

Income Growth and the Distributional Effects of Urban Spatial Sorting*

Victor Couture Cecile Gaubert Jessie Handbury Erik Hurst

August 2021

Abstract

We explore the impact of rising incomes at the top of the distribution on spatial sorting patterns within large U.S. cities. We develop and quantify a spatial model of a city with heterogeneous agents and non-homothetic preferences for neighborhoods with endogenous amenity quality. As the rich get richer, demand increases for the high-quality amenities available in downtown neighborhoods. Rising demand drives up house prices and spurs the development of higher quality neighborhoods downtown. This gentrification of downtowns makes poor incumbents worse off, as they are either displaced to the suburbs or pay higher rents for amenities that they do not value as much. We quantify the corresponding impact on well-being inequality. Through the lens of the quantified model, the change in the income distribution between 1990 and 2014 led to neighborhood change and spatial resorting within urban areas that increased the welfare of richer households relative to that of poorer households, above and beyond rising nominal income inequality.

*We thank Hero Ashman, Allison Green, Serena Xu, and Xianya Zhou for outstanding research assistance. We thank Treb Allen, David Atkin, Don Davis, Jonathan Dingel, Ingrid Ellen, Alessandra Fogli, Ed Glaeser, Gabriel Kreindler, Christopher Palmer, Steve Redding, Andres Rodriguez-Clare, Jung Sakong, Yichen Su, Jacob Vigdor, and Jonathan Vogel for many helpful comments as well as numerous seminar and conference participants. Victor Couture thanks the Fisher Center for Real Estate and Urban Economics at UC Berkeley for financial support. Cecile Gaubert thanks the Fisher Center for Real Estate and Urban Economics and the Clausen Center for International Business and Policy at UC Berkeley for financial support. Jessie Handbury thanks the Zell/Lurie Real Estate Center at Wharton for financial support. Authors contact information: victor.couture@ubc.ca, cecile.gaubert@berkeley.edu, handbury@wharton.upenn.edu, and erik.hurst@chicagobooth.edu.

1 Introduction

Over the last three decades income inequality in the United States has grown sharply, with income growth concentrated at the top of the earnings distribution. At the same time, higher income individuals have been moving back into urban cores, transforming downtown neighborhoods and sparking a policy debate on gentrification within many US cities. We posit that these trends are linked. Rich households are more likely to live downtown than middle-income households in part because these areas afford them access to local amenities like restaurants or entertainment venues.¹ As the incomes of the rich increase, aggregate demand for neighborhoods with these luxury urban amenities rises and more of them choose to reside downtown, triggering the redevelopment of previously low-income neighborhoods.

In this paper, we develop a model to formalize this mechanism. We use the model to quantify how top income growth contributes to neighborhood change, measure the associated welfare effects, and guide policy designed to curtail the resulting gentrification. The key model features include (i) non-homothetic preferences for location: rich and poor households make systematically different location choices within a given city and (ii) endogenous neighborhood development: the quality of amenities in city neighborhoods responds endogeneously to demand. The macro and trade literature have highlighted that a rise in nominal income inequality can induce an even stronger increase in real income inequality in the presence of such non-homothetic preferences and endogenous supply responses. We apply this logic to the endogenous development of neighborhoods within cities.

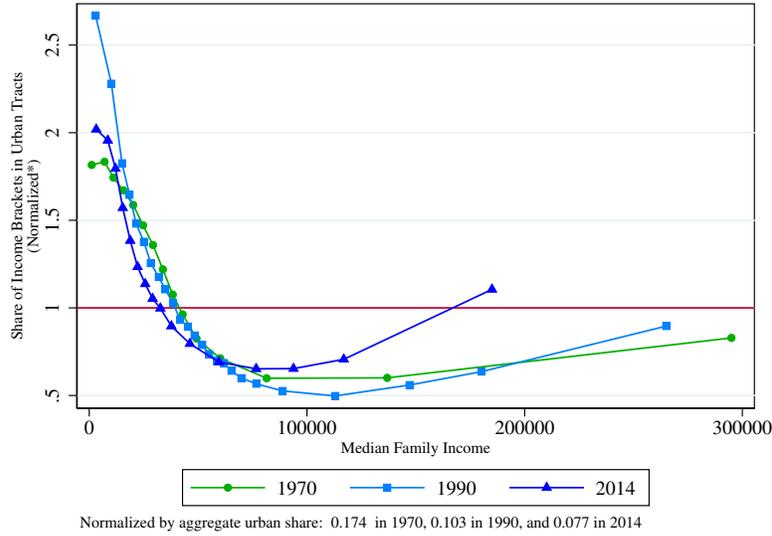
Figure 1 motivates our analysis by summarizing the changes in spatial sorting by income in large U.S. cities from 1970 to today. Specifically, the figure plots the propensity of households in each Census income bracket to live “downtown” in 1970, 1990, and 2014 relative to the average household in each period. Downtown is defined as the set of constant boundary Census tracts closest to the city center that account for 10% of each CBSA’s population in 2000.² We restrict our analysis to the 100 largest CBSAs based on 1990 population.³ The figure documents that the propensity to live downtown is U-shaped in household income, and has been so since at least 1970. As is well known, poorer households are over-represented in downtown areas. A perhaps lesser known fact is that richer households are also systematically over-represented downtown: above \$100,000, the propensity to live downtown increases with income. Notably, Figure 1 also shows that the U-shape has shifted in the recent period. Between 1990 and 2014, the rich have become more likely to reside downtown, and the poor less so.

¹For example, Aguiar and Bills (2015) estimate that restaurant meals and non-durable entertainment are among the goods with the highest income elasticities. Couture and Handbury (2017) document that downtown areas of major cities have a higher density of such amenities.

²CBSAs are Core-Based Statistical Areas defined by the U.S. Census Bureau. CBSAs consist of a core area with substantial population, together with adjacent communities with a high degree of economic and social integration with the core. CBSAs with population above 50,000 are also referred to as Metropolitan Statistical Areas (MSAs).

³See Appendix A for a detailed discussion of the construction and robustness of Figure 1. The patterns in Figure 1 hold for reasonable variation in our spatial definition of downtowns, as well as within detailed demographic groups. These patterns are robust to controlling for many demographic characteristics, thus alleviating concerns that Figure 1 reflects demographic characteristics that are correlated with income, or changes in demographic characteristics that are correlated with changes in income.

Figure 1: Downtown Residential Income Propensity by Income



Note: This figure shows sorting by income in 100 large CBSAs in 1970, 1990, and 2014. The 100 CBSAs are those with the largest population in 1990. Each point plots the share of families in a given Census income bracket who reside downtown in a given year – normalized by the share of all families who reside downtown that year – against the median family income (in 1999 dollars) of that Census income bracket in that year. The downtown area of each CBSA is defined as the set of tracts closest to the center of each CBSA that account for 10% of that CBSA’s population in 2000. The number of points on the graph is limited by the number of income brackets reported by the Census for tract-level data. We compute the median income in each bracket using IPUMS microdata for the corresponding year in the 100 largest CBSAs. The IPUMS microdata is adjusted for top-coding using the generalized Pareto method, as described in Appendix C.

We develop a model of a stylized city that accommodates such U-shape sorting patterns and then use the model to investigate how much of the recent change in the U-shape can be traced back to changes in the income distribution over time. In the model, households are heterogeneous in incomes and choose where to live among neighborhoods that offer different qualities of amenities and housing. Households trade off neighborhood attractiveness against cost of living. This cost depends on local housing prices and on the cost of commuting to work. Preferences for neighborhoods are non-homothetic: households with higher incomes are more willing to pay the higher cost of living in desirable neighborhoods. On the supply side, neighborhoods are built by private developers who compete for land within each area of the city. As top incomes grow, demand for high-quality neighborhoods downtown rises leading to an increase in prices throughout downtown, including in low-quality neighborhoods where the poor live. An increase in the supply of high-quality neighborhoods, and an associated decrease in the variety of low-quality neighborhoods, amplifies this price mechanism. Through these mechanisms, an influx of richer households downtown unambiguously hurts the lower income renters residing there.

The model also builds in mitigating forces through which an influx of higher incomes downtown could benefit incumbent poor households. First, local governments build public amenities like parks or schools that benefit all households in a given location. The provision of public amenities

increases as the tax base downtown increases. Second, some low-income downtown residents own their homes, and hence reap the benefits of house price appreciation. Given these mitigating forces, how an influx of richer households into downtown areas affects the well-being of lower income households on net is a quantitative question.

Before fully estimating the model, we use micro data to provide evidence for the model’s key sorting mechanism. We exploit changes in spatial sorting patterns in response to a CBSA-wide income shock, driven by plausibly exogenous variation in labor demand across cities. Our instrumental variable estimation shows that an income shock raises house prices and amenity quality downtown more than in the suburbs and induces the rich to re-sort into downtowns. These results suggest that income growth triggers within-city spatial sorting disproportionately drawing the rich downtown, consistent with our model of non-homothetic location choices.

This micro evidence yields an estimate for the elasticity that governs how *income-dependent* within-city location choices are, a key parameter in our quantitative analysis. After fully quantifying the model, we show that it is able to replicate the fact that downtown areas are disproportionately populated by both very low and very high earners, mimicking the U-shape in Figure 1. In the model, low-income households minimize the costs of housing and commuting by residing in low-quality neighborhoods downtown. At the same time, higher income households are attracted downtown by the density of high-quality, high-amenity luxury neighborhoods offered there.

We use the quantified model to back out how much of the change in spatial sorting between 1990 and 2014 comes from changes in the income distribution. We find that (i) the increased incomes of the rich since 1990 are causing a phenomenon that looks like urban gentrification – the in-migration of higher income residents downtown causes the amenity mix of neighborhoods to change – and (ii) this mechanism can explain roughly forty percent of the urbanization of the rich (top income decile) and roughly sixty percent of the suburbanization of the poor (bottom income decile).⁴ These findings highlight that, while other forces outside of the model are also arguably contributing to neighborhood change over the last few decades, the rising incomes of the rich play a quantitatively important role in the recent rise in gentrification of urban centers.

To further validate the model, we present additional counterfactual analyses. First, we show that a similar procedure applied to both the 1950-1970 and the 1970-1990 change in the income distribution lead to spatial sorting responses that are qualitatively different to 1990-2014, but similar to those observed in the data over the same earlier periods. While incomes increased during these periods, our analysis suggests that there was not a sufficiently large increase in households at very high income levels to trigger urbanization of the higher-income households. Second, we

⁴Throughout the paper, we often use “neighborhood change” of low-income neighborhoods and “gentrification” interchangeably. We realize that gentrification is a complex process with many potential definitions and drivers. Our interpretation is closest to the definition in the Merriam-Webster dictionary that defines gentrification as “the process of renewal and rebuilding accompanying the influx of middle-class or affluent people into deteriorating areas that often displaces poorer residents.” Our paper is not intended to explore all potential underlying causes of neighborhood gentrification. Rather, we wish to focus on the dimension of gentrification that follows the rise in top incomes. Specifically, we focus on the interaction of rising top incomes, non-homothetic preferences for urban amenities, and endogenous spatial responses.

show that the model also performs quite well at explaining cross-CBSA variation in spatial sorting during the 1990 to 2014 period. We feed into our model changes in the income distribution that are CBSA specific. The resulting model predictions match well the actual patterns of neighborhood change observed empirically within each CBSA. These additional counterfactuals further highlight the ability of our model to match empirical patterns linking changes in the income distribution and changes in the propensity of high-income individuals to locate downtown.

We next use the quantified model to study the normative implications of changes in spatial sorting. We find that the 1990-2014 increase in income inequality triggered an even larger increase in well-being inequality once accounting for changes in neighborhood quality and spatial sorting. Overall, we find that accounting for endogenous spatial sorting responses increased well-being inequality between the top and bottom deciles of the income distribution by 3.6 percentage points (on a base of 19 percentage points) during the past three decades. We further estimate that the welfare of low-income renters was actually reduced by 0.75% from the resulting gentrification stemming from top income growth during this period. Quantitatively, mitigating forces like the public provision of amenities are not strong enough to overcome the base mechanism, which hurts the poor primarily through higher rents.

Finally, we use the model to study the effect of stylized ‘anti-gentrification’ policies, for instance one that taxes housing in high-quality neighborhoods downtown to subsidize housing in low-quality neighborhoods downtown. We find that such policies can be effective in maintaining a diverse income mix downtown. However, these policies do not overturn the increase in well-being inequality that we find for 1990-2014, as price and quality upgrades are to a large extent pushed to the suburbs. In contrast, a policy that relieves housing supply constraints downtown mitigates the negative welfare impact of neighborhood change on the poor, but does not curb the influx of the rich into downtowns.

Related Literature This paper contributes to four main literatures. First, a growing literature studies how nominal income inequality growth can induce even stronger real income inequality growth in models with non-homothetic preferences and endogeneous supply responses, especially in the context of international trade.⁵ We apply this logic to the endogenous development of neighborhoods within cities.

Second, we contribute to the quantitative spatial economics literature reviewed in Redding and Rossi-Hansberg (2017), more specifically to the strand that studies the internal structure of cities (Ahlfeldt et al. (2015); Allen et al. (2015); Redding and Sturm (2016)). Different from our approach, these papers feature homogeneous workers, homothetic preferences, and model a specific city with no extensive margin of within-city locations. We propose a stylized model of a representative city that allows us to model the extensive margin: the number and quality of neighborhoods in a city

⁵Faber and Fally (2021) show that more productive firms target wealthier households; Jaravel (2018) shows that innovation is skewed towards the growing top income market segment; Fajgelbaum et al. (2011), Fajgelbaum and Khandelwal (2016) and Faber (2014) study the welfare consequences of trade across the income distribution. Dingel (2016) provides evidence that the higher income residents generate endogenous supply of higher quality products in U.S. cities.

is endogenous. Our approach also uniquely studies spatial sorting and well-being across the full distribution of incomes, with a common non-homothetic preference structure across incomes.⁶ The model’s core mechanisms are drawn from Fajgelbaum et al. (2011) and relate to the assignment model of Davis and Dingel (2020). Our framework retains the tractability of quantitative spatial models, which allows us to take it to the data and quantify the impact of policies on neighborhood change and welfare.⁷ Our paper therefore also complements recent work examining the welfare implications of urban policies by Diamond et al. (2019), Eriksen and Rosenthal (2010), Baum-Snow and Marion (2009), Diamond and McQuade (2019), and Hsieh and Moretti (2019).

Third, we contribute to the literature that studies changes in spatial sorting over time. In an early contribution, Gyourko et al. (2013) show that the increase in high incomes nationally can explain the upward co-movement of incomes and house prices observed in “superstar cities.” Diamond (2016) shows that homophily among the college educated amplifies sorting behavior across cities. Different from Diamond (2016), households have identical non-homothetic preferences in our model. Changes in consumption and location choices stem from changes in income. Furthermore, we model the economics behind the endogenous supply of neighborhoods within a city that fuels increases in welfare inequality. Contemporaneous work also studies welfare inequality within cities. For example, Fogli and Guerrieri (2019) focus on educational outcomes while Su (2018) emphasizes the role of rising value of time for high skilled workers. Our focus on urban amenities as an important dimension of neighborhood heterogeneity follows the early insights of Brueckner et al. (1999) and Glaeser et al. (2001) on the “consumer city.” Lee (2010) studies the role of luxury urban amenities in the sorting of the high-skilled into large cities. Recent work by Almagro and Dominguez-Iino (2019) and Hoelzlein (2019) also study how endogenous amenities reinforce sorting by income within cities.

Fourth, our approach complements a flourishing literature that highlights various causes and consequences of gentrification and neighborhood change.⁸ Within this literature, our paper builds on a growing strand that concludes that amenities play an important role in explaining demographic shifts downtown, relative to changing job locations (Glaeser et al. (2001), Baum-Snow and Hartley (2020), Couture and Handbury (2017), Su (2018)). Couture and Handbury (2017), for example, document rising average commute distance for high-wage workers from 2002 to 2011 despite their moving into downtown areas, and rising propensity to reverse-commute among the rich, i.e., to live downtown but work in the suburbs. These findings illustrate that changing job location or changing

⁶Our paper complements recent work by Tsivanidis (2019) who uses Stone-Geary preferences to study the distributional effects of infrastructure investment in Bogota across two skill groups. In country-wide spatial equilibrium models, Peters et al. (2018) use PIGL preferences of a representative agent to study structural change across U.S. counties and Fajgelbaum and Gaubert (2020) study optimal spatial policies in models with heterogeneous workers who have heterogeneous but homothetic preferences over endogenous city amenities.

⁷Gaigne et al. (2017) theoretically analyze an extension of a classic linear city model with jobs and amenities exogenously given at different locations on the line, in which non-homothetic preferences generates heterogeneous spatial sorting.

⁸See, for example, Guerrieri et al. (2013), Edlund et al. (2019), Ellen et al. (2019), Berkes and Gaetani (2018), Vigdor et al. (2002), Lance Freeman (2005), McKinnish et al. (2010); Ellen and ORegan (2010), Ding et al. (2016); Brummet and Reed (2019), Meltzer and Ghorbani (2017), Lester and Hartley (2014), and Autor et al. (2017).

taste for commutes alone are unlikely to rationalize the rising propensity of high-skilled workers to live downtown. We contribute to this literature by documenting and quantifying a novel channel (rising top incomes, coupled with income effects on location choice), as well as methodologically, by providing a quantitative model that allows for policy assessment.

The rest of the paper proceeds as follows. Section 2 lays out the model and its properties. Section 3 presents the data and provides empirical evidence for the main mechanism of the model, while the model is fully quantified in Section 4. The impact of the 1990-2014 change in income distribution on within-city spatial sorting is presented in Section 5. Section 6 discusses the corresponding normative and policy implications. Section 7 discusses the robustness of our findings to alternative specification.

2 A model of the residential choice of heterogeneous households

2.1 Benchmark Model

We propose a model of residential location in a city. On the demand side, the model resembles a conventional discrete choice of location framework, widely used in the quantitative spatial economics literature, except that it features location choices that are *income-dependent*. On the supply side, the model acknowledges that the urban landscape changes in response to shifts in demand, and in particular to shifts in incomes, by featuring endogenous development of neighborhoods of heterogeneous qualities. Derivations and proofs are given in Appendix B.

2.1.1 Demand for Neighborhoods

The city is populated by households who have heterogeneous income, with a continuum of households at each level of income w . The CDF of the distribution of incomes is denoted $F(\cdot)$ and is taken as a primitive of the model. A key object of interest will be how the city equilibrium changes as the income distribution $F(w)$ changes.

Each household ω makes a discrete choice of a neighborhood n to reside in. Neighborhoods are grouped into four broad types. They are characterized, first, by their geographic region r . Specifically, neighborhoods can be located either downtown ($r = D$) or in the suburbs ($r = S$). Second, they differ in their intrinsic quality q , which captures for instance the attractiveness of their amenities and the quality of their housing stock. Specifically, neighborhoods can be of High ($q = H$) or Low ($q = L$) quality. Within each of the four types (r, q) , there are N_{rq} neighborhoods to choose from. They are assumed for simplicity to be symmetric. Importantly, the number of neighborhoods of each type, N_{rq} , is an endogenous variable in the model, as supply of neighborhoods responds to demand.

Household ω makes a discrete choice of a neighborhood n where to live, trading-off quality of life in a given neighborhood against the corresponding cost-of-living. Specifically, they maximize the following indirect utility:

$$v_n^\omega = (w^\omega - p_n) B_n b_n^\omega. \tag{1}$$

In this expression, w^ω is the income of household ω , p_n the cost of living in neighborhood n , B_n summarizes the quality of amenities in neighborhood n , and b_n^ω is household ω 's idiosyncratic preference for neighborhood n , detailed below. We think of p_n as the cost of housing, although the model can be easily extended to encompass other costs like commuting costs.⁹ Formulation (1) assumes that households have a unit consumption of housing. This simple assumption ensures that lower-income households are more impacted by the high cost of living in attractive neighborhoods than are higher-income households, which leads to spatial sorting by income, as we establish below. Note that, despite this assumption, not all households spend the same amount on housing. Indeed, households with different incomes choose systematically different neighborhoods, which have different house prices.

Households have idiosyncratic preferences for neighborhoods denoted b_n^ω . They are drawn iid from a Generalized Extreme Value (GEV) distribution:

$$G(\{b_n\}) = \exp \left(- \left[\sum_{r,q} \left(\sum_{n \in \mathcal{R}_{r,q}} b_n^{-\gamma} \right)^{-\frac{\rho}{\gamma}} \right] \right), \quad (2)$$

with $\rho \leq \gamma$, where $\mathcal{R}_{r,q}$ is the set of neighborhoods of type (r, q) . The parameter ρ captures the substitutability of neighborhoods of different quality and location types, while γ governs the higher substitutability across neighborhoods of the same type. This structure gives rise to a demand system similar to a nested-logit, popular in the empirical literature on discrete-choice modeling. Households first make a choice over quality and location of neighborhoods (the ‘‘upper nest’’), then over horizontally differentiated neighborhoods within this category (the ‘‘lower nest’’).¹⁰

The share of households who locate in a neighborhood of type (r, q) among households with income w takes the familiar discrete-choice form:

$$\lambda_{r,q}(w) = \sum_{r \in \mathcal{R}_{r,q}} \lambda_r(w) = \frac{V_{r,q}^\rho(w)}{\sum_{r',q'} V_{r',q'}^\rho(w)}. \quad (3)$$

In this expression,

$$V_{r,q}(w) = N_{r,q}^{\frac{1}{\gamma}} B_{r,q} (w - p_{r,q}) \quad (4)$$

⁹We make the implicit assumption that, in a first step that is not modeled, workers find jobs with income w in the city. In a second step, they choose where to live within the city. Heterogeneous commuting costs of neighborhoods are encompassed in expression 1, which can be derived from the more general specification

$$v_n(\omega) = ((1 - \tau_n)w^\omega - \tilde{p}_n) \tilde{B}_n b_n^\omega,$$

where \tilde{p}_n is the cost of housing in neighborhood n and τ_n is the commuting cost to work. Denoting $p_n = \frac{\tilde{p}_n}{1 - \tau_n}$ and $B_n \equiv \tilde{B}_n (1 - \tau_n)$ yields expression (1).

¹⁰See e.g. Fajgelbaum et al. (2011) for a detailed discussion of the properties of the resulting demand system. Note that formulation (2) corresponds to a Fréchet distribution when $\gamma = \rho$, in which case the independence of irrelevant alternatives applies and all neighborhoods are equally substitutes for all households. The nested specification in (2) departs from these restrictions, capturing more realistic patterns of individual preferences. Furthermore, as will be clear from the properties of the model, the parameters ρ and γ govern two distinct economic forces in this framework.

is the inclusive value of neighborhoods of type (r, q) for a household with income w , summing across all the possible choices of neighborhood r of type (r, q) . Importantly, this inclusive value increases with the number of neighborhoods N_{rq} to choose from. This is because of a love-of-variety effect: given the idiosyncratic preferences (2), more neighborhoods to choose from leads to better matches between households and neighborhoods on average, yielding higher utility. Conveniently, this love of variety for neighborhoods enters as an amenity shifter, similar to B_{rq} , in (4). Consequently, we will refer thereafter to $N_{rq}^{\frac{1}{\gamma}} B_{rq}$ as an “amenity composite” for neighborhood type (r, q) . This amenity composite is an endogenous variable in the model, because the provision of neighborhoods is itself endogenous, as we turn to next. This feature is important as it captures the fact that the quality of the urban landscape, beyond housing prices alone, responds to income composition in the city. The intensity of this effect is governed by the parameter γ .

Taking stock, we see from (3) and (4) that, at all levels of incomes, the propensity of a household to reside in a given type of neighborhood depends positively on the quality of its amenities, positively on the variety of neighborhoods of that type the household can choose from, and negatively on its housing cost.

2.1.2 Supply of Neighborhoods

Neighborhoods are supplied by private developers, who choose the quality and location (r, q) of the neighborhood they develop. To develop a neighborhood of type (r, q) , developers pay a fixed cost f_{rq} , and then rent land K_r from local landowners to build $H_{rq} = K_r/k_{rq}$ housing units of quality q in location r . There is free entry of developers into each segment (r, q) .

There are two geographically segmented markets for land: one in D and one in S . In each market, there is perfect competition between atomistic landowners who rent out their land to developers, with the following aggregate supply of land in market r :

$$K_r = K_r^0 (R_r)^{\epsilon_r}, \text{ for } r \in \{D, S\} \quad (5)$$

where R_r is rent in region r , K_r is land supply, and K_r^0 is an exogenous size shifter. The parameter ϵ_r captures the elasticity of land supply in location r . It is allowed to be higher in the suburbs ($\epsilon_S > \epsilon_D$), capturing the notion that it is easier to expand land at the outskirts of the city - through greenfield development and sprawl - than in densely built downtowns.

We allow for housing in high-quality neighborhoods to have a higher area requirement than in low-quality neighborhoods, capturing the fact that housing units are larger in higher quality neighborhoods, and/or that developers devote more space to residential amenities there. Formally, we denote k_{rq} is the unit size of housing in a neighborhood of type (r, q) , and allow for $k_{rH} > k_{rL}$.

The pricing and entry behavior of developers is summarized here, and described in further detail in Appendix B.2. Given unit housing demand and monopolistic competition, equations (3) and (4)

lead to the following housing pricing formula:

$$p_{rq} = \frac{\gamma}{\gamma + 1} k_{rq} R_r + \frac{1}{\gamma + 1} \bar{W}_{rq} \quad (6)$$

where $k_{rq} R_r$ is the unit cost of housing in (r, q) and \bar{W}_{rq} is a measure of demand for a neighborhood of type (r, q) .¹¹ The parameter γ governs the size of markups ($p_{rq} - k_{rq} R_r$) that developers can extract - the lower the γ , the more differentiated neighborhoods are within a type, and the higher the market power of developers.

Finally, through the free entry condition, the number N_{rq} of neighborhoods of type (r, q) adjusts so that the operating profit of a developer in neighborhood type (r, q) just offsets the fixed cost, that is:

$$N_{rq} = \frac{\int_w \lambda_{rq}(w) (p_{rq} - k_{rq} R_r) dF(w)}{f_{rq}}. \quad (7)$$

This equation determines how the supply of neighborhoods of each quality responds to the city-wide income distribution $F(\cdot)$.

Definition 1. An equilibrium of the model is a distribution of location choices by income $\lambda_{rq}(w)$, housing prices p_{rq} , land rents R_r , and number of neighborhoods N_{rq} such that (i) households maximize their utility; (ii) developers and landowners maximize profits; (iii) developers make zero profits; and (iv) the markets for land and housing clear.

Given the structure of the model, it is straightforward to show that an equilibrium can be expressed in terms of changes relative to another reference equilibrium that has different primitives, such as a different city-level distribution of income. We leverage this approach in section 5.¹²

2.2 Equilibrium Properties

We now highlight the main properties of the model. We first show that, in a given equilibrium, the model can capture rich income-based locational sorting patterns. We then turn to a comparative statics exercise and ask: How does the spatial equilibrium change following a change in the city-wide income distribution $F(w)$?

2.2.1 Sorting by income

Our first result is that the model yields residential sorting of households by income.

Proposition 1. *High income households are over-represented in high cost-of-living neighborhoods.*

¹¹ Specifically, $\bar{W}_{rq} = \frac{\int_w (w - p_{rq})^{-1} \delta_{rq}(w) w dF(w)}{\int_w (w - p_{rq})^{-1} \delta_{rq}(w) dF(w)}$ where $\delta_{rq}(w) = 1\{w - p_{rq} > 0\}$.

¹²The model may give rise to multiple equilibria if the agglomeration effects at play are too strong compared to the dispersion forces, driven by the housing supply (in-)elasticity ϵ_r and the idiosyncratic preference for locations-quality types driven by ρ and γ . Around our estimated parameter values, we have not found evidence for such multiple equilibria, suggesting that the calibrated parameters are low enough for equilibrium uniqueness.

Formally, Proposition 1 simply stems from the fact that $\frac{\partial^2 \log V_{rq}(w)}{\partial w \partial p_{rq}} > 0$, or equivalently $\frac{\partial^2 \log \lambda_{rq}(w)}{\partial w \partial p_{rq}} > 0$. Higher income households are less sensitive to house prices than lower income households so that, all else equal, they are over-represented in expensive neighborhoods. This is obtained here by assuming that each household consumes one unit of housing, which yields simple non-homotheticity in consumption, but we would obtain a similar result with any non-homothetic preferences for c and h in which housing is a necessity (see Gaubert and Robert-Nicoud (2021)).¹³

Second, in our setup, the intensity of this sorting by income in the city is crucially shaped by the preference parameter ρ :

Lemma 1. *The intensity of income-based residential sorting increases with ρ , all else equal.*

Indeed, the relative propensity to live in various neighborhood types by income can be written as:

$$\frac{\lambda_{rq}(w)/\lambda_{rq}(w')}{\lambda_{r'q'}(w)/\lambda_{r'q'}(w')} = \left[\frac{(w - p_{rq}) / (w' - p_{rq})}{(w - p_{r'q'}) / (w' - p_{r'q'})} \right]^\rho, \quad (8)$$

so that $\frac{\partial}{\partial \rho} \left(\frac{\partial^2 \log \lambda_{rq}(w)}{\partial w \partial p_{rq}} \right) > 0$: the higher is ρ , the more richer households are over-represented in expensive neighborhoods, all else equal. In that sense, the parameter ρ ends up governing the strength of income effects in location choice. This makes ρ a particularly important parameter to estimate in our quantitative exercise.

Finally, we go back to a key stylized fact that our model aims to capture: the U-shaped propensity to live downtown by income, established empirically in Figure 1. The corresponding object of interest in the model is share of households that lives in $r = D$ at each level of income, that is:

$$\lambda_D(w) = \sum_q \lambda_{Dq}(w).$$

Recalling that, in the model, there are four neighborhood types in the city (DH , DL , SH , and SL), we derive the following result:

Lemma 2. *The share of households living downtown is a U-shaped function of income if the following condition holds:*

$$p_{DL} < p_{SL} < p_{SH} < p_{DH}. \quad (9)$$

The formal proof of Lemma 2 establishes that the share of households living in the most expensive neighborhoods in the city necessarily increases with income, while the share of households

¹³In contrast to what we do here, the literature in economic geography frequently models housing consumption assuming Cobb-Douglas preferences, which deliver a constant expenditure share of housing and a unit income elasticity of housing. This assumption is well suited for models of location choice *across cities with homogeneous workers*, as shown by Davis and Ortalo-Magne (2011). Davis and Ortalo-Magne (2011) compute, city by city, the distribution of expenditure on housing divided by income, and show that the *median* of this distribution is stable *across* cities. This approach is silent on how housing shares vary by income within cities, which our model focuses on. Aguiar and Bils (2015) show using CEX data that housing consumption has an income elasticity that is lower than 1. Our model assumption is better suited than one relying on Cobb-Douglas preferences to capture the empirical fact that the expenditure share on housing decreases with income, as we discuss below.

living in the least costly neighborhoods of the city necessarily decreases with income. When both the most expensive and the least expensive neighborhoods are downtown, as in (9), the propensity of households to live in downtown is a U-shaped function of income. Intuitively, high-income households are attracted downtown by high-quality amenities provided in expensive neighborhoods, while low-income households are attracted downtown by a relatively low cost of living in low-quality neighborhoods.

2.2.2 Effect of a change in income distribution

We have established in equations (3) and (4) that location choices depend on two endogenous neighborhood characteristics: an “amenity composite” ($N_{rq}^{\frac{1}{\gamma}} B_{rq}$), which is endogeneous through N_{rq} - the supply of neighborhoods of a given quality-location - and the cost-of-living (p_{rq}), determined in equation (6). How do these characteristics change following a change in the income distribution in the city?

Suppose that initially, the spatial equilibrium displays the empirically-relevant U-shape pattern of sorting as in Lemma 2. We now ask: what is the effect of a small increase in the relative number of high-income households? Formally, we study a small increase in the relative number of high-income households in the sense of first-order stochastic dominance (FOSD). We consider a class of income distributions indexed by ι , $F_{\iota}(w)$, ordered in the sense of FOSD, that is:

$$\iota > \iota' \Rightarrow F_{\iota}(w) \leq F_{\iota'}(w).$$

Starting at income distribution $F_{\iota}(w)$, we consider an infinitesimal increase in ι . The key sorting mechanism in the model in response to a shift in the income distribution acts through changes in housing prices, themselves driven by changes in land prices (equation (6)), as follows:

Lemma 3. *Given a small increase ι in the relative number of high-income households, downtown land prices increase: $\frac{\partial R_D}{\partial \iota} \geq 0$.*

Through the market for urban land, a demand shock for housing in high-quality neighborhoods pushes up downtown land prices, transmitting the shock to the entire urban area, including in low-quality downtown neighborhoods. The intensity of the price increase is shaped by the housing supply elasticity ϵ_D . A more inelastic supply downtown leads to steeper price increases there, all else equal. Given the impact of neighborhood prices on sorting (see Proposition 1), increased prices downtown tend to increase the share of high income residents and decrease the share of low-income residents there.

This price mechanism is further reinforced by an amplification mechanism, as follows:

Lemma 4. *Given a small increase ι in the relative number of high-income households, the perceived quality of the most expensive neighborhood option increases ($\frac{\partial N_{DH}}{\partial \iota} > 1$), while the perceived quality of the least expensive neighborhood option decreases ($\frac{\partial N_{DL}}{\partial \iota} < 1$).*

This result is quite intuitive: as the number of high-income households increases in the city, there is increased demand for the luxury option DH , which results in an increased supply of such high-quality neighborhoods downtown. Importantly, this endogenous increase in the supply of DH neighborhood makes the DH option even more attractive, because preferences for neighborhoods embed a love-of-variety effect (the strength of which is governed by γ). This fuels further the sorting of richer households downtown, amplifying the price mechanism described above. Because γ governs this love of variety effect, it is a key parameter for our quantitative exercise.

To conclude, shifts in the city-wide income distribution impact not only housing prices but also the urban landscape more generally: neighborhoods endogenously becomes higher quality - a phenomenon that looks like urban gentrification. Our framework therefore formalizes a model for the gentrification of poorer neighborhoods in a city, as a result of demand and supply shifts for these neighborhoods.

2.3 Extensions and Empirical Implementation

Our benchmark model is stylized along several dimensions. Notably, the benchmark model's main mechanisms increase welfare inequality unambiguously in response to top income growth: endogenous supply responses (lemma 4) benefit the rich while hurting the poor through price effects (lemma 3). In this section, we develop model extensions that add nuance to this benchmark model, including mechanisms through which the inflow of richer households downtown may benefit the poor. For simplicity, we review these extensions one at a time, although we incorporate them jointly in a unified quantitative framework in sections 4 and 5.

2.3.1 Publicly-financed amenities

A limitation of the benchmark model is that it does not account for the fact that higher income households moving downtown increase the tax base and hence the provision of public amenities in urban municipalities, whereby benefiting poor incumbent households. To capture this notion, we assume that part of the attractiveness of neighborhoods is driven by public amenities such as parks, public schools, or policing, financed by local governments. They respond to taxes according to:

$$B_{rq} = B_{rq}^o (G_r)^\Omega, \quad (10)$$

where G_r is local government spending in location r , Ω is the supply elasticity of public amenities, and B_{rq}^o captures the part of amenities not determined by local government spending. Governments can levy taxes on local households, summarized by $T_r(w)$ for households with income w . Spending is equal to taxes levied in the location r :

$$G_r = \int \lambda_r(w) T_r(w) dF(w). \quad (11)$$

As the tax base downtown increases, government revenues G_r increase, which raises amenities

for all households downtown, irrespective of the quality of their neighborhood.

2.3.2 Home ownership

The benchmark model is static, hence it cannot capture the important notion that, as neighborhoods downtown gentrify, homeowners in gentrifying neighborhoods reap the price appreciation of their home. We embed this notion in the quantitative model in a simple way, allowing for household income to depend not only on their labor income, but also on the capital appreciation of their real estate portfolio, denoted by $\chi(w)$. Combined with taxes described above, net household income becomes:

$$w + \chi(w) - T_r(w). \quad (12)$$

We allow the real estate portfolio $\chi(w)$ to depend on household type w to capture the corresponding heterogeneity in homeownership rates and initial locations in the data. As land prices in downtown neighborhoods increase, incumbent homeowners receive a capital gain on their housing stock making them better off compared to renters. The extent of this mitigating force is governed by $\chi(w)$ which we discipline empirically by matching homeownership rates by income in different locations.

2.3.3 Change in Job Location

Our focus is to isolate the effect of changes in income on changes in spatial sorting, through non-homothetic preferences for neighborhoods. We have kept the model minimal on other dimensions that likely contribute, in part, to changes in sorting. Change in jobs location is one such dimension. The framework can be extended – at the cost of a more cumbersome exposition – to model the joint choice of workplace and residential location as in Ahlfeldt et al. (2015). This extension would allow us to account for the possibility that changes in job location triggers skill-dependent re-sorting, a different channel for changes in spatial sorting by income. However, we note that current evidence in the literature indicates that spatial job sorting plays little to no role in explaining recent downtown gentrification. First, Su (2018) investigates job urbanization by skill from 1994 to 2010, and finds no difference in job location trends between high- and low-skilled jobs: they both suburbanize somewhat over that time period, at the same rate. Second, Baum-Snow and Hartley (2020) decompose the role of residential demand vs. labor demand in explaining 2000-2010 downtown gentrification. They conclude that labor demand plays little role. Third, Couture and Handbury (2017) conclude that high wage jobs - in the upper third of the wage distribution - do not urbanize fast enough from 2002-2011 to be an important driver of urban gentrification.¹⁴ Given the general conclusion in the urban literature that job location is not an important driver of recent

¹⁴Moreover, data from the U.S. Census shows that commuting times *increased* the most for downtown CBSA residents in the top income deciles between 1990 and 2014 (authors' calculations). This is consistent with the recent literature documenting the increased propensity of high-income households living in urban centers to reverse commute - i.e., to live downtown and work in the suburbs.

urban gentrification, and given also the limited impact that commuting costs ultimately have on our quantification, we have refrained from complexifying the model in that dimension.

The mechanisms described in this subsection mitigate the adverse effect of an increase in income inequality on welfare inequality, while endogenous neighborhood supply amplify the baseline price effect. The net effect of increased incomes of the rich on welfare inequality through spatial sorting is therefore ambiguous. We turn next to its quantification.

3 Evidence on Spatial Sorting by Income

In this section, we provide empirical evidence in support of the income-based sorting mechanism at the heart of our model. The strength of this sorting mechanism pins down the parameter ρ , which will be important for the quantitative exercise that follows.

3.1 Mapping Model to Data

Our empirical work requires data on household location decisions and housing costs mapped into the spatial units employed in the model. Herein, we describe this mapping and summarize our main data sources. The full exposition of the data work is detailed in Appendix C.

Spatial Units and Classifications We equate the notion of a city to a Core-Based Statistical Area (CBSA), and that of a neighborhood to a census tract. All data are interpolated to constant 2010-boundary tracts and 2014-boundary CBSAs using the Longitudinal Tract Data Base (LTBD).

Within each CBSA, we define downtown as the set of tracts closest to the center of the CBSA that accounted for 10 percent of the CBSA’s population in 2000.¹⁵ This spatial boundary of downtown is constant across all years. We refer to all remaining non-downtown tracts in each CBSA as *suburban*, so that tracts are either classified as downtown (*D*) or suburban (*S*). Appendix A.1 features maps showing which tracts in New York, Chicago, Philadelphia, San Francisco, Boston and Las Vegas are classified as downtown and suburban.¹⁶

We define neighborhood quality using residential demographic composition. We draw from Diamond (2016), who shows that the college-educated share can proxy for endogenous amenities. Specifically, we classify a neighborhood to be high quality if at least 40 percent of residents between the ages of 25 and 65 have at least a bachelor’s degree. Under this definition, 15, 22, and 32 percent of census tracts in the top 100 CBSAs are respectively classified as high quality in 1990, 2000, and 2014.

Income and House Prices Our data on location choice by household income is from the 1990 Census and the 2012-2016 American Community Surveys (henceforth referred to as the “2014

¹⁵We use the CBSA centers from Holian and Kahn (2012), who identify the coordinates returned by Google Earth for a search of each CBSA’s principal city.

¹⁶We explore the robustness of our main results to alternative definitions of the downtown area in the appendix.

ACS”). For the analysis in this section, we use census tables that report the number of households residing in each census tract by income bracket. The boundaries of census brackets change over time so, to make inter-temporal comparisons of residential choices by income, we define alternative income brackets that are constant in real terms over time. We summarize each constant bracket m by the median income w_m of households within the bracket.¹⁷ We drop brackets with median annual income smaller than \$25,000; given the presence of public housing, such households are not well represented by the model. In the CBSA-level analysis herein, $\lambda_{rq,ct}(w_m)$, denotes the share of households in income bracket m that reside in tracts of area-quality (rq) , out of all households in that CBSA c at time t .

We use data from the National Historical Geographic Information System (NHGIS) 1990 Census tables to measure 1990 house prices by (r, q) -pairs. Specifically, we start with data on median house prices within each census tract, and compute a population-weighted median over all tracts within a given neighborhood type (i.e., within a given (r, q) -pair). We use these 1990 housing prices to calibrate our baseline model.

Our estimation of ρ , meanwhile, relies on changes in house prices between 1990 and 2014. The cross-sectional nature of the Census data does not allow us to measure the change in housing costs for the same set of housing units within a census tract over time. So we instead use *Zillow’s 2 Bedroom Home Value Index* when measuring changes in housing prices over time within an area-quality pair (r, q) . The Zillow neighborhood house price indices are specifically designed to capture house price changes over time for a constant set of housing units within a given neighborhood. Since house prices from Zillow are not available prior to 1996, we proxy for Zillow house prices in our initial period (1990) using data for years 1996-1998.¹⁸ We pool prices for years 2012-2016 for our end period (2014). Finally, the Zillow indices measure neighborhood house price changes at the level of zipcodes which are larger than census tracts. We use crosswalks provided by HUD to map zipcode to census tracts. We then aggregate tracts to area-quality pairs.¹⁹ In the CBSA-level analysis herein, for example, $p_{rq,ct}^h$ denotes the annual user cost of a median priced 2 bedroom house in area-quality pair (r, q) within CBSA c at time t .

All nominal variables including income and house prices are deflated to 1999 dollars using the urban CPI. House prices are converted into an annual user cost using ratios of 5.0 percent in 1990 and 4.6 percent in 2014 from the Lincoln Institute of Land Policy.

¹⁷We re-allocate households to these constant brackets assuming a uniform income distribution within each original bracket. For each interior bracket, w_m is equal to its mid-point. We assign the top bracket the median income of that bracket in the 2000 IPUMS microdata. A full discussion of how we deal with topcoding in this data can be found in Appendix C.

¹⁸House prices were relatively flat over the 1990 to 1995 period suggesting that this measurement issue is unlikely to bias our results in any meaningful way.

¹⁹Given the inherent difficulty of measuring house prices at a small spatial scale in both levels and changes over time, we perform an extensive set of robustness exercises exploring the sensitivity of our key parameter estimates to alternate house price data. In particular, our online appendix discusses the replication of estimates of ρ using house price data from both Census and Zillow, as well as variants of these indices. Our estimates of ρ are similar across a variety of different house price series.

3.2 Changes in the Income Distribution and Residential Sorting: Evidence

We now provide empirical support for the sorting mechanism in the model. Our analysis uses data for 1990 and 2014 and focuses on the 100 CBSAs with the largest 1990 population.

3.2.1 Descriptive Evidence

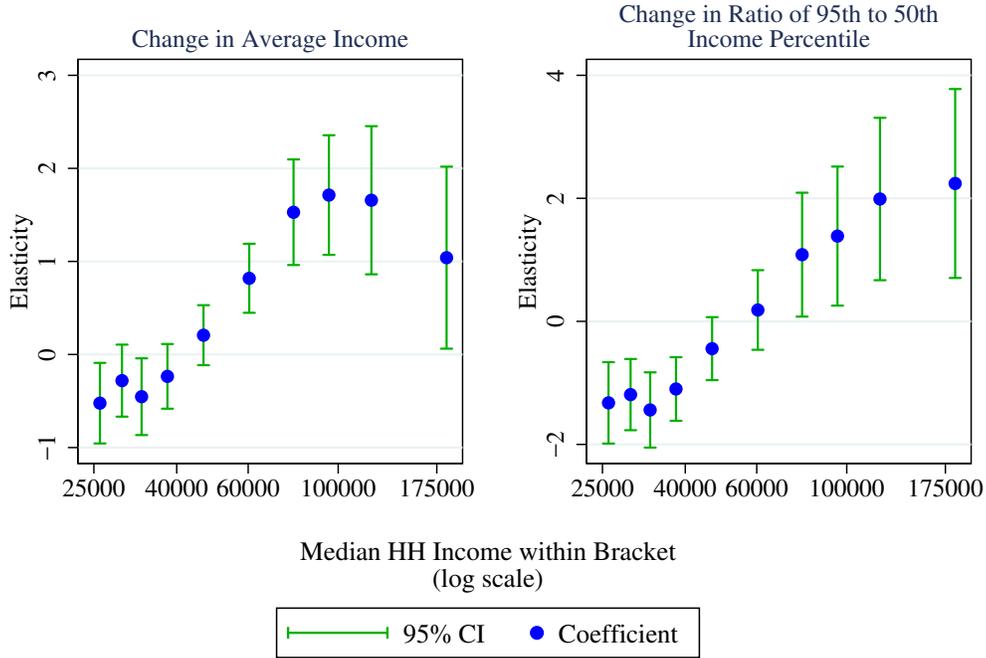
This analysis aims to test whether changes in a city’s income distribution generate different spatial sorting patterns for different income groups. As the city gets richer, our model predicts that richer households move downtown (“urbanize”) relative to the average household, while poorer households move out (“suburbanize”). A first look at the data supports this prediction. We show that different income groups display systematically different spatial re-sorting patterns in cities that exhibit different city-wide income growth patterns. We use cross-city variation to estimate the correlation between city-wide income growth and spatial sorting for each income bracket m with the following regression:

$$\Delta \ln \left(\frac{\lambda_{D,c}(w_m)/\lambda_{D,c}}{\lambda_{S,c}(w_m)/\lambda_{S,c}} \right) = \alpha^m + \beta^m \Delta \text{Income}_c + \nu_c^m, \quad (13)$$

where $\Delta \left(\frac{\lambda_{D,c}(w_m)/\lambda_{D,c}}{\lambda_{S,c}(w_m)/\lambda_{S,c}} \right)$ is the change in the propensity of households at income level w_m to reside downtown relative to all households in CBSA c between 1990 and 2014. ΔIncome_c is one of two measures of income growth in c during the same time period. The first measure is simply CBSA average income growth. We could observe average income growth from a neutral rise in income. The second measure is designed to specifically capture top income growth, as the change in the ratio between the 95th percentile and the median of the CBSA income distribution. Note that $\beta^m > 0$ implies that income group m becomes more likely to reside downtown (relative to other income groups) in CBSAs that experienced more income growth across all areas.

Figure 2 plots point estimates and 95 percent confidence bands for the β^m coefficients from regressions (13). The left-hand plot shows that richer households urbanized more, and poorer households less, in CBSAs that experienced more average income growth. The right-hand plot shows even stronger patterns of urbanization of rich households and suburbanization of poor households in CBSAs with more top income growth. A 10 percent increase in the 95/50 ratio of incomes is associated with a 20 percent increase in the propensity of households earning more than \$150,000 to live downtown relative to the average household, and a 10 percent decrease in the corresponding propensity for households earning less than \$40,000. These cross-CBSA correlations suggest that shifts in the income distribution may be a quantitatively important factor in explaining the evolution of sorting by income from 1990 to 2014 shown in Figure 1, i.e., in explaining the relative urbanization of the rich and suburbanization of the poor.

Figure 2: Correlation between Change in Propensity to Live Downtown and Change in CBSA Income Distribution by Household Income between 1990 and 2014



Note: This figure shows income bracket-specific coefficients, along with 95 percent confidence intervals, for regressions of the propensity of households in that income bracket to reside downtown in a CBSA against a measure of CBSA income growth. Each income bracket specific regression comes from exploiting cross-CBSA variation with observations weighted by CBSA population. The independent variable in the left panel is the change in average CBSA income while the independent variable in the right panel is the change in the ratio of the 95th to the the 50th percentile of household income in the CBSA. The y-axis shows the income bracket specific regression coefficients and the x-axis shows median income within each income bracket. All changes are from 1990 to 2014.

3.2.2 Model-Consistent Estimation

The descriptive results above suggest that income growth and sorting by income are associated across CBSAs, but it still remains to show that this link is consistent with the mechanism proposed in Section 2, and plausibly causal. We now fill this gap, testing the key sorting mechanism in the model with an instrumental variable strategy that addresses factors that confound the link between income growth and spatial re-sorting.

Estimating Equation Equation (8) in the model is key for what we do next. It shows that household location choice is a simple function of disposable income (income net of housing prices). In particular, equation (8) implies that any changes in sorting by income are driven solely by changes in housing prices p_{rq} , i.e., by the higher willingness of richer households to pay the higher cost of locating in more attractive neighborhoods. Put differently, changes in house prices p_{rq} capture all there is to know about how a change in the income distribution in the city will impact

sorting. Our empirical strategy is therefore to measure how, across cities, changes in house prices, driven by changes in income growth, trigger spatial re-sorting by income. To do so, we derive the following estimating equation from (8), interpreting different time periods and different cities as different equilibria of the model and taking log differences across two time periods and two pairs of income groups m and m' for each CBSA c :

$$\Delta \ln \left(\frac{\lambda_{Dq,c}(w_m)}{\lambda_{Sq,c}(w_m)} \right) - \Delta \ln \left(\frac{\lambda_{Dq,c}(w_{m'})}{\lambda_{Sq,c}(w_{m'})} \right) = \rho \left[\Delta \ln \left(\frac{w_m - p_{Dq,c}}{w_m - p_{Sq,c}} \right) - \Delta \ln \left(\frac{w_{m'} - p_{Dq,c}}{w_{m'} - p_{Sq,c}} \right) \right] + \varepsilon_{c,q,(m,m')} \quad (14)$$

Normalizing brackets m and m' such that $w_m > w_{m'}$, the dependent variable is positive if the richer income bracket urbanizes faster than the poorer income bracket. The independent variable is positive when house prices rise faster downtown than in the suburbs. Therefore, if the rich urbanize faster than the poor in response to house price growth downtown relative to the suburbs, then ρ will be positive, consistent with the model. In contrast, if sorting patterns were unrelated to income, we would estimate $\rho = 0$. Estimating (14) therefore provides a test of our income sorting mechanism.²⁰ Finally, note that equation (14) is triple differenced - by time, area, and income. So our estimate of ρ is robust to omitted variables that are time-invariant, or that affects D and S equally, or that affect every income group equally. There may be, however, confounding factors that are time-varying, income-biased, downtown-specific, and not captured by our measure of quality. We now turn to describing these threats to identification in more detail.

Identification Beyond the sorting mechanism that we aim to estimate, there may be other local shocks that drive both sorting and housing price growth across cities, confounding identification. Note that any local shocks that are valued equally by all incomes, or whose impact is captured by a switch in quality level q , will not compromise identification. These shocks are captured by B_{rq} in the model and controlled for by differencing across income groups m and m' within a quality tier in a CBSA. Similarly, income-specific shocks that affect the attractiveness of both downtown and suburban neighborhoods of a given quality tier in a CBSA will not compromise identification. These shocks are captured by the common utility shifter ($\sum_{r',q'} V_{r',q'}^\rho(w)$) in the model and controlled for by differencing between the downtown D and suburban S areas of each CBSA.²¹ To confound identification, shocks need to be both biased towards either downtown or the

²⁰Without micro panel data we cannot observe how the location choice of a given household changes as their income w changes. We instead estimate (14) using 45 distinct income-bracket pairs of w_m and $w_{m'}$ for each quality tier j . We pool our estimation across income pairs and organize the data such that $w_m > w_{m'}$. As a result, the maximum number of observations in each of our regressions is 9000, though in many specifications, we have less given missing data at the CBSA-area-quality triplet. We also remove any observation with $w - p_{crq}^h < 0$ (1.4% of our sample). We then censor the top and bottom 1 percent of $\ln \left(\frac{w - p_{Dq}^h}{w - p_{Sq}^h} \right)$ in each year.

²¹Note that taking a difference across two time periods (1990 to 2014) is not necessary from the perspective of the model. Unlike the other two differences (across income groups and across D, S), the time difference does not remove a systematic unobserved model component. We prefer to estimate the model in changes rather than in the cross-section as this estimation strategy illustrates the main sorting mechanism in our model.

suburbs and income-specific, and not controlled for by our quality levels. Using the terminology of the model, such a shock would make the attractiveness of neighborhood of type rq income specific, i.e., equal to $B_{rq}(w_m)$ instead of B_{rq} , thus introducing a systematic error term (beyond measurement error) into equation (14). Denoting the unobserved income-specific component of neighborhood attractiveness as $\epsilon_{rq}(w_m) = B_{rq}(w_m) - B_{rq}$, the error in equation (14) is equal to $\Delta \ln \left(\frac{\epsilon_{Dq,c}(w_{m'})}{\epsilon_{Sq,c}(w_{m'})} \right) - \Delta \ln \left(\frac{\epsilon_{Dq,c}(w_m)}{\epsilon_{Sq,c}(w_m)} \right)$. Such shocks could include downtown-biased growth in amenities by local city planners that are valued more by the high-skilled, like private schooling, luxury retail, and proximity to high-skilled jobs, or the decline in central city violent crime since 1990 (e.g., Levitt 2004), which may also be valued more by the high-skilled (Ellen et al., 2019). If these factors make downtowns more attractive to high-income households and drive house prices up downtown relative to the suburbs, they would bias our estimate of the coefficient ρ upwards.

To overcome this identification challenge, we use an instrumental variable strategy. We instrument for relative changes in house prices using an idea closely related to our theory. First note that the housing supply elasticity is lower downtown (both by assumption in our model and in the recent estimates from Baum-Snow and Han 2019). So, CBSA-level income growth will generate more house price growth downtown than in the suburbs; we verify that this is the case empirically below. This suggests instrumenting the difference in house price growth between downtown and the suburbs in equation (14) using a plausibly exogenous CBSA-level income shock. We implement this idea using a shift-share (Bartik) shock to CBSA per capita income. The Bartik shock predicts the change in CBSA average earnings by projecting trends in industry-level average earnings observed elsewhere in the country on each CBSA’s initial industry mix.

The exclusion restriction is that proposed in Borusyak et al. (2021): industry shocks need to be conditionally exogenous, in the sense that they are uncorrelated with the income- and downtown-biased error term described above. Specifically, we assume that industries that experienced higher national wage growth were not initially disproportionately located in CBSAs where downtowns gained skill-biased amenities relative to the suburbs.

This exclusion restriction may be violated if the industries that experienced higher national wage growth are themselves both downtown- and skill-biased. For example, if tech firms employ high-skilled individuals and are initially over-represented downtown, then national wage growth in tech could attract high-income individuals downtown. To address this concern, we show that our results are robust to excluding various sets of industries that are either downtown- and/or skill-biased. First, we exclude the top quartile of downtown-biased industries from the computation of our Bartik shock. Specifically, we remove industries in which residents of urban areas are most likely to work.²² This isolates wage growth in suburbanized industries to instrument for relative house price growth downtown. Second, we recompute a Bartik instrument leaving out tech industries and then separately finance, insurance, and real estate (FIRE) industries. These industries disproport-

²²To do so, we first rank industries by the share of their workers that lives downtown. Then, starting from the industry with the most urbanized workers, we remove industries entirely from our Bartik computation until 25 percent of all workers have been removed. We then renormalize CBSA-level industry shares so they are relative to total CBSA employment excluding these downtown-biased industries.

tionately employ higher skilled workers. As we highlight below, all three alternative instruments yield similar results to our main instrument. We interpret this as further evidence, consistent with recent research, that access to jobs is not driving the recent in-migration of high-income individuals into urban centers.

Finally, we conduct pre-trend and balance tests similar to those proposed by Borusyak et al. (2021), and find no evidence of pre-trends or balance violation. A full discussion of these results can be found in Appendix D.

Identifying Variation Before estimating ρ using equation (14), we illustrate the variation in the data that allows for identification. We first verify that, in line with the logic above, the Bartik income shock raises house prices more downtown than in the suburbs. To illustrate this variation, we plot our Bartik shock between 1990 and 2014 for each CBSA (on the x-axis) against $\Delta \ln(p_{Dq,c}^h/p_{Sq,c}^h)$ (on the y-axis) in the left panel of Figure 3. There are 200 observations in the figure: 2 quality tiers within each of our 100 CBSAs. We find that within each quality tier, a more positive income shock raises housing prices downtown relative to the suburbs. This variation underlies the significant first-stage statistic in our estimation of ρ below.

Next, we report the reduced-form regression of change in spatial sorting directly on the Bartik shock. To simplify the presentation, we pool quality tiers and show the results of the following regression for each of our 10 income brackets:

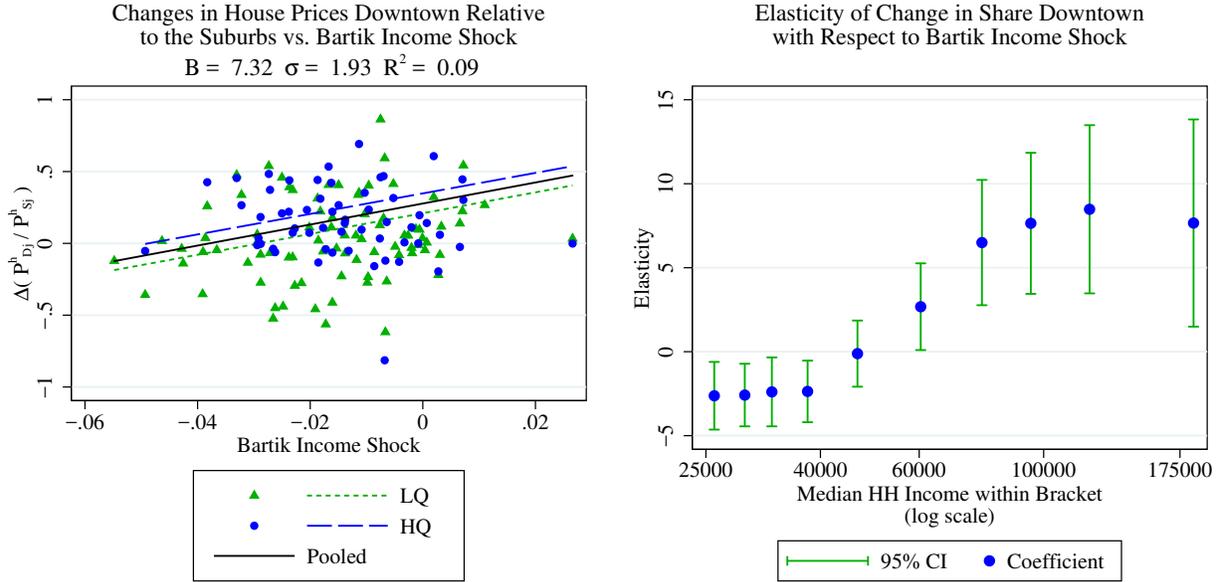
$$\Delta \ln \left(\frac{\lambda_{D,c}(w)/\lambda_{D,c}}{\lambda_{S,c}(w)/\lambda_{S,c}} \right) = \alpha^w + \beta^w \widehat{\Delta \text{Income}_c}^{\text{Bartik}} + \nu_c^w. \quad (15)$$

This regression is exactly the same as our descriptive regression (13), except that the Bartik income shock replaces actual income growth. In equation (15), $\beta^w > 0$ implies that following a positive CBSA Bartik shock, the propensity of income group w to live downtown rises relative to that of the average CBSA resident. The right panel of Figure 3 reports estimates from equation (15), along with their 95 percent confidence bounds. We find that a CBSA income shock causes differential spatial sorting responses from the rich vs. the poor. In particular, rich households are more likely than poor households to move downtown in response to an income shock. For all the top five income groups, $\beta^m > 0$ and all estimates are statistically significant at the 5 percent level. Conversely, all the bottom five income groups have estimates of $\beta^m < 0$, with all but the middle income group estimate being statistically significant.

To summarize, Figure 3 provides reduced-form evidence consistent with the key mechanism in our model. As CBSA income increases, house prices grow faster downtown, and richer households are more likely to re-sort downtown relative to poorer households.

Estimates of ρ Having clarified the variation that identifies (14), we are now ready to estimate ρ – the parameter that governs the intensity of income-based sorting. Our baseline results are reported in the first two columns of Table 1. Column 1 shows an OLS estimate, and column 2 shows an IV estimate using the Bartik instrument. We weight both OLS and IV regressions

Figure 3: Identifying Variation for ρ



Note: On the left, we plot changes in downtown relative to suburban house prices within each quality tier on the y-axis, against the Bartik income shock between 1990 and 2014 on the x-axis for each of the largest 100 CBSAs. The house price data is from the Zillow 2 Bedroom Index in 1996-1998 and 2012-2016. We drop the top and bottom 1% of $\Delta(p_{Dq,c}^h/p_{Dq,c}^h)$ from the plot. On the right, we show income bracket-specific coefficients, along with 95 percent confidence intervals, from equation (15) on the y-axis (regression of Bartik income shock on changes in normalized urban share from 1990 to 2014), against median income within each income bracket on the x-axis. Both panels show CBSA population weighted-regression coefficients

by the number of observations in each cell. This downweights cells with fewer individuals where measurement error may be higher.

Our OLS estimate is somewhat lower than our IV estimate (2.48 vs. 3.04), but we cannot reject that they are the same. Our instrument has strong first stage predictive power with an F-stat of 26. Columns (3)-(5) show IV estimates for alternative Bartik shocks that exclude urbanized, FIRE, and Tech industries. These estimates are similar to our base results, ranging from 2.3 to 3.2. This suggests that our instrument is not correlated with labor demand shocks that are concentrated in urban centers of CBSAs and that disproportionately affect high-income households. In Appendix D we show further robustness of our estimates of ρ to many different specifications, including different time periods, different house price measures, different definitions of downtown area, and different quality cut-offs. These robustness estimates are almost all within two standard error bands of our preferred estimate.

To summarize, we use our preferred IV estimate in column 2 and set $\rho = 3.0$ in our model calibration. As we show later, ρ is an important parameter determining our welfare results. In our counterfactual exercises we show the sensitivity of our results to alternate values of ρ between 1.5 and 4.5 which encompass roughly the two standard deviation bands of our estimate in column 2.

As an additional robustness exercise, we show that residential amenities – as measured by restaurant quality – also increased more in downtown areas relative to the suburbs in response to CBSA income growth. The endogenous response of amenities to the changing income distribution

Table 1: Estimation of elasticity ρ

	(1)	(2)	(3)	(4)	(5)
$\hat{\rho}$	2.48 (0.34)	3.04 (0.70)	2.69 (0.76)	2.26 (0.76)	3.20 (0.87)
Instrument	None	Base	Omit Top Urban Industries	Omit FIRE Industries	Omit Hi-Tech Industries
R ²	0.21				
KP F-Stat		26.3	16.0	15.6	19.0
Obs	5,878	5,878	5,878	5,878	5,878

Notes: This table shows estimates from equation (14). Data from 100 largest CBSAs where neighborhood quality defined from education mix of residents. Each observation is weighted by the number of households in the income bracket. Column 3 to 5 also control for share of omitted industries. KP F-Stat = Kleibergen-Paap Wald F statistic. Standard errors clustered at the CBSA-quality level are in parentheses for column (1)-(5).

is a key amplification mechanism in our model. To conserve on the space needed to introduce our restaurant quality index, we relegate these results to Appendix E.

4 Model Quantification

Having established the empirical relevance of the key model mechanism and provided a micro-based estimate of ρ , we now turn to quantifying the remaining structural parameters of the model. We do so in two stages. In a first stage, we quantify the model elasticities. In a second stage, we use the method of moments to calibrate the parameters of the model that govern the levels of prices and amenities, conditional on model elasticities. We refer the reader to the online appendix for the full details of the procedure.

4.1 Model Parametrization

Table 2 lists the model’s seven elasticities that we estimate or calibrate in the first stage. This includes ρ , whose estimation was described above. We discuss the calibration of each of the remaining parameters briefly here. The role played by these parameters in driving sorting patterns and welfare results is discussed in detail in Section 5. All of the parameters selected here are for our baseline calibration. We explore robustness over a range of values for each parameter in Section 7.

Land Supply Elasticities (ϵ_S and ϵ_D) In the model, the area-specific elasticity of land supply (ϵ_r) directly translates into an elasticity of housing supply. We calibrate ϵ_D and ϵ_S to match the Saiz (2010) housing supply elasticity estimates for cities that have an average household density similar to that in our representative downtown and suburban areas. This yields $\hat{\epsilon}_D = 0.60$ and

Table 2: Key Model Parameters

Parameter	Description	Value	Source
<u>Non-homotheticity</u>			
ρ	Between-type neighborhood substitution elasticity	3.0	Estimation
<u>Land Price Responses</u>			
ϵ_D	Downtown land supply elasticity	0.6	Calibrated to Saiz (2010)
ϵ_S	Suburban land supply elasticity	1.3	Calibrated to Saiz (2010)
<u>Amplification</u>			
γ	Within-type neighborhood substitution elasticity	6.8	Literature
<u>Other</u>			
T_D	Downtown local property tax	0.2	IPUMS 2000
T_S	Suburban local property tax	0.3	IPUMS 2000
Ω	Public amenity supply elasticity	0.05	Literature

$\hat{\epsilon}_S = 1.33$. These numbers are roughly similar to the recent within CBSA housing supply elasticities estimated in Baum-Snow and Han (2019).

Key Amplification Parameter (γ) Income sorting (governed by ρ) is amplified in the model by endogenous neighborhood development. The intensity of this effect is governed by the shape parameter γ , which controls gains from variety in residential neighborhood choice. Ahlfeldt et al. (2015) estimate a within-type neighborhood substitution elasticity of 6.8 using detailed micro data from Germany. For our baseline calibration, we use the Ahlfeldt et al. (2015) estimate and set $\gamma = 6.8$. Given that the elasticity of substitution across neighborhoods of a given quality estimated in Germany may not map exactly to the preferences of Americans, we show the sensitivity of our results to alternate values of γ .

Public Amenities (T_D, T_S, Ω) We calibrate local taxes (T_n) to match the unit-level average real estate taxes paid as a share of annualized housing costs in 2000, using tract-level data from the 2000 Census. This implies a local property tax rate as a fraction of the annual user cost of housing of 30% in the suburbs and 20% downtown. We set the elasticity of the endogenous component of the public amenity with respect to these tax revenues (Ω) to 0.05 (Fajgelbaum et al., 2018).

Homeownership ($\chi(w)$) In our benchmark model, we assume that all housing rents in the city (land rents and fixed costs of development) accrue to an absentee landlord and none are transferred to the city residents, i.e., that $\chi(w) = 0$ for all w . In our counterfactual analysis, however, we want to be able to account for the heterogeneous rate of home ownership in contributing to spatial sorting responses, in order to allow households who own their home to reap the benefits of rising

house prices. To do so, we discipline $\chi(w)$ by transferring to households at each labor income level capital gains corresponding to their average real estate portfolio. This transfer equals the average house price growth in the neighborhoods where households of that income lived in the previous period, which is then scaled by the share of households who were homeowners according to the 2000 IPUMS data. Empirically, this share of home ownership increases systematically with labor income.

4.2 Second Stage: Method of Moments

Armed with estimates for the key elasticities of the model, we conclude the calibration of the model using a method of moments to estimate two key vectors of composite model parameters: (i) the relative amenity composite of each neighborhood type $N_{rq}^{\frac{1}{\gamma}} B_{rq}$, and (ii) the price of housing in each neighborhood type p_{rq} , which together pin down the calibrated values for location choices $\{\lambda_{rq}(w)\}$ at all levels of income w in the baseline equilibrium. The procedure does not separately identify all of the structural parameters of the model that shape these composites. But these composite parameters are just sufficient to compute any counterfactual equilibrium of the model.

We target two sets of moments that summarize the key economic concepts we aim to capture: (i) the 1990 distribution, by income level, of the share of workers living downtown (i.e., the U-shape sorting patterns presented in the introduction), and (ii) the 1990 level of house prices by neighborhood type. To accurately capture the location choices of higher-income households, we target the downtown share of households at a finer income grid than the Census income brackets represented in the introduction. To this end, we construct the same plot as Figure 1 but for finer \$5,000 income brackets (in 1999 dollars) using the micro IPUMS data. The additional detail in the income dimension comes at the expense of precision in the spatial dimension and, as a result, we are limited to studying 27 CBSAs of our original 100 in the calibration and counterfactual exercises.²³ We perform this calibration and counterfactuals for a representative city that is an average of these 27 CBSAs. The U-shape patterns of residential sorting for these 27 CBSAs are very similar to the U-shape patterns documented in Figure 1.

The identification of the model in this second stage is quite straightforward. First, house price moments directly inform the calibration of p_{rq} . Then, conditional on prices, the U-shape pattern of the location choice data helps identify the relative attractiveness $N_{rq}^{\frac{1}{\gamma}} B_{rq}$ of different types of neighborhoods, by a revealed preference approach applied to our non-homothetic demand function: the same level of price and quality of a neighborhood generates different demand patterns at

²³The IPUMS data identifies the locations of respondents at the PUMA (Public Use Microdata Area), each of which contains approximately 100,000 individuals, relative to the 4,000 contained in each Census tract. To replicate the urban share for each fine income brackets, we first construct a cross-walk between PUMAs and our tract-based downtown areas. There are 27 CBSAs in which PUMAs are small enough relative to the downtown definition so as to allow for useful inference here. See Appendix C for more details.

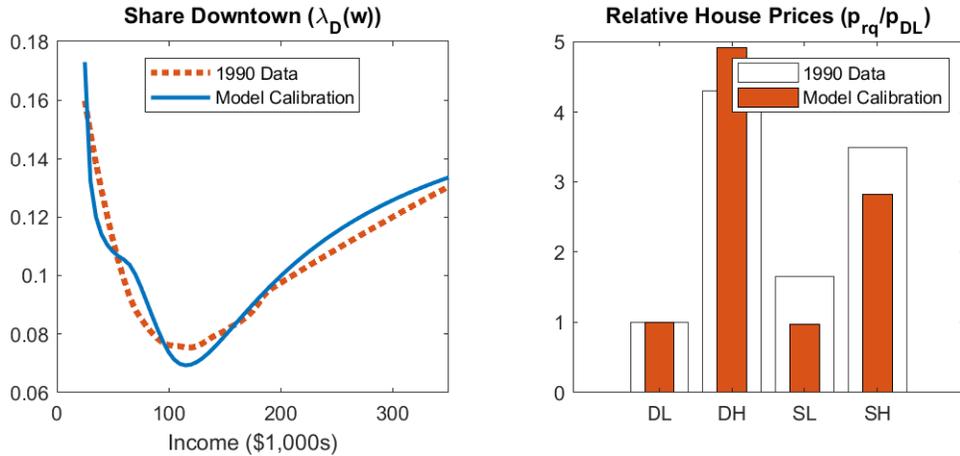
different levels of income. Concretely, the identification relies on the equation:

$$\frac{\lambda_D(w)}{\lambda_S(w)} = \frac{\sum_{q=L,H} N_{Dq}^{\frac{\rho}{\gamma}} B_{D,q}^{\rho} [w - p_{D,q}]^{\rho}}{\sum_{q=L,H} N_{Sq}^{\frac{\rho}{\gamma}} B_{S,q}^{\rho} [w - p_{S,q}]^{\rho}}.$$

Given w , the calibration backs out the $N_{rq}^{\frac{1}{\gamma}} B_{rq}$ and p_{rq} that allow to best match the distribution of location choices in the data. The vectors are pinned down up to a normalization level, whose value does not impact the counterfactuals done in the following section.

Moment Fit. The moment fit is presented in Figure 4. Since the model is over-identified, neither moment can be matched perfectly. The procedure trades-off a better fit of the U-shape for location choices against a better fit for housing prices. Despite a sparse specification, the calibrated model is able to match the rich non-monotonic U-shape patterns of location choice by households of various incomes remarkably well. The model also matches the relative housing prices between downtown and suburban high and low-quality neighborhoods. In 1990, the model and Census data house prices in low-quality downtown neighborhoods and in the suburbs are close. Both the model and data have high-quality suburban neighborhoods having housing prices being about 3 times higher than low-quality suburban neighborhoods. Importantly, both the model and the data have downtown high-quality neighborhoods being between 4 and 5 times higher than low-quality downtown neighborhoods.

Figure 4: Calibration to 1990 Urban Shares and Neighborhood Prices



Notes: These figures show the fit of the calibrated model to the two targeted moments. The left-hand plot shows the share of households in each \$5,000 income bracket that reside downtown in 1990. The dashed line shows the data, while the solid line shows the prediction of the calibrated model. The data are constructed from micro IPUMS data and reflect the propensity to reside downtown by income in the 27 CBSAs in which PUMAs (the finest spatial unit the IPUMS data) are small enough relative to the downtown definition to make useful inference here. The curve is interpolated to address top-coding in the IPUMS data. See Appendix C for more details. The clear bars in the right-hand plot shows the average Census house price in tracts of each location-quality type, normalized by the average index in low-quality tracts downtown, in 1996. The solid red bars show the predicted relative housing costs predicted by the calibrated model.

5 Income Growth and Changing Spatial Sorting

Armed with our quantified model, we now turn to our main question of interest. Using counterfactual analysis, we gauge the extent to which a change in the income distribution in the city ($F(w)$) can help rationalize the observed changes in spatial sorting within the city. We start by analyzing the 1990-2014 period before turning to 1970-1990 and 1950-1970.

Given the structure of the model, a counterfactual equilibrium can be computed with the following parameters on hand, as detailed in Appendix G: (i) the model elasticities $\{\rho, \gamma, \epsilon_r, \Omega\}$, and (ii) the initial equilibrium values for population in each neighborhood type and house prices as calibrated above. Given that the model is over-identified, the baseline model matches the 1990 data imperfectly. We treat the log-differences between data and model as measurement error, and hold it constant across periods when we conduct counterfactuals.

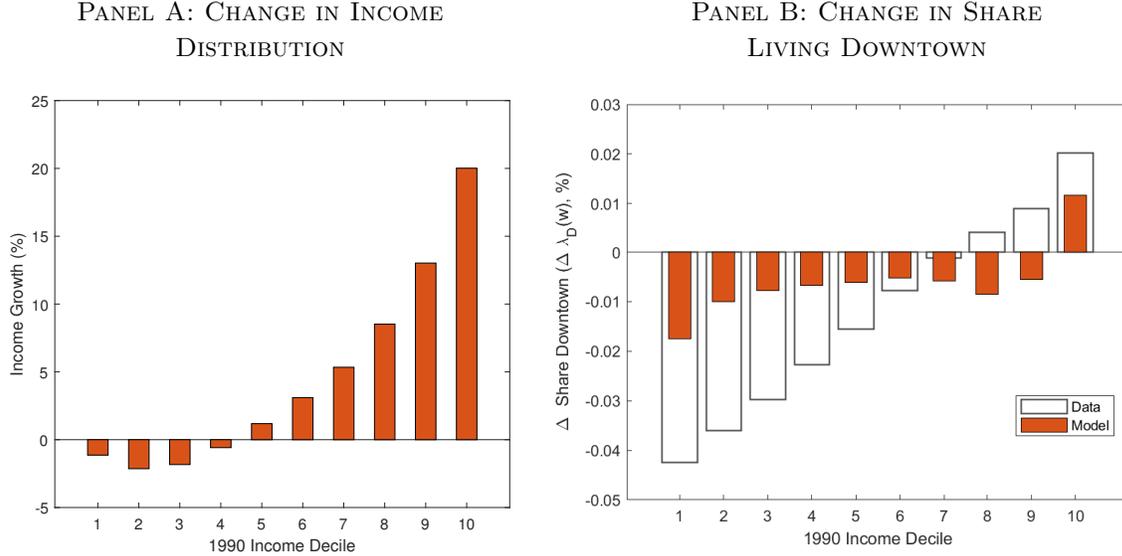
5.1 Baseline counterfactual: 1990-2014 change in income distribution

Between 1990 and 2014, the income distribution of the largest CBSAs became more unequal, mirroring the patterns documented for the economy as a whole. Panel A of Figure 5 summarizes this change plotting the percentage change in income between 1990 and 2014 for each income decile in the representative city made of 27 large CBSAs that we used to calibrate the model above. Inflation-adjusted income per capita grew on average by 10%. For the bottom decile, however, income actually fell slightly by approximately 1 percent, while for the top decile, income increased by about 18 percent. Overall, the 90-10 income gap widened by 19 percentage points.

How much did this change in income distribution, in isolation, contribute to changes in spatial sorting within cities? We use the quantified model to answer this question. We compute the counterfactual spatial equilibrium that corresponds to the 2014 income distribution, without changing any other parameters of the model. We then compare sorting in this model-based counterfactual to that in the actual spatial equilibrium in 2014. In Panel B of Figure 5, the clear wide bars show the actual empirical change in the propensity to live in downtown areas between 1990 and 2014 for each decile of the income distribution, summarizing the shift in the U-shape of Figure 1. The skinnier solid red bars show the changes predicted by the model in response to the shift in the income distribution.

In the model, the 1990-2014 change in the income distribution generates a shift in location choices that matches the general trend we observe in the data. High income households move downtown, while low-income households move out of downtown. The predictive power of the income shock alone is substantive: it explains about 40 percent of the suburbanization of the bottom decile of the income distribution, and about 60 percent of the urbanization of the top income decile. The income shock does less well at explaining the changing location choices of individuals in upper-middle income deciles. This suggests that factors aside from the changing income distribution are also quantitatively important in determining the changing location choices of residents of large

Figure 5: Counterfactual change in sorting from shift in income distribution



Notes: Panel A shows income growth between 1990 and 2014 by income decile for the 27 CBSAs used to calibrate our model. This panel summarizes the shift in the income distribution that we feed into the model. Panel B shows the change in the propensity to live downtown resulting from the change in the income distribution by income decile (solid red bars). The clear bars show the change in the propensity to live downtown between 1990 and 2014 by income decile in the data.

cities.²⁴

5.2 Tests of Model Predictions

We perform two model validation exercises. First, we go further back in time and ask whether shifts in the income distribution in 1950-1970 and 1970-1990 can speak to changes in spatial sorting patterns during these periods, as they do in our 1990-2014 baseline counterfactual. Second, we replicate our baseline calibration and 1990-2014 counterfactual one-by-one for each of the CBSAs that make up the representative city in our baseline analysis and use the results to study whether the model can reproduce salient differences in spatial sorting *across* CBSAs from 1990 to 2014.

5.2.1 Predictions Going Backwards in Time: 1970 and 1950 Counterfactual

In this exercise, we feed the 1950 and 1970 income distributions into the baseline model, calibrated to 1990, and compute the model predictions for the effect of changes in income inequality on spatial sorting between 1950 and 1970 and then 1970 and 1990.²⁵ Incomes unambiguously grew in the

²⁴Note that the predicted urbanization of the highest income decile reflects both a shift along the calibrated U-shape of Figure 4 as well as an endogenous uptick in the U-shape, generated by the change in the income distribution. In the appendix we also show that the model also explains a significant portion of the observed uptick itself.

²⁵The 1950 and 1970 income distributions we feed into the baseline calibration are based on the IPUMS microdata for the set of 27 cities included in our calibration. As we did with our base specification, we interpolate the aggregate income distribution across these cities above the respective top-codes for each year using the generalized Pareto

1950-70 and 1970-90 periods, as the top panel of Figure 6 shows. Specifically, during these periods, the fraction of households with low income uniformly decreased while the fraction of households with higher incomes uniformly increased. The orange bars in the lower panel of Figure 6, meanwhile, show that, through the lens of the model, this income growth generated different changes in sorting from 1950-1970 and 1970-1990 than that predicted in our main counterfactual for 1990-2014. For example, our model predicts lower income individuals *urbanized* and higher income households *suburbanized* during the 1950-1970 period. Additionally, the model predicts little change in spatial sorting patterns by income during the 1970-1990 period. These patterns contrast with the large amount of higher income households moving into downtown areas predicted by our model during the 1990-2014 period.

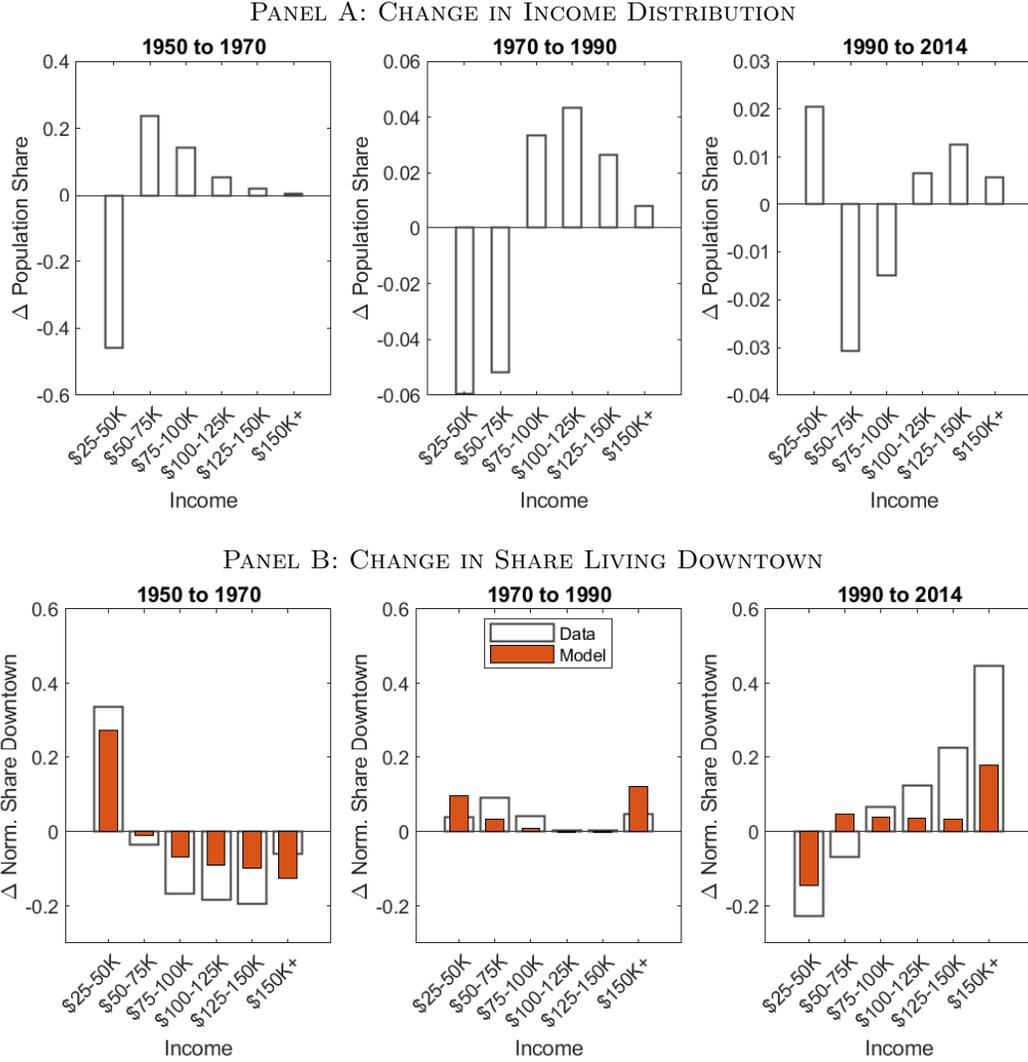
These results reflect that income growth during a given period is not a sufficient condition to cause high-income households to disproportionately move downtown. What then does generate the change in sorting patterns in response to the income growth during the different time periods? The difference stems from where the income growth takes place in the income distribution. Specifically, the shifts in the income distribution in earlier decades were less skewed towards the very rich (in absolute levels) than the 1990-2014 shift. The high-income bracket seeing the largest growth in population share was \$50,000-75,000 from 1950-1970, \$100,000-125,000 from 1970-1990, and then \$125,000-150,000 from 1990-2014. In the 1950s and 1960s, the income growth of higher income households primarily occurred on the downward portion of the U-shape; as a result, our model predicts the suburbanization of high-income households during this period. This prediction is consistent with the data over the same period (shown in the clear bars). Conversely, in the 1970s and 1980s, the income growth for high-income households occurred around the bottom of the U-shape implying only a small change in urbanization rates for these households during this period. Again, this is roughly consistent with actual empirical spatial sorting patterns by income during this period. Finally, between 1990 and 2014, the top income growth shifted households away from the suburban middle of the U-shape and towards the urbanized upward sloping portion of the U-shape and, accordingly, the model predicts the shift of higher-income households downtown. Collectively, these results show that the model can successfully predict different dynamics of changes in spatial sorting patterns by income, depending on where in the income distribution income growth takes place.

5.2.2 Cross-CBSA Predictions

As a second validation exercise, we assess whether the model can match the salient heterogeneity in the changes in residential sorting patterns *across cities*. To that end, we re-calibrate the model separately for individual CBSAs (rather than for a representative city as in the baseline). We allow CBSAs to differ from each other in their initial 1990 income distribution and initial spatial sorting patterns, in the change in their income distribution between 1990 and 2014, and in their land supply elasticities (ϵ_D and ϵ_S). The other parameters in Table 2 are assumed to be identical across

method, as described in the Appendix C.

Figure 6: 1950 and 1970 Counterfactuals



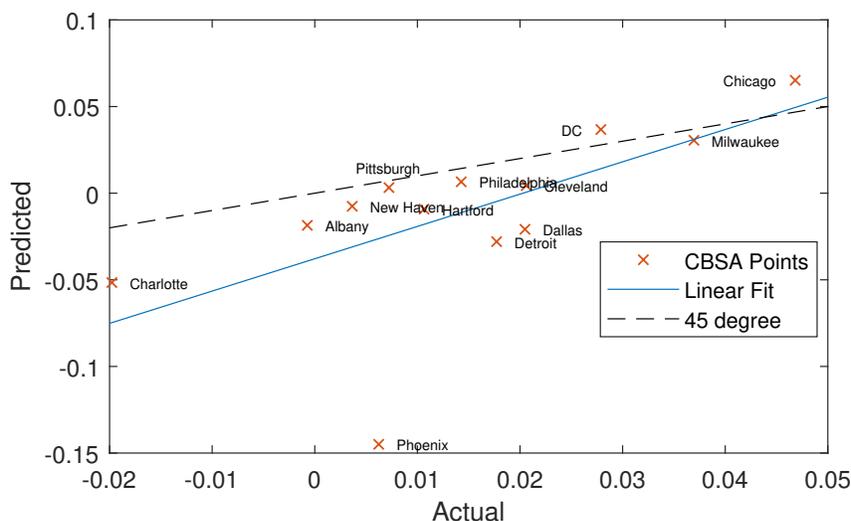
Notes: Panel A shows the change in the share of the population in each income bracket in the 27 CBSAs modeled in our calibration between 1950 and 1970 (left), 1970 and 1990 (middle), and 1990 and 2014 (right). The corresponding plots in Panel B show the change in the share living downtown for each income bracket as observed in the data (clear bars) and as predicted by the model (orange bars). Income brackets are reported in 1999 dollars.

CBSAs. For each CBSA, we calibrate the model by targeting the 1990 distribution of location choice by income within the CBSA, and then compute the model’s prediction for how the CBSA’s spatial sorting patterns change in response to the actual change in the CBSA income distribution. We then compare the predictions of the model to the empirical changes in residential sorting within each city.

Figure 7 compares the cross-CBSA heterogeneity in spatial sorting predicted by the model with that observed in the data using a simple summary statistic: the propensity of households with incomes higher than \$100,000 to reside downtown relative to the average household. The plot compares the *change* in the share of households with incomes above \$100,000 that reside downtown

between 1990 and 2014 in the data to the corresponding change predicted by the model in response to the CBSA-specific shock to the income distribution over the same time period.²⁶ The results show that, through the lens of the model, CBSA-level changes in the income distribution explain CBSA-level changes in spatial sorting of high-income individuals quite well. The CBSAs predicted by the model to have a large relative increase in high-income individuals residing downtown are actually the ones where we observe such an increase empirically.

Figure 7: 1990-2014 Change in the Urban Share of Households Earning Above \$100,000 Less the Average Urban Share



Notes: This figure plots the change in the share of households earning above \$100,000 (in 1999 dollars) that reside downtown between 1990 and 2014, as predicted by the model for each CBSA vs. as observed in the data over the same period.

We conclude from these out-of-sample analyses that the model does quite well at matching time series changes for a representative CBSA as well as cross-CBSA heterogeneity. We view this as a strong test of the model’s implications linking the growth in income at the top of the income distribution with the influx of the rich into downtown neighborhoods within a CBSA. In particular, many national stories that could be confounding our baseline results get differenced out in the cross-CBSA analysis.

6 Welfare and Policy Implications

Having established the model’s ability to reproduce salient empirical sorting patterns, we turn to using the model to analyze the normative implications of changes in urban spatial sorting. We first use the model to assess the well-being consequences, for different income groups, of the neighborhood change and spatial re-sorting triggered by top income growth between 1990 and 2014.

²⁶For this analysis, we can only use 13 of 27 CBSAs. These are the CBSAs for which there is sufficient coverage above the state-specific IPUMS income topcodes to implement the generalized Pareto interpolation procedure that we use to measure CBSA-level income distributions and U-shapes.

In doing so, we highlight the economic forces within the model that drive welfare differences across groups. We end this section by discussing the effectiveness of policies aimed at mitigating these changes in spatial sorting.

6.1 Changes in Welfare Inequality

The framework in Section 2 delivers the following function for the representative utility of a household with income w :

$$V(w) = \left(\sum_{r',q'} V_{r',q'}^\rho(w) \right)^{1/\rho}. \quad (16)$$

where $V_{r,q}$ is the inclusive value of all neighborhoods of type (r,q) defined in (4). We quantify the change in welfare between 1990 and 2014 by income decile, using a related dollar-denominated measure - compensating variation - as follows:

$$CV(i) = m_2(i) - m_2(V_2^{-1}(V_1(m_1(i)))) ,$$

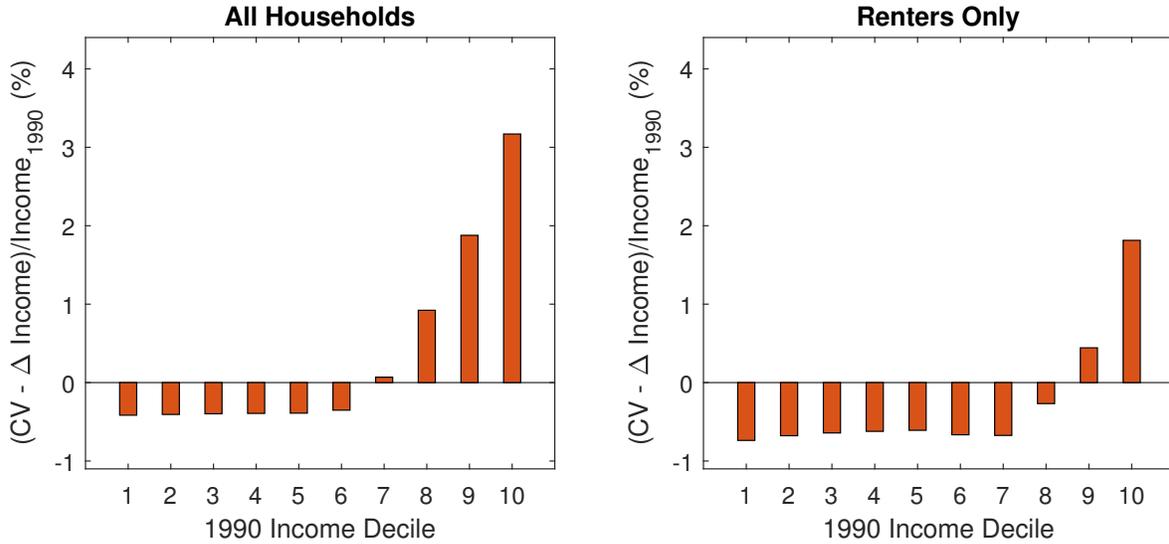
where $m_t(i)$ is income of percentile i in equilibrium t . $CV(i)$ reflects changes in well-being associated with not only changing income, but also changing housing costs and changing endogenous amenity quality. To isolate the welfare gains due to changing housing costs and amenity quality alone, we simply subtract the income growth of a given percentile i from their welfare (i.e., CV) growth:

$$\Delta \mathcal{W}^c(i) = \frac{CV(i) - (m_2(i) - m_1(i))}{m_1(i)}.$$

Whenever $\Delta \mathcal{W}^c(i) > 0$, income growth *understates* the increase in well-being at percentile i .

Figure 8 reports our headline welfare results for each income decile. It shows the welfare gains from within-city spatial sorting triggered by the 1990-2014 income shift. The left panel averages results between homeowners and renters, while the right panel isolates the effects on renters. Focusing first on the average results by decile, we find that the spatial sorting response *amplifies* the differences in well-being between the rich and the poor during this time period. In the top decile of the income distribution, well-being grew more than income, by an additional 3.2 percentage points. As high earners move downtown, the supply of high-quality neighborhoods that they value rises endogenously, making them better off. House prices increase as well, but for high earners the amenity benefit of neighborhood change dominates the price effect. In contrast, at the bottom of the income distribution, households' well-being is hurt by the same house price increases without the same compensatory neighborhood variety benefits. As a result, well-being changes are even more negative than income changes in the bottom decile, by an additional 0.4 percentage points. Overall, the well-being gap between the top and bottom deciles of the income distribution grew by an additional 3.6 percentage points, compared to the 19 percentage point growth in the nominal income gap. That is, within city spatial sorting amplifies the growing welfare gap between

Figure 8: Welfare Changes From Spatial Sorting Response to Changing Income Distribution



Notes: This figure shows the percent welfare growth that households in each income decile are predicted to receive, above and beyond income growth, between the 1990 and 2014 model equilibria. The left-hand panel shows results averaged across homeowners and renters; the right-hand panel focuses on renters only.

the rich and the poor from rising income inequality by almost 20 percent (3.6/19).

Comparing the welfare results for all households in the left panel of Figure 8 with those for renters only in the right panel, we see that capital gains from house price appreciation benefits households at all income levels, to some extent. For example, about 30 percent of individuals from the lowest income decile who resided downtown in 1990 owned their home, limiting the negative welfare effects of spatial re-sorting. Without this effect – i.e., for renters only – welfare losses of low-income households are much larger. Renters in the bottom decile experienced a 0.7 percentage point reduction in their welfare stemming from the changing spatial sorting that resulted from the shift in the income distribution. At the top of the income distribution, the amenity benefit of neighborhood change is strong enough that high-income renters still gain from gentrification, in spite of facing the full brunt of the housing cost growth. They see a 1.8 percentage point growth in welfare.

Before analyzing the main mechanisms at play behind these findings, it is worth commenting on the magnitudes of these welfare effects. Figure 8 implies that a renter in the first decile of the income distribution - earning on average \$30,000 per year - is made roughly \$210 worse off per year in consumption equivalent terms. There are two reasons for this relatively small overall welfare impact. First, the largest welfare losses from an influx of rich households are concentrated on downtown residents, i.e., only 15 percent of individuals earning \$30,000 per year live downtown. If we isolate the most impacted group, low-income renters who remain downtown, we find a welfare loss that is three times larger, at \$630.²⁷ To put that number in perspective, it represents roughly

²⁷We only take into account changes in amenities and prices for these households, holding constant their idiosyncratic preference shocks for location.

one month’s rent for these households. Second, note that we isolate the effect of a change in the income distribution, holding population constant. In reality, population grew a lot in these large CBSAs between 1990 and 2014. Including this population growth along with the changes in the income distribution further amplifies the welfare losses of low-income renters by a factor of five.²⁸. This large magnitude is commensