older adults (Greaney, 2025). The increasing employment of workers approaching retirement age manifests as more local employment of these workers in cities with older populations. This likely crowded out the employment of younger adults in these cities, causing them to move to cities that employ young workers more intensively. As a consequence, cities with more young workers in 2010 grew the most and experienced higher house price growth.

hoods differ in the value of a minimal lot R_{ic} as they do in the Fact above:

$$R_{ic} = \alpha_c i \tag{17}$$

where α_c is a function such that the ratio α_c/ι_c is strictly increasing in c. Within cities, neighborhoods are ordered by average skill and inversely ordered by density because expensive minimal lots attract higher-skill households that consume more land per household. This variation also causes skill sorting in high-productivity cities. This logic is summarized in Proposition 1.

Proposition 1. *Suppose $\bar{A} = 0$, $\Omega(z) = 0$ for all z, and perfect mobility ($\rho, \theta \to \infty$). Under the parameterization in (17) with α_c/ι_c strictly increasing in c, the following must hold:*

1. *City skill is increasing in c*
2. *In each city c, the density of housing units is decreasing in i*
3. *Every city has a negative skill-density gradient. The magnitude of the slope of this gradient is increasing in c.*

Proof. See Appendix B.4.[26]

Having built a connection between the model and the data, I proceed to discuss two aspects of the model and their importance in determining the consequences of national deregulation. The first aspect concerns how aggregate productivity responds to deregulation in the presence of skill sorting. The second pertains to the efficacy of regulation in correcting the neighborhood choice externality. I conclude with a welfare analysis via the Social Planner's Problem, which incorporates each aspect.

3.3 The importance of skill sorting for aggregate productivity

The literature emphasizes that housing regulation depresses aggregate labor productivity because it limits the size and output of productive cities. However, these papers abstract from modeling household skill heterogeneity. I argue that skill heterogeneity is important for assessing the productivity benefits of deregulation. With skill heterogeneity, city size depends on both the number of households and the average household skill. If regulation causes productive cities to have higher average skill, as suggested by Proposition 1, then the productivity gains from deregulation could be offset by the out-migration of high-skill households and the in-migration of low-skill households.

To understand better, consider the thought experiment of increasing the stock of housing units in San Francisco by 15%. In a world with no skill heterogeneity across households, San Francisco's output would increase by 15%, reflecting equal population growth. However, in a world with skill heterogeneity, migration toward San Francisco would comprise households with below-average skills. Consequently, San Francisco's output would increase by less than 15%. High-skill households might leave as a result of declining neighborhood amenities and increased residential congestion, attenuating output growth even further.

[26] I assume $\bar{A} = 0$, $\Omega(z) = 0$, and perfect mobility to improve the exposition of the proof. In doing so, I abstract from other motives for skill sorting that could be generated by non-homothetic preferences or self-reinforcing sorting on endogenous neighborhood amenities. The goal is to isolate the effect of regulation itself.

I connect this logic to the model. Recall that productivity is described by the total output of non-housing production per capita (Equation 12). Totally differentiating this expression with respect to the equilibrium population of skill z households:

$$\frac{\partial \tilde{g}}{\tilde{g}} = \frac{1}{\bar{\bar{L}}} \frac{\sum_{c \in C} \sum_{i \in N(c)} \iota_c \times z \times \partial L_i(z)}{\tilde{g}} \tag{18}$$

holding location choices of other household types fixed. This expression indicates that the marginal effect of a household's migration on aggregate productivity increases in z. This means that aggregate productivity gains are smaller when population flows $\partial L_i(z)$ toward high-productivity cities comprise of disproportionately low-z households. In the deregulation exercise of Section 6, I find comparable magnitudes of population flows $\sum_{z,i} |\partial L_i(z)|$ after deregulation to the literature, but the effect on aggregate productivity is significantly lower. This attenuation effect is compounded by the deterioration of neighborhood amenity values caused by the migration of low-z households. To better understand this quantitative result, I present a simple two-city model in Appendix B.5 and prove that the incorporation of skill heterogeneity into the model causes lower productivity growth after deregulation.

3.4 Externalities, skill sorting, and neighborhood quality

Under the simplified model of Proposition 1, all skill sorting across neighborhoods and cities is induced only by regulation, and there are no externalities arising from neighborhood choice because $\Omega(z) = 0$. Regulated, high-income neighborhoods in productive cities offer cheaper prices per unit of housing services to compensate for the excess housing consumption required to live there. Additional welfare inequality is induced by the creation of rich neighborhoods.

In this section, I contrast this regressive outcome with one where regulation, when targeted to neighborhoods that would be rich absent regulation, can increase the amenity values of all neighborhoods and facilitate efficient Tiebout sorting. This is achieved by reallocating the poorest households in rich neighborhoods to become the richest households in poor neighborhoods, thereby increasing the average income everywhere. The neighborhoods that would be rich in the absence of regulation, and thus should be targeted, depend crucially on the distribution of fundamental amenity values $\nu_i(z)$. Endogenous amenities that appear in these neighborhoods create both the incentive for lower-skill households to free-ride and the scope for regulation to limit this free-riding. The idea that sorting absent regulation creates an externality has been used to argue that fiscal centralization across neighborhoods is more efficient than decentralization (Calabrese et al., 2011, 2007). The reasoning behind this argument is the same as that used here.

To build a better understanding, I detail a simple model. Suppose there are three income types $z_l < z_m < z_h$ (low, medium, and high), one closed city, and neighborhoods $i \in \{0, 1\}$ with an identical production technology for housing. These neighborhoods will be ordered in equilibrium by affluence. I consider a set of *economies* that are parameterized by $t \in \{0, 1\}$. These economies vary in terms of the structure of their fundamental amenities. The $t = 0$ economy is characterized by no variation in income across neighborhoods, which can be achieved by setting $\nu_{t=0,i}(z) = 1$ across all neighborhoods and skill types. In contrast, the $t = 1$ economy is associated with the largest variance in neighborhood income in the absence of regulation. This is achieved by restricting $\nu_{t=1,1}(z_h) = 1$, $\nu_{t=1,0}(z_h) = 0$, $\nu_{t=1,0}(z_l) = 1$ and $\nu_{t=1,1}(z_l) = 0$; along with $\nu_{t=1,i}(z_m) = 1$ for every i. Under this parameterization, income sorting is at its strongest because there will be complete segregation of z_l and z_h types in either neighborhood.

Proposition 2. *Assume that fundamental amenities take the described form for a set of economies parameterized by $t \in \{0, 1\}$.*
Then, the following is true about an equilibrium where regulation does not bind:

1. ***Inclusionary Zoning**: In the $t = 1$ economy, a marginal increase in regulation in $i = 1$ increases average income in all locations.*

2. ***Exclusionary Zoning**: In the $t = 0$ economy, a marginal increase in regulation in any neighborhood(s) does not increase average income in all locations. Instead, the average income across neighborhoods weighted by the population of z_h types increases.*

Proof. See Appendix B.6 □

Proposition 2 suggests that the efficacy of regulation depends in part on how well it targets neighborhoods that high-income households fundamentally value. An important question is which, if any, of the economies considered in Proposition 2 characterizes the actual world in which we live. The model can distinguish between these hypotheses because it can fit variation in fundamental amenities informed by neighborhood choice data.

In the main deregulation exercise of Section 6, I find that the contribution of changing amenities to overall welfare is negligible for the average household, but there are large changes in the neighborhood income distribution and desegregation overall. Other quantitative applications of the model predict that fundamental amenities matter for the welfare effects of deregulation. I show that broadly shared welfare gains can be achieved by accurately targeting regulation to neighborhoods with high $\nu_i(z)$ for high-skill households, as Proposition 2 suggests.[27]

In the following section, I complete the theoretical analysis by solving the Social Planner's Problem. This problem weighs the benefits of the regulation outlined in this section against the accompanying distortions in housing consumption.

3.5 The Social Planner's Problem

Regulation distorts housing consumption: How does its efficiency compare with other policy tools? It is well known that head taxation and transfers can correct fiscal externalities in local public finance (Calabrese et al., 2011) and other salient externalities in urban economics (Fajgelbaum and Gaubert, 2020; Donald et al., 2024). In this section, I consider how a social planner would optimally design a system of spatial taxes and transfers (equivalently, head taxation) to correct the neighborhood choice externality. I also compare this problem with what regulation achieves. This problem is nested within the framework of Fajgelbaum and Gaubert (2020). For the simplicity of the exposition, I assume that the within- and across-city migration elasticities are equal, or $\theta = \rho$. I also assume that there is no distinction between zones within neighborhoods.

[27] I want to stress that variation in fundamental amenities across neighborhoods is *not* required for minimum lot size regulation to be desirable. Net welfare gains from regulation (with redistribution) when fundamental amenities do not differ across space are achievable through classical Tiebout sorting – a better allocation of congested public goods to income levels – as demonstrated in Calabrese et al. (2007).

The **Social planner's problem** for a set of household weights $\{\alpha(z)\}_{z\in Z}$ and landowner weights α^L is defined as the choice of numeraire consumption allocations $g_i^C(z)$ and $g_i^L(z)$ for households and landowners respectively, housing consumption allocations $A_i(z)$, and total capital inputs into housing production g_i^A to solve

$$\max \sum_{z\in \boldsymbol{Z}} \alpha(z) \log \boldsymbol{W}(z) + \alpha^L \Pi \tag{19}$$

where $\log \boldsymbol{W}(z)$ is the average household welfare (8), and $\Pi := \sum_{i\in N} g_i^L$ is the total amount of the numeraire paid to landowners, subject to the following resource and free mobility constraints:

$$\sum_{i\in \cup_{c\in \mathbb{C}} N(c)} \Big[\sum_{z\in Z} g_i^C(z) L_i(z)\Big] + g_i^A + g_i^L = \underbrace{\sum_{c\in C}\left[\iota_c \sum_{i\in N(c), z\in Z} z L_i(z)\right]}_{\text{Total production of numeraire}} \tag{20}$$

$$\forall i, \ \sum_{z\in Z} A_i(z) L_i(z) = \underbrace{\lambda_i^{\frac{1}{1+\epsilon_i}} {g_i^A}^{\frac{\epsilon_i}{1+\epsilon_i}} T_i^{\frac{1}{1+\epsilon_i}}}_{\text{Local production of housing services}} \tag{21}$$

$$\forall i, z, \ \ V_i(z) - \frac{1}{\theta}\log\left[\frac{L_i(z)}{\bar{L}(z)}\right] = \log \boldsymbol{W}(z) \tag{22}$$

$$\forall z, \ \sum_{i\in N} L_i(z) = \bar{L}(z) \tag{23}$$

These constraints are intuitive. Equation (20) requires that the total production of the numeraire balances with the social planner's allocations across space. Equation (21) represents the same condition, but for housing consumption, which cannot be traded across neighborhoods. Equation (22) is an algebraic manipulation of the household's utility-maximizing neighborhood choice when $\theta = \rho$ (Equation 7). The final constraint ensures the population is balanced for each skill level. In Proposition 3, I derive two necessary conditions for a solution to this problem that hold irrespective of welfare weights.

Proposition 3. *Let $V_i^g(z)$ and $V_i^A(z)$ be the marginal utility of numeraire and housing consumption, respectively, for a skill z household living in i. Suppose $\theta = \rho$. The following must be true in a social planner's allocation:*

1. ***Efficiency of the housing sector**. For all i and z,*

$$\underbrace{\frac{V_i^g(z)}{V_i^A(z)}}_{\text{-MRS of housing for numeraire}} = \underbrace{\frac{\epsilon_i}{1+\epsilon_i}\left[\frac{g_i^A}{\lambda_i T_i}\right]^{-\frac{1}{1+\epsilon_i}}}_{\text{Marginal product of capital in } i\text{'s housing sector}}$$

2. ***Spatial efficiency**. For all i and z*

$$\underbrace{\sum_{z'\in Z} \Omega(z') \frac{\partial \log Inc_i}{\partial \log L_i(z)} V_i^g(z')^{-1} L_i(z')}_{\text{Total amenity spillovers from skill } z \text{ households}} + \underbrace{\iota_c z}_{\text{Productivity gains}} - \underbrace{\frac{1}{\theta} V_i^g(z)^{-1}}_{\text{Spatial redistribution}} =$$

$$\Xi(z) + \underbrace{g_i^C(z) + \frac{V_i^A(z)}{V_i^g(z)} A_i(z)}_{\text{Social planner's costs}}$$

for some constants $\Xi(z)$ that depend only on household skill z.

Proof. See Appendix B.7 □

These two conditions are informative. The first expresses a standard condition for consumption and production efficiency: the local marginal rate of substitution for housing and the numeraire good equals the local marginal rate of technical substitution. This condition does not hold for any allocation in which regulation R_i is binding because it distorts housing consumption. Lower-skilled households observe larger deviations from this efficiency condition.

The second concerns itself with an efficient spatial allocation induced by these chosen transfers. When designing an allocation that induces a household to move to another neighborhood, a social planner must balance the **benefits** minus the **costs** of doing so *across all neighborhoods*. There are three benefits to doing so, represented by each term on the left-hand side of Condition 2. The first term is the sum of the amenity spillovers arising from the reallocation of a skill-z household. In some allocations, the amenity spillover term from middle-skill households can be simultaneously negative in rich neighborhoods and positive in poor neighborhoods, implying that there are benefits to the redistribution of these households (as discussed in Proposition 2). The second benefit arises from reallocating a household to a higher productivity city, which increases the production of the numeraire good; this is the channel emphasized by Hsieh and Moretti (2019) and others. This marginal benefit is larger when reallocating a higher-skilled household. The third arises from a well-known redistributive motive caused by idiosyncratic taste shocks.[28] Each of these benefits is expressed in units of the numeraire good. In contrast, the *costs* of household reallocation are the total value of what the social planner allocates to them in the destination neighborhood. These are represented by two terms on the right-hand side of Condition 2. These include both the allocation of housing services expressed in units of the numeraire good and the numeraire itself.

A key consideration is that regulation does not attain the social planner's ideal. Housing regulation may induce movement that brings the market allocation closer to spatial optimality, but it must necessarily do so at the expense of distorting housing consumption. Counterfactual exercises in Section 6 suggest that distortions to housing consumption are more costly.

3.6 Model extensions

Before using the model to quantify the effects of deregulation, I discuss three extensions used to probe the robustness of the quantitative results.

Endogenous productivity The production technology in Equation (11) is restrictive in three important ways. First, population flows across cities cause changes in city productivity via agglomeration effects (Combes and Gobillon, 2015). I extend the baseline production technology so that wages respond to city size with elasticity α, or

$$\iota_c = \tilde{\iota}_c L_c^{\alpha} \tag{24}$$

where $\tilde{\iota}_c$ is an exogenous component of city productivity, and L_c is the population of city c. Second, the relative flows of high- and low-income workers alter the relative wages they earn

[28] It is well known that idiosyncratic preferences over locations introduce redistribution motives if there are differences in the marginal utility of consumption across space (Fajgelbaum and Gaubert, 2020; Donald et al., 2024). Setting $\theta \to \infty$ (removing idiosyncratic preferences) eliminates this concern.

if these workers are not perfect substitutes in production (Card, 2009). I extend the model to allow for households to differ on both skill z and education $s \in S := \{\text{College}, \text{NonCollege}\}$ and modify the technology in (11) as follows:

$$g_c = \left[\sum_{s \in S} \left[\iota_{cs} \sum_{z \in Z} z L_{cs}(z) \right]^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{25}$$

where σ is the elasticity of substitution between education levels, $\iota(c,s)$ is an education-augmenting productivity term, and $L_{cs}(z)$ is the population of (s,z) types in city c. Third, these agglomeration effects may be biased in their benefit to college workers, as suggested by the evidence in Diamond (2016) and Baum-Snow et al. (2018). I allow education-augmenting productivity ι_{cs} to respond to a city's educational composition as follows:

$$\iota_{cs} = \tilde{\iota}_{cs} \prod_{s' \in S} L_{cs'}^{\alpha(s',s)} \tag{26}$$

where L_{cs} is the population of s types in c. Because deregulation changes both the population and skill composition of cities, each of these extensions may be quantitatively important. I show that the conclusions of this paper are robust to each of these extensions.[29]

Incorporating capital gains and losses So far, I have abstracted from households' land ownership. Deregulation will have a large impact on equilibrium land values. This affects the lifetime consumption of many households because most of them are heavily invested in housing, especially low- and medium-wealth homeowners (Greaney, 2025). To this end, I allow households to differ by a designation of *renter* or *homeowner*. Homeowners are endowed with shares in a national land portfolio with value Π. Total household income now includes dividends on this portfolio

$$w_c z + s^{\text{Nat}}(z)\Pi \tag{27}$$

where $s^{\text{Nat}}(z)$ is the exogenous fraction of total land wealth owned by a z homeowner. Renters do not observe capital gains or losses and are equivalent to the definition of a household in the baseline model.

Preferences for residential density Another common justification for minimum lot size regulation are household preferences for low density neighborhoods. The model abstracts from these preferences, attributing preference-based sorting across neighborhoods entirely to income. This is almost surely incorrect. Gyourko and McCulloch (2024) document that suburban homeowners exhibit strong and heterogeneous preferences for low housing density. However, their evidence also reveals why directly identifying this preference separately from preferences over neighbors' income composition is difficult. When they augment their hedonic specification with demographic controls for neighbors' race, income, education, and family structure, these controls account for nearly two-thirds of the welfare loss associated with increased rental-unit density exposure. The Facts in Section 3.2 reflect the difficulty of

[29]However, some controversies remain. Despite allowing for endogenous labor productivity, I find that most of the observed variation in productivity across cities is driven by exogenous differences, as implied by virtually all estimates of the strength of agglomeration forces α. Recent work argues that the endogenous location choices of heterogeneous firms explain practically all of this leftover variation (Hong, 2025). The counterfactuals I perform should be interpreted as holding firm locations fixed. Output effects of deregulation are likely even lower with heterogeneous and mobile firms.

identification in this context: income and density are negatively and tightly correlated within cities.

To verify the robustness, I extend the model to allow for density to affect neighborhood amenities:

$$\log b_i(z) \;=\; \Omega(z) \log \text{Inc}_i - \alpha_L \log L_i + \log \nu_i(z)$$

where α_L is the elasticity of preferences to density L_i, which is set to 0.08. This is an upper bound of estimates of the elasticity of congestion to city size in Combes, Duranton, and Gobillon (2019). The conclusions of this paper are robust to this extension.

Durable housing Recent work shows that housing construction is inherently slow to respond to economic shocks, especially redevelopment (Baum-Snow and Han, 2024). My model abstracts from housing supply dynamics. Consequently, counterfactuals should be interpreted as comparing long-run outcomes after structures become sufficiently dilapidated to be replaced. However, transition dynamics are important for policy conclusions (Greaney, Parkhomenko, and Van Nieuwerburgh, 2025). I introduce a major extension to the model that can capture the effects of deregulation in the short run, where new housing construction occurs only on undeveloped land. The derivation is long and is relegated to Appendix B.8.

4 Calibration

In this section, I calibrate the model parameters to rationalize the observed regulatory stringency R_i, wages, per-unit housing prices, and the spatial distribution of household income as a model equilibrium. Table 6 in Appendix C summarizes all parameter choices. Appendix C contains additional details. I begin with the key sources of data that serve as externally calibrated parameters.

4.1 Additional Data

City wages. Central to the model is the sorting of high-skilled households into productive cities. To measure city wages in terms of efficiency units of labor, I follow Baum-Snow et al. (2018) and regress log hourly wages in a set of occupation, sex, race, ancestry, year, quadratic in age, and years of education, including MSA fixed effects using the 2015-2019 ACS individual sample. The MSA fixed effects in this regression define w_c. I normalize w_c so that it is, on average, one across cities, noting that the relevant interpretation of household efficiency units z is then the total labor income that can be earned in an average city.

Local skill distributions. Amenity values $b_i(z)$ depend crucially on the local mass of households by skill level $L_i(z)$. I construct the mass of households in two steps. First, I take block-group household income distributions from the 2016–2020 ACS, reported in 17 nominal income bins and aggregated to the 7 bins used throughout. Bins are defined by yearly household income: \$0–25k, \$25k–50k, \$50k–75k, \$75k–100k, \$100k–150k, \$150k–200k, \$200k+. I scale these bin shares by the 2020 Census housing unit counts to obtain population counts. Second, I construct the support of the type distribution Z from the 2015–2019 ACS individual sample. I deflate total household income by the corresponding city wage w_c to arrive at a measure of household efficiency units and set the support point z_k for bin k to the population-weighted mean of this measure across all sample households in that bin.

A key issue is that the ACS reports income distributions at the block group level in dollar bins, which do not correspond directly to skill-unit bins because higher-productivity cities command higher wages per unit of household skill. I adjust for this by applying a cumulative density function (CDF) interpolation to block group income shares. For each block group in a city with productivity w_c, I evaluate the piecewise-linear income CDF at bin boundaries scaled by w_c and take differences to obtain shares expressed in efficiency unit terms.[30]

Per-unit housing prices in regulated zones. I recover per-unit housing prices P_{iR} from the hedonic regression

$$\log[\text{Value}_{iht}] \;=\; \log P_{iR} \;+\; \text{Controls}_{iht} \;+\; \text{FE}_t^{\text{year}} + \; \text{FE}_t^{\text{month}} \tag{28}$$

where Value_{iht} is the arm's-length transaction price of property h in block group i at time t, estimated on 2016–2022 CoreLogic transactions linked to 2022 property assessments and restricted to regulated structures (1–4 housing units per lot). The block group fixed effects identify P_{iR}, conditional on structure type, floorspace, lot size, room and bathroom counts, HVAC system quality, and foundation type, along with year- and month-fixed effects.[31]

Converting the value of a minimal lot to user cost. Constructing the regulatory stringency measure R_i requires expressing the minimum lot value as an annual cost of housing consumption. I set the user cost rate to be 4.2% of home values, which is the national median user cost rate in 2017 measured by Bishop, Dowling, Kuminoff, and Murphy (2025). However, Bishop et al. (2025) documents substantial heterogeneity in user cost rates across US counties. Importantly, highly regulated cities tend to have low rent-to-price ratios and thus user cost rates (Gyourko et al., 2013). All results in this paper are qualitatively robust to the incorporation of the heterogeneous Bishop et al. (2025) estimates.

In what follows, I describe how I internally calibrate the remaining parameters, allowing the model to match the data above. These parameters can be separated into two groups: those pertaining to housing demand and those pertaining to housing supply. I describe each of these in sequence.

[30] Concretely, the nominal income cutpoints $\{0, \$25\text{k}, \dots, \$200\text{k}\}$ defining the ability bins are multiplied by w_c before interpolation. A city with above-average wages has a higher nominal income for a given level of efficiency units; therefore, some mass shifts from higher bins to lower bins relative to the unadjusted shares. In practice, this adjustment has very modest effects on the income distribution of the average block group.

Moreover, I address the presence of zero population counts because they have been shown to cause instability in counterfactual outcomes (Dingel and Tintelnot, 2026). This is done via a smoothing procedure applied to population counts within three broad groups (bins 1–2, 3–4, and 5–7). Zero-count cells in a group are imputed to half the average count of the other bins in that group. Reassuringly, all counterfactual results are insensitive to this procedure.

[31] I impute missing categorical variables to use as much of the transactions data as possible. Prices are winsorized at the 2.5% tails within each MSA. This yields indices P_{iR} for approximately 95% of block groups, where CoreLogic transactions are sufficiently populated. To achieve complete coverage, I construct lower-quality indices using residuals of a hedonic regression on block group aggregates using the 2016-2020 ACS. To ensure that these prices are of a similar scale, I adjust the ACS prices so that the log mean and variance are identical to the prices derived from the transactions and assessments. The ACS block group tabulations do not report housing characteristics that are specific to regulated structures. Hedonic indices derived from ACS data pool across observations of regulated and unregulated structures.

4.2 Housing Demand Parameters

Consumption values and amenities define the value of living in a neighborhood. Their construction requires three sets of parameters. The first is the price per unit of housing services. The second is the measure of regulatory stringency – the value of a minimal lot – to understand by how much regulation has distorted housing consumption. Third, preference parameters $\bar{A}$ and β and the migration elasticities ρ and θ, which inform the importance of housing consumption absent regulation and the ability of households to substitute between zones, cities, and neighborhoods. I directly inferred the prices per unit of housing services in regulated zones and the value of a minimal lot from the data. In what follows, I *choose* prices in unregulated zones and housing preference parameters to target 1) the observed share of households that choose regulated structures, 2) the observed aggregate spending on housing services by household income, and 3) the causal effect of regulation on income sorting estimated by Song (2025). I describe each procedure below.

Calibrating preferences and the within-city migration elasticity ρ. Three parameters govern housing demand and the intensity of within-city sorting: the Cobb-Douglas weight β, the Stone–Geary minimum housing expenditure $\bar{A}$, and the within-city migration elasticity ρ. These parameters are jointly determined in the sense that all three affect the model-implied fraction of households living in regulated structures and the overall stringency of regulation. Therefore, I calibrate all three simultaneously.[32] The calibration targets two sets of moments:

Housing expenditure shares. The first set of targets consists of the median housing expenditure share by skill group from the 2015–2019 ACS. To calculate this expenditure share, I condition the data on renters for whom housing costs are directly observed.[33] The calibrated values $\beta = 0.09$ and $\bar{A} = 3{,}850$ per year reproduce the observed expenditure gradient: modeled spending shares decline monotonically from 39% for the lowest skill group to 9% for the highest skill group. With only two parameters, the model matches data expenditure shares for 7 skill groups, with a maximum error of 3 percentage points.

Song (2025) border discontinuity. The second moment disciplines ρ using quasi-experimental evidence. Song (2025) estimates the causal effect of minimum lot size regulation on income sorting using a border discontinuity design: neighborhoods just inside a binding minimum lot size restriction are compared to observationally similar neighborhoods just outside, and the income gap between them is attributed solely to regulation. Song (2025) finds an elasticity of neighborhood income to minimum lot size regulation of 0.16 (Table 6, Column 5).

I replicate this experiment within the model. For each regulated neighborhood, I simulate the equilibrium effect of a 10% reduction in the minimum lot size on the neighborhood's average income. To do so, I split each neighborhood into two equal-sized partitions separated by a "border". I reduce regulation only on one side of this border and allow households to re-sort across it. For this experiment, I hold city wages, amenities, and the populations of all neighborhoods fixed.[34] The calibrated ρ equates the percentage difference in income at the left

[32] For computational tractability, the parameter space is an interpolation over a coarse grid and is evaluated on a large random sample of 1,500 block groups. Targeted moments from the sample are within less than one percent of the targeted moments of the entire population.

[33] To construct spending shares on housing by skill group, I follow the procedure of Finlay and Williams (2022). They define expenditure shares using income, net of taxes and transfers. This yields following spending shares by household income: 0.389 for households that make between 0-25k annually, 0.232 for 25-50k, 0.177 for 50-75k, 0.153 for 75k-100k, 0.138 for 100-150k, 0.122 for 150k-200k, and 0.089 for 200k+.

[34] These assumptions are required to interpret the border discontinuity estimate absent violations to the Stable Unit Value Treatment assumption.

and right of the border to target an elasticity of 0.16; specifically, a 1.6% change in income for a 10% reduction in the minimum lot size. This yields $\rho = 5.5$. This is smaller but comparable to the estimate of 8.5 in Baum-Snow and Han (2024).[35]

Per-unit prices in unregulated zones. Given $(\beta, \bar{A})$ and ρ, I calibrate unregulated per-unit prices P_{iU} so that the share of households choosing regulated structures in the model matches the data in every neighborhood. The identification argument is intuitive: a neighborhood with an expensive minimal lot but many households choosing regulated structures reveals that P_{iU} must be high—households must find unregulated structures costly. Conversely, an abundance of unregulated structures in a neighborhood implies a low P_{iU}.

Amenity values $b_i(z)$ and the across-city migration elasticity θ. Given the consumption values and ρ, the amenity values $b_i(z)$ are identified through the Rosen-Roback spatial equilibrium condition:

$$e^{V_i(z)} L_i(z)^{-\frac{1}{\rho}} L(\mathbb{C}(i), z)^{\frac{1}{\rho}-\frac{1}{\theta}} \bar{L}(z)^{\frac{1}{\theta}} = \boldsymbol{W}(z), \tag{29}$$

where $L(\mathbb{C}(i), z)$ is the city-level population of type z in the city containing the neighborhood i. This is a direct algebraic manipulation of the neighborhood choice Equation (7). Identification follows the standard Rosen-Roback logic: amenities must be high in neighborhoods where housing is expensive relative to wages, serving as a compensating differential. For the across-city migration elasticity, I use $\theta = 4$ from Hornbeck and Moretti (2018), who identify it through the city-level employment response to exogenous wage shocks across US labor markets.[36]

4.3 Housing supply parameters

Parameters that govern the productivity of the housing sector remain. Four objects characterize housing supply: the supply elasticity ϵ_i, housing productivity λ_i, and zone-specific land areas $T_{i\boldsymbol{o}}$ for $\boldsymbol{o} \in \{R, U\}$. I take supply elasticities ϵ_i directly from Baum-Snow and Han (2024), who causally estimate tract-level variation in these elasticities using a shift-share labor demand shock. Housing productivity λ_i is then identified by the developers' profit-maximizing level of construction on a minimal lot:

$$R_i \quad = \quad \lambda_i \, P_{iR}^{1+\epsilon_i} \, l_i$$

where l_i is the unit density restriction (Equation 3). Land areas for each zone $T_{i\boldsymbol{o}}$ are then chosen to clear housing markets, equating supply with the demand in each zone.[37]

[35] These estimates use very different sources of variation. Baum-Snow and Han (2024) regress population growth on 7-year changes in wages and housing prices. By contrast, the estimate in Song (2025) uses cross-sectional variation in regulation on a small spatial scale.

[36] I report summary statistics for calibrated demand parameters in Table 7 of Appendix C.3, additionally broken down for cities in and out of the top quartile of the productivity distribution (superstars and non-superstars). I make three observations. First, the cross-sectional dispersion of amenity values is larger than that reported in Hsieh and Moretti (2019) because this model captures within-city dispersion in amenity values. Between-city variance in amenity values is comparable to Hsieh and Moretti (2019), especially for households with average skills. Low-skill households derive relatively greater amenity value in productive cities than high-skill households. Second, Unregulated prices P_{iU} average roughly 50% above regulated prices in the same block group. This reflects the logic in Section 3: higher prices in unregulated zones are a compensating differential for regulation. In high-wage cities, the differences between regulated and unregulated prices are even greater. Third, consumption values are greater in high-wage cities, but only for households with skills that earn them more than \$50k annually in an average city.

[37] For block groups with no regulated structures, I set T_{iU} equal to the measured land area of the block group, which identifies λ_i from an unregulated housing market equilibrium.

Untargeted Moments and Model Fit The calibration exactly targets a rich set of empirical moments. The most important of these are local neighborhood income distributions and regulation R_i. The model's untargeted moment is the distribution of developed land mass implied by the sum of the calibrated land areas T_{iR} and T_{iU}. Because developed and total land area differ — the latter includes undevelopable terrain, water, and vacant land — one would not expect exact correspondence between the model's implied land use and actual geographic land masses. Nevertheless, a regression of model-implied developed land mass on actual block group land area yields an R^2 of 36%. This is a strong correlation given these conceptual differences, providing reassurance that the model captures genuine cross-sectional variation in land use intensity across the United States.

5 Estimating $\Omega(z)$

How much do households value neighborhood affluence? Equation (9) specifies the causal effect of neighborhood affluence on amenities, and the strength of this relationship is governed by the set of elasticities $\Omega(z)$. These are crucial for gauging the welfare effects of regulations that prevent free-riding in rich neighborhoods, as demonstrated in Proposition 2.

The identification of $\Omega(z)$ is challenging because of at least two endogeneity issues. The first is reverse causality bias. Large unobserved amenities imply high housing prices, and thus a high price for the minimal lot, putting a disproportionate penalty on local consumption for low-income households and driving sorting patterns. Moreover, non-homothetic preferences cause income sorting on housing prices irrespective of regulation (Lee and Lin, 2017; Couture et al., 2023). The second arises because high-amenity locations may be disproportionately valued by the rich, irrespective of outcomes in the local housing market. Alternatively, the opposite might be true, as suggested by the negative relationship between income and density within cities (Section 3.2).

To address endogeneity concerns, I propose a donut strategy in which the characteristics of other neighborhoods are used as instruments for income in the focal neighborhood. I use the degree of sloped terrain as this characteristic, as it is a strong predictor of neighborhood income (Lee and Lin, 2017; Saiz, 2010). The reason for this is not well understood, but is likely driven by a combination of supply- and demand-side factors. Slopes make residential development costly, but also may be directly demanded by households because they are associated with desirable homefront views.[38] The donut design is meant to address the possibility that localized slopes are correlated with local demand factors.

Let $S(d)$ be the average slope over a set d of block groups. The estimating equation is

$$\log b_i(z) \quad = \quad \Omega(z) \log \text{Inc}_i + \beta_1(z) S(d_{1i}) + \log \tilde{\nu}_i(z) \tag{30}$$

where Inc_i is average income and d_{1i} is the set of block groups whose centroids are within distance d_1 of the boundary of i.[39] I use average slopes within a ring of length d_1 to d_2 with $d_1 < d_2$ as an instrument for average income in (30). There are two identification assumptions. First, households do not demand sloped terrain outside the buffer of length d_1. Second, slopes may be correlated with excluded demand factors in ν (such as lakefront views) insofar

[38] The idea that housing supply constraints are not orthogonal to neighborhood amenities has been used to argue that they are invalid instruments to estimate neighborhood demand parameters (Davidoff, 2016).

[39] Here, I define $\log \tilde{\nu}_i(z) = \log \nu_i(z) - \beta_1(z) S\left[d_{1i}\right]$; these are the fundamental amenities of a neighborhood orthogonal to slopes within the buffer region d_1.

as slopes in the set of instrument block groups are uncorrelated with these demand factors *conditional on slopes in* d_{1i}. In other words,

$$S(d_{2i}) \perp \log \tilde{\nu}_i(z) \mid S(d_{1i}) \quad \text{for all } z \in Z. \tag{31}$$

Where d_{2i} is the set of block groups whose centroids lie within distance d_1 and d_2 of the boundary of i. In Appendix D.1, I present an econometric model that incorporates latent and spatially correlated demand factors that may themselves be correlated with slopes. I use this to explain the meaning of the identification assumption in mathematical terms.

Instrument relevance follows from the idea that neighborhoods outside the control buffer d_1 are substitutes (Bayer et al., 2007). Therefore, the characteristics of substitute locations serve as instruments for the characteristics of the focal neighborhood. Steeply sloped neighborhoods attract disproportionately high-income residents because sloped terrain may be directly valued (homefront views) and because it constrains residential supply. Sloped substitute neighborhoods attract high-income households *away* from the focal neighborhood, depressing the average income of the focal neighborhood. This substitution effect implies a negative first stage: slopes outside the control buffer should be negatively associated with income in the focal neighborhood, conditional on slopes in the control buffer that account for the focal demand and supply effects of steep terrain.

Results. Table 2 presents the main results using a control buffer of $d_1 = 10$km and an instrument buffer of $d_2 = 16$km. Panel A reports IV estimates of $\Omega(z)$ for three aggregated skill groups: low ($0–50k yearly income), medium ($50k–100k), and high ($100k+), and Panel B reports the OLS counterparts.[40] All specifications absorb MSA fixed effects and include log block-group land area, mean and standard deviation of elevation, average commuting time, building age, public transit usage share, and CBD distance ranking as controls. Standard errors are clustered using a Bartlett kernel reaching 35 kilometres. The Kleibergen-Paap F-statistics of approximately 19 confirm instrument relevance. IV estimates of $\Omega(z)$ are positive for all skill groups and are precisely estimated for medium- and high-income households. The estimates are $\hat{\Omega}(\text{low}) = 0.2$, $\hat{\Omega}(\text{med}) = 0.23$, and $\hat{\Omega}(\text{high}) = 0.44$.

How should the magnitude of these estimates be interpreted? Equation (10) shows that the consumption value a household of skill z is willing to forgo to live in a neighborhood with one percent greater income is exactly $\Omega(z)$. However, these consumption values are not quoted in money-metric units. Expressing $\Omega(z)$ relative to the average marginal utility of income reveals that a medium-skill household is willing to pay 0.2% of wages to live in a neighborhood with 1% greater income. For low- and high-skill households, this is 0.16% and 0.31% of wages, respectively. These results provide two insights. First, the amenity value of income is high for the average household. Second, high-income households have the highest willingness to pay for income, suggesting that the amenities created by neighborhood affluence are disproportionately valued by these households. This means that high-income households sort into expensive, high-amenity neighborhoods – the sorting condition that underlies Tiebout models (Epple and Platt, 1998). This also means that income changes brought on by deregulation will be made larger by self-reinforcing sorting (Diamond, 2016; Su, 2022).

[40] I aggregate the 7 income bins into three groups only for estimation, to address zero-population cells by income type in some block groups. Amenity value for each aggregated group is a population-weighted average of constituent types.

Table 2: Baseline IV Estimates by skill group.

Dependent variable: $\log b_i(z)$						
	Panel A: IV (Donut: 10-16km)			*Panel B: OLS*		
Skill group:	(1) Low	(2) Med	(3) High	(4) Low	(5) Med	(6) High
ln Income	0.2043*	0.2318***	0.4355***	0.0076	0.0835***	0.3064***
	(0.1146)	(0.0718)	(0.0679)	(0.0074)	(0.0055)	(0.0058)
Slope Control	0.0130***	0.0064***	0.0074***	0.0092***	0.0035	0.0051*
	(0.0029)	(0.0020)	(0.0020)	(0.0032)	(0.0031)	(0.0029)
Local Slope Control	-0.0030	-0.0054	-0.0075**	0.0063***	0.0016*	-0.0016***
	(0.0055)	(0.0035)	(0.0031)	(0.0014)	(0.0010)	(0.0005)
Observations	185,490	185,667	181,561	185,490	185,667	181,561
KP F-stat	19.9	19.6	19.9	—	—	—
MSA Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Base Controls	Yes	Yes	Yes	Yes	Yes	Yes
Amen/Topo Controls	No	No	No	No	No	No
Density Control	No	No	No	No	No	No

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Panel A: IV standard errors from Bartlett spatial kernel (35 km bandwidth). Panel B: OLS standard errors clustered at the MSA level (standard errors are smaller when clustering using the 35km Bartlett kernel). Estimates disaggregated for 3 skill groups: low ($0–50k yearly income), medium ($50k–100k), and high ($100k+), measured in yearly household income if living in an average productivity city. Columns are ordered by skill groups: low (0-50k annual income), medium (50-100k) and high(100k+). All specifications include MSA fixed effects and include own log block-group land area, as well as mean and standard deviation of elevation. "Slope Control" is the average slope within the block group plus a buffer with a length equal to d_1. "Local Slope Control" is the average slope within the block group. ln Income is instrumented with the average slopes of block groups that have centroids within buffers d_1 and d_2. The "Base Controls" include travel time, building age, public transport and bus shares in commuting, and CBD distance."Amen/Topo" controls include various amenities (density of coffee shops, parks, and restaurants) and various topographic features (cover of different types of forest, such as deciduous or evergreen, wetlands, and perennial snow cover)."Density Control" is the within-MSA density ranking of the block group.

Panel B of Table 2 shows that the OLS estimates of $\Omega(z)$ are substantially smaller than the IV estimates for all skill groups. The IV corrects for downward bias because dense and central neighborhoods tend to have high fundamental amenity values but low average incomes.[41] This is suggested by the empirical facts presented in Section 3.2. For low-income households, this bias is strong enough to render the OLS coefficient indistinguishable from zero, even though the IV estimate is positive and economically significant.

The estimates of the first stage regression have the theoretically correct sign. Table 8 in Appendix D reports first stage estimates for a pooled regression across skill groups. The columns correspond to progressively richer control sets. In the baseline specification (Column 2), the 10–16 km slope instrument has a negative coefficient. The slopes in the focal block group are positively associated with income across all specifications, consistent with sloped terrain serving as local supply constraints that make housing expensive. The negative and significant coefficient on the donut ring confirms the substitution mechanism: sloped neighborhoods 10–16 km away attract high-income households away from the focal neighborhood, reducing its income.[42]

Robustness. Appendix D.3 reports a battery of robustness checks on the estimates. Table 9 shows that the pooled estimates of $\Omega(z)$ are stable across the full range of different controls

[41] Note that the measure of amenities includes commuting time and within-city access to high paying jobs. The model abstracts from commuting and job access (Ahlfeldt et al., 2015).

[42] Block groups with no neighboring block group centroids in the 10–16 km ring are dropped from the estimation sample; these correspond to block groups with unusually large land mass and low population density.

for local amenities. Most importantly, estimates are stable when including residential density rankings as controls, suggesting density is not confounding preferences for affluent neighbors (Gyourko and McCulloch, 2024). Table 10 reports estimates across alternative donut radii from 2–4 km to 16–25 km. As expected, shorter-range donuts (closer to the focal neighborhood) produce weaker instruments because nearby slopes are more likely to confer demand spillovers that offset the negative substitution effect. A larger donut of 16–25 km yields even larger point estimates, and the counterfactual results are quantitatively robust to them. Instrument buffers at greater distances are too weak to be reliable.

I further probe the validity of these estimates in two ways. First, I construct a placebo test that exploits the timing of neighborhood income changes. It relies on the assumption of *no retrocausality* – that future income levels cannot cause changes to the fundamental amenity values of neighborhoods in the past. If a correlation exists, this likely reflects a violation of the exclusion restriction: incomes are correlated with predetermined fundamental amenity values of the past, and likely also with those of the present. In Appendix D.4, I describe this placebo test in more detail and show that this strategy survives the test for each skill group.

Second, I test an alternative instrument that uses local topography and metro-level income variation to predict local income distributions. This approach is inspired by that of Davis, Gregory, and Hartley (2023), who applied it to measure racial homophily in location preferences. This approach yields smaller but qualitatively similar estimates, which are discussed in Appendix D.4. All the conclusions of this paper are robust to these estimates.

In the following section, I use these estimates to assess the welfare impacts of complete deregulation and other spatially targeted policy schemes. A large willingness to pay for income, especially for high-skill households, will matter for the conclusions drawn.

6 Evaluation of Counterfactual Policies

How would deregulation affect the aggregate economy, and how would households of varying affluence realize these gains? In this section, I study the impact of deregulation on both welfare outcomes and income sorting across neighborhoods and cities. I pay special attention to three channels:

1. Gains from the expansion of productive cities (Hsieh and Moretti, 2019; Duranton and Puga, 2023)
2. Losses from the neighborhood choice externality (Hamilton, 1976; Calabrese et al., 2007)
3. Increased housing affordability through smaller lots (Song, 2025; Kulka, 2019)

The main message of this paper is that the gains from the expansion of productive cities and losses from the neighborhood choice externality are negligible relative to the increase in housing affordability. I perform three counterfactual exercises. The first is a nationwide deregulation exercise. The second halves the stringency of regulation in San Francisco, mirroring the recent elimination of single-family zoning in California. The third studies a spatial permutation of the observed distribution of regulation. This scheme targets regulation to neighborhoods that provide high amenity values to rich households more accurately.

6.1 Complete deregulation

I start by studying the impact of complete deregulation, amounting to an equilibrium where $l_i = 0$ in all neighborhoods. To define household welfare, I take the ordinal measure of utility $\log \boldsymbol{W}(z)$ from Equation (8). I express this measure using the equivalent variation reported in annual per-household dollar terms. I provide more details on the construction of this welfare measure in Appendix E.2. Figure 3 reports the equivalent variation in the "total" column, broken out by household skill. Social welfare is calculated by taking a population-weighted average across skill levels.

Figure 3:
Equivalent variation for deregulation by household skill.

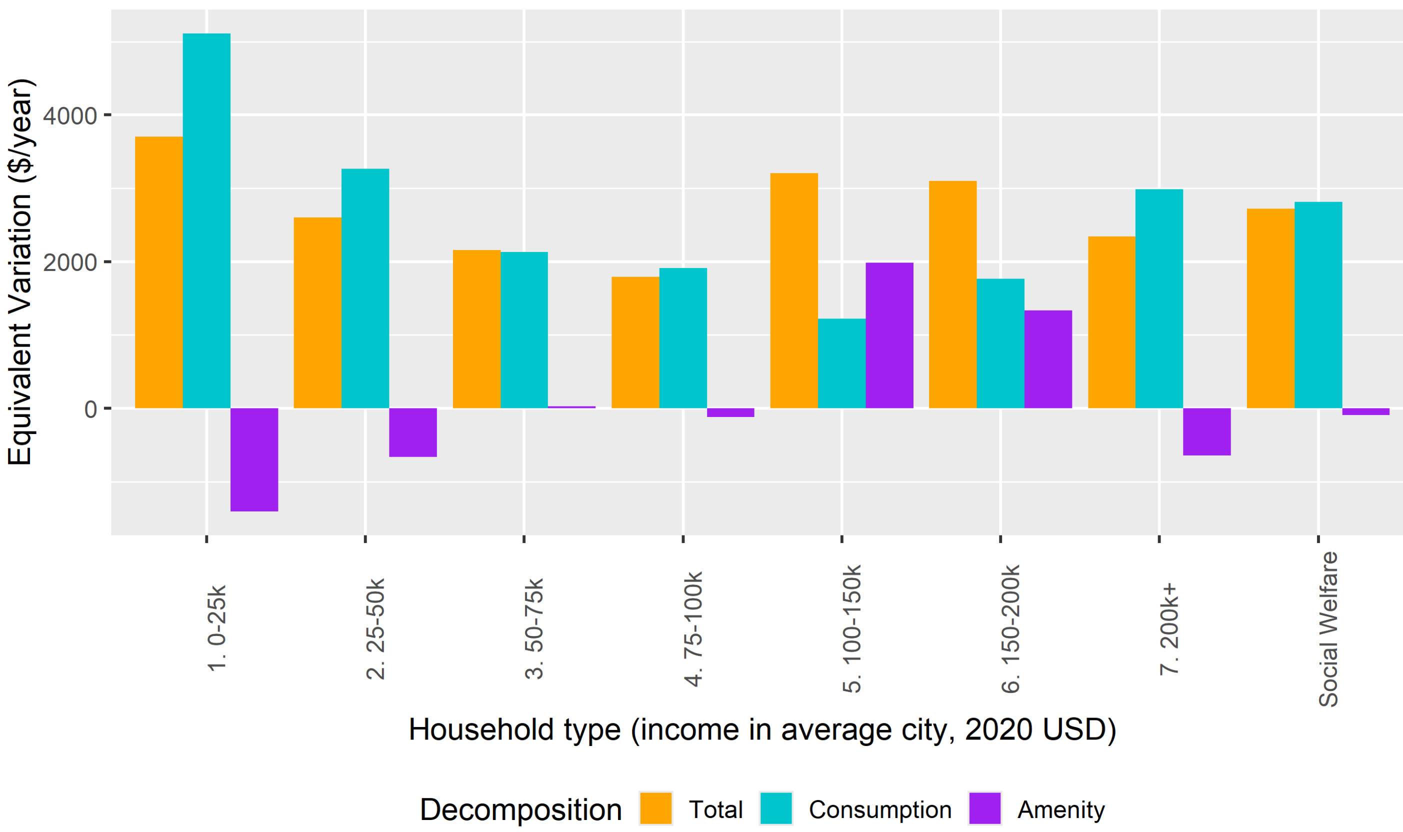

Welfare is measured as the equivalent variation and expressed in dollar terms, reported in the "Total" column. Higher values indicate higher welfare gains. Social welfare is the population-weighted average of welfare by type. "Amenity" and "Consumption" components are constructed using Shapley values, with a procedure outlined in Appendix E.4. The components add to the total column.

Viewing the "Total" column, Figure 3 shows that the social welfare of households increases by approximately $\$2{,}600$ per year. Households earning less than $\$25{,}000$ annually gain approximately $\$3{,}600$ per year, the largest absolute gain of any income group, reflecting the fact that regulation distorts housing consumption the most for this group. However, welfare gains are also substantial across the income distribution, with middle-income households ($\$50{,}000$–$\$75{,}000$) gaining roughly $\$2{,}000$ per year and the highest income households ($\$200{,}000$+) gaining approximately $\$2{,}100$ per year. Deregulation has widespread benefits for all households that own no land wealth.

Neighborhood choice externalities How has the neighborhood choice externality affected the gains from deregulation observed in Figure 3? Answering this question requires parsing the welfare consequences of changing amenities $b_i(z)$ and neighborhood consumption values $C_i(z)$. To this end, I decompose welfare changes using a Shapely value decomposition, with details of its construction in Appendix E.4. I report the results of this decomposition for each type in Figure 3 under the "Consumption" and "Amenity" columns, which sum to the "Total" column.

The results reveal that households along the skill distribution benefit from deregulation for markedly different reasons. For the lowest income households (below $\$50{,}000$), the amenity contribution is negative — deregulation reduces the average neighborhood income in the regulated neighborhoods they now find affordable — while consumption gains are very large, exceeding $\$4{,}000$ per year. For middle-income households ($\$50{,}000$–$\$100{,}000$), the amenity contribution is approximately zero, with nearly all gains flowing from reduced housing costs. Strikingly, upper-middle-income households ($\$100{,}000$–$\$200{,}000$) experience *positive* amenity contributions alongside consumption gains: freed from regulation, they sort into high-amenity neighborhoods that were otherwise unaffordable, as predicted by the simple model in Section 3.4. The highest-income households lose as deregulation renders their neighborhoods nonexclusive and less amenitized.

Most importantly, the social welfare loss ascribed to the amenity component is negative, but essentially negligible.[43] The unimportance of amenity value changes stands in stark contrast to other papers that study minimum lot size reforms quantitatively (Calabrese et al., 2007). The observed spatial distribution of minimum lot size regulation appears unable to facilitate efficient Tiebout sorting despite its potential to do so (Proposition 3).

Alternatively, the importance of endogenous amenities can be assessed by comparing the welfare impacts of deregulation to those that hold amenity value fixed. This is equivalent to setting $\Omega(z) = 0$ for each z and re-calibrating the model. This alternative delivers essentially identical welfare results: the neighborhood choice externality matters relatively little in the aggregate, while increased housing affordability drives welfare gains.

While amenity changes appear unimportant for policy evaluation, the effect of deregulation on land value is crucial for policy evaluation. I find that the aggregate value of land decreases substantially after deregulation, at 7%. This has a simple explanation through the lens of the model. Regulation makes housing consumption artificially large, which means that housing prices and land values must be adjusted upward to clear housing markets. An equivalent view is that deregulation increases land availability, which is accompanied by a decrease in land values nationally.

Incorporating capital losses on land values Housing is one of the most widely held assets, so capital losses on land induced by deregulation affect the average household. The baseline model makes no attempt to weigh welfare benefits against capital losses on land. To this end, I allow for households to differ along the dimension of *renter* and *homeowner*, with the income of homeowners now including the rental income earned from a national land portfolio. The total homeowner income is given by Equation (27). Household portfolio weights $s(z)$ are calculated such that the total share of the portfolio owned by skill-z households is proportional to the total homeowner spending on housing services for that skill level. Renters retain the same welfare measures as in the baseline model. For computational reasons, the model is not re-calibrated with this new dimension of household heterogeneity. Instead, I assume that prices and land values change in accordance with the baseline model without homeownership and use these inputs in the welfare calculation. I provide additional details of this procedure in Appendix E.3.

[43] In fact, certain parameter choices on ρ and $\Omega(z)$ used in robustness checks yield small and *positive* amenity contributions, so I do not emphasize the importance of the negative value.

Figure 4 reports the welfare effects by skill level, broken down by homeowners and a pooled average over renters and homeowners.[44] The pooled social welfare gain from deregulation is approximately $\$1{,}700$ per year. This is driven by large gains for lower-income households, of which renters are the majority. By contrast, homeowners earning over $\$200{,}000$ uniquely experience welfare *losses* of roughly $\$1{,}200$ per year. This reflects that these households hold a disproportionate share of land wealth, and the capital losses from deregulation outweigh the affordability gains for this group. This pattern underscores that deregulation is redistributive: it generates large net gains in the aggregate but shifts resources from landowners to renters, a point highlighted extensively in the literature (Parkhomenko, 2023; Greaney et al., 2025).[45]

Figure 4:
Equivalent variation for deregulation by household type,
pooled over renters and homeowners and including capital losses.

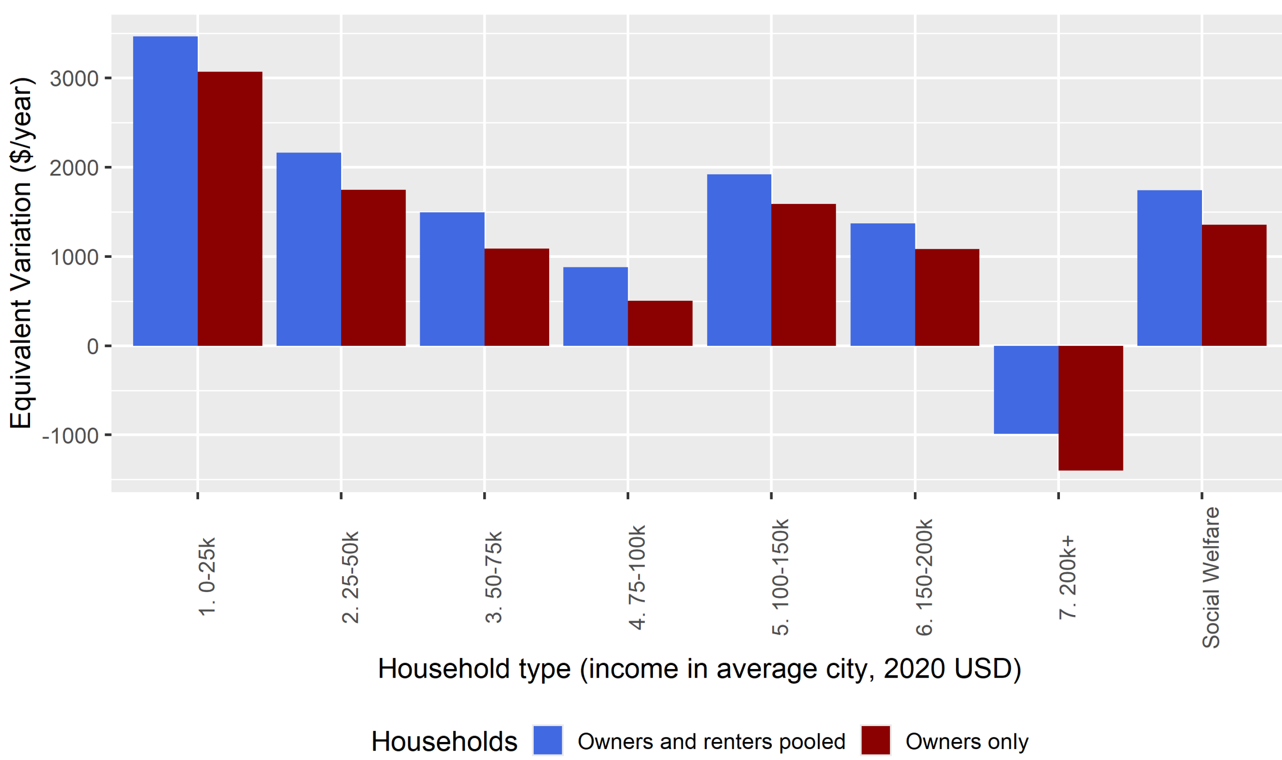

Welfare is measured as the equivalent variation and expressed in dollar terms. Higher values indicate higher welfare gains. Social welfare is the population-weighted average of welfare by type. The results incorporate capital losses for homeowners by income type, using the procedure outlined in Appendix E.3. The results include equilibrium adjustments to neighborhood amenity values, as in the baseline counterfactual. Renters and homeowners are pooled by type using a population-weighted average of welfare changes.

Gentrification Deregulation appears to only slightly reduce neighborhood amenity values, but are there distributional consequences across neighborhoods? Evidence from Section 3.2 suggests that deregulation increases income in high-density neighborhoods relative to low-density neighborhoods, especially within superstar (high-wage) cities. The model confirms this prediction. In Figure 5, I compare the income density gradient for both the observed and counterfactual equilibrium separately for the superstar and non-superstar city samples

[44] Averages are taken with respect to the share of owner-occupiers by skill level from the 2015-2019 ACS. These are: 0.375 for 0-25k, 0.511 for 25-50k, 0.618 for 50-75k, 0.708 for 75-100k, 0.793 for 100-150k, 0.857 for 150-200k and 0.89 for 200k+. More details are in Appendix E.3.

[45] The welfare consequences for each skill group are robust to a multitude of robustness checks pertaining to the measurement of regulation R_i. These include: using properties built after 1970 for the calculation of the Unit Density Restriction, assuming all lots below the minimum are unregulated structures when calibrating the model, and using user cost estimates from (Bishop et al., 2025). Deviations in welfare changes for each of these robustness checks average less than 2 percentage points for each skill group.

(Section 3.2).[46] Each panel corresponds to a city sample. Within each panel, I plot flexible regressions of income (demeaned at the MSA level) on the baseline density ranking. The yellow and purple lines correspond to the counterfactual and initial equilibria, respectively.

Figure 5 shows the stark prediction made by the model.[47] The income-density gradient in superstar cities is almost entirely eliminated under complete deregulation: the initial gradient of approximately -0.61 log points (from the empirical estimation) collapses to near-zero in the counterfactual, erasing virtually all differences in sorting patterns between superstar and non-superstar cities. This suggests that variation in regulation across space can account for nearly *all* differences in income sorting on density in expensive cities. Low-density neighborhoods, which currently command the highest incomes, experience significant income declines, while formerly dense and lower-income neighborhoods gentrify. Non-superstar cities see more modest changes, reflecting the smaller within-city variation in regulation that is characteristic of these markets. The migration patterns within productive cities closely resemble recent US gentrification, wherein affluent households move toward central, high-density neighborhoods.

Figure 5:
Gentrification post deregulation, especially in superstar cities.

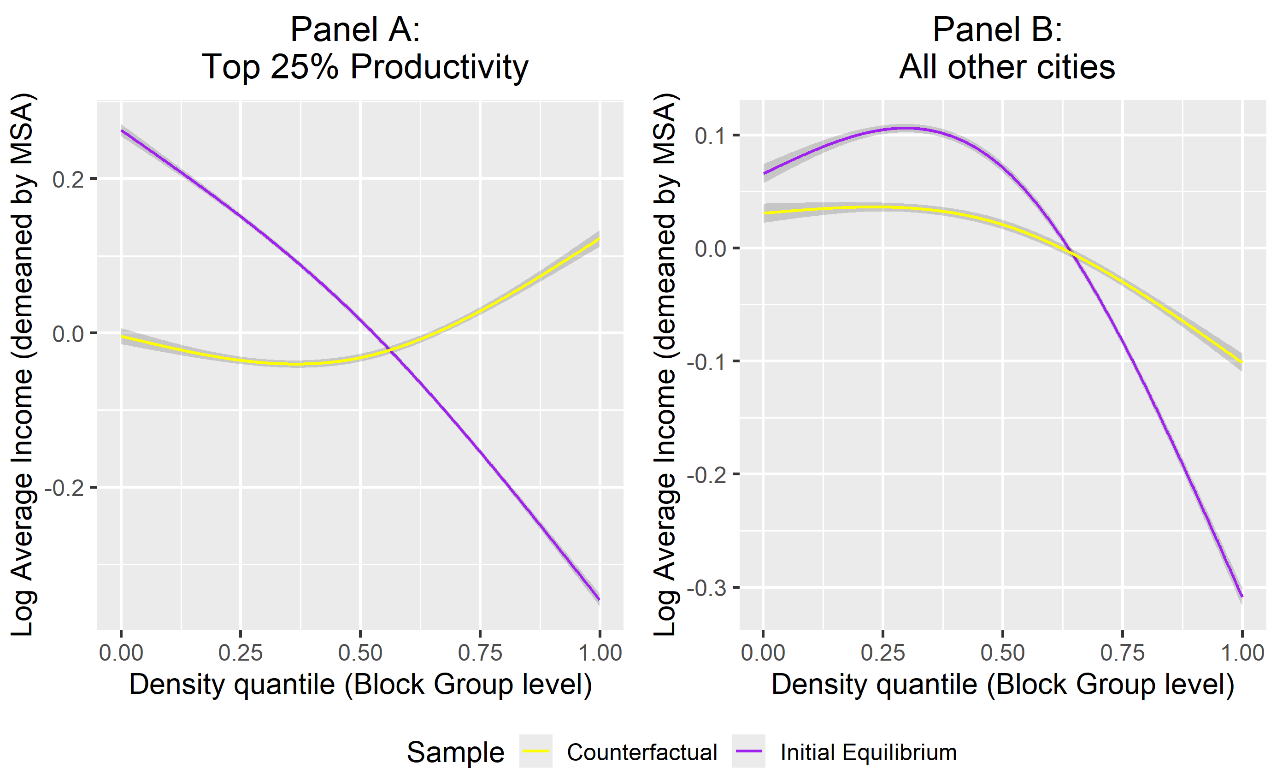

For each panel, demeaned log average income at the neighborhood level is regressed against the observed density ranking of neighborhoods at baseline. The purple regression uses data from the baseline equilibrium that matches the data, as shown in Figure . The yellow regression uses data generated from the counterfactual with no minimum lot size requirement. Each panel corresponds to "superstar" sample cities (Top 25% wages) and "non-superstar" sample cities (Bottom 75% wages), as in Figure 2.

Aggregate labour productivity In Section 3, I showed that skill heterogeneity and sorting matters for the productivity benefits of deregulation. I find that complete deregulation increases aggregate labor productivity by only 0.11%, equivalent to roughly a twentieth of a typical year of growth in the US. In contrast, Duranton and Puga (2023) find that eliminating housing supply restrictions in superstar cities would increase output per person by around

[46]Note that the model closely matches the average income of all neighborhoods in the ACS, so these differences under the counterfactual are directly comparable to observed neighborhood income distributions. Slight differences arise due to aggregation bias in calibration.

[47]In Figure 12 of Appendix E.5, I additionally reconstruct Panel A of Figure 2 treating counterfactual incomes as data.

8.2%. Hsieh and Moretti (2019) estimate the number to be 3.2% in an exercise where they increase the housing supply elasticity in San Jose, San Francisco and New York.

I show that the skill sorting responses drive this low value in two ways. First, I plot the growth rate in the number of households against the growth rate in average household skill across all cities in Figure 6, with each city represented by a circle that is colored by productivity and has a radius proportional to the city population. In general, we observe large net household inflows in cities with higher wages and housing prices. [48] These cities also observe large decreases in the average skill level. Second, I calculate what aggregate productivity would be if 1) cities had household growth as predicted by deregulation and 2) the average skill in each city remained at observed levels. This exercise yields an aggregate productivity of 1.57%, which is more than ten times higher than what is predicted when income sorting occurs. The endogenous amenity response also contributes to low productivity growth, which is driven by the deteriorating amenity values offered in productive cities. Counterfactual productivity growth when amenities are assumed exogenous would be almost four times as high, at 0.39%.

Figure 6: City income sorting post deregulation

The y axis is defined as the change in the average income that a household could earn in an average city from baseline to counterfactual. The correlation between the growth rate in the number of households and productivity is approximately 50%, and this correlation is 70% for unadjusted housing prices. Cities on the top left (high income, negative population growth) tend to have low productivity, and cities with higher productivity tend to be in the bottom right.

These results are robust to multiple assumptions about externalities and production technologies. First, I allow city productivity ι_c to respond endogenously to the city population, as in Equation (24), with an elasticity $\alpha = 0.05$ (Combes and Gobillon, 2015). Aggregate productivity growth under the counterfactual increases from 0.11% to 0.38%. Second, I introduce education as an additional dimension by which households differ, as in Equation (25); households with varying education levels are substitutes with elasticity $\sigma = 1.3$ (Card, 2009). Under the same counterfactual, aggregate productivity growth is 0.21%. Finally, I allow households of differing education levels to arbitrarily affect labour productivity by education, as in Equa-

[48]Correlations between population flows with wages and housing prices are 60% and 83%, respectively. For changes in average city skill, these correlations are instead −50% and −67%.

tion (26) (Diamond, 2016).[49] This counterfactual yields an aggregate productivity growth of 0.47%. In Table 15 of Appendix E.5, I organize counterfactual productivity changes by each of these model extensions. Moreover, performing the same counterfactual in a model without any skill heterogeneity implies that aggregate productivity growth would be 8 times higher, at 0.8%.[50]

Landlords in regulated neighborhoods Given that all renters are better off, how can the model rationalize why minimum lot sizes are frequently imposed? I find that land values, on average, decrease by a larger amount in initially stringent neighborhoods. This is commensurate with the idea that homeowners impose land use restrictions to maximize their land values (Parkhomenko, 2023; Hilber and Robert-Nicoud, 2013; Fischel, 2001). Theoretically, two competing effects determine the relationship between land value growth and initial regulatory stringency. On the one hand, regulation distorts housing consumption choices, thus decreasing neighborhood demand. However, neighborhood demand increases when stringent regulations increase neighborhood affluence. I find that in an equilibrium where amenities are endogenous, changes in land value are negatively related to initial levels of regulatory stringency $\tilde{R}_{ic}$, but *positively* related when amenities are assumed exogenous. Comparing this result to the literature, Parkhomenko (2023) provides analytical conditions on parameters to predict when land use regulation will increase or decrease land values. An important message is that larger congestion externalities contribute to the positive relationship between land use regulation and rents. In this paper, the amenity value of income has an identical effect: it mediates this relationship.

6.2 Unilateral deregulation in San Francisco

Next, I simulate the effects of halving the value of a minimal lot R_i in the San Francisco Metropolitan Area.[51] This policy change is designed to mirror the recent elimination of single-family zoning in California. I study welfare effects, as well as how income, land values, and productivity change both in San Francisco and nationally.

Deregulation causes San Francisco to densify through an influx of primarily lower-income households. The number of households in the city grows by 12%, but the average household skill falls by 15%, eroding all of the aggregate productivity gains that would have been achieved as a result of city expansion. Assuming away all income-composition effects of city migration implies that aggregate labor productivity would have increased by 0.15%. This is sizable, especially for a policy change affecting an area that represents a small share of nationwide economic activity.

Despite deteriorating neighborhood amenities, households of all skill levels benefit from

[49] For this counterfactual, I use the calibrated estimates of $\alpha_{s's}$ from Fajgelbaum and Gaubert (2020). Let C index college households and N index non-college households. Using estimates from Diamond (2016), they calibrate $\alpha(N,N) = 0.003$, $\alpha(N,C) = 0.02$, $\alpha(C,C) = 0.053$ and $\alpha(C,N) = 0.044$. These estimates suggest that productivity spillovers mostly come from the presence of college educated workers.

[50] Aggregate productivity growth after deregulation in the model without skill heterogeneity is still considerably smaller than those estimated by Hsieh and Moretti (2019) and Duranton and Puga (2023). Differences stem from two sources. First, they consider deregulation only in high productivity cities. Second, their measures of regulation are broader and summarized in the Wharton Regulation Index (Gyourko et al., 2021). The important fact I emphasize is that, in this model, aggregate productivity gains are attenuated tenfold by the presence of skill heterogeneity and endogenous neighborhood amenities.

[51] The value of a minimal lot R_i is an endogenous object (Equation 3). It is not directly manipulated by policymakers, whereas the unit density restriction l_i is. In principle, any value of R_i can be achieved by policymakers with a correctly chosen density restriction, l_i. In practice, halving R_i yields almost identical effects under this counterfactual as halving l_i.

the unilateral deregulation of San Francisco (on average, 0.1% of income). Land values in San Francisco also increase slightly in the aggregate, at 1.5%. Given these city-wide benefits, what explains the widespread political challenge to deregulation, particularly by local residents who are directly affected? The discussion of city-wide outcomes ignores significant heterogeneity in land value changes across neighborhoods. Figure 13 in Appendix E.5 maps this heterogeneity. Panel A plots the initial income distribution before deregulation across the metro, and Panel B shows the observed measure of regulatory stringency from Equation (13). High-density neighborhoods in San Francisco, Berkeley, Oakland, and Hayward have the lowest income and relatively lower levels of regulation, while low-density neighborhoods farther inland from these municipalities have the opposite. Panels C and D plot the changes in neighborhood income and land values after deregulation. Both income and land values tend to increase in these high-density municipalities as high-income households leave previously stringent neighborhoods to relocate there. That is, these neighborhoods gentrify, as suggested by the facts in Section 3.2.

I show that the endogenous amenity response explains the massive heterogeneity in outcomes across neighborhoods. Figure 14 in Appendix E.5 plots the distribution of land value changes in when amenities are exogenous and endogenous, respectively. When amenity values are endogenous, there is a large tail of neighborhoods in which land values drop substantially. These correspond to initially stringent, affluent, and low-density neighborhoods. By contrast, when amenity values are exogenous, land values increase by 17% city-wide, and no such tail exists. Instead, land values increase in initially stringent neighborhoods, reflecting the option value of building at a higher density. This heterogeneity provides a natural explanation for why state-level deregulation faces entrenched local opposition: landowners in initially stringent neighborhoods bear concentrated losses, giving them strong incentives to block the development of higher-density housing. The model suggests that the endogenous amenity value of neighborhoods mediates these incentives.

6.3 A better spatial allocation of regulation

These counterfactual outcomes suggest a dismal view of the minimum lot size: complete deregulation is considerably more efficient than the status quo, and cities are disincentivized from unilaterally deregulating to protect the land and amenity values of certain neighborhoods. Does this imply that regulation cannot be used to improve social welfare? In this section, I show that there is a significant opportunity for an improved spatial distribution of regulation.

The idea behind a better spatial allocation echoes the logic of Proposition 2. A social planner wants to impose stringent regulation in neighborhoods that would be rich in the absence of regulation, and reduce or eliminate regulation in poor neighborhoods. In Section 6.1, I provided suggestive evidence that minimum lot sizes were targeting neighborhoods valued by rich households, but this correlation is far from perfect, especially within productive cities (Figure 5). I consider a policy change that perfects this correlation. To this end, I define a neighborhood's *high-skill amenity score* ν_{Si} as follows:

$$\nu_{Si} = \sum_{z \in Z} \frac{\nu_i(z)}{\sum_{z' \in Z} \nu_i(z')} z \tag{32}$$

ν_{Si} is the fundamental amenity-weighted average skill level of i, which measures how much high-skill households value a neighborhood's amenities relative to low-skill households. I consider a counterfactual policy that orders the observed values of a minimal lot R_i and calibrated shares of regulated land $\frac{T_{iR}}{T_{iU}+T_{iR}}$ by the ranking of ν_{Si}.

Which cities and neighborhoods score highly? First, I find that the correlation between the amenity score and city productivity is negative, albeit somewhat weak. This means that low-skill households value amenities in high-productivity cities more than high-skill households do, on average. Second, there is a strong negative correlation between the score and neighborhood density within cities. This means that the within-city distribution of regulation is higher in neighborhoods that high-income households value. These empirical relationships explain how the policy change affects the neighborhood income distribution.

This policy change delivers large welfare benefits averaging $\$2{,}200$ per year per household, and, crucially, generates an increase in national land values of 4.1%. This stands in stark contrast to complete deregulation, where households incur capital losses. In Figure 7, I break down these welfare effects by skill level and find that both renters and owners of *all* skill levels are better off. Strikingly, homeowners benefit more than the pooled average at nearly every income level: the capital appreciation from land compounds the direct welfare gains, particularly for high-income households that hold large land portfolios. The welfare gains are also considerably larger than those from complete deregulation: the better spatial allocation yields $\$2{,}200$ per household compared to approximately $\$1{,}700$ under complete deregulation (after accounting for capital losses in the latter case). This means that not only are the costs of nationwide deregulation avoidable, but also that a well-targeted spatial reallocation of regulation can generate positive returns to existing landowners, removing the very tension between renters and homeowners that makes deregulation politically difficult.

Figure 7:
Effects of permuted regulation scheme
pooled over renters and homeowners and including capital losses

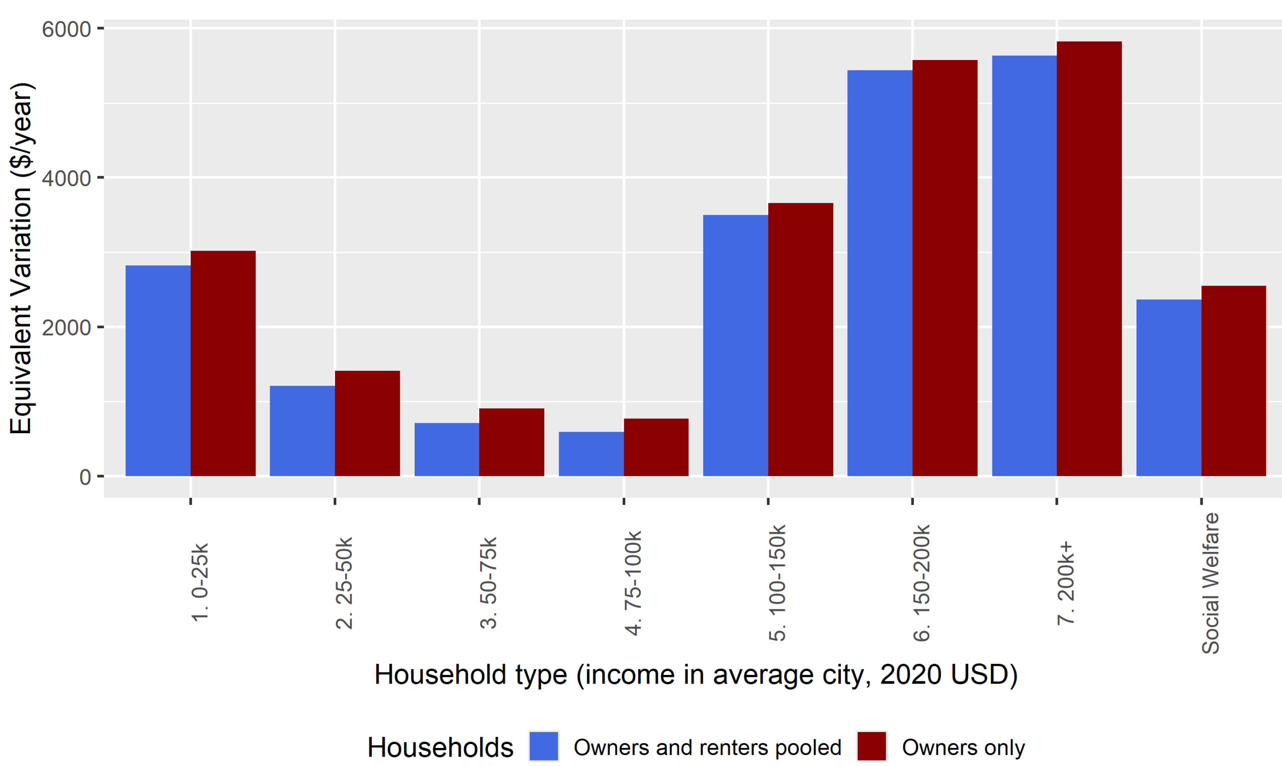

Welfare is measured as the equivalent variation and expressed in dollar terms. Higher values indicate higher welfare gains. Social welfare is the population-weighted average of welfare by type. The results incorporate capital losses for homeowners by income type, using the procedure outlined in Appendix E.3. The results include equilibrium adjustments to neighborhood amenity values, as in the baseline counterfactual. Renters and homeowners are pooled by type using a population-weighted average of welfare changes.

Importantly, these welfare gains are driven by reasons that differ from those of complete deregulation. To see this, I decompose households' welfare into changes in neighborhood consumption and amenity values in Figure 15 of Appendix E.5. For the lowest income households, the amenity contribution is negative; they lose access to desirable neighborhoods that were originally unregulated. However, their consumption gains from newly affordable housing in neighborhoods they value more than compensate, yielding net gains of approximately $\$2{,}850$ per year. For high-income households, the pattern reverses: they gain enormously from amenity improvements (rising neighborhood incomes in their now-protected preferred locations), while paying somewhat more for housing. The gains for upper-income households ($\$100{,}000$–$\$200{,}000$) in dollar terms exceed $\$3{,}500$ per year, driven almost entirely by the amenity channel. This suggests that the permuted allocation does a substantially better job of correcting the neighborhood choice externality than the status quo does. The figure shows that the policy facilitates efficient Tiebout sorting relative to both the observed equilibrium and that without any regulation. This echoes the message that regulation can be an efficient substitute for head taxation to correct externalities Calabrese et al. (2007). However, this model suggests that regulation is misallocated and needs to target the right neighborhoods.

Income sorting within and across cities This policy change shifts regulation toward neighborhoods that are fundamentally valued by the rich. Which neighborhoods are these? How does this new policy change the skill distribution across neighborhoods and cities? In Figure 16 of Appendix E.5, I plot how the average household skill changes with productivity across cities. I find a large negative relationship, with the most productive cities observing substantial losses in average skill. This means that the model predicts that productive cities are too stringent relative to their fundamental amenity value for low-skill households. In other words, cities like San Francisco and San Jose are too desirable for lower-income households to justify their exclusion at the observed levels of regulation.

There are also large changes in the neighborhood income distribution within cities. In Figure 17 of Appendix E.5, I plot the income-density gradient at the baseline level of regulation and after the policy change for both superstar and non-superstar cities. This is a replication of Figure 2 for the new policy. Within productive cities, the figure shows that this new policy causes the gentrification of high-density neighborhoods and an attenuation of the income-density gradient very similar in magnitude to that generated by complete deregulation. The model predicts that low-density neighborhoods in these cities are too stringent when benchmarked against their fundamental amenity value for high-skill households. Low income households benefit from the policy delivering affordable housing in the low-density suburbs of productive cities, like San Francisco.

7 Conclusion

Recent work studying housing regulation asserts that these regulations have implications that extend beyond the issue of housing affordability. First, they are thought to decrease aggregate labor productivity by restricting the size of America's most productive cities (Glaeser and Gyourko, 2018; Hsieh and Moretti, 2019; Duranton and Puga, 2023; Parkhomenko, 2023). Second, they are argued to limit the externalities associated with lower income households free riding off amenities in rich neighborhoods (Hamilton, 1976; Calabrese et al., 2011; Brueckner, 2021) or general congestion externalities (Parkhomenko, 2023; Glaeser and Gyourko, 2018).

In this paper, I asked how a particular type of regulation – the minimum lot size – has shaped welfare, inequality, and the spatial distribution of income. To do so, I built a model of minimum lot sizes that accommodates rich heterogeneity across cities, households, and neighborhoods, featuring direct preferences for neighborhood affluence. This model generates a novel measure of regulatory stringency – the value of a minimal lot – whose spatial variation explains two facts I document: low-density neighborhoods are more affluent within cities, and this pattern is steeper in productive cities. I provided evidence that these facts are caused by regulation by exploiting a shift-share instrument based on the national rise in employment among older workers. The theory shows that patterns of income sorting matter crucially for the welfare impacts of deregulation, which motivates the use of a model that can fit variation in stringency both within and across cities.

I used this model to assess how the implications of complete deregulation – aggregate labor productivity gains and the correction of neighborhood choice externalities – compare in magnitude to the gains in housing affordability. I found very little aggregate productivity gains from deregulation because the ensuing inflow of lower-skill households to productive cities offsets the productivity gains from city expansion. I also found that the observed distribution of minimum lot sizes does little to increase the average neighborhood amenity value. Instead, households of all income levels benefit primarily from the opportunity to consume smaller, more affordable homes. The neighborhoods that gain in income and amenity value are the high-density neighborhoods of expensive cities, mirroring recent gentrification patterns in US downtowns. Taken together, housing affordability is the most important consequence of large-scale deregulation. However, complete deregulation imposes large losses on landowners, and I showed that these losses are concentrated in initially stringent, affluent neighborhoods.

I also studied the case of unilateral deregulation in San Francisco. All households benefit on average, and city-wide land values increase slightly. However, the endogenous amenity response concentrates land value losses in initially stringent, low-density neighborhoods, giving their landowners strong incentives to block densification. This provides a new explanation for the entrenched local opposition to state-level deregulation: it need not reflect a lack of city-wide incentive to reform, but rather the concentrated losses borne by politically influential landowners in affluent neighborhoods. Strong preferences for affluent neighbors are necessary and sufficient to shape these maligned incentives.

Despite these challenges, I showed that minimum lot sizes *can* be a powerful policy tool when well-targeted. A spatial reallocation that concentrates regulation in neighborhoods that rich households value absent regulation delivers welfare gains for all skill levels – both renters and owners – and generates capital gains on national land values. This removes the tension between renters and homeowners that makes nationwide deregulation politically difficult, and suggests a path to housing reform that minimizes the negative consequences for certain groups.

This paper focuses on the role of housing regulation for correcting the externality arising from strong preferences for affluent neighbors. Future work in urban and macroeconomics should look beyond housing regulations as the only relevant policy tool. Is there room for place-based redistribution that can prevent free-riding in neighborhoods while limiting harm to low-income households? Can other well-designed regulatory schemes approximate this ideal? Among an exciting and dynamic literature, I hope to push this line of inquiry as far as it can go.

References

Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015). The economics of density: Evidence from the berlin wall. *Econometrica 83*(6), 2127–2189.

Allen, T., C. Arkolakis, and X. Li (2024, December). On the equilibrium properties of spatial models. *American Economic Review: Insights 6*(4), 472–89.

Amalgro, M. and T. Dominguez-Iino (2021, 12). Location sorting and endogenous amenities: Evidence from amsterdam. *Working paper*.

Anagol, S., F. V. Ferreira, and J. M. Rexer (2021, October). Estimating the economic value of zoning reform. Working Paper 29440, National Bureau of Economic Research.

Barseghyan, L. and S. Coate (2016, August). Property taxation, zoning, and efficiency in a dynamic tiebout model. *American Economic Journal: Economic Policy 8*(3), 1–38.

Bartik, A., A. Gupta, and D. Milo (2025, August). The costs of housing regulation: Evidence from generative regulatory measurement. SSRN Working Paper.

Baum-Snow, N., M. Freedman, and R. Pavan (2018, October). Why has urban inequality increased? *American Economic Journal: Applied Economics 10*(4), 1–42.

Baum-Snow, N. and L. Han (2024). The microgeography of housing supply. *Journal of Political Economy 132*(6), 1897–1946.

Baum-Snow, N. and D. Hartley (2020). Accounting for central neighborhood change, 1980–2010. *Journal of Urban Economics 117*, 103228.

Baum-Snow, N. and R. Pavan (2011, 08). Understanding the City Size Wage Gap. *The Review of Economic Studies 79*(1), 88–127.

Bayer, P., F. Ferreira, and R. McMillan (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy 115*(4), 588–638.

Bayer, P., R. McMillan, A. Murphy, and C. Timmins (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica 84*(3), 893–942.

Bertaud, A. and J. K. Brueckner (2005). Analyzing building-height restrictions: predicted impacts and welfare costs. *Regional Science and Urban Economics 35*(2), 109–125.

Bishop, K. C., J. Dowling, N. V. Kuminoff, and A. D. Murphy (2025). Tax policy and the heterogeneous costs of homeownership. *Journal of Urban Economics 147*, 103750.

Brueckner, J. K. (2021). Jue insight: Zoning and property taxation revisited—was hamilton right? *Journal of Urban Economics*, 103393.

Brueckner, J. K. and S. S. Rosenthal (2009, 11). Gentrification and Neighborhood Housing Cycles: Will America's Future Downtowns Be Rich? *The Review of Economics and Statistics 91*(4), 725–743.

Brueckner, J. K. and R. Singh (2020). Stringency of land-use regulation: Building heights in us cities. *Journal of Urban Economics 116*, 103239.

Brueckner, J. K., J.-F. Thisse, and Y. Zenou (1999). Why is central paris rich and downtown detroit poor?: An amenity-based theory. *European Economic Review 43*(1), 91–107.

Bunten, D. M. (2017, 6). Is the rent too high? aggregate implications of local land-use regulation. Working Paper 2017-064, Board of Governors of the Federal Reserve System.

Calabrese, S., D. Epple, and R. Romano (2007). On the political economy of zoning. *Journal of Public Economics 91*(1-2), 25–49.

Calabrese, S. M., D. N. Epple, and R. E. Romano (2011, 12). Inefficiencies from Metropolitan Political and Fiscal Decentralization: Failures of Tiebout Competition. *The Review of Economic Studies 79*(3), 1081–1111.

Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica 87*(3), 741–835.

Card, D. (2009, May). Immigration and inequality. *American Economic Review 99*(2), 1–21.

Chavent, M., V. Kuentz-Simonet, A. Labenne, and J. Saracco (2018, Dec). Clustgeo: an r package for hierarchical clustering with spatial constraints. *Computational Statistics 33*(4), 1799–1822.

Chetty, R. and N. Hendren (2018, 02). The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects*. *The Quarterly Journal of Economics 133*(3), 1107–1162.

Combes, P.-P., G. Duranton, and L. Gobillon (2008). Spatial wage disparities: Sorting matters! *Journal of Urban Economics 63*(2), 723–742.

Combes, P.-P., G. Duranton, and L. Gobillon (2019, 07). The costs of agglomeration: House and land prices in french cities. *The Review of Economic Studies 86*(4), 1556–1589.

Combes, P.-P., G. Duranton, and L. Gobillon (2021). The production function for housing: Evidence from france. *Journal of Political Economy 129*(10), 2766–2816.

Combes, P.-P. and L. Gobillon (2015). Chapter 5 - the empirics of agglomeration economies. In G. Duranton, J. V. Henderson, and W. C. Strange (Eds.), *Handbook of Regional and Urban Economics*, Volume 5 of *Handbook of Regional and Urban Economics*, pp. 247–348. Elsevier.

Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2023, 04). Income growth and the distributional effects of urban spatial sorting. *The Review of Economic Studies 91*(2), 858–898.

Couture, V. and J. Handbury (2020). Urban revival in america. *Journal of Urban Economics 119*, 103267.

Cui, T. (2023, 11). The emergence of exclusionary zoning across american cities. *Working paper*.

Davidoff, T. (2016). Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors. *Critical Finance Review 5*(2), 177–206.

Davidoff, T., A. Pavlov, and T. Somerville (2022). Not in my neighbour's back yard? laneway homes and neighbours' property values. *Journal of Urban Economics 128*, 103405.

Davis, M. A., J. Gregory, and D. A. Hartley (2023, 06). Preferences over the racial composition of neighborhoods: Estimates and implications. *Working paper*.

Diamond, R. (2016, March). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review 106*(3), 479–524.

Dingel, J. I. and F. Tintelnot (2026). Spatial economics for granular settings. *Econometrica 94*(2), 407–464.

Donald, E., M. Fukui, and Y. Miyauchi (2024, 03). Unpacking aggregate welfare in a spatial economy. *Working paper*.

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

Eeckhout, J., R. Pinheiro, and K. Schmidheiny (2014). Spatial sorting. *Journal of Political Economy 122*(3), 554–620.

Epple, D., B. Gordon, and H. Sieg (2010, June). A new approach to estimating the production function for housing. *American Economic Review 100*(3), 905–24.

Epple, D. and G. J. Platt (1998). Equilibrium and local redistribution in an urban economy when households differ in both preferences and incomes. *Journal of Urban Economics 43*(1), 23–51.

Fajgelbaum, P. D. and C. Gaubert (2020, 01). Optimal Spatial Policies, Geography, and Sorting*. *The Quarterly Journal of Economics 135*(2), 959–1036.

Fernandez, R. and R. Rogerson (1996, 02). Income Distribution, Communities, and the Quality of Public Education*. *The Quarterly Journal of Economics 111*(1), 135–164.

Fernandez, R. and R. Rogerson (1997). Keeping people out: Income distribution, zoning, and the quality of public education. *International Economic Review 38*(1), 23–42.

Finlay, J. and T. Williams (2022, 09). Housing demand, inequality, and spatial sorting. *Working paper*.

Fischel, W. A. (2001). *The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies*. Harvard University Press.

Fogli, A. and V. Guerrieri (2019, 08). The end of the american dream? inequality and segregation in us cities. Working Paper 26143, National Bureau of Economic Research.

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

Glaeser, E. L., M. E. Kahn, and J. Rappaport (2008). Why do the poor live in cities? the role of public transportation. *Journal of Urban Economics 63*(1), 1–24.

Greaney, B. (2025, 10). Homeownership and the distributional effects of uneven regional growth*. *Working paper*.

Greaney, B., A. Parkhomenko, and S. Van Nieuwerburgh (2025, February). Dynamic Urban Economics. NBER Working Papers 33512, National Bureau of Economic Research, Inc.

Grieson, R. E. and J. R. White (1981). The effects of zoning on structure and land markets. *Journal of Urban Economics 10*(3), 271–285.

Guerrieri, V., D. Hartley, and E. Hurst (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics 100*, 45–60.

Gyourko, J., J. S. Hartley, and J. Krimmel (2021). The local residential land use regulatory environment across u.s. housing markets: Evidence from a new wharton index. *Journal of Urban Economics 124*, 103337.

Gyourko, J., C. Mayer, and T. Sinai (2013, 11). Superstar cities. *American Economic Journal: Economic Policy 5*(4), 167–99.

Gyourko, J. and S. E. McCulloch (2024, October). The distaste for housing density. Working Paper 33078, National Bureau of Economic Research.

Gyourko, J. and R. Molloy (2015). Chapter 19 - regulation and housing supply. In G. Duranton, J. V. Henderson, and W. C. Strange (Eds.), *Handbook of Regional and Urban Economics*, Volume 5 of *Handbook of Regional and Urban Economics*, pp. 1289–1337. Elsevier.

Hamilton, B. W. (1975). Zoning and property taxation in a system of local governments. *Urban Studies 12*(2), 205–211.

Hamilton, B. W. (1976). Capitalization of intrajurisdictional differences in local tax prices. *The American Economic Review 66*(5), 743–753.

Helsley, R. W. and W. C. Strange (1995). Strategic growth controls. *Regional Science and Urban Economics 25*(4), 435–460. Recent Advances in Urban Economics and Land Use: A Special Issue in Honour of Hiroyuki Yamada.

Henderson, J. V. and J. Thisse (2001). On strategic community development. *Journal of Political Economy 109*(3), 546–568.

Herkenhoff, K. F., L. E. Ohanian, and E. C. Prescott (2018). Tarnishing the golden and empire states: Land-use restrictions and the u.s. economic slowdown. *Journal of Monetary Economics 93*, 89–109. Carnegie-Rochester-NYU Conference on Public Policy held at the Stern School of Business at New York University.

Hilber, C. A. and F. Robert-Nicoud (2013). On the origins of land use regulations: Theory and evidence from us metro areas. *Journal of Urban Economics 75*, 29–43.

Hong, G. (2025, 03). Two-sided sorting of workers and firms: Implications for spatial inequality and welfare. *Working paper*.

Hornbeck, R. and E. Moretti (2018, May). Who benefits from productivity growth? direct and indirect effects of local tfp growth on wages, rents, and inequality. Working Paper 24661, National Bureau of Economic Research.

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

Kulka, A. (2019, 12). Sorting into neighborhoods: The role of minimum lot sizes. *Working paper*.

Kulka, A., A. Sood, and N. Chiumenti (2026, 03). Under the (neighbor)hood: Understanding interactions among zoning regulations. *The Review of Economics and Statistics*, 1–45.

Lee, S. and J. Lin (2017, 03). Natural Amenities, Neighbourhood Dynamics, and Persistence in the Spatial Distribution of Income. *The Review of Economic Studies 85*(1), 663–694.

Martellini, P. (2022, 2). Local labor markets and aggregate productivity. *Working paper*.

Mei, M. (2022, 10). House size and household size: The distributional effects of the minimum lot size regulation. *Working paper*.

Mills, E. S. (2005). Why do we have urban density controls? *Real Estate Economics 33*(3), 571–585.

Molloy, R. (2020). The effect of housing supply regulation on housing affordability: A review. *Regional Science and Urban Economics 80*(C).

Molloy, R., C. G. Nathanson, and A. Paciorek (2022). Housing supply and affordability: Evidence from rents, housing consumption and household location. *Journal of Urban Economics 129*, 103427.

Ortalo-Magné, F. and A. Prat (2014, February). On the political economy of urban growth: Homeownership versus affordability. *American Economic Journal: Microeconomics 6*(1), 154–81.

Parkhomenko, A. (2023). Local causes and aggregate implications of land use regulation. *Journal of Urban Economics 138*, 103605.

Parkhomenko, A. (2025, 09). Homeownership, polarization, and inequality. *The Review of Economic Studies*, rdaf068.

Saiz, A. (2010, 08). The Geographic Determinants of Housing Supply*. *The Quarterly Journal of Economics 125*(3), 1253–1296.

Song, J. (2025, 01). The effects of residential zoning in u.s. housing markets. *Working paper*.

Su, Y. (2022, February). The rising value of time and the origin of urban gentrification. *American Economic Journal: Economic Policy 14*(1), 402–39.

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

White, M. J. (1975). The effect of zoning on the size of metropolitan areas. *Journal of Urban Economics 2*(4), 279–290.

Wood, S. (2012, October). mgcv: Mixed gam computation vehicle with gcv/aic/reml smoothness estimation.

Zabel, J. and M. Dalton (2011). The impact of minimum lot size regulations on house prices in eastern massachusetts. *Regional Science and Urban Economics 41*(6), 571–583.

A Appendix: Data and Facts Continued

A.1 Data Construction

In this section, I detail the data used in both the empirical work and the estimation of the structural model. This paper uses two broad sources of data:

1. CoreLogic:
   (a) Universe of property assessments (most current assessments as of December 2022)
   (b) 2008-2012 and 2015-2022 universe of transactions
2. Other data:
   (a) 2008-2012 and 2016-2020 American Community Survey (ACS) tabulations, harmonized to 2020 block group definitions
   (b) 2010 and 2020 census housing counts harmonized to 2020 block groups.
   (c) 2007-2010 and 2016-2017 National Neighborhood Data Archive (NaNDA) data at the tract level.
   (d) The US Geological survey's EDNA database (2003 version)

where the separated data ranges are used to construct two panels over block groups. These two sources are used to construct a block-group level panel:

1. **Minimum lot sizes and housing value density**. I defer a discussion of how minimum lot sizes and housing value densities are constructed to Appendices A.2 and A.3, as this requires detail.
2. **The share of housing units in regulated structures**. These enter into the stringency measure in (13). The ACS reports housing counts for structures between 1, 2, and 3-4 housing units, labeled as **Units in Structure**. These are aggregated to arrive at household counts in regulated structures. I do not use CoreLogic data to calculate these shares as many assessments in large multifamily structures do not have information on the number of housing units.
3. **Block group level densities of housing units and incomes**. Average incomes and shares of housing units in regulated structures are directly calculated from the ACS tabulations for each panel. Average incomes in each block group are calculated by dividing tabulated aggregate neighborhood income by the number of sampled households.[52] Housing counts used to construct density measures are from the 2020 Census.
4. **Various controls used in estimation**. These include (from the ACS) median building age, household size; share of cars in commuting, buses and other public transport; shares of households that are families, average travel time, white and college shares; (from NaNDA) density of performing arts facilities, sports, casinos, recreational, transit stops, fast food, restaurants, coffee shops and bars; fraction of land area in parks, covered by perennial snow, deciduous forest, evergreen, mixed forest, shrubs, herbaceous, woody and herbaceous wetlands.

[52]These tabulations are labeled **Aggregate Household Income in the Past 12 Months** and **Housing Units**

5. **Slopes**. I use the USGS EDNA database at 30×30m resolution to create average slope measures at the block group level. This data is used solely in the estimation of the effect of income on neighborhood amenity value in Section 5.

There are roughly $196,000$ block groups in 2013-definition Metropolitan Statistical Areas (MSAs). To estimate the regression (14), I remove block groups with less than 25 housing units per square mile, representing about 8000 block groups (4% of the sample). The facts are robust to censoring at a wide range of densities, including not censoring at all. Summary statistics for reported variables above given in Table 4.

A.2 Constructing Zoning Districts

Block groups generally do not correspond to areas by which local governments assign uniform land use regulation. To measure "bunching" of minimum lot sizes around regulatory levels, I need to take a stand on the geographic unit by which to construct lot size distributions. I call these geographic units "Zoning Districts". More often than not, zoning districts can be inferred from zoning codes reported in the CoreLogic data.[53] I assign a zoning code to a block group by taking the modal code across parcels in the block group. Block groups with no populated zoning code data are assigned missing.

For roughly $66,000$ block groups (one third of the sample), there is no data on zoning codes to apply regulation. To extend coverage of zoning districts, I cluster the remaining block groups. The fundamental challenge in doing so is a trade-off when choosing the size of these zoning districts: on one hand, large zoning districts may pool together multiple levels of regulation, and there is no way to distinguish between them. On the other hand, zoning districts that are too small may result in spurious measurements due to the lack of observations. I perform the clustering algorithm within each US municipality, as these are typically responsible for setting regulation. I assign at most one municipality to a block group by matching municipalities to geocoded parcels, and taking the modal municipality across parcels. I test the algorithm on two different definitions of a municipality: the municipality reported by CoreLogic for assessment purposes, and the municipality from the US Place Shapefiles.

To aggregate block groups into zoning districts, I employ a highly flexible clustering algorithm, which I describe here. This algorithm nests the clustering algorithm of Chavent et al. (2018) and is implemented via their R package. This is useful because it allows for the weighing of two types of variables to assign clusters: geographic proximity and other non-geographic variables. I use block-group modal lot size and first decile of the lot size distribution as these variables. The algorithm allows me to consider how important geographic proximity needs to be relative to other location based characteristics that might signal similar minimum lot sizes.

The number of clusters to assign is a parameter in Chavent et al. (2018)'s algorithm. This governs how large zoning districts are. I provide a maximum level of a *targeted average number of block groups per cluster* to as a hyperparameter in the algorithm. This parameter defines the minimum number of clusters that must appear in a municipality. For example, in a municipality with 100 block groups without zoning codes, and if I impose that the maximum cluster size is roughly 25 block groups, then the minimum number of clusters is roughly 4. Subject to this minimum number of clusters, I select the actual number of clusters to minimize the *silhouette score* on non-geographic variables, which is a metric that assesses both within-cluster similarity and across-cluster dissimilarity and is typically used to assess clustering performance. I consider a range of targeted maximum cluster sizes of 5, 15, 25, 100 and 250 block groups, respectively.

Summarizing, my algorithm incorporates three important hyperparameters:

[53] However, no information on specific regulation can be extracted from these codes.

1. Two definitions of a municipality. The first is the municipality reported in the assessments. The second is the incorporated city taken from the US Place shapefiles. If block groups are not in a reported municipality, I treat the county as a municipality instead.[54]

2. The weight of geographic proximity relative to non-geographic variables when assigning clusters (Chavent et al., 2018).

3. Targeted maximum sizes (in number of block groups) for the average cluster.

The algorithm can be roughly described as follows:

Algorithm 1 Clustering

For all hyperparameters do:

For all municipalities do:

1. Assign block groups with same non-missing zoning code to district
2. Cluster remaining block groups choosing number of clusters between 2 and $N - 1$, where $N - 1$ is the number of block groups to be clustered
3. Select desired number of clusters to maximize the silhouette score subject to the minimum number of clusters constraint. All block groups in a cluster are assigned a different zoning code from other clusters.

Figure 8 shows the optimal set of zoning districts for Hayward in the San Francisco metropolitan area (optimal will be defined shortly). The optimal algorithm places largest weight on geographic proximity when assigning clusters. The algorithm also uses a maximum cluster size of 5 block groups, as well as CoreLogic's municipality definition. The average zoning district (including block groups that have missing and non-missing zoning codes) has 3.3 block groups, with a standard deviation of 11. The median size of a zoning district is one block group. Clusters are not sensitive to choices of other hyperparameters; I provide more discussion when assessing the performance of the algorithm in Appendix A.4.

[54]Unincorporated locations typically have land use regulation set by the county they are in.

Figure 8: Zoning Districts in Hayward, California

Hayward, CA Optimal Zoning Districts

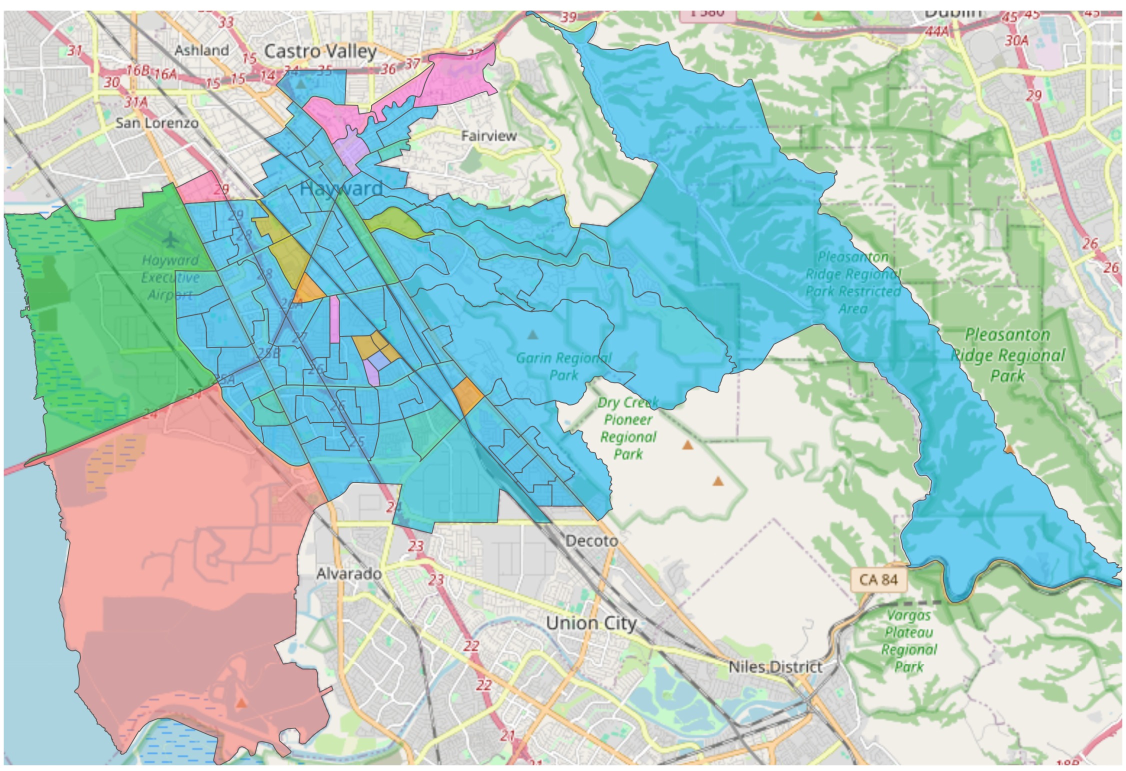

A.3 Measuring Minimum Lot Sizes and Stringency

The zoning districts give a set of geographic boundaries for which to measure "bunching" in the lot size distribution. I use the mode of the distribution as a measure of this bunching. Within each zoning district, I take the modal lot size for single family homes, duplexes, triplexes and fourplexes, and adjust the lot size by the implied number of housing units per lot (for example, dividing the minimum lot size by 2 for duplexes). This is useful because municipalities often assign different lot size restrictions to these four different types of structures. I define the *baseline unit density restriction* as the minimum of these to be as conservative as possible, as detecting bunching alone cannot measure regulations pertaining to which of these four structures are allowed to be built. This minimum is the unit density restriction that is used both in the facts and in the calibration of the model.

To build intuition behind why the mode is an accurate measure of bunching, consider the distribution of lot sizes in Hayward, California (Figure 1). The official minimum lot size for most of Hayward's zoning districts is 5000 square feet. The mode of the observed distribution of lots in Hayward is also 5000 square feet, as shown below. Mostly all block groups in Hayward are assigned the 5000 square foot minimum lot size with this measurement procedure.

Housing Value density Next, I outline how I measure the density of housing values, which enters into the measure of regulatory stringency in Equation (13). I limit the sample to *regulated structures* (single family, duplex, triplex and fourplexes) as the only use for housing value density pertains to the stringency of regulation (in both the model and the empirical work). I deviate from directly calculating the density of house values (for example, by taking the median sales value across regulated structures and dividing by the median lot size for each block group). This is because housing value density varies significantly within block groups between houses with small lots and those with large lots; this leads to large measurement error in regulatory stringency when the measured minimum lot size deviates from the median lot size within a neighborhood.

Instead, I predict housing value densities *at the minimal lot* using a linear model. Within each MSA, I estimate the following regression over assessments indexed by a in neighborhood i and city c

$$\text{HouseValue}_{aic} = \beta_{0ic} + \beta_{1c}\text{LotSize}_{aic} + \epsilon_{aic} \tag{33}$$

The slope parameter of this model informs about the relative price of houses on small lots relative to large lots within the same MSA. I then use the β_{1c} parameters from these regressions to predict the house value of the minimal lot in every neighborhood i using the formula

$$\text{PredHouseValue}_{ic} = \text{MedianHouseValue}_{ic} - \beta_{1c}\text{MedianLotSize}_{ic} + \beta_{1c}\text{MinLotSize}_{ic} \tag{34}$$

where MinLotSize_{ic} is the measured minimum lot size (ignoring implicit unit density restrictions, such as dividing by 2 if duplexes are allowed). This formula ensures that if the minimum lot size is also at the median in the lot size distribution, then its house value is also predicted to be the median house value of the neighborhood.[55] Finally, I measure housing

[55]This often occurs in neighborhoods with extreme "bunching" – where most lots are built at the minimum.

value densities as

$$\text{HousingValueDensity}_{ic} = \text{PredHouseValue}_{ic}/\text{MinLotSize}_{ic} \tag{35}$$

which enter directly into the stringency measure, Equation (13).

Additional cleaning I perform additional cleaning to ensure my measure of regulatory stringency is not wrought with measurement error and other pathologies. I list each procedure below:

1. Transactions data are particularly sparse in small and inexpensive cities. This means that there are select block groups in small cities that report massive home values relative to nearby block groups, which makes the measure of regulatory stringency large. This is driven by measurement error when using observed transactions to estimate the transaction value of a median home in these block groups. In cities that report a median transaction below 300k USD (below the US average), I impute block groups median transaction value to the MSA median if it is above 1 million USD. This imputation has the effect of limiting measurement error that compounds in model counterfactuals, making some cities look artificially more stringent then they plausibly are; driven by only a handful of block groups that appear extremely stringent for spurious reasons. Conclusions of counterfactual exercises remain the same whether or not I do this cleaning.

2. Sometimes, the algorithm fails in the sense that the size of the measured minimum lot size is well above the average lot size in some block groups. This usually happens in remote locations where parcels are sparse and their size idiosyncratic. For model counterfactuals and to estimate the facts, I set the observed minimum lot size to zero if the assigned minimum lot size is twice the average. The facts and model counterfactuals are robust to different thresholds for which to perform this cleaning; including no cleaning at all.

In Table 4, I report summary statistics of the measured unit density restrictions, housing value densities, and the accompanying measure of regulatory stringency from the Facts in Section 3.2. The average unit density restriction across block groups is 0.272 acres per housing unit with a standard deviation of 3.74. Table 3 additionally reports boxplots of stringency for select cities. Figure 9 shows the measured unit density restrictions for Hayward, California by zoning district.

Figure 9: Hayward Unit density restrictions

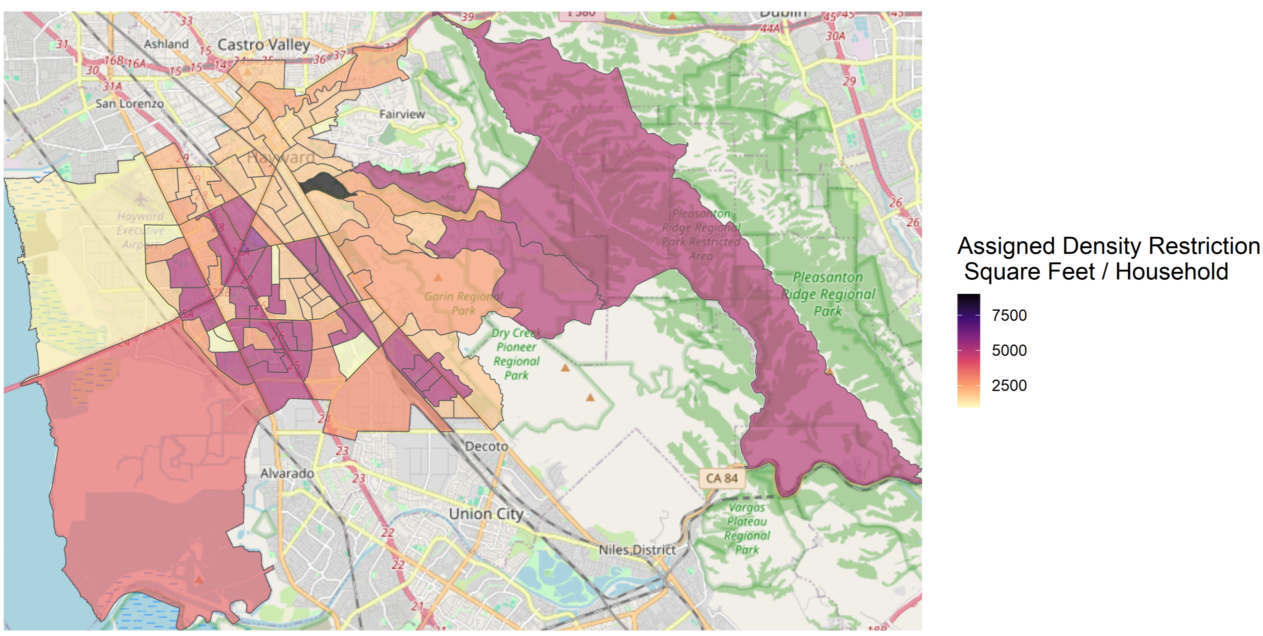

A.4 Validating Minimum Lot Sizes

To validate the algorithm and to select clustering hyperparameters, I introduce two sources of official data on minimum lot sizes (or more generally, unit density restrictions). The first source, which mirrors validation exercises in Song (2025) and Cui (2023), is the Terner California Land Use Survey (TCLUS). This source of data is aggregated to the municipality level and cannot be used to test how well the algorithm performs on microgeographic variation in regulation. The second data source – the MAPC Zoning Atlas – addresses this limitation by offering zoning-district level regulation data. I address how the validation procedure works for each source below.

TCLUS The TCLUS reports minimum lot sizes for both single family homes and ”multifamily homes” at the municipality level for over 200 Californian municipalities. As recognized by Song (2025), performing validation when the constructed measure of regulation varies within municipalities is difficult. Notwithstanding, I assume the TCLUS data reports the median minimum lot size (weighted by population) for each of these structure types (single family and multifamily).

I take, at the municipality level, the median modal lot across block groups separately for single family homes, duplexes, triplexes and fourplexes (recall that the measured unit density restriction in a neighborhood is the minimum of each of these modes after adjusting for the implied number of units per lot). In each municipality, I match the TCLUS-reported minimum lot size for single family homes with the modal single family lot. I also match the TCLUS-reported minimum lot size for multifamily structures with the closest modal lot amongst duplexes, triplexes, and fourplexes. I then calculate the *error rate* for each municipality for both single family and multifamily homes. On this sample of municipalities, the median error on the best set of clustering hyperparameters is 7% of the official minimum lot size for single family homes and 16% for multifamily homes.

Note that the unit density restriction actually used in the paper is the minimum across single family homes, duplexes, triplexes and fourplexes, as described in Appendix A.3. By definition, this is conservative relative to the minimum lot sizes used to compare to the TCLUS.

MAPC Zoning Atlas The MAPC Zoning Atlas is a publicly available GIS database compiled by the Metropolitan Area Planning Council that records the official minimum lot size and permitted use type for every zoning district in Massachusetts. I restrict to residential zones with positive, non-missing minimum lot sizes and convert the official minimum lot size from square feet to acres. I then spatially intersect the MAPC zoning polygons with block groups in the Greater Boston CBSA, retaining only those block groups where the share of housing units in regulated structures exceeds 25%. This yields 504 zoning districts available for comparison.

The fundamental challenge is that my estimated zoning districts (clusters of block groups) do not coincide with official MAPC district boundaries. To address this, I aggregate to the MAPC zoning district level: for each district, I compute the area-weighted median of the estimated single-family modal lot size across all intersected block groups, where weights are proportional to the intersection area between the MAPC polygon and each block group. I then calculate the absolute deviation of this estimate from the official minimum lot size, expressed as a fraction of the official lot size, and take the area-weighted median of these deviations across all districts. This error rate is 18.4% of the official unit density restriction at the optimal set of clustering hyperparameters. Although larger than the TCLUS error, this reflects the additional challenge of recovering intra-municipal variation in regulation, and the estimate remains quite accurate. Reassuringly, performance of the algorithm does not vary much across hyperparameters, which implies that I am not over-fitting the data used to validate the algorithm.

A.5 Summary Statistics

Table 3: Regulatory stringency boxplot for select cities.

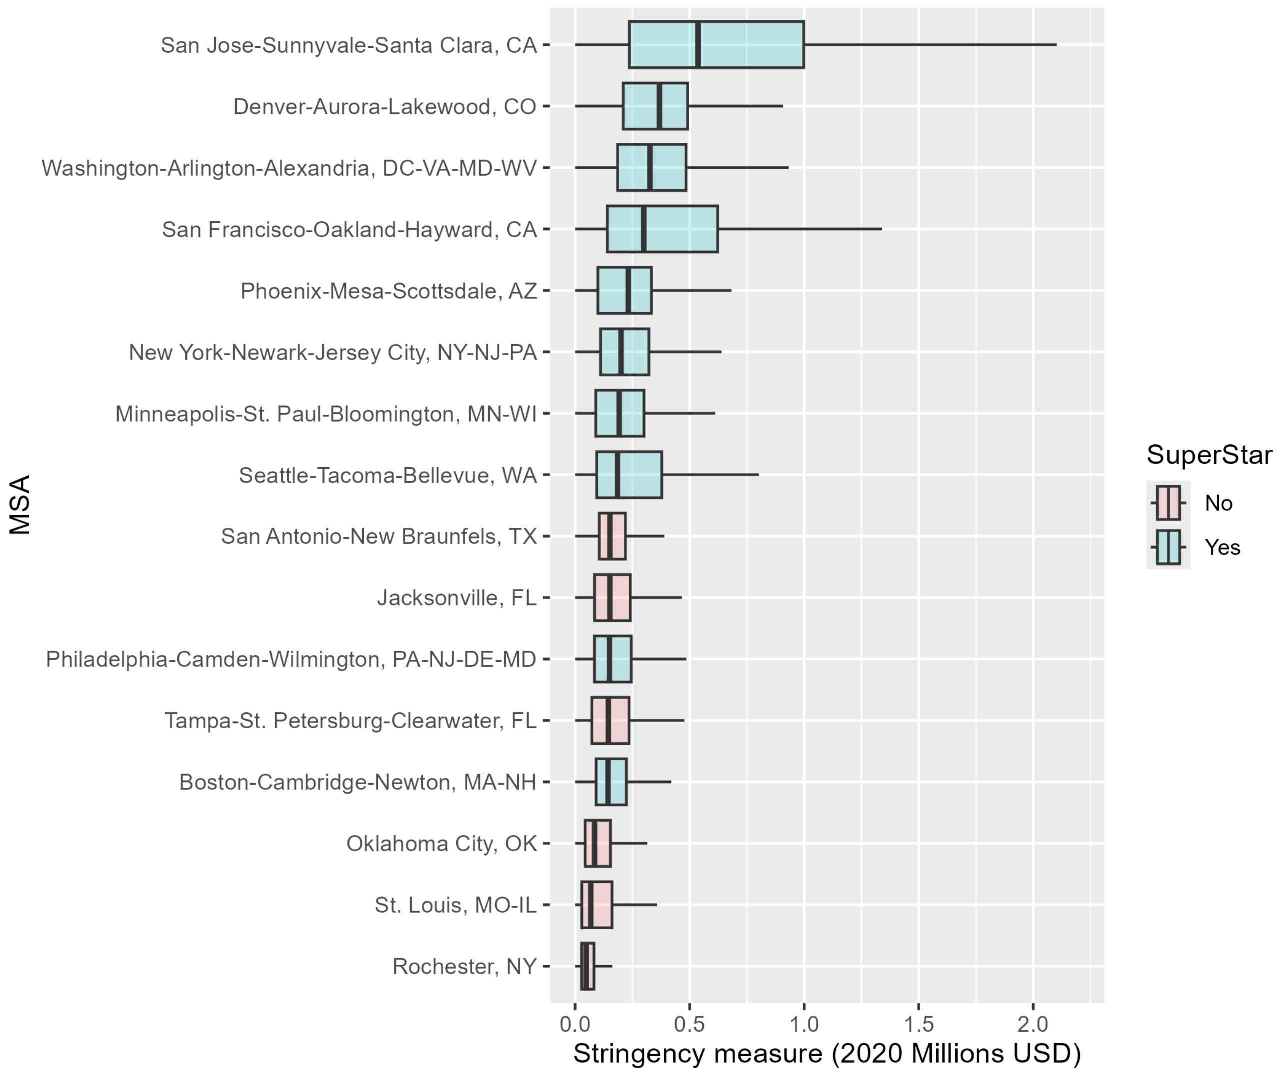

Regulatory stringency boxplot for select cities. Reports stringency measure introduced in Equation (13) for select cities. Cities are colored based on whether they are included in the superstar sample used to construct the Facts in Section 3.2.

Table 4: Summary Statistics for Key Variables,
disaggregated by superstar city status

Superstar?	Aggregate				No				Yes			
Variable	N	Mean	Sd	Median	N	Mean	Sd	Median	N	Mean	Sd	Median
ln Average Income	194533	11.2	0.567	11.2	83861	11	0.524	11	110672	11.4	0.556	11.4
Unit Density Restriction (acres)	190685	0.302	3.76	0.115	83141	0.386	5.2	0.14	107544	0.237	2.04	0.0869
Stringency measure $\tilde{R}_{ic}$ (2020 millions USD)	187690	0.195	0.479	0.126	81759	0.123	0.11	0.0897	105931	0.251	0.624	0.169
Housing Value Density (2020 millions USD/acre)	187531	4.56	133	1.49	81764	1.56	4.89	0.839	105767	6.88	177	2.47
Regulated housing units (share)	190183	0.794	0.252	0.899	82971	0.811	0.222	0.894	107212	0.78	0.273	0.905

Summary statistics. ”Unit Density Restriction” refers to the measured physical unit density restriction that enters into Equation (13), and is measured in Appendix A.3. ”Regulated Housing Units” refers to the share of housing units in single family homes, duplexes, triplexes and fourplexes. housing value density is measured in Appendix A.3 and enters into (13). ”Stringency measure” is the empirical regulatory stringency measure introduced in Equation (13). Variation in nonmissing data are a result of either 1) incomplete transactions coverage or 2) additional cleaning procedures described in Appendix A.3.

Table 5: Relationship between regulatory stringency and productivity.

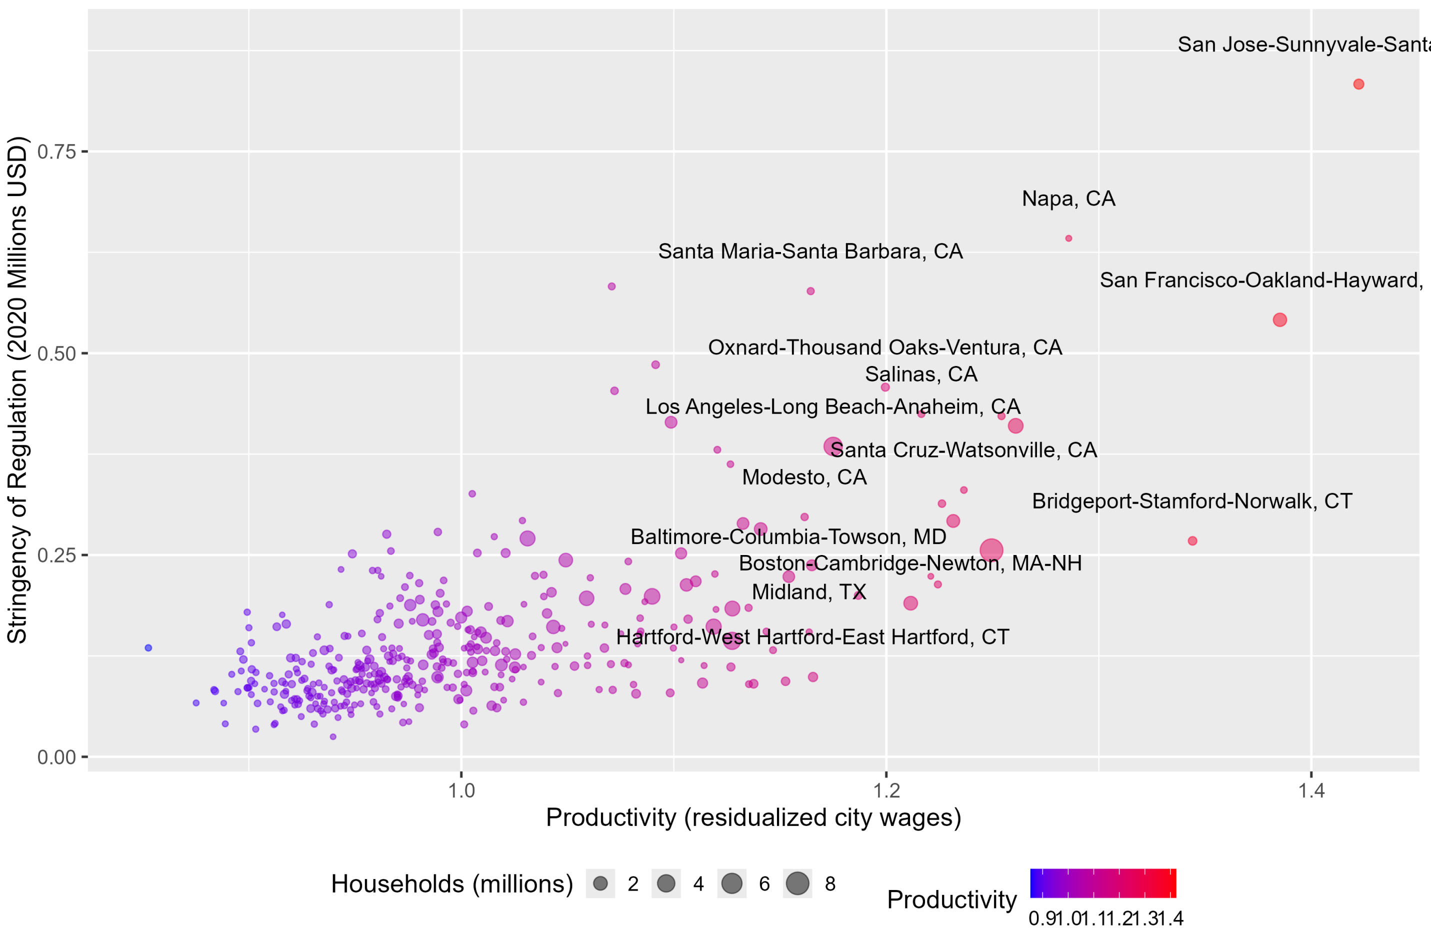

Reports stringency measure introduced in Equation (13) plotted against residualized city wages with some additional city characteristics. For methodology behind construction of wages, see Section 4. Productivity is normalized to be on average one across cities.

A.6 Discussion of Alternative Specifications

In this section, I provide additional robustness checks for the facts. I also discuss instances where the Facts in Section 3.2 do not hold.

Distance to CBD The Facts in section 3.2 are not robust when considering distance to the CBD instead of the density ranking. In Figure 10, I reproduce regressions of demeaned income (Panel A) and regulatory stringency (Panel B) on the distance to CBD. We observe a generally increasing relationship between income and CBD distance. Panel B shows that there is evidence that high density neighborhoods close to the CBD are less stringent in superstar cities. However, a key difference is that the relationship is not monotone for each city sample.

Figure 10: Facts in Section 3.2 when considering distance to CBD

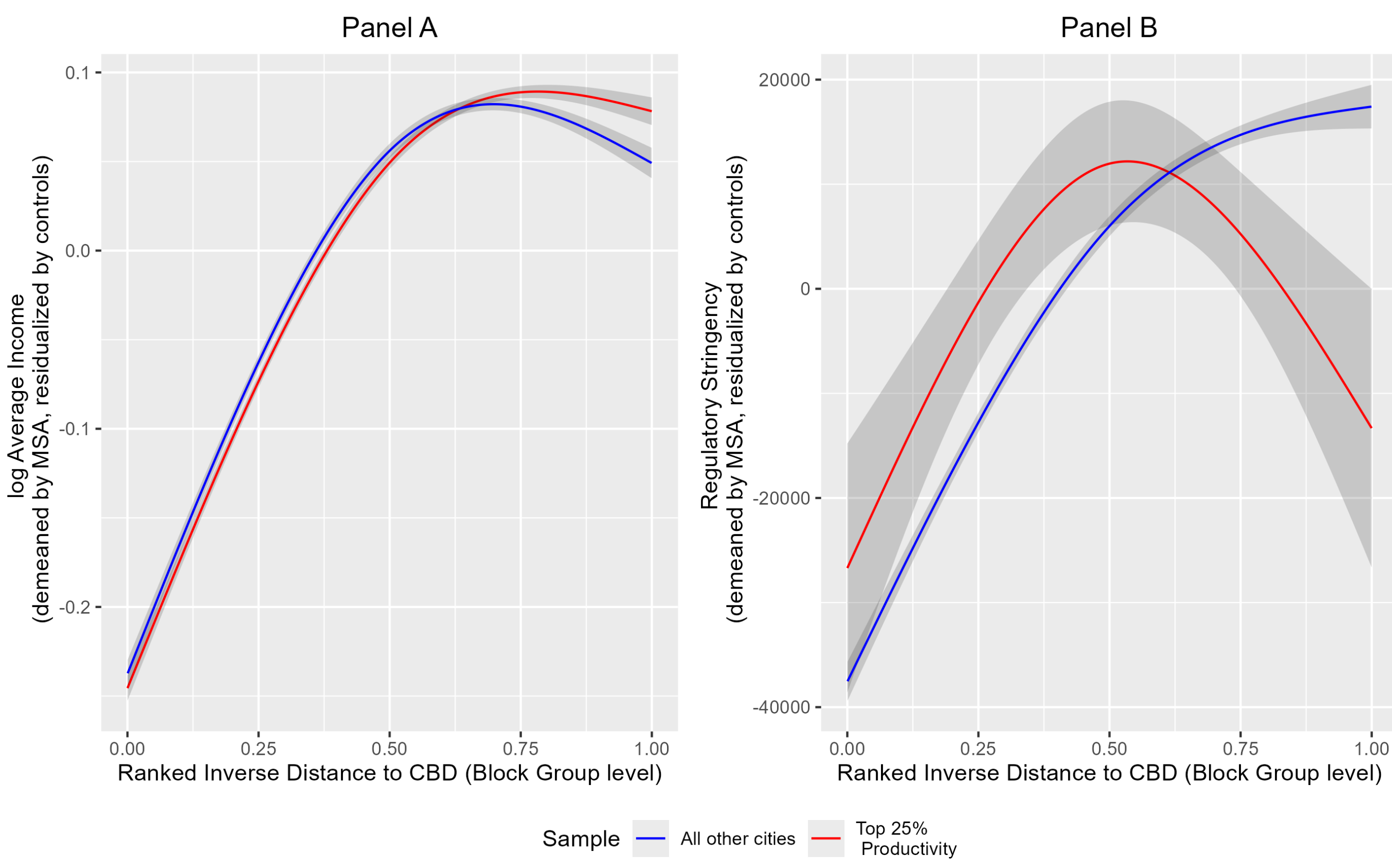

Relationship over time The Facts are robust to various time periods. I repeat the same exercise using 2008-2012 ACS data, with results reported in Figure 11. The difference in the negative income-density gradient across samples is somewhat stronger in this sample relative to the 2016-2020 ACS. This likely reflects the recent gentrification of high density neighborhoods nationwide. Results also hold when residualizing regressions by the standard set of controls.

Alternative measures of superstars The Facts in Section 3.2 look quantitatively identical when adopting three different measures of "superstar" cities: the top 25% and bottom 75% of density, housing prices, and a measure of city productivity alone (see Section 4 for details on productivity measurement). Moreover, these facts are robust to alternative measures of superstars that use 10% and 50% thresholds in their definition (rather than the baseline 25%).

Weights Facts in Section 3.2 use regressions that weigh block groups evenly, roughly corresponding to a population-weighted regression. The facts are not driven by larger weights

Figure 11: Fact in 2008-2012 ACS sample

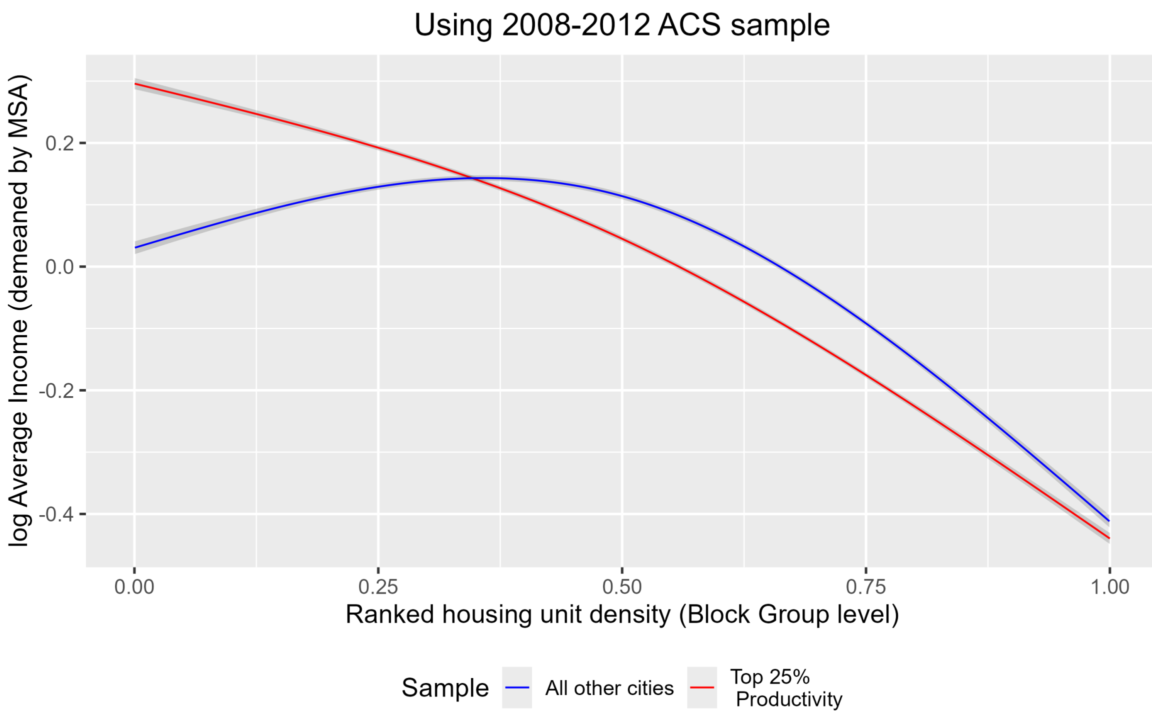

ascribed to bigger cities; they hold when each city receives equal weight. Moreover, each figure is identical when weighting the regression by the number of households in each block group.

Various hyperparameters There is a large number of hyperparameters that characterize zoning districts, and thus the measure of regulation used to show the Facts in Section 3.2. Reassuringly, these facts look quantitatively identical for the entire hyperparameter space considered. This is likely driven by the fact that there are only minimal differences in clusters across hyperparameters (after all, two-thirds of all block groups are assigned regulation by populated zoning codes).

B Appendix: Theory

B.1 Derivation of the distortion factor

For brevity, I drop zone and neighborhood subscripts for this derivation. I start with solving the households maximization problem, restated here:

$$V(z) := \max_{A,g} \underbrace{z^{-1}\beta^{-\beta}(1-\beta)^{-(1-\beta)}(A-\bar{A})^{\beta}g^{1-\beta}}_{\text{Consumption value}} + \underbrace{\log b(z)}_{\text{Amenity value}} \tag{36}$$

subject to

$$PA \geq R \text{ and}$$
$$PA + g \leq wz$$

It is instructive to first consider the case where $R = 0$ (or equivalently, when regulation is not binding). This problem has the well known solution (ignoring amenities as they do not matter for consumption decisions)

$$V(z) = w\frac{(1-\frac{P\bar{A}}{wz})}{P^{\beta}}$$

with spending shares equal to

$$\beta + (1-\beta)\frac{P\bar{A}}{wz}$$

This means that regulation R is binding if and only if

$$\frac{R}{wz} > \beta + (1-\beta)\frac{P\bar{A}}{wz}$$

If regulation is binding, consumption g is $wz - R$ and housing consumption A is $\frac{R}{P}$. Substituting this into the maximization problem above yields

$$V(z) = z^{-1}\beta^{-\beta}(1-\beta)^{-(1-\beta)}(R-P\bar{A})^{\beta}P^{-\beta}(wz-R)^{1-\beta}$$

Multiplying and dividing $V(z)$ by the value $wz(1-\frac{P\bar{A}}{wz})$ and rearranging allows us to express this indirect utility as a product of the utility if minimum lot sizes were not binding and a distortion factor, as is suggested in the text:

$$V(z) = \underbrace{w\left[\frac{1-\frac{P\bar{A}}{wz}}{P^{\beta}}\right]}_{\text{Undistorted utility}} \underbrace{\left[\frac{\left(1-\frac{R}{wz}\right)\left(1-\frac{P\bar{A}}{wz}\right)^{-1}}{1-\beta}\right]^{1-\beta}\left[\frac{\left(R-P\bar{A}\right)\left(wz-P\bar{A}\right)^{-1}}{\beta}\right]^{\beta}}_{\text{Distortion Factor}} \tag{37}$$

Assuming $\bar{A} = 0$ (Cobb-Douglas preferences), Equation (37) reduces down to what is reported in the text:

$$\underbrace{\frac{w_i}{P^{\beta}}}_{\text{Undistorted utility}} \times \underbrace{\left[\frac{\frac{R_i}{w_i z}}{\beta}\right]^{\beta}\left[\frac{1-\frac{R_i}{w_i z}}{1-\beta}\right]^{1-\beta}}_{\text{Distortion factor}}$$

where neighborhood and zone indicators are reintroduced in the equation for comparison.

B.2 Derivation of Equation (6)

Recall the household's problem (4), dropping neighborhood and zone subscripts (noting that $R > 0$ only applies to regulated zones by assumption). The objective is to prove that

$$\frac{\partial}{\partial z}\left[\frac{\partial V(z)}{\partial R} / \frac{\partial V(z)}{\partial P}\right] < 0 \tag{38}$$

Whenever regulation is binding at z. To do this, we use the formula $V(z) = z^{-1}\beta^{-\beta}(1-\beta)^{-(1-\beta)}(R - P\bar{A})^{\beta}P^{-\beta}(wz - R)^{1-\beta} + \log b(z)$ from Appendix B.1. It can be easily shown that

$$\frac{\partial V(z)}{\partial R} = V(z)\left(\beta\frac{1}{R - P\bar{A}} - (1-\beta)\frac{1}{wz - R}\right)$$

and

$$\frac{\partial V(z)}{\partial P} = -V(z)\beta\left(\frac{1}{P} + \frac{\bar{A}}{R - P\bar{A}}\right)$$

so that

$$\frac{\partial V(z)}{\partial R} / \frac{\partial V(z)}{\partial P} = -\frac{\beta\frac{1}{R-P\bar{A}} - (1-\beta)\frac{1}{wz-R}}{\beta\left(\frac{1}{P} + \frac{\bar{A}}{R-P\bar{A}}\right)} \tag{39}$$

Differentiating this expression with respect to z yields

$$\frac{\partial}{\partial z}\left[\frac{\partial V(z)}{\partial R} / \frac{\partial V(z)}{\partial P}\right] = -\left[\beta\left(\frac{1}{P} + \frac{\bar{A}}{R - P\bar{A}}\right)\right]^{-1}(1-\beta)\frac{w}{(wz - R)^2} \tag{40}$$

Since regulation is binding ($\frac{R}{wz} > \beta + (1-\beta)\frac{P\bar{A}}{wz}$ from B.1) we know that $R > P\bar{A}$. Then, for all $wz > R$, the above expression is strictly negative. This proves the result.

B.3 Microfoundations for endogenous amenities channel

In this section, I provide two approximate microfoundations for the relationship in Equation (9). I abstract away from different zones within a neighborhood i.

Local public goods financed through property taxes Suppose each neighborhood implements a property tax rate t_i on the value of a home (in terms of numeraire consumption). The revenue from this tax is rebated equally to all residents. Let S_i be total spending on housing in neighborhood i by all residents. The average value of a home (which is equal to average spending on a home) from the housing market clearing condition is

$$\frac{S_i}{L_i}$$

where L_i is the neighborhood population and so the consumption rebate is $t_i \frac{S_i}{L_i}$ for each household irrespective of income. Suppose households amenity value $b_i(z)$ is a composite of the public good K and fundamential amenities $\nu_i(z)$, or $b_i(z) = K^{\Omega(z)} \nu_i(z)$. Then provision of the public good K is $t_i \frac{S_i}{L_i}$. Then, amenity values are $b_i(z) = t_i^{\Omega(z)} \left[\frac{S_i}{L_i}\right]^{\Omega(z)} \nu_i(z)$. When housing consumption is Cobb-Douglas and no minimum lot sizes are binding, this means that $S_i = \beta Y_i$, where Y_i is total income in neighborhood i. Putting this all together,

$$b_i(z) = \beta^{\Omega(z)} t_i^{\Omega(z)} \left[\frac{Y_i}{L_i}\right]^{\Omega(z)} \nu_i(z)$$

where $\frac{Y_i}{L_i}$ property tax rate terms $t_i^{\Omega(z)}$ and spending shares $\beta^{\Omega(z)}$ are absorbed into the fundamental amenities term $\nu_i(z)$ (if they can reasonably be taken as exogenous). Obviously, property tax rates are endogenous in the local public finance literature, and respond to the changing income composition of the neighborhood to facilitate Tiebout sorting (Calabrese et al., 2007, 2011). Since I take a broad view of this neighborhood choice externality and include many reasons for why it occurs, I abstract away from these specific concerns.

B.4 Proof of Proposition 1

Proposition 1 has three parts. First, it says that cities with higher productivity are relatively more affluent because of variation in regulatory stringency. Second, it says that, within cities, neighborhoods that have higher regulatory stringency in equilibrium also have lower density. Thirdly, regulation generates a negative income-density gradient that is stronger in more productive cities. For this proof only, I adopt a separate index for cities and neighborhoods – c and i, respectively. In what follows, note that we assume a perfect mobility equilibrium for this proof and thus a spatial equilibrium is a set of allocations across neighborhoods $L_{ic}(z)$ such that $V_{ic}(z) \leq V(z)$ for some $V(z)$ and $V_{ic}(z) < V(z) \implies L_{ic}(z) = 0$ for all (i, c) pairs. Also note that preferences are Cobb-Douglas, $\bar{A} = 0$ and $\Omega(z) = 0$ for every z, each neighborhood provides unit amenity value and has an identical production technology for housing. Lastly, recall that I make no distinction between regulated and unregulated zones to make the exposition as simple as necessary.

For simplicity, I assume i and c and z are continuous which will directly imply all results on discrete space. For a location with regulatory stringency R_{ic} that offers wage w_c, let $s_{ic}(z) := \frac{R_{ic}}{w_c z}$. I start with two lemmas. The first lemma characterizes the bid rent curve – the maximal rent willing to be paid to live in an arbitrary neighborhood offering wages w and regulation levels R.

Lemma 1. *The bid rent curve θ satisfies $\theta_{ic}(z) = V(z)^{-\frac{1}{\beta}} w_c^{\frac{1}{\beta}} \frac{s_{ic}(z)}{\beta} \left(\frac{1-s_{ic}(z)}{1-\beta}\right)^{\frac{1-\beta}{\beta}}$ if $s_{ic}(z) \in [\beta, 1)$, $\theta_{ic}(z) = V(z)^{-\frac{1}{\beta}} w_c^{\frac{1}{\beta}}$ if $s_{ic}(z) < \beta$. If $s_{ic}(z) \geq 1$ then $\theta_{ic}(z) = 0$*

Proof. This follows directly from rearranging Equation (5) to determine what the price of housing services must be to rationalize utility $V(z)$ at wages w and regulation R. In addition, recall that we set utility to zero if the price of a minimal lot exceeds income. □

The second lemma says that the slope of the bid rent curve across cities is strictly increasing income, all else equal, which will eventually imply income sorting on regulation. This result is essentially implied by the supermodularity property of (6) combined with regulatory stringency increasing faster than wages in c.

Lemma 2. *If $V_{ic}(z) = V(z)$ then $\frac{\partial \theta}{\partial c}$ is strictly increasing in z when $s_{ic}(z) \in (\beta, 1)$ or is constant when $s_{ic}(z) < \beta$ or $s_{ic}(z) \geq 1$.*

Proof. The slope of the bid rent curve when $\beta < s < 1$ can be expressed as

$$\frac{\partial \theta}{\partial c} = \theta_{ic}(z) \left[\frac{1}{\beta} \frac{\iota_{c'}}{\iota_c} + \frac{\partial s}{\partial c} s_{ic}(z)^{-1} - \frac{1-\beta}{\beta} \frac{\partial s}{\partial c} (1 - s_{ic}(z))^{-1} \right]$$

Recall that $s_{ic}(z) = \frac{R_{ic}}{\iota_c z} = \frac{\alpha(c) i}{\iota_c z}$ which is strictly increasing in c by one of the Proposition's assumptions. This can be used to show $\frac{\partial \theta}{\partial c}$ is increasing in z using similar arguments from Appendix B.2.

Lastly, if $s_{ic}(z) < \beta$ then $\theta_{ic}(z) = V(z)^{-\frac{1}{\beta}} \iota_c^{\frac{1}{\beta}}$ which is constant in z because of the assumption $V(z) = \max_{i,c} V_{ic}(z) = \frac{\iota_c}{P_i^{\beta}}$. The case where $s_{ic}(z) \geq 1$ is trivial.

□

The third lemma deals with across-neighborhood comparisons of the bid rent function.

Lemma 3. *If $V_{ic}(z) = V(z)$ then $\frac{\partial \theta}{\partial i}$ is strictly increasing in z when $s_{ic}(z) \in (\beta, 1)$ or is constant when $s_{ic}(z) < \beta$ or $s_{ic}(z) \geq 1$.*

Proof. Follows directly from the derivation in Appendix B.2. □

Using these three lemmas will help with proving the three statements in Proposition 1.

1. High productivity cities are more affluent:

   Proof. Recall that there are two income types $z_1 > z_0$. Fix any two cities c_0 and c_1 with $c_0 < c_1$. This result can be shown if one can prove that, for all neighborhoods i where the minimum lot size in either city is binding for at least one agent

   $$\text{If } L_{i,c_0}(z_1) > 0 \text{ then } L_{i,c_1}(z_1) > 0 \text{ and } L_{i,c_1}(z_0) = 0$$

   I will prove this. $L_{i,c_0}(z_1) > 0$ implies that type z_1 can successfully bid in the neighborhood, so that $\theta_{i,c_0}(z_1) = \max_z \theta(i, c_0, z)$. Suppose for contradiction that $L_{i,c_1}(z_1) = 0$ or $L_{i,c_1}(z_0) > 0$. Then, in either case, $\theta_{i,c_1}(z_0) = \max_z \theta_{i,c_1}(z)$. But this is impossible, since Lemma 2[56] and $\theta_{i,c_0}(z_1) = \max_z \theta_{i,c_0}(z)$ implies that $\theta_{i,c_1}(z_1) > \theta_{i,c_1}(z_0)$.

   □

2. Negative income density gradient in productive cities.

   Proof. Fix two cities with $c_0 < c_1$. It is sufficient to show that, If $L(i_1, c_0, z_1) > 0$ and the minimum lot size is binding for at least one type in at least one city, then the following hold

   (a) $L_{i_1,c_1}(z_1) > 0,$ and $L_{i_1,c_1}(z_0) = 0$ with

   (b) $L_{i_0,c_1}(z_1) = 0,$ or $L_{i_0,c_1}(z_0) > 0$

   I prove this. a) $L_{i,c_1}(z_1) > 0$ and $L_{i,c_1}(z_0) = 0$ follows from the same argument in the proof of Statement 1. It remains to show b). Suppose, for contradiction, that $L_{i_0,c_1}(z_1) > 0$ and $L_{i_0,c_1}(z_0) = 0$; that is, high income types strictly outbid low income in the low regulation neighborhood of the productive city. Then $\theta_{i_0,c_1}(z_1) > \theta_{i_0,c_1}(z_0)$. However, this immediately implies that $\theta(i_0, c, z_1) > \theta(i_0, c, z_0)$ for every c because i_0 neighborhoods are unregulated and thus the bid rents have the same slopes across income types (Lemma 2). However, this also then implies $\theta(i_1, c, z_1) > \theta(i_1, c, z_0)$ for all c because Lemma 3 ensures that the slope of the bid rent across neighborhoods is nondecreasing in z. Hence, z_0 types are outbid from all neighborhoods and cities, which cannot be an equilibrium.

   □

3. The density of housing units is decreasing in i.

   Proof. It is sufficient to fix a city c and show for any i_0 and i_1 such that $i_0 < i_1$, $\sum_z L_{i_1,c}(z) \leq \sum_z L_{i_0,c}(z)$.[57] Since $\theta_{ic}(z)$ is decreasing in i (regulation decreases utility when binding, thus maximal rent willing to pay), we must have that the equilibrium rent curve $\theta_{ic} = \max_z \theta_{ic}(z)$ is also decreasing in i.

[56] Since we assume the minimum lot size is binding for at least one agent *and* not all agents observe the value of a minimal that exceeding income, Lemma 2 implies strict supermodularity.

[57] Recall that all neighborhoods have identical land mass.

I prove that this is only commesurate with a population that is decreasing in i. Lemma 3 ensures that there are (weakly) relatively more z_1 types in i_1, and, for both z types, housing expenditure per capita is weakly larger in i_1 than i_0 (strict if regulation is binding). Suppose for contradiction that the population is higher in i_1. This *must* imply that total housing expenditure in i_1 is strictly higher than i_0. This implies strictly higher rents in i_1 than i_0 because we assume they have the same production technology for housing, contradicting the fact that $\theta_{ic} = \max_z \theta_{ic}(z)$ is decreasing in i.

□

B.5 A simple model showing why income heterogeneity matters for aggregate productivity

In this section, I provide a simple model to explain why the presence of income heterogeneity can attenuate the productivity benefits of deregulating productive cities. This model does not depend on assumptions about endogenous amenities, though they may exacerbate this attenuation effect. To this end, consider a *productive city* indexed by c and an *outside option* indexed by o. The outside option (representing the entire US economy not including the city c) offers an exogenous wage ι_o per unit of skill and utility level V_0. The city c offers an exogenous wage $\iota_c > \iota_o$, and households pay endogenous rents P_c to live in the city. Preferences are Cobb-Douglas with parameter β. Assume total land supply in the city is fixed at $T(c)$, and there is an exogenous density of one housing unit per unit of land (i.e. the housing supply elasticity is zero).

First, consider the case where this city has no minimum lot size regulation. Housing market clearing in city c implies a unique price P_c such that

$$\beta \sum_{z \in Z} \frac{\exp[\frac{\iota_c}{P_c^{\beta}}]^{\theta}}{\exp[\frac{\iota_c}{P_c^{\beta}}]^{\theta} + [V_0]^{\theta}} zL(z) = P_c T(c) \tag{41}$$

where there are Gumbel idiosyncratic preference shocks with migration semi-elasticity θ and $L(z)$ is the mass of z households in the economy. The assumptions about Cobb-Douglas preferences mean that location choice fractions $\frac{\exp[\frac{\iota_c}{P_c^{\beta}}]^{\theta}}{\exp[\frac{\iota_c}{P_c^{\beta}}]^{\theta}+[V_0]^{\theta}}$ are independent of skill z. That is, all skills choose the productive city in the same proportion. This means that we can compare two models where the support of the skill distribution Z differs in an unregulated equilibrium. In fact, equilibria under any skill distributions $L(z)$ with the same population $\sum_z L(z)$ and total skill levels $\sum_z zL(z)$ will have equivalent prices P_c and aggregate productivity (population-weighted average wages) given in equilibrium

$$f_c \iota_c \sum_z zL(z) + (1 - f_c)\iota_0 \sum_z zL(z) \tag{42}$$

where f_c is the z-independent share of households who choose city c. This logic is synthesized by the lemma:

Lemma 4. *Suppose the above assumptions hold and there is no regulation. Then, with any arbitrary skill distribution with the same total income $\sum_z zL(z)$ results in the same equilibrium prices and relative skill levels across cities. This implies that aggregate productivity is invariant to any of these distributions.*

Having established this, I seek to show that a regulated equilibrium in a model with skill dispersion exhibits higher aggregate productivity than a model with only one skill level. With the logic above, this implies that complete deregulation leads to smaller productivity growth with skill heterogeneity relative to a model with homogeneity, which is the goal of the exercise. To this end, compare two distinct aggregate skill distributions indexed by L_d for $d \in \{0, 1\}$ with the following properties:

$$\sum_{z \in Z} z L_0(z) = \sum_{z \in Z} z L_1(z) \quad \text{and} \tag{43}$$

$$\sum_{z \in Z} L_0(z) = \sum_{z \in Z} L_1(z) \tag{44}$$

Next, assume $L_0(z)$ has a support over only one z value (no income heterogeneity), and that $L_1(z)$ has skill heterogeneity (a support of size 2 or more). The idea is to compare $L_1(z)$ with $L_2(z)$. Let l_c be the physical minimum lot size in the city. This means that there is an effective cap on the population of $\frac{T(c)}{l_c}$, with this cap being reached exactly if every household faces binding regulation in conjunction with housing market clearing. This leads to the following statement:

Lemma 5. *Suppose all assumptions above hold and l_c is sufficiently large. Then, aggregate productivity with under regulation l_c is higher in the economy with the skill distribution L_1 relative to L_0.*

The proof relies on constructing an equilibrium with l_c large enough so that regulation binds for all skill levels in the support of both L_0 and L_1.[58] In a high-regulation environment, this means that equilibrium populations are identical under these skill distributions and equal to this population cap. However, because of income sorting driven by regulation differences, there will be relatively more high skill people in city c relative to low-skill people under distribution L_1. With the assumed similarities between distributions above, numeraire good output in city c will be higher with skill dispersion relative to no skill dispersion, and thus the same holds for aggregate productivity.

Setting l_c large enough to bind all households is unlikely to be borne in the data. However, this is a sufficient condition and not a necessary one. The main takeaway is that, if regulation is highly binding, increased skill dispersion manifests through skill sorting into productive cities and does not cause changes to equilibrium city size (in terms of the number of households).

Below, I present a formal and complete argument.

Proposition 4 (Skill Heterogeneity Attenuates Productivity Gains from Deregulation)**.** *Consider the two-city model given above, with a productive city c offering wage $\iota_c > \iota_o$ and an outside option o offering wage ι_o. Households have Cobb-Douglas preferences ($\bar{A} = 0$) with expenditure share β, idiosyncratic Gumbel location preference shocks with elasticity θ, and skill $z \in Z$. City c has fixed land supply $T(c)$ with an exogenous housing density of one unit per unit of land, and a minimum lot size l_c.*

Let L_0 and L_1 be two aggregate skill distributions satisfying

$$\sum_{z \in Z} z\, L_0(z) \;=\; \sum_{z \in Z} z\, L_1(z) \;=:\; \bar{Y}, \tag{45}$$

$$\sum_{z \in Z} L_0(z) \;=\; \sum_{z \in Z} L_1(z) \;=:\; \bar{L}, \tag{46}$$

where L_0 has support on the single skill level $\bar{z} := \bar{Y}/\bar{L}$ (no skill heterogeneity), and L_1 has support on at least two distinct skill levels (skill heterogeneity). Suppose l_c is sufficiently large that the minimum lot size constraint binds in city c for every skill level in the support of $L_0 \cup L_1$.

[58] Such a finite regulation level must exist for finite support of a skill distribution as an equilibrium outcome.

Then aggregate productivity $\tilde{g}_d$, defined as the population-weighted average wage across both locations under distribution L_d ($d \in \{0, 1\}$), satisfies

$$\tilde{g}_1 \; > \; \tilde{g}_0.$$

That is, the regulated economy with skill heterogeneity achieves strictly higher aggregate productivity. By Lemma 4 (invariance under no regulation), this implies that complete deregulation yields strictly smaller productivity gains under L_1 than under L_0.

Proof. The proof proceeds in four steps.

Step 1: Indirect utility under binding regulation.

Appendix B.1 gives the indirect utility of a type-z household when the minimum lot size constraint binds in city c:

$$V_c(z) \;\propto\; \frac{l_c^{\beta}\left(\iota_c\, z - P_c\, l_c\right)^{1-\beta}}{z}, \tag{47}$$

where P_c is the equilibrium price of housing services and $P_c l_c$ is the cost of the minimum-lot bundle. The minimum lot size constraint binds for type z if and only if $P_c l_c > \beta\, \iota_c\, z$, which holds for all z in the support of L_d by assumption.

Step 2: Strict supermodularity of utility in z.

Differentiating the difference $V_c(z) - V_0$ with respect to z:

$$\begin{aligned}
\frac{\partial}{\partial z}(V_c(z) - V_0) &= V_c(z)\left[\frac{(1-\beta)\,\iota_c}{\iota_c\, z - P_c\, l_c} - \frac{1}{z}\right] \\
&= V_c(z)\frac{(1-\beta)\,\iota_c\, z \;-\; \left(\iota_c\, z - P_c\, l_c\right)}{z\left(\iota_c\, z - P_c\, l_c\right)} \\
&= V_c(z)\frac{P_c\, l_c \;-\; \beta\, \iota_c\, z}{z\left(\iota_c\, z - P_c\, l_c\right)}.
\end{aligned} \tag{48}$$

Since regulation binds, the numerator satisfies $P_c l_c - \beta \iota_c z > 0$, and the denominator satisfies $z(\iota_c z - P_c l_c) > 0$ (as $wz = \iota_c z > R = P_c l_c$ is required for the household to afford the minimum lot). Hence (48) is *strictly positive* for every z in the support, establishing that $V_c(z) - V_0$ is strictly increasing in z.

Step 3: High-skill sorting into city c under L_1.

With Gumbel idiosyncratic preference shocks and migration elasticity θ, the equilibrium fraction of type-z households residing in city c is

$$f_c(z) \;:=\; \frac{\exp\left[V_c(z)\right]^{\theta}}{\exp\left[V_c(z)\right]^{\theta} + \exp\left[V_0\right]^{\theta}}, \tag{49}$$

which is strictly increasing in the difference $V_c(z) - V_0$, and therefore *strictly increasing in z* by Step 2.

Let $L_c^d(z) := f_c^d(z)\, L_d(z)$ denote the equilibrium mass of type-z residents in city c under distribution L_d, where $f_c^d(z) := f_c(z;\, P_d(c))$ is evaluated at the equilibrium price $P_d(c)$. Housing market clearing (with land supply $T(c)$ and unit density) requires

$$\sum_{z \in Z} f_c^d(z)\, L_d(z) \;=\; \frac{T(c)}{l_c} \quad \text{for } d \in \{0, 1\}, \tag{50}$$

determining $P_d(c)$ under each distribution. The binding-regulation assumption ensures that $f_c^d(z)$ is strictly increasing in z at both equilibrium prices.

Define $\Sigma_d := \sum_z z\, f_c^d(z)\, L_d(z)$, the total skill allocated to city c under L_d.

Under L_0: Since all mass $\bar{L}$ is concentrated at $\bar{z} = \bar{Y}/\bar{L}$, market clearing gives $f_c^0(\bar{z})\,\bar{L} = T(c)/l_c$, so

$$\Sigma_0 \;=\; \bar{z} \;\cdot\; \frac{T(c)}{l_c}. \tag{51}$$

Under L_1: Decompose Σ_1 using the identity $\sum_z z\, h(z)\, \mu(z) = \mathrm{Cov}_\mu(z,\, h(z)) + E_\mu[z] \cdot E_\mu[h(z)]$ with $\mu = L_1/\bar{L}$ and $h(z) = f_c^1(z)$:

$$\Sigma_1 \;=\; \bar{L}\; \mathrm{Cov}_{L_1/\bar{L}}\big(z,\; f_c^1(z)\big) \;+\; \bar{z} \;\cdot\; \frac{T(c)}{l_c}, \tag{52}$$

where we used $E_{L_1/\bar{L}}[z] = \bar{z}$ and $E_{L_1/\bar{L}}[f_c^1(z)] = T(c)/(l_c\bar{L})$ from (50).

Since $z \mapsto f_c^1(z)$ is *strictly* increasing and L_1 places positive mass on *at least two distinct* values of z, the covariance in (52) is strictly positive:

$$\mathrm{Cov}_{L_1/\bar{L}}\big(z,\; f_c^1(z)\big) \;>\; 0.$$

Combining with (51) and (52):

$$\Sigma_1 \;>\; \Sigma_0. \tag{53}$$

Step 4: Aggregate productivity comparison.

Aggregate productivity under L_d is the population-weighted average wage:

$$\tilde{g}_d \;=\; \frac{\iota_c\, \Sigma_d \;+\; \iota_o\, \big(\bar{Y} - \Sigma_d\big)}{\bar{L}} \;=\; \frac{\iota_o\, \bar{Y} \;+\; \big(\iota_c - \iota_o\big)\, \Sigma_d}{\bar{L}},$$

where we used $\sum_z L_o^d(z) = \bar{Y} - \Sigma_d$ from (45). Since $\iota_c > \iota_o$ and $\Sigma_1 > \Sigma_0$ by (53), it follows immediately that

$$\tilde{g}_1 \;>\; \tilde{g}_0.$$

which is what we wanted to show.

□

B.6 Proof of Proposition 2

Setup and notation. Throughout this proof the model is a single closed city with two neighborhoods $i \in \{0,1\}$ and three skill types $z_l < z_m < z_h$. I write $s_i(z) := R_i/(wz)$ for the regulation-to-income ratio of a type-z household in neighborhood i (so regulation binds iff $s_i(z) > \beta$), and $\text{Inc}_i := \sum_z wzL_i(z)/\sum_z L_i(z)$ for average income. Let $Y_i := w\sum_z zL_i(z)$ denote total income and $N_i := \sum_z L_i(z)$ total population in i. Define the following objects for any neighborhood i and skill z:

- $o_i(z) := wzL_i(z)/Y_i$ — share of i's total income earned by type-z households;
- $f_i(z) := L_i(z)/N_i$ — fraction of i's residents who are type z;
- $\tilde{f}_i(z) := L_i(z)/\bar{L}(z)$ — fraction of all type-z households who reside in i;
- $\psi_i(z) := o_i(z) - f_i(z)$ — deviation of income share from population share.

Note that $\sum_z o_i(z) = \sum_z f_i(z) = 1$, so $\sum_z \psi_i(z) = 0$. Also, $\psi_i(z)$ has the same sign as $wz - \text{Inc}_i$: types earning above (below) average income in i have $\psi_i(z) > 0$ (< 0). Recall the distortion factor $D_i(z) := \left[s_i(z)/\beta\right]^{\beta}\left[(1-s_i(z))/(1-\beta)\right]^{1-\beta}$ from Equation (5), which equals 1 when $s_i(z) = \beta$ (just-binding) and is strictly decreasing in $s_i(z)$ for $s_i(z) > \beta$.

Indirect utility is

$$V_i(z) \;=\; \frac{w}{P_i^{\beta}} D_i(z) \;+\; \Omega(z)\log \text{Inc}_i \;+\; \log \nu_i(z), \tag{54}$$

where $D_i(z) = 1$ whenever $s_i(z) \le \beta$ (regulation is not binding).

Equilibrium and linearization. I assume $\Omega(z)$ is increasing in z and sufficiently small to guarantee a unique interior equilibrium.[59] An equilibrium with within-city mobility elasticity ρ satisfies

$$V_i(z) - \frac{1}{\rho}\log L_i(z) \le \xi(z) \tag{55}$$

for some $\xi(z)$ that depends only on z, for every i and z; where the condition holds with equality whenever $L_i(z) > 0$. This follows directly from Equation (7) in the single-city, single-zone case, where $\log \boldsymbol{W}(z)$ plays the role of $\xi(z)$. Totally differentiating (55) and taking differences across neighborhoods yields the *population adjustment formula* for neighborhood $i = 1$:

$$\frac{\partial L_1(z)}{L_1(z)} \;=\; \rho\big[1 - \tilde{f}_1(z)\big]\Big[\Delta C(z) \;+\; \Omega(z)\,\Delta \text{Inc} \;+\; \Delta\nu(z)\Big], \tag{56}$$

where $\Delta C(z) := \partial V_1^C(z) - \partial V_0^C(z)$ is the differential change in consumption value (the first term of (54)), $\Delta\text{Inc} := \partial\text{Inc}_1/\text{Inc}_1 - \partial\text{Inc}_0/\text{Inc}_0$, and $\Delta\nu(z) := \partial\nu_1(z)/\nu_1(z) - \partial\nu_0(z)/\nu_0(z)$.

With Cobb-Douglas preferences and housing supply elasticity ϵ, housing prices respond to total income as $\partial P_i/P_i = (1+\epsilon)^{-1}\partial Y_i/Y_i$. Using $\partial Y_0 = -\partial Y_1$ (fixed city income), the consumption value difference becomes

$$\Delta C(z) \;=\; \tilde{c}\,\Delta_1 Y \;+\; \frac{w}{P_1^{\beta}}\,\partial_{R_1} D_1(z), \tag{57}$$

[59] For Ω sufficiently small, the equilibrium is stable because the disutility of a marginal rent increase from population growth dominates any endogenous amenity response. Uniqueness allows well-defined comparative statics across parameterizations. Uniqueness can be proven using standard methods (Allen, Arkolakis, and Li, 2024).

where $\Delta_1 Y := \partial Y_1/Y_1$, $\tilde{c} := -\frac{\beta}{1+\epsilon}\big[w/P_1^\beta + (w/P_0^\beta)(Y_1/Y_0)\big] < 0$, and $\partial_{R_1} D_1(z)$ is the derivative of $D_1(z)$ with respect to R_1 (which is zero when regulation does not bind for type z in neighborhood 1, and strictly negative when it does).[60] Finally, average income in neighborhood i satisfies

$$\frac{\partial \text{Inc}_i}{\text{Inc}_i} \;=\; \sum_z \psi_i(z)\,\frac{\partial L_i(z)}{L_i(z)}, \tag{58}$$

which follows from the definition of $\text{Inc}_i = Y_i/N_i$.

1. **Inclusionary Zoning**: In the $t = 1$ economy, a small increase in regulation in $i = 1$ increases average income in all locations.

   Proof. At $t = 1$, fundamental amenities satisfy $\nu_1(z_l) = 0$ and $\nu_0(z_h) = 0$, so z_l types reside only in $i = 0$ and z_h types reside only in $i = 1$. Only z_m types inhabit both neighborhoods. So, the spatial equilibrium equation above only holds with equality for the respective neighborhoods these types occupy. I show that a marginal increase in R_1 that is *just binding* for z_m types — i.e., beginning from $s_1(z_m) = \beta$ — causes z_m households to leave neighborhood $i = 1$. Because z_m types earn below-average income in $i = 1$ (which hosts z_h types) and above-average income in $i = 0$ (which hosts z_l types), their departure raises average income in both neighborhoods.

   Since $\partial L_i(z_h) = \partial L_i(z_l) = 0$ (these types are pinned to their respective neighborhoods), $\Delta_1 Y = o_1(z_m)\Delta L_m$ and $\Delta\text{Inc} = \tilde{I}_m\,\Delta L_m$, where I define $\Delta L_m := \partial L_1(z_m)/L_1(z_m)$ and

   $$\tilde{I}_m \;:=\; \psi_1(z_m) \;+\; \psi_0(z_m)\,\frac{L_1(z_m)}{L_0(z_m)}.$$

   Substituting into (56) and using (57) (with $\partial_{R_1} D_1(z_m) = 0$ initially since $s_1(z_m) = \beta$ means the constraint is just binding, and including the first-order response to R_1) yields the following linear equation in ΔL_m:

   $$\big[1 - \rho\big(1 - \tilde{f}_1(z_m)\big)\big(\tilde{c}\,o_1(z_m) + \Omega(z_m)\tilde{I}_m\big)\big]\Delta L_m \;=\; \rho\big(1 - \tilde{f}_1(z_m)\big)\,\frac{w}{P_1^\beta}\,\partial_{R_1} D_1(z_m). \tag{59}$$

   The right-hand side is strictly negative because $\partial_{R_1} D_1(z_m) < 0$. The coefficient on the left is strictly positive for $\Omega(z_m)$ sufficiently small: since $\tilde{c} < 0$, the term $-\rho(1 - \tilde{f}_1(z_m))\tilde{c}\,o_1(z_m) > 0$ already exceeds 1 for moderate ρ, and adding the $\Omega(z_m)$ term does not reverse the sign for small Ω. Hence $\Delta L_m < 0$.

   Finally, $\partial\text{Inc}_i/\text{Inc}_i = \psi_i(z_m)\Delta L_m$ from (58). Since $\psi_1(z_m) < 0$ (middle-skill earns below average in $i = 1$) and $\psi_0(z_m) > 0$ (above average in $i = 0$), both income changes $\psi_i(z_m)\Delta L_m$ are positive. Income rises in both neighborhoods. □

2. **Exclusionary Zoning**: In the $t = 0$ economy, increasing regulation in any neighborhood(s) does not increase average income in all locations. Instead, average income across neighborhoods weighted by the population of z_h types increases.

   Proof. At $t = 0$, $\nu_i(z) = 1$ for all i and z, so the equilibrium is symmetric: $L_i(z) = \bar{L}(z)/2$, $P_0 = P_1 =: P$, and $\tilde{f}_i(z) = 1/2$ for all i, z. I consider a marginal increase in R_1 starting from the point where $s_1(z_l) = \beta$ (just-binding for z_l).

[60] Since $D_1(z)$ attains its maximum of 1 at $s_1(z) = \beta$ and is strictly concave for $s_1(z) > \beta$, any marginal increase in R_1 that pushes $s_1(z)$ above β decreases $D_1(z)$, so $\partial_{R_1} D_1(z) < 0$.

Step 1: Income divergence ($\Delta\mathrm{Inc} > 0$). At the symmetric equilibrium $o_i(z) = o(z)$, $f_i(z) = f(z)$, so $\psi_i(z) = \psi(z) := o(z) - f(z)$ and $\sum_z \psi(z) = 0$. Using (56), (57), and (58), the definition of $\Delta\mathrm{Inc}$ reduces to

$$\Delta\mathrm{Inc} \;=\; \rho \sum_z \psi(z)\Big[\tilde{c}\,\Delta_1 Y + \Omega(z)\,\Delta\mathrm{Inc} + \frac{w}{P^\beta}\,\partial_{R_1} D_1(z_l)\,\mathbf{1}_{z=z_l}\Big].$$

The $\tilde{c}\,\Delta_1 Y$ term vanishes upon summation because $\sum_z \psi(z) = 0$. Solving for $\Delta\mathrm{Inc}$:

$$\Delta\mathrm{Inc} \;=\; \frac{\rho\,\psi(z_l)\,\frac{w}{P^\beta}\,\partial_{R_1} D_1(z_l)}{1 - \rho\sum_z \psi(z)\Omega(z)}. \tag{60}$$

The denominator is strictly positive for Ω small. The numerator is positive because $\psi(z_l) = o(z_l) - f(z_l) < 0$ (low-skill types earn a smaller share of income than their population share) and $\partial_{R_1} D_1(z_l) < 0$. Hence $\Delta\mathrm{Inc} > 0$: neighborhood incomes diverge in response to regulation.

Step 2: Inc_1 *rises and* Inc_0 *falls.* From (58) and (56) at the symmetric equilibrium:

$$\frac{\partial\mathrm{Inc}_1}{\mathrm{Inc}_1} \;=\; \frac{\rho}{2}\sum_z \psi(z)\Big[\tilde{c}\,\Delta_1 Y + \Omega(z)\,\Delta\mathrm{Inc} + \frac{w}{P^\beta}\,\partial_{R_1} D_1(z_l)\,\mathbf{1}_{z=z_l}\Big].$$

Using $\sum_z \psi(z) = 0$ to eliminate the $\Delta_1 Y$ term and substituting (60):

$$\frac{\partial\mathrm{Inc}_1}{\mathrm{Inc}_1} \;=\; \frac{\rho}{2}\,\psi(z_l)\,\frac{w}{P^\beta}\,\partial_{R_1} D_1(z_l)\cdot\frac{1}{1 - \rho\sum_z \psi(z)\Omega(z)} \;>\; 0.$$

The sign holds because $\psi(z_l) < 0$, $\partial_{R_1} D_1(z_l) < 0$, and the denominator is positive. By symmetry, $\partial\mathrm{Inc}_0/\mathrm{Inc}_0 = -\partial\mathrm{Inc}_1/\mathrm{Inc}_1 < 0$. Neighborhood 1 becomes richer and neighborhood 0 poorer. Since z_h types are driven toward the richer regulated neighborhood, the average income weighted by z_h-type populations strictly increases. No such unambiguous statement can be made for the unweighted average. □

B.7 Social planning problem and proof of Proposition 3

I start by restating the problem for convenience. The **Social planner's problem** for a set of household weights $\{\alpha(z)\}_{z\in Z}$ and absentee landowner weights α^L is defined as the choice of numeraire consumption allocations $g_i^c(z)$ and g_i^L for households and landowners respectively, housing consumption allocations $A_i(z)$, and total capital inputs into housing production g_i^A, and welfare values $\boldsymbol{W}(z)$ to solve

$$\max \sum_{z\in \boldsymbol{Z}} \alpha(z) \log \boldsymbol{W}(z) + \alpha^L \Pi \tag{61}$$

where $\Pi := \sum_{i\in N} g_i^L$ is the total numeraire consumption paid to landowners. Maximization is done subject to the following resource, free mobility and population balancing constraints:

$$\sum_{i\in\cup_{c\in\mathbb{C}}N(c)} \Big[\sum_{z\in Z} g_i^c(z)L_i(z)\Big] + g_i^A + g_i^L = \underbrace{\sum_{c\in C}\left[\iota(c)\sum_{i\in N(c),z\in Z} zL_i(z)\right]}_{\text{Total production of numeraire}} \tag{62}$$

$$\forall i,\ \sum_{z\in Z} A_i(z)L_i(z) = \underbrace{\lambda_i^{\frac{1}{1+\epsilon_i}} g_i^{A\,\frac{\epsilon_i}{1+\epsilon_i}} T_i^{\frac{1}{1+\epsilon_i}}}_{\text{Local production of housing services}} \tag{63}$$

$$\forall i,z,\ \ V_i(z) - \frac{1}{\theta}\log\left[\frac{L_i(z)}{L(z)}\right] = \log \boldsymbol{W}(z) \tag{64}$$

$$\forall z,\ \sum_{i\in\cup_{c\in C}N(c)} L_i(z) = L(z) \tag{65}$$

where

$$V_i(z) = (A_i(z)-\bar{A})^{\beta}\big[g(i,z)\big]^{1-\beta} + \Omega(z)\log \text{Inc}_i + \log \nu_i(z)$$

In what follows, let Λ^C be the lagrange multiplier for the constraint Equation (62), $\Lambda^A(i)$ for (63), $\Lambda^{FM}(i,z)$ for (64) and $\Lambda^L(z)$ for (65). In addition, let $V_i^g(z)$ and $V_i^A(z)$ be the marginal utility of numeraire and housing consumption, respectively, for a household of skill z in neighborhood i.

For use with derivations, I list important first order conditions associated with this problem:

1. FOC w.r.t $g_i^c(z)$ for fixed z, i

$$-\underbrace{\Lambda^{FM}(i,z)V_i^g(z)}_{\text{Weighted marginal utility of numeraire}} = \Lambda^C L_i(z) \tag{66}$$

2. FOC w.r.t. $L_i(z)$

$$\sum_{z'\in Z} -\Lambda^{FM}(i,z')\Omega(z')\frac{\partial \log \text{Inc}_i}{\partial L_i(z)} + \Lambda^{FM}(i,z)\frac{1}{\theta}\frac{1}{L_i(z)} + \Lambda^C\iota_c z = \Lambda^L(z) + \Lambda^C g_i^c(z) + \Lambda^A(i)A_i(z) \tag{67}$$

3. FOC w.r.t $A_i(z)$ for fixed z, i

$$-\underbrace{\Lambda^{FM}(i,z)V_i^A(z)}_{\text{Weighted marginal utility of housing services}} = \Lambda^A(i)L_i(z) \tag{68}$$

4. FOC w.r.t. g_i^A

$$\underbrace{\tilde{\lambda}(i)\frac{\epsilon_i}{1+\epsilon_i}\left[\frac{g_i^A}{T_i}\right]^{-\frac{1}{1+\epsilon_i}}}_{\text{Marginal product of capital in } i\text{'s housing sector}} = \frac{\Lambda^C}{\Lambda^A(i)} \tag{69}$$

I emphasized two conditions that partially characterize the solution to this problem.

Condition 1 The first condition concerns how a social planner trades off numeraire and housing consumption in each location:

$$\underbrace{\frac{V_i^g(z)}{V_i^A(z)}}_{\text{-MRS of housing for numeraire}} = \underbrace{\frac{\epsilon_i}{1+\epsilon_i}\left[\frac{g_i^A}{\lambda_i T_i}\right]^{-\frac{1}{1+\epsilon_i}}}_{\text{Marginal product of capital in } i\text{'s housing sector}} \tag{70}$$

This is the standard MRS = MRTS condition that characterizes general equilibrium in the neoclassical model. This cannot be achieved with minimum lot size regulation because it necessarily distorts housing consumption for a given set of equilibrium prices, and more so for low income households.

I use the first order conditions above to derive this equation. Rearranging and substituting Equation (68) into Equation (66) yields the following expression

$$\frac{V_i^g(z)}{V_i^A(z)} = \frac{\Lambda^C}{\Lambda^A(i)} \tag{71}$$

and substituting Equation (69) into this yields the desired Condition 1.

Condition 2 The second concerns an efficient spatial allocation of households:

$$\sum_{z'\in Z}\Omega(z')\frac{\partial \log \text{Inc}_i}{\partial L_i(z)}V_i^g(z')^{-1} + \iota(c)z - \frac{1}{\theta}V_i^g(z)^{-1} = \Xi(z) + g_i^c(z) + \frac{V_i^A(z)}{V_i^g(z)}A_i(z) \tag{72}$$

for some constants $\Xi(z)$ I derive below. This equation ensures that the positive and negative externalities from amenity spillovers are balanced with aggregate productivity benefits of packing households in productive cities as well as the costs to the social planner to induce this movement.

I use the first order conditions to derive this. Substitute the expression for $\Lambda^{FM}(i,z)$ from Equation (66) into Equation (67), as well as the expression for $\Lambda^A(i)$ from (71) to get

$$\begin{aligned}\Lambda^C\sum_{z'\in Z}\Omega(z')\frac{\partial \log \text{Inc}_i}{\partial L_i(z)}V_i^g(z')^{-1}L_i(z') - \Lambda^C\frac{1}{\theta}V_i^g(z)^{-1} + \Lambda^C\iota_c z = \\ \Lambda^L(z) + \Lambda^C g_i^c(z) + \Lambda^C\frac{V_i^A(z)}{V_i^g(z)}A_i(z)\end{aligned} \tag{73}$$

Divide both sides of the equality by Λ^C to obtain

$$\sum_{z' \in Z} \Omega(z') \frac{\partial \log \text{Inc}_i}{\partial L_i(z)} V_i^g(z')^{-1} L_i(z') - \frac{1}{\theta} V_i^g(z)^{-1} + \iota_c z = \frac{\Lambda^L(z)}{\Lambda^C} + g_i^c(z) + \frac{V_i^A(z)}{V_i^g(z)} A_i(z) \tag{74}$$

and define $\Xi(z) := \frac{\Lambda^L(z)}{\Lambda^C}$, which completes the derivation.

B.8 Extending the model to allow for durable housing

In this section of the Appendix, I detail a major extension to the model to allow for durable housing and new development that occurs only on undeveloped land.

Site selection into development Each neighborhood and zone consists of a continuum of *sites* on the unit interval $[0,1]$, each site x containing a mass T_{io} of land. Developers pay a fixed cost $F_i(x)$ per unit of land to develop any given site x, and these fixed costs are heterogeneous across sites (Baum-Snow and Han, 2024). Sites are ordered in terms of increasing fixed costs, and are parameterized by

$$F_{io}(x) = \mu_i \Big[-\log x \Big]^{-\frac{1}{\sigma}} \tag{75}$$

where μ_i governs the mean fixed cost over sites, and σ is inversely related to the variance of fixed costs over sites. Sites are developed if the profits of doing so are positive: when the value of developed land $\frac{1}{1+\epsilon_i}\lambda_i P_{io}^{1+\epsilon_i}$ exceeds the fixed cost. This means that the mass of developed land X_{io} is given by

$$X_{io} = \exp\left(- \frac{\lambda_i^{-\sigma} P_{io}^{-\sigma(1+\epsilon_i)}}{(1+\epsilon_i)^{-\sigma}\mu_i^{-\sigma}} \right) \times T_{io} \tag{76}$$

There are three important quantities that determine the marginal price elasticity of land supply $\frac{\partial \log X_{io}}{\partial \log P_{io}}$. The first is σ: the less heterogeneous the site preparation costs, the more sites there are at the margin of being developed in response to land value growth. The second is the elasticity of (developed) land value with respect to the price of housing services, $1+\epsilon_i$. The third, and most important, is the fraction of land developed itself: if all land in a neighborhood is developed, there is no land left to select into development.[61]

Durable housing and short-run equilibria I define a short run equilibrium of the model as one where 1) deregulation only occurs on undeveloped sites and 2) total floorspace on developed sites is fixed at pre-equilibrium values (this can be achieved by imposing $\epsilon_i = 0$ on already-developed land). While this short run equilibrium is not implemented in the draft of this paper, this extension can be done in a relatively straightforward way.

[61] Baum-Snow and Han (2024) estimate that the lion's share of the variation in the housing supply elasticity within cities is driven by variation in the fraction of land developed. In this model, stringently regulated (e.g. high R_i) neighborhoods can cope with positive housing demand shocks by building low density housing on vacant land.

C Appendix: Calibration

C.1 The complete calibration algorithm

In this section, I detail the full algorithm that calibrates the model. I omit details of the calculations of hedonic prices, residualized wages, and other parameters because they are straightforward and explained with detail in Section 4.

Algorithm 2 Calibration

(Consumption values and supply parameters)

For all block groups do

If Partially Regulated[62]

Uniquely unregulated prices P_{iU} given all other parameters, including regulated prices P_{iR} and observed value of a minimal lot R_i to set the supply equal to the demand for housing in every zone such that the fraction of households choosing regulated structures equals ACS data. Defines consumption value $C_i(z)$ and aggregate housing expenditure for each zone.

Else

Directly calculate $C_i(z)$ and housing expenditure given some housing price P_i and the value of a minimal lot R_i.

If Fully or Partially Regulated

Calculate λ_i to uniquely solve $\lambda_i P_{iR}^{1+\epsilon_i} l_i = R_i$ for all partially or fully regulated block groups.

Choose $T_{i\boldsymbol{o}}$ for each $\boldsymbol{o}$ to clear housing markets by zone. If Fully Regulated, $T_{iU} = 0$.

Else

Set T_{iU} to some observed land mass of the block group and $T_{iR} = 0$.

Calculate λ_i to clear housing markets.

(Amenities)

For all block groups do

1. Choose amenities to uniquely solve (29) given $C_i(z)$ and populations $L_i(z)$ (up to a normalization) for each type.

---

[62] A neighborhood is partially regulated if a block group is 1) assigned regulation previously, 2) has a positive fraction of housing units in regulated structures observed in the ACS, but this fraction is less than one. Fully regulated neighborhoods are identical to partially regulated neighborhoods except that the fraction of housing units in regulated structures is 1. Unregulated neighborhoods are all neighborhoods not satisfying one of the definitions above.

[62] In an unregulated neighborhood, T_{iR} and λ_i are not separately identifiable parameters, so this allocation has no impact on counterfactual outcomes.

C.2 Calibration when households differ on education

Extending the model to allow for households to differ by education requires some additional calibration methodology, which I detail here.

Wages by Education Since households of differing education levels are not perfect substitutes in production (Equation 25), they require a separate measure of wages. I regress log hourly wages in a set of occupation, sex, race, ancestry, year, quadratic in age and years of education, including MSA fixed effects separately for both college and non-college educated workers. These MSA fixed effects for both education groups form city and education specific wages per unit of effective labour, each normalized to be one in an average city. Productivity $\iota(c, s)$ by education s is then derived to uniquely rationalize employment by education in each city and this city wage under the technology in Equation (25) after specifying a value of $\sigma = 1.3$ (Card, 2009).

Local education-skill distributions The ACS block group tabulations do not report the joint distribution of households by skill z and education status. I impute the joint distribution of neighborhood income and education by calculating the share of college workers by skill type at the city level using the ACS microdata sample. I assume that this share by type applies uniformly over all neighborhoods within a given city.

All other calibration methodology when households differ on education is precisely the same as the baseline model where they do not (a type is redefined to be an income-education pair). I also assume that the support of the type distribution and spending shares on housing are the same across education levels within household skill types.

C.3 Summary Statistics for key calibrated parameters

Table 6: Table of calibrated parameters.

Parameter	Description	Value	Source or Target
Housing Demand			
P_{iR}	Housing prices (regulated zone)		Hedonic regression using CoreLogic transactions (Section 4.1)
P_{iU}	Housing prices (unregulated zone)		ACS fraction of households living in regulated structures
$(\beta, \bar{A})$	Housing preference parameters	(0.09, 3850)	ACS aggregate spending on housing services by income; max error of 3 percentage points across 7 income groups
Neighborhood Demand			
θ	Across-city migration elasticity	4	Hornbeck and Moretti (2018)
ρ	Within-city migration elasticity	5.5	Targets Song (2025) border discontinuity
$b_i(z)$	Amenity values by skill		ACS block-group income distribution
w_c	City productivity		Mincer regression with city fixed effects on ACS household data (Section 4.1)
Housing Supply			
R_i	Value of a minimal lot		Data (empirical facts, Section 3.2)
λ_i	Construction productivity		Clear housing markets in each zone (Section 4.3)
T_{io}	Land available for construction by zone		Clear housing markets in each zone (Section 4.3)
ϵ_i	Housing supply elasticity		Baum-Snow and Han (2024)

Table 7: Summary Statistics for Key Calibrated Parameters,
disaggregated by SuperStar city sample

Superstar?	Aggregate			No			Yes		
Variable	N	Mean	Sd	N	Mean	Sd	N	Mean	Sd
Housing price (Regulated zone)	189012	1	1	81801	0.52	0.27	107211	1.4	1.2
Housing price (Unregulated zone)	189012	1.5	1.9	81801	0.82	0.77	107211	2.1	2.2
ln Amenity value (0-25k)	188309	-0.07	0.38	81508	-0.27	0.28	106801	0.082	0.37
ln Amenity value (25-50k)	188309	-0.036	0.27	81508	-0.17	0.19	106801	0.066	0.28
ln Amenity value (50-75k)	188511	-0.026	0.23	81626	-0.13	0.18	106885	0.058	0.24
ln Amenity value (75-100k)	188511	-0.022	0.22	81626	-0.12	0.19	106885	0.056	0.22
ln Amenity value (100-150k)	184337	-0.0021	0.24	79601	-0.12	0.21	104736	0.085	0.23
ln Amenity value (150-200k)	184337	-0.0042	0.25	79601	-0.073	0.24	104736	0.048	0.25
ln Amenity value (200k+)	184337	-0.018	0.32	79601	-0.19	0.29	104736	0.11	0.27
Consumption value (0-25k)	189012	0.56	0.31	81801	0.64	0.26	107211	0.5	0.34
Consumption value (25-50k)	189012	0.8	0.2	81801	0.82	0.13	107211	0.79	0.24
Consumption value (50-75k)	189012	0.89	0.14	81801	0.87	0.08	107211	0.9	0.17
Consumption value (75-100k)	189012	0.92	0.11	81801	0.89	0.061	107211	0.95	0.13
Consumption value (100-150k)	189012	0.95	0.088	81801	0.91	0.051	107211	0.98	0.097
Consumption value (150-200k)	189012	0.96	0.08	81801	0.91	0.046	107211	1	0.079
Consumption value (200k+)	189012	0.98	0.076	81801	0.92	0.043	107211	1	0.067

Summary statistics for calibrated (and estimated) housing prices, consumption values, and amenity values by skill type. These statistics are broken down by the SuperStar and non-SuperStar sample used to construct the Facts in section 3 (top quartile of wages per unit of skill). They are also reported for the entire sample. Amenity values $b_i(z)$ are normalized to have a mean of 1 (in levels, not logs). Differences in observation counts for amenity values are in block groups with no counts of the corresponding type.

D Appendix: Estimating $\Omega(z)$

D.1 Econometric model

In this section, I specify a data generating process that justifies the use of the donut instrument to estimate each $\Omega(z)$. Consider a closed city with N neighborhoods; these constitute the observed data. The structural equation for amenity value is, as in Equation (30),

$$\log b_i(z) = \Omega(z) \log \text{Inc}_i + \beta_1(z) S\left[d_{1i}\right] + \log \nu_i(z)$$

where neighborhood income is determined by the neighborhood choice equation (**??**) subject to prices clearing housing markets. I specify that fundamental amenities $\log \nu_i(z)$ are a sum of two components:

$$\log \nu_i(z) = \underbrace{\Upsilon_i(z)}_{\text{Unobserved demand factors}} + \underbrace{\xi_i(z)}_{\text{Noise}}$$

where $\xi_i(z)$ is random noise that may be spatially correlated. $\Upsilon_i(z)$ contain unobserved demand factors that are not themselves slopes but can be correlated with slopes, i.e. weather, climate, views and ecetera. Note that if there were no unobserved demand factors[63], an instrument is still required because Inc_i will be endogenously determined by the shocks $\xi_i(z)$.

I allow for slopes in the buffer region to cause unobserved local demand factors (or be correlated through some underlying latent variable). Structurally,

$$\Upsilon_i(z) = \tilde{\gamma}(z) S\left[d_{1i}\right] + \sigma_i(z) \tag{77}$$

for some $\sigma_i(z)$ capturing other demand unobservables orthogonal to slopes in the buffer region d_1. This term may have arbitrary spatial correlation. The identification assumption is that slopes outside the buffer region cannot be correlated with any unobserved demand factors, or

$$S\left[d_{2i}\right] \perp \sigma_i(z) \tag{78}$$

which in turn requires that $S\left[d_{2i}\right]$ be "excluded" as a demand factor in Equation (77). That is, slopes outside of the buffer cannot provide any additional information to predict unobserved demand factors beyond its ability to predict slopes within the buffer region d_1. Note that this allows for the possibility that slopes have an arbitrary spatial correlation. This identification assumption is equivalent to saying

$$S\left[d_{2i}\right] \perp \log \tilde{\nu}_i(z) \mid S\left[d_{1i}\right]$$

which is the assumption reported in the paper.

[63] That is, $\Upsilon_i(z) = 0$ with probability 1 for every i and z.

D.2 First Stage Regression

Table 8: First Stage Regressions pooled across income types.

VARIABLES	(1) ln Income	(2) ln Income	(3) ln Income	(4) ln Income
Slope Donut (10-16km)	-0.0240***	-0.0210***	-0.0207***	-0.0193***
	(0.0053)	(0.0048)	(0.0048)	(0.0048)
Slope Control	-0.0141**	-0.0076	-0.0068	-0.0032
	(0.0061)	(0.0052)	(0.0054)	(0.0053)
Local Slope Control	0.0519***	0.0463***	0.0464***	0.0401***
	(0.0025)	(0.0021)	(0.0022)	(0.0021)
Observations	194,203	191,668	191,371	191,371
Specification	FirstStage	FirstStage	FirstStage	FirstStage
Donut	10-16km	10-16km	10-16km	10-16km
Base Controls	No	Yes	Yes	Yes
Amen/Topo Controls	No	No	Yes	Yes
Density Control	No	No	No	Yes

Standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

All specifications include MSA fixed effects and standard errors are clustered using a 35km Bartlett kernel. "Slope Donut" is the instrumental variable – average slopes within block groups between a set of distance buffers. "Local Slope Control" is the average slope within the block group. $\ln$ Income is instrumented with the average slopes of block groups that have centroids within buffer d_1 and d_2. "Base Controls" include travel time, building age, public transport and bus shares in commuting and CBD distance. "Amen/Topo" controls include various amenities (density of coffee shops, parks, restaurants) and various topographic features (cover of different types of forest such as deciduous or evergreen, wetlands, perennial snow cover). "Density Control" is the within-MSA density ranking of the block group.

D.3 Robustness

Various controls Table 9 compares estimates pooled over each income type with various controls. Column (1) is the OLS estimate controlling only for block group land mass. Column (2) introduces the IV estimate with the same set of controls as (1), showing downward bias. Column (3) adds an additional a set of ”base” controls: commuting time, median building age, share of public and bus transport in commuting, and CBD distance rankings. Column (4) adds additional topographic characteristics (such as forest cover) and observed amenities (density of coffee shops, bars, parks) from NaNDA with little changes to the estimate. This is the pooled version of the baseline IV estimates in Table 2. Lastly, Column (5) additionally controls for housing unit density.

Table 9: Pooled IV Specification with various controls

VARIABLES	(1) ln Amenity	(2) ln Amenity	(3) ln Amenity	(4) ln Amenity	(5) ln Amenity
ln Income	0.1379***	0.3020***	0.2857***	0.2758***	0.3060***
	(0.0043)	(0.0835)	(0.0771)	(0.0720)	(0.0763)
Slope Control	0.0019	0.0065**	0.0089***	0.0089***	0.0062***
	(0.0030)	(0.0027)	(0.0021)	(0.0020)	(0.0018)
Local Slope Control	0.0031***	-0.0058	-0.0053	-0.0030	0.0009
	(0.0007)	(0.0045)	(0.0037)	(0.0035)	(0.0032)
Observations	188,361	188,361	185,979	185,691	185,691
Specification	OLS	IV	IV	IV	IV
Donut	.	10-16km	10-16km	10-16km	10-16km
Base Controls	No	No	Yes	Yes	Yes
Amen/Topo Controls	No	No	No	Yes	Yes
Density Control	No	No	No	No	Yes
FStat Bart c 35 km	.	22.1	19.8	18.9	16.4

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Pooled IV Specification with various controls, using the average amenity across low, medium and high income groups. ”Slope Control” is the average slope within the block group plus a buffer with length equal to d_1. All specifications include MSA fixed effects and standard errors are clustered using a 35km Bartlett kernel. ”Local Slope Control” is the average slope within the block group. $\ln$ Income is instrumented with the average slopes of block groups that have centroids within buffer d_1 and d_2. ”Base Controls” include travel time, building age, public transport and bus shares in commuting and CBD distance. ”Amen/Topo” controls include various amenities (density of coffee shops, parks, restaurants) and various topographic features (cover of different types of forest such as deciduous or evergreen, wetlands, perennial snow cover). ”Density Control” is the within-MSA density ranking of the block group.

Different donuts To test how estimates vary by donut size, I compare various pooled estimates in Table 10. I find that donut radii reaching between 2-10km produce either small point estimates or those that do not differ from OLS, consistent with the idea that slopes in these donut regions do not satisfy the exclusion restriction. Column 1 corresponds to a very weak instrument, so these point estimates accompany very large standard errors. Point estimates rise sharply for donuts with larger control buffers. Column 4 reports the baseline donut estimates. Column 5 reports estimates for a larger 16-25km buffer, and finds larger point estimates. Counterfactual results quantitatively robust to these larger point estimates, though they predict stronger income sorting across neighborhoods in response to deregulation.

Table 11 tests the first stage regression under different donut radii. For all radii reaching from 2-25km, we find negative coefficients associated with slopes in the donut region, supporting the rationale behind instrument relevancy. To test if the positive first stage coefficients are associated with differences in the price of housing services, I repeat regressions in Table 11 but with the independent variable being the hedonic housing price index. This is reported in Table 12. For the preferred specifications reaching beyond 10km, instrumented slopes are negatively correlated with housing prices in the focal neighborhood.

Table 10: Pooled IV Specification with various donuts

VARIABLES	(1) ln Amenity	(2) ln Amenity	(3) ln Amenity	(4) ln Amenity	(5) ln Amenity	(6) ln Amenity
ln Income	2.7062	-0.0091	0.0929	0.2857***	0.3727***	0.5505***
	(3.7476)	(0.0855)	(0.0721)	(0.0771)	(0.0639)	(0.1830)
Slope Control	-0.0327	0.0066***	0.0055***	0.0089***	0.0113***	0.0204**
	(0.0493)	(0.0020)	(0.0015)	(0.0021)	(0.0028)	(0.0086)
Local Slope Control	-0.0890	0.0062*	0.0032	-0.0053	-0.0088***	-0.0166*
	(0.1306)	(0.0037)	(0.0035)	(0.0037)	(0.0031)	(0.0086)
Observations	178,685	183,846	185,845	185,979	186,149	185,062
Specification	IV	IV	IV	IV	IV	IV
Donut	1-2km	2-4km	5-10km	10-16km	16-25km	25-35km
Base Controls	Yes	Yes	Yes	Yes	Yes	Yes
Amen/Topo Controls	No	No	No	No	No	No
Density Control	No	No	No	No	No	No
FStat Bart c 35 km	0.5	29.1	27.2	19.7	39.8	3.3

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Pooled IV Specification with various donuts, using the average amenity across low, medium and high income groups. "Slope Control" is the average slope within the block group plus a buffer with length equal to d_1. All specifications include MSA fixed effects and standard errors are clustered using a 35km Bartlett kernel. "Local Slope Control" is the average slope within the block group. ln Income is instrumented with the average slopes of block groups that have centroids within buffer d_1 and d_2. "Base Controls" include travel time, building age, public transport and bus shares in commuting and CBD distance. "Amen/Topo" controls include various amenities (density of coffee shops, parks, restaurants) and various topographic features (cover of different types of forest such as deciduous or evergreen, wetlands, perennial snow cover). "Density Control" is the within-MSA density ranking of the block group.

Additional checks The calibrated amenity values used in estimation depend on the calibration of consumption values $C_i(z)$ and the choice of within-city migration elasticities ρ. The estimates are robust to a host of alternative calibrations. First, are similar for low and high values of ρ based on confidence intervals in Baum-Snow and Han (2024). Second, controlling for the hedonic price index derived from Equation (28) only changes estimates by a few percentage points. Third, estimation results are similar when preferences are assumed Cobb-Douglas

Table 11: Pooled First Stage with various donuts

VARIABLES	(1) ln Income	(2) ln Income	(3) ln Income	(4) ln Income	(5) ln Income	(6) ln Income
Slope Donut	0.0016	-0.0171***	-0.0218***	-0.0210***	-0.0291***	-0.0082
	(0.0034)	(0.0041)	(0.0051)	(0.0052)	(0.0054)	(0.0058)
Slope Control	0.0116***	0.0165***	0.0040	-0.0076	-0.0180*	-0.0442***
	(0.0041)	(0.0047)	(0.0046)	(0.0065)	(0.0093)	(0.0118)
Local Slope Control	0.0356***	0.0401***	0.0450***	0.0463***	0.0461***	0.0468***
	(0.0022)	(0.0021)	(0.0023)	(0.0025)	(0.0027)	(0.0029)
Observations	184,176	189,465	191,532	191,668	191,857	190,731
Specification	IV	IV	IV	IV	IV	IV
Donut	1-2km	2-4km	5-10km	10-16km	16-25km	25-35km
Base Controls	Yes	Yes	Yes	Yes	Yes	Yes
Amen/Topo Controls	No	No	No	No	No	No
Density Control	No	No	No	No	No	No
Cluster	MSA	MSA	MSA	MSA	MSA	MSA

Robust standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Pooled First Stage Specification with various donuts, using neighborhood log average income as an outcome variable. "Slope Donut" is the instrumental variable – average slopes within block groups between a set of distance buffers. "Slope Control" is the average slope within the block group plus a buffer with length equal to d_1. All specifications include MSA fixed effects and standard errors are clustered at the MSA level. "Local Slope Control" is the average slope within the block group. $\ln$ Income is instrumented with the average slopes of block groups that have centroids within buffer d_1 and d_2. "Base Controls" include travel time, building age, public transport and bus shares in commuting and CBD distance. "Amen/Topo" controls include various amenities (density of coffee shops, parks, restaurants) and various topographic features (cover of different types of forest such as deciduous or evergreen, wetlands, perennial snow cover). "Density Control" is the within-MSA density ranking of the block group.

rather than Stone-Geary.

Table 12: Pooled First Stage with various donuts. Hedonic housing price index is the dependent variable.

VARIABLES	(1) hedonicPrice	(2) hedonicPrice	(3) hedonicPrice	(4) hedonicPrice	(5) hedonicPrice	(6) hedonicPrice
Slope Donut	0.0187*	0.0058	0.0093	-0.0218*	-0.0503***	-0.0055
	(0.0106)	(0.0114)	(0.0148)	(0.0116)	(0.0149)	(0.0164)
Slope Control	0.0270**	0.0320**	0.0228**	0.0438**	0.0418*	0.0004
	(0.0112)	(0.0130)	(0.0112)	(0.0193)	(0.0216)	(0.0269)
Local Slope Control	0.0184***	0.0236***	0.0319***	0.0329***	0.0353***	0.0418***
	(0.0047)	(0.0063)	(0.0092)	(0.0097)	(0.0108)	(0.0129)
Observations	178,753	183,919	185,919	186,052	186,224	185,136
Specification	IV	IV	IV	IV	IV	IV
Donut	1-2km	2-4km	5-10km	10-16km	16-25km	25-35km
Base Controls	Yes	Yes	Yes	Yes	Yes	Yes
Amen/Topo Controls	No	No	No	No	No	No
Density Control	No	No	No	No	No	No
Cluster	MSA	MSA	MSA	MSA	MSA	MSA

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Pooled First Stage Specification with various donuts using the hedonic housing price index as an outcome variable. ”Slope Donut” is the instrumental variable – average slopes within block groups between a set of distance buffers. ”Slope Control” is the average slope within the block group plus a buffer with length equal to d_1. All specifications include MSA fixed effects and standard errors are clustered at the MSA level. ”Local Slope Control” is the average slope within the block group. $\ln$ Income is instrumented with the average slopes of block groups that have centroids within buffer d_1 and d_2. ”Base Controls” include travel time, building age, public transport and bus shares in commuting and CBD distance. ”Amen/Topo” controls include various amenities (density of coffee shops, parks, restaurants) and various topographic features (cover of different types of forest such as deciduous or evergreen, wetlands, perennial snow cover). ”Density Control” is the within-MSA density ranking of the block group.

D.4 Alternative Instruments and Tests

Placebo test For further validation, I construct a placebo test that exploits the timing of neighborhood income changes. It relies on the assumption that future incomes cannot cause past amenity values — a plausible assumption if either moving costs are low enough that location decisions do not depend strongly on expectations of future amenity values,[64] or if households cannot accurately forecast future neighborhood income growth.[65]

Let t index years. The placebo centers around the specification:

$$\log \nu_{t=0,i}(z) = \tilde{\beta}(z) \log \mathrm{Inc}_{t=1,i} + \beta_1(z)\, S[d_{1i}] + \log \tilde{\nu}_{t=0,i}(z) \tag{79}$$

where $\nu_{t=0,i}(z)$ are fundamental amenities in the earlier period $t = 0$, $\mathrm{Inc}_{t=1,i}$ is average income in the later period $t = 1$, and $\tilde{\nu}_{t=0,i}(z)$ is the residual. Here $t = 0$ corresponds to the 2008–2012 ACS and 2010 Census cross-section, while $t = 1$ corresponds to the baseline 2016–2020 sample. Under the main identification assumption and the timing assumptions above, $\tilde{\beta}(z) = 0$; if instead $\tilde{\beta}(z) \neq 0$, the instrument is correlated with unobserved $t = 0$ amenity factors and likely with $t = 1$ amenity factors as well. To estimate (79), I invert the spatial equilibrium condition to recover $\log \nu_{t=0,i}(z) = \log b_{t=0,i}(z) - \hat{\Omega}(z) \log \mathrm{Inc}_{t=0,i}$, where estimated values of $\Omega(z)$ are treated as known, and then instrument for $\log \mathrm{Inc}_{t=1,i}$ using the same donut slope.

Table 13 reports the results. The pooled IV estimate of $\tilde{\beta}(z)$ (Column 1) is small and statistically indistinguishable from zero. The OLS counterpart (Column 2) is negative and significant, confirming that the placebo test has statistical power to detect violations. Estimates by income group (Columns 3–5) lead to the same conclusion, except for the highest skill groups on the cusp of statistical significance. Reported standard errors represent a lower bound on the true standard errors because they treat $\hat{\Omega}(z)$ as exact rather than estimated in a previous stage. The placebo therefore provides additional support for the exclusion restriction for all skill groups.

[64] That is, one-time moving costs as in dynamic models of location choice (Caliendo et al., 2019; Bayer et al., 2016).

[65] A weaker sufficient condition is that households can predict future income only from past income, so that future income after conditioning on past income is unpredictable. The placebo also implicitly assumes that the effect of income on amenities is not dynamic in the sense of depending on lagged income flows.

Table 13: Placebo tests, pooled and by income type

VARIABLES	(1) ln Fund. Amenity	(2) ln Fund. Amenity	(3) ln Fund. Amenity (low)	(4) ln Fund. Amenity (med)	(5) ln Fund. Amenity (high)
ln Income (2020)	0.0469	-0.0455***	0.1058	0.0478	-0.0082
	(0.0913)	(0.0045)	(0.1239)	(0.0784)	(0.0837)
Slope Control	0.0123***	0.0101***	0.0183***	0.0092***	0.0087***
	(0.0026)	(0.0022)	(0.0035)	(0.0022)	(0.0024)
Local Slope Control	0.0009	0.0031***	0.0012	-0.0008	0.0019
	(0.0022)	(0.0006)	(0.0029)	(0.0019)	(0.0021)
Observations	187,816	187,816	187,764	187,543	180,673
Specification	IV	OLS	IV	IV	IV
Donut	10-16km	10-16km	10-16km	10-16km	10-16km
Base Controls	Yes	Yes	Yes	Yes	Yes
Amen/Topo Controls	No	No	No	No	No
Density Control	No	No	No	No	No
FStat Bart c 35 km	23.2		23.0	23.2	23.4

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Placebo tests. Column (1) reports the estimate of a pooled specification of (79). Column (2) estimates (79) with OLS under the assumption that the pooled baseline estimate of $\Omega(z)$ is identified. The remaining columns estimate (79) with IV disaggregated by 3 skill groups: low ($0–50k yearly income), medium ($50k–100k), and high ($100k+), measured in yearly household income if living in an average productivity city. All specifications include MSA fixed effects and standard errors are clustered using a 35km Bartlett kernel. "Local Slope Control" is the average slope within the block group. $\ln$ Income is instrumented with the average slopes of block groups that have centroids within buffer d_1 and d_2. "Base Controls" include travel time, building age, public transport and bus shares in commuting and CBD distance. "Amen/Topo" controls include various amenities (density of coffee shops, parks, restaurants) and various topographic features (cover of different types of forest such as deciduous or evergreen, wetlands, perennial snow cover). "Density Control" is the within-MSA density ranking of the block group.

Davis, Gregory, and Hartley (2023) Instrument As an additional robustness check, I instrument for neighborhood income using a shift-share instrument adapted from Davis et al. (2023). The key idea is to predict neighborhood income using metro-level population shares by skill level and only topographic features of the land—specifically average slope, average elevation, the standard deviation of elevation, and the fraction of land with slopes exceeding 15 degrees.

The instrument is constructed in three steps. In the first step, I compute the within-metro choice probability for each of seven income bins $k \in \{1, \ldots, 7\}$:

$$\pi_{ik} = \frac{L_{ik}}{\sum_{i' \in \mathbb{C}_i} L_{i'k}} \tag{80}$$

where L_ik is the number of residents in income bin k living in block group i, the sum is taken over all block groups i' in the same metropolitan area $\mathbb{C}_i$, and the counts are drawn from the 2016–2020 ACS and 2020 Census. Block groups with zero residents in bin k are excluded from estimation for that bin.

In the second step, I regress $\log \pi_{ik}$ on the topographic variables and metropolitan area fixed effects:

$$\log \pi_{ik} = \alpha_k \overline{\text{Slope}}_i + \gamma_k \overline{\text{Elev}}_i + \delta_k \text{SD\_Elev}_i + \lambda_k \text{SteepSlope}_i + \mu_{\mathbb{C}_i} + \varepsilon_{ik} \tag{81}$$

where $\overline{\text{Slope}}_i$ is mean slope, $\overline{\text{Elev}}_i$ is mean elevation, SD_Elev_i is the standard deviation of elevation, SteepSlope_i is the fraction of land with slope exceeding 15 degrees, and μ_c is a metropolitan area fixed effect. Coefficients are estimated separately for each bin k. Predicted probabilities $\hat{\pi}_{ik}$ are then computed for *all* block groups in the sample, including those with zero observed counts, by exponentiating the fitted values from the full model matrix.

In the third step, I recover predicted neighborhood income shares via a shift-share formula. Let $\Pi_k(c) = L_{kc}/L_c$ denote the metropolitan-level share of residents in income bin k in city c. The predicted share of bin-k residents in neighborhood i is

$$\widehat{s}_{ik} = \frac{\Pi_k(c)\, \hat{\pi}_{ik}}{\sum_{k'=1}^{7} \Pi_{k'}(c)\, \hat{\pi}_{ik'}} \tag{82}$$

where c is the city associated with i. Predicted average neighborhood income is then

$$\widehat{\text{Inc}}_i = \sum_{k=1}^{7} \widehat{s}_{ik}\, \bar{w}_k(c) \tag{83}$$

where $\bar{w}_k(c)$ is the estimated ability (wage) measure for income group k in metro c from the city productivity estimation in Section 4. This instrument thus combines city-level income composition (the "shares") with topography-driven predicted sorting within the city (the "shifts"), and is constructed entirely from variables that are predetermined with respect to contemporary amenity shocks. As in Davis et al. (2023), I control for the direct effect of each of these topographic variables on location choices nationwide. Identification relies on local topography being uncorrelated with the difference in unobserved amenities by skill group, conditional on the direct effect of topography on neighborhood choice itself.

The results are reported in Table 14. Estimates are slightly smaller for medium- and low-skill households relative to the preferred specification. High income households have virtually identical estimates. All model results are quantatively robust to these estimates.

Table 14: Alternative IV specification, disaggregated by skill group

VARIABLES	(1) ln Amenity (Low)	(2) ln Amenity (Med)	(3) ln Amenity (High)
log_Average_Income	0.1869***	0.1634***	0.4190***
	(0.0618)	(0.0348)	(0.0348)
Avg_Slope	0.0031	0.0006	-0.0035**
	(0.0025)	(0.0015)	(0.0015)
Observations	185,485	185,661	181,552
Specification	IV	IV	IV
Base Controls	Yes	Yes	Yes
Amen/Topo Controls	Yes	Yes	Yes
Density Control	No	No	No

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

IV Specification by skill group using the Davis et al. (2023)-style instrument. Estimates disaggregated for 3 skill groups: low ($0–50k yearly income), medium ($50k–100k), and high ($100k+), measured in yearly household income if living in an average productivity city. Avg Slope is the block group level average slope. Additional controls include average elevation, variance in elevation, and the fraction of land with steep slopes exceeding 15 degrees (omitted). "Base Controls" include travel time, building age, public transport and bus shares in commuting and CBD distance. "Amen/Topo" controls include various amenities (density of coffee shops, parks, restaurants) and various topographic features (cover of different types of forest such as deciduous or evergreen, wetlands, perennial snow cover). "Density Control" is the within-MSA density ranking of the block group.

E Appendix: Counterfactuals

E.1 Algorithm for computing counterfactuals

In this section, I detail the algorithm used to compute counterfactual outcomes. There are two main sets of state variables in this algorithm: 1) Populations (by neighborhood, zone and type) and 2) Spending shares on housing (by neighborhood, zone and type) that are iterated on to solve for an equilibrium both across neighborhoods and for housing markets within neighborhoods.

Algorithm 3 Counterfactuals

Initialize error on populations (with tolerance ϵ_L) and spending shares (tolerance ϵ_S)

While Population Error $> \epsilon_L$ or Spending Share Error $> \epsilon_S$ do

For each neighborhood i, zone $\boldsymbol{o}$

1. Solve for housing prices $P_{i\boldsymbol{o}}$ to clear markets given current neighborhood-zone allocation $L_{i\boldsymbol{o}}(z)$ and spending shares $\beta_{i\boldsymbol{o}}(z)$
2. Update wages, amenities and/or productivity given current neighborhood allocation (when applicable)
3. Given wages, amenities, productivity and housing prices from above, do for each z:
   (a) Calculate what the desired neighborhood allocation $L_{i\boldsymbol{o}}(z)$ would be (Equation 7)
   (b) Calculate what the desired spending share on housing $\tilde{\beta}_{i\boldsymbol{o}}(z)$ would be (using the solution to Equation 4)

Update Population Error $= \max_{i,\boldsymbol{o},z} |\tilde{L}_{i\boldsymbol{o}}(z) - L_{i\boldsymbol{o}}(z)|$ and Spending Share Error $= \max_{i,\boldsymbol{o},z} |\tilde{\beta}_{i\boldsymbol{o}}(z) - \beta_{i\boldsymbol{o}}(z)|$

Set $L_{i\boldsymbol{o}}(z) = L_{i\boldsymbol{o}}(z) + \kappa[\tilde{L}_{i\boldsymbol{o}}(z) - L_{i\boldsymbol{o}}(z)]$ and $\beta_{i\boldsymbol{o}}(z) = \beta_{i\boldsymbol{o}}(z) + \kappa[\tilde{\beta}_{i\boldsymbol{o}}(z) - \beta_{i\boldsymbol{o}}(z)]$ for some parameter $\kappa < 1$.

E.2 Defining the equivalent variation

This section defines the equivalent variation measure of welfare: defined as the percentage change in income received uniformly across locations that makes the average household as well off under baseline prices and amenities when compared to the counterfactual equilibrium. Let $\boldsymbol{W}_t(z)$ be the welfare measure (85) evaluated for model t a household of skill z. $t = 1$ denotes the welfare under the counterfactual equilibrium; and $t = 0$ is the baseline.

Next, let $V_0(i, z, \tau)$ be the indirect utility associated with a neighborhood i for skill z under the baseline calibration prices and amenity values, and for a level of *equivalent variation* τ that scales wages in the indirect utility equation (4). That is,

$$V_0(i, z, \tau) := \max_{A,g,\boldsymbol{o}} \underbrace{z^{-1}\beta^{-\beta}(1-\beta)^{-(1-\beta)}(A-\bar{A})^{\beta}g^{1-\beta}}_{\text{Consumption value}} + \underbrace{\log b_i(z)}_{\text{Amenity value}} \tag{84}$$

subject to

$$P_{i\boldsymbol{o}}A \geq R_i \text{ if } \boldsymbol{o} = R \text{ and}$$
$$P_{i\boldsymbol{o}}A + g \leq \underbrace{w_c(1+\tau)}_{\text{Cash injection}} z$$

$V_0(i, z, \tau)$ defines the location utility after a cash injection of $100 \times \tau\%$ of income in that location. I define this equivalent variation measure τ that uniquely solves the equation

$$\log \left[\sum_{c \in \mathbb{C}} \left[\sum_{i' \in N(c)} \left[e^{V_0(i',z,\tau)} \right]^{\rho} \right]^{\frac{\theta}{\rho}} \right]^{\frac{1}{\theta}} = \log \boldsymbol{W}_1(z) \tag{85}$$

The interpretation of τ is the amount of cash (expressed as a percentage of income in location i) delivered to z households in all locations simultaneously to make them as well off in the baseline equilibrium with regulation as in the counterfactual equilibrium with complete deregulation. Importantly, this measure ignores the general equilibrium effects of this cash injection on housing prices and amenities arising from new location choices.

E.3 Weighing renters and landowners welfare: continued

In this appendix, I outline exactly how portfolio weights $s(z)$ are calculated and how welfare measurement is done without re-calibrating or re-simulating the model and incorporating homeownership status.

Welfare weights In the baseline model, households have an implied amount spending on housing services, which we denote by $S_{i\boldsymbol{o}}(z)$ and depends on the household type (spending shares by household are targeted in calibration), as well as the neighborhood and the zone of residence. I define total spending by type z homeowners as

$$\boldsymbol{S}(z) := H(z) \sum_{i \in \cup_{c \in \mathbb{C}} N(c);\ \boldsymbol{o} \in \{R,U\}} S_{i\boldsymbol{o}}(z) L_{i\boldsymbol{o}}(z)$$

where $H(z)$ is the national share of homeowners conditional on being a type z household (from the data), and $L_{i\boldsymbol{o}}(z)$ is the model-based population choosing neighborhood i and zone $\boldsymbol{o}$. The normalized share of total spending

$$\frac{\boldsymbol{S}(z)}{\sum_{z' \in Z} \boldsymbol{S}(z')}$$

from this procedure is 0.04 for households making between $0 - 25k$, 0.09 for $25 - 50k$, 0.12 for households making $50 - 75$ and $75 - 100k$, 0.2 for households making $100 - 150k$, 0.14 for $150 - 200k$, and 0.29 for households making $200k+$. Finally, I define the weight $s(z)$ as

$$s(z) = \frac{1}{H(z)L(z)} \frac{\boldsymbol{S}(z)}{\sum_{z \in Z} \boldsymbol{S}(z)}.$$

where $L(z)$ is the total mass of z households.

Calculating welfare for homeowners and renters Next, I detail how welfare calculation is performed. Importantly, I do not re-calibrate or re-simulate the model with this new dimension of household heterogeneity. This means I assume prices, land values and amenities change after deregulation in the same way as the baseline model with absentee landowners.

To define the equivalent variation of a homeowner, I consider the definition from Section E.2 and make two changes. First, I add income on land to the measure of location indirect utility for homeowners of type z. In the counterfactual equilibrium, I enter land income after complete deregulation into the calculation of indirect utility. Then, I calculate the equivalent variation in the exact same way as in Section E.2, with the exception of baseline and counterfactual land incomes for homeowners.

E.4 Shapley value decomposition

This appendix briefly explains the Shapley decomposition to isolate the roles of changing consumption and amenity values performed in Section 6. Consider any the two equilibrium outcome vectors $C_t(z)$ (consumption value) and $b_t(z)$ (amenity value), where $t = 0$ denotes pre-counterfactual and $t = 1$ counterfactual outcomes, and vector components represent individual neighborhoods. Let $\log \boldsymbol{W}(C_t, b_t)$ be the renter welfare measure from Equation (85) evaluated at some vector of consumption values C_t and amenity values b_t.

I define an *effect* of changing consumption as

$$\tilde{\boldsymbol{W}}_{C,t}(z) = \log \boldsymbol{W}(C_1(z), b_t(z)) - \log \boldsymbol{W}(C_0(z), b_t(z)) \tag{86}$$

and similarly an *effect* of changing amenities as

$$\tilde{\boldsymbol{W}}_{b,t}(z) = \log \boldsymbol{W}(C_t(z), b_1(z)) - \log \boldsymbol{W}(C_t(z), b_0(z))$$

for some $t \in \{0, 1\}$. A *Shapley effect* $S_C(z)$ for the consumption is the unweighted average of consumption effects over t. That is,

$$S_C(z) = \frac{1}{2} \sum_{t \in \{0,1\}} \tilde{\boldsymbol{W}}_{C,t}(z) \tag{87}$$

with an analogous definition for amenities. By definition, $S_C(z) + S_b(z) = \log \boldsymbol{W}(C_1(z), b_1(z)) - \log \boldsymbol{W}(C_0(z), b_0(z))$, which is the total welfare effect of a counterfactual. In Table 3 (and other tables) I report this Shapley decomposition with the total effect $S_C(z) + S_b(z)$ rescaled to match the equivalent variation associated with the counterfactual.

E.5 Supplementary figures and tables

Figure 12: Income-Density Gradients in baseline and counterfactual after complete deregulation.

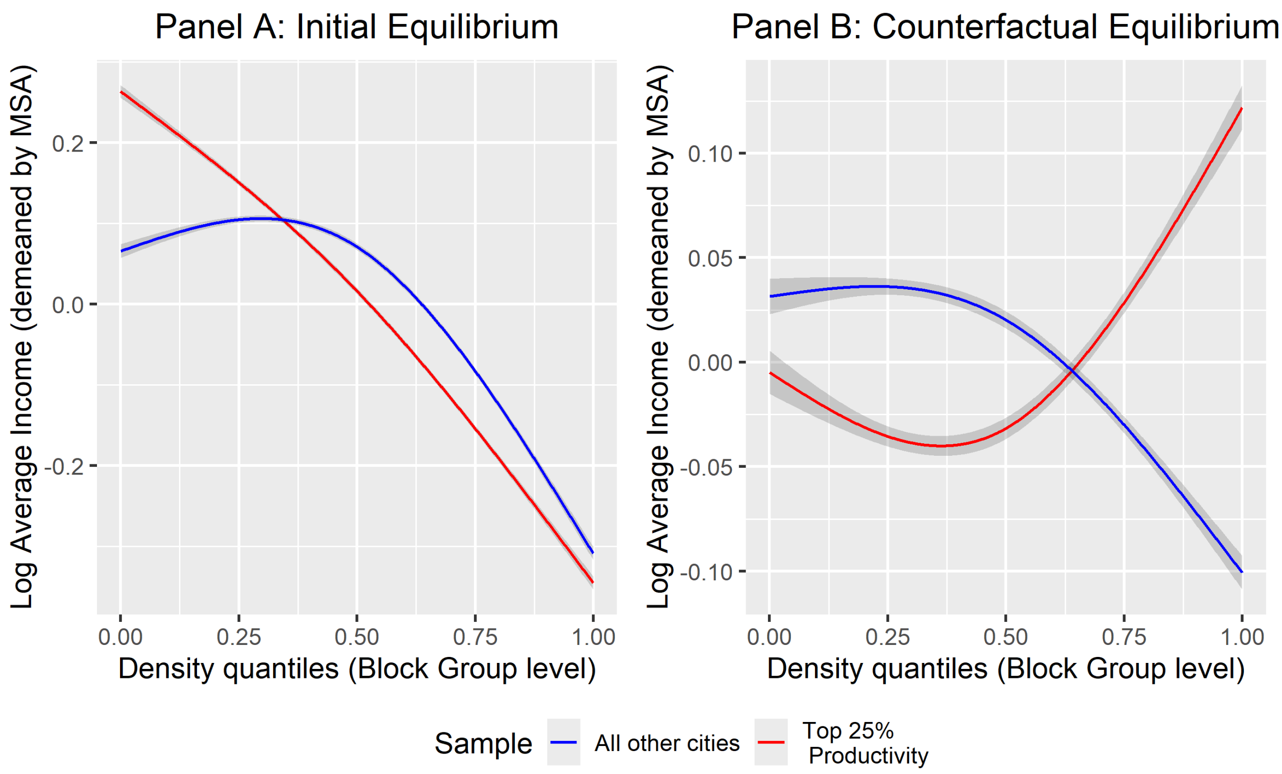

Panel A corresponds to the same estimates of Panel A of Figure 2. Panel B compares the income density gradient across "superstar" sample cities and "non-superstar" cities that is generated in an equilibrium without minimum lot sizes. The gentrification of high density neighborhoods is apparent. Differences in the income density gradient across samples disappear when transitioning from the initial equilibrium to the counterfactual equilibrium.

Table 15: Productivity changes from counterfactual by model assumptions.

End. Amenities	End. Productivity	Education	Prod. Growth	No Income Sorting
✓			0.12%	1.65%
			0.41%	2.68%
✓	✓		0.4%	2.11%
✓		✓	0.22%	1.52%
✓	✓	✓	0.49%	1.84%

Productivity changes from counterfactual by model assumptions. "Prod. Growth" refers to productivity changes from counterfactual. "No income sorting" refers to calculating what productivity growth would be if cities grew at levels determined by counterfactual, but with no compositional changes to city income or education. For the full model (last row), I ignore how city wages change when calculating the no-income-sorting counterfactual. This is because uniform city growth when productivity is endogenous by education causes income sorting, which, in turn, is due to agglomeration being education-augmenting, as in Baum-Snow et al. (2018).

Figure 13:
San Francisco, initial data and counterfactual outcomes.

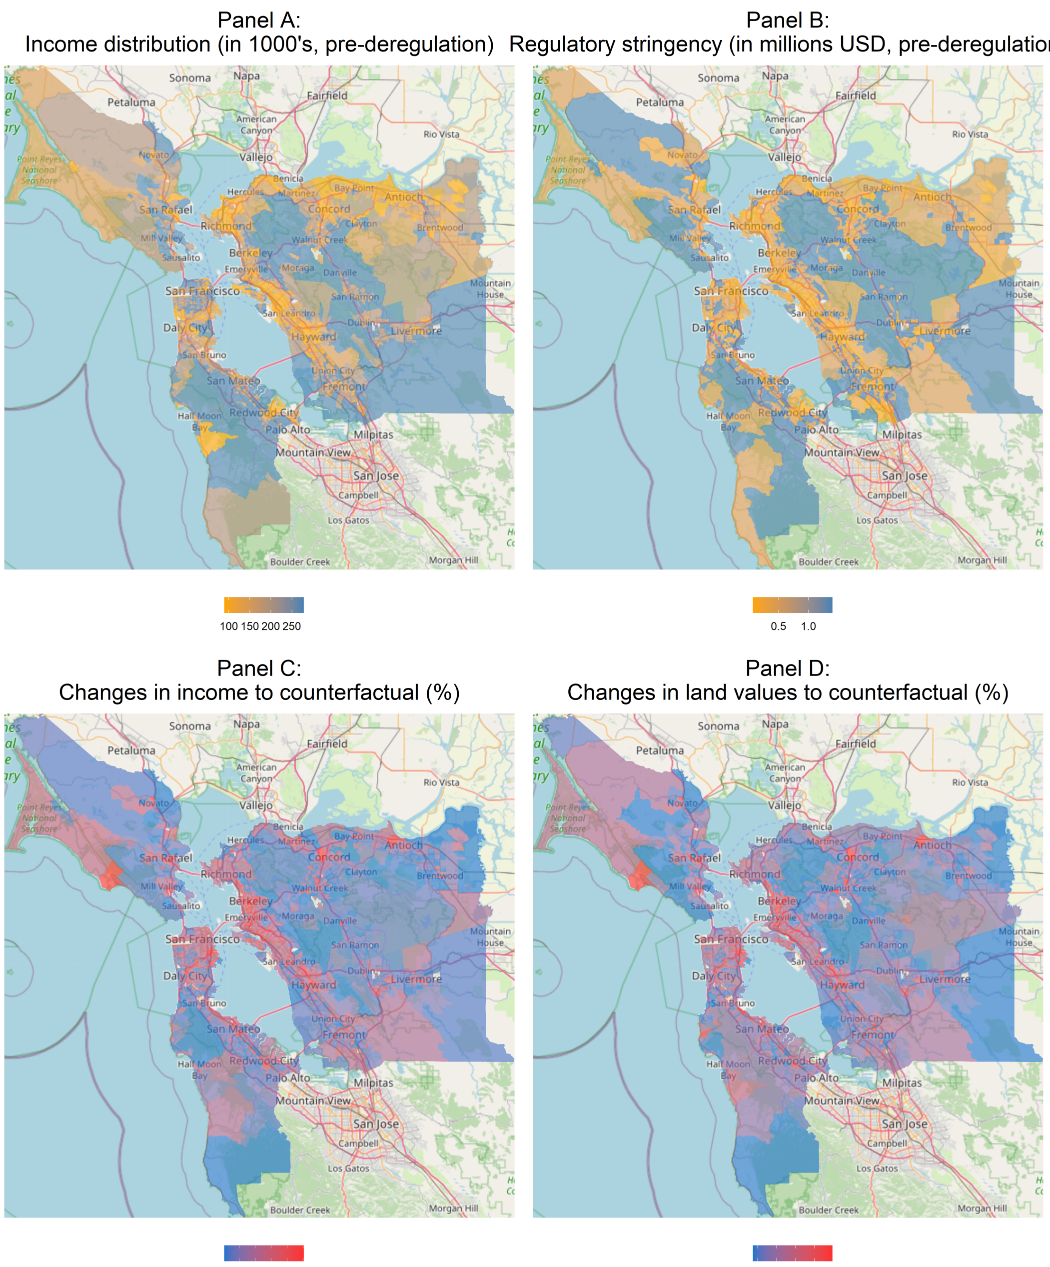

San Francisco's initial data on income and regulatory stringency, along with changes in incomes and land values after halving the minimum lot size. Values are binned by quartile for improved readability of the graph. Regulatory stringency is the empirical measure introduced in Equation (13).

Figure 14: Distribution of Land Value Changes in San Francisco after unilateral deregulation

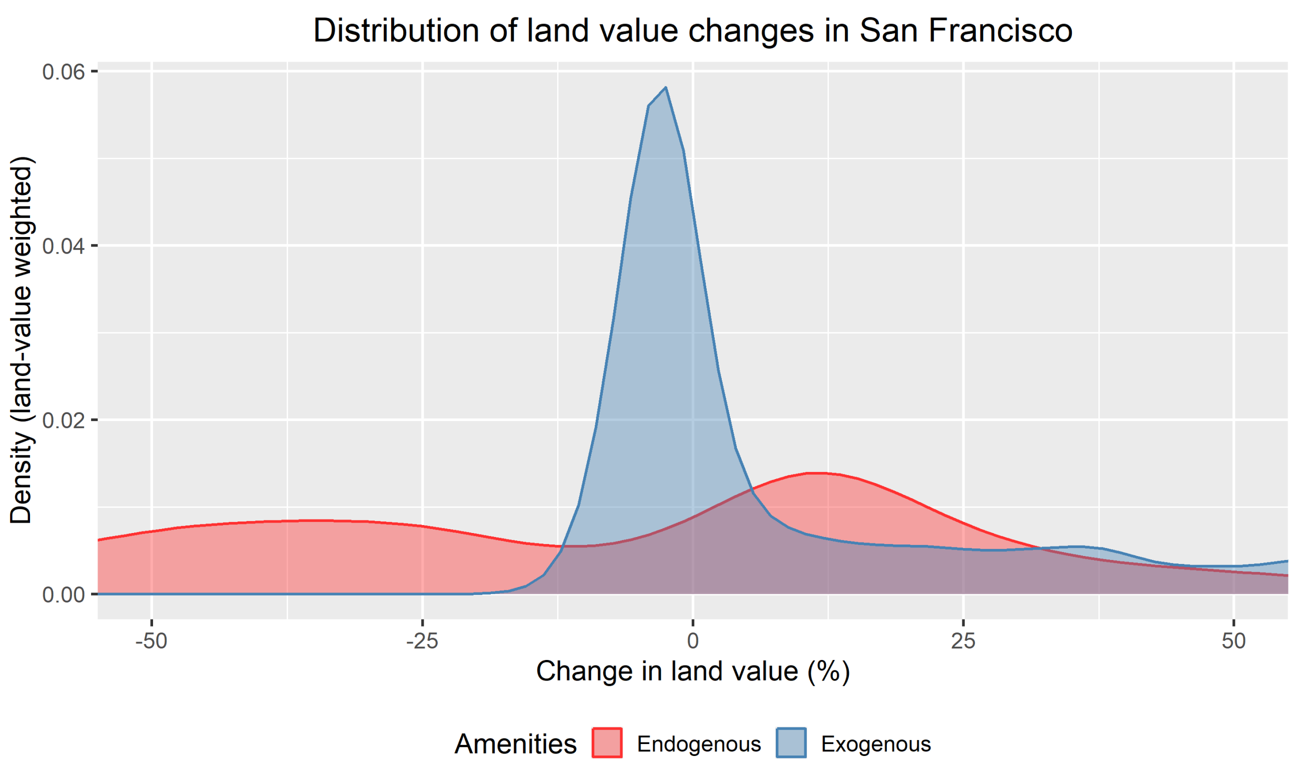

Distributions are weighted by land value in the initial equilibrium and are shaded for counterfactuals where neighborhood amenities are endogenous and exogenous.

Figure 15: Shapley decomposition of welfare effects from permuted policy, in $

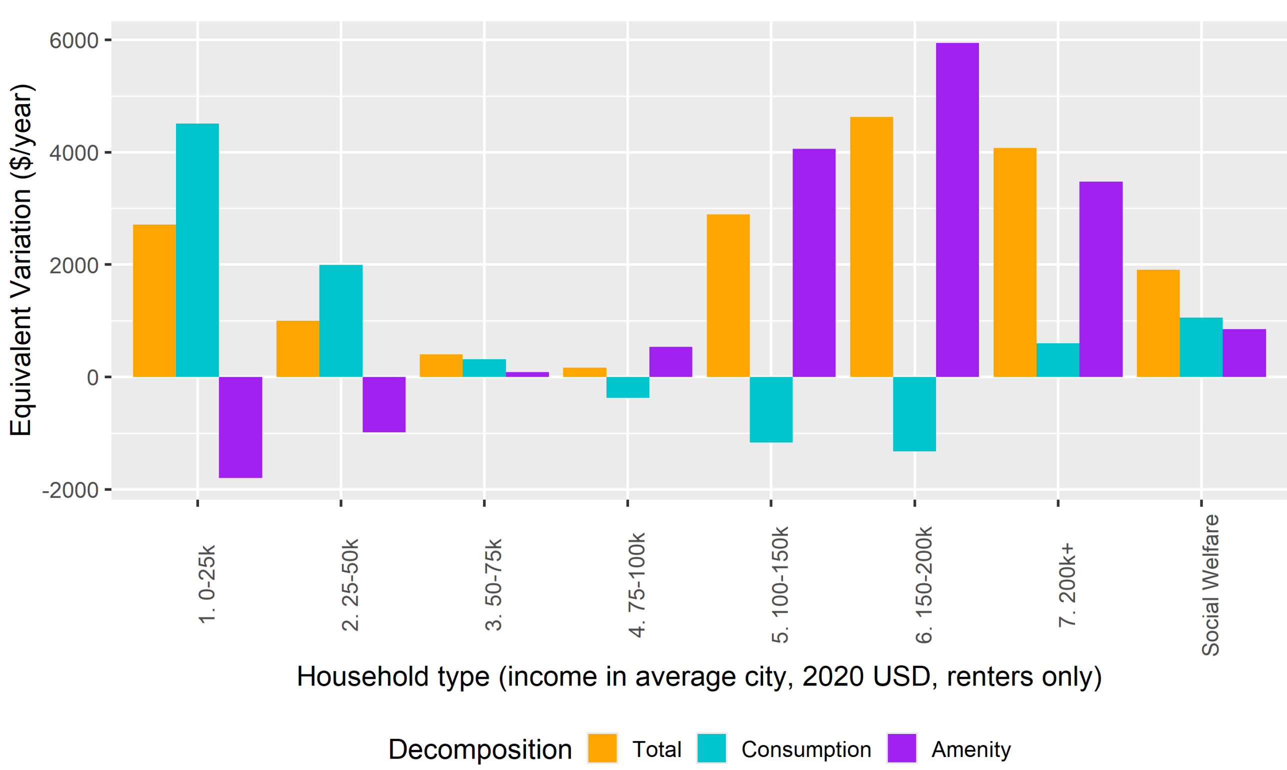

Welfare is measured as the equivalent variation. Higher values mean higher welfare gains. Social welfare is the population weighted average of welfare by type. "Amenity" and "Consumption" components are constructed using Shapley values, with a procedure outlined in Appendix E.4. The sum of components add to the equivalent variation associated with the policy that permutes regulation to target rich neighborhoods.

Figure 16: City income sorting after permuting regulation.

The y axis is defined as the change in the average income that a household could earn in an average city from baseline to counterfactual. The x axis measures city productivity (residualized wages from the data). The permuted policy disproportionately deregulates productive cities, suggesting that the amenity valuation of productive cities for high skill people is not high enough to justify regulation.

Figure 17:
Gentrification in superstar cities after permuting regulation.

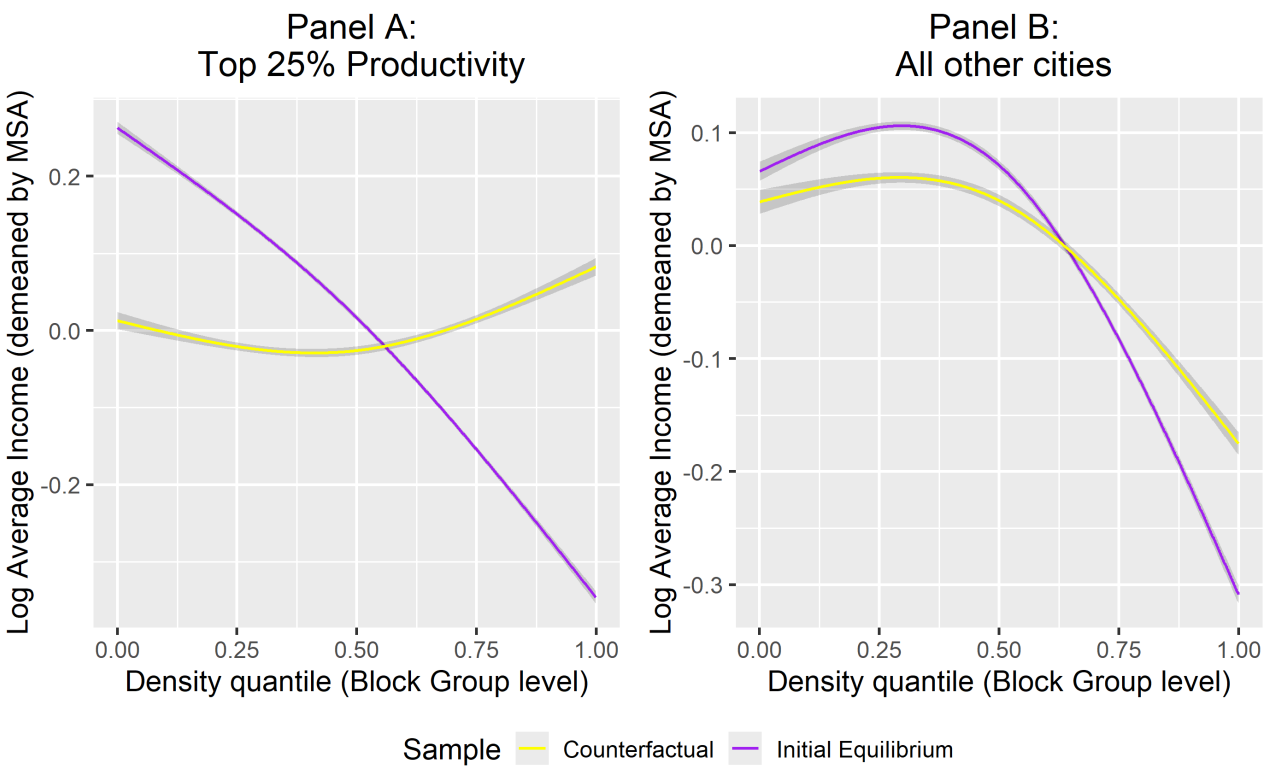

For each panel, demeaned log average income at the neighborhood level is regressed against the observed density ranking of neighborhoods at baseline. The purple regression uses data from the baseline equilibrium that matches data, as in Figure . The yellow regression uses data generated from the counterfactual where regulation is permuted to target rich neighborhoods based on the high-skill amenity score. Each panel corresponds to ”superstar” sample cities (Top 25%) and ”non-superstar” sample cities (Bottom 75%), as in Figure . This policy change suggests that low-skill households somewhat value low density neighborhoods in productive cities, but regulation otherwise makes them too exclusive.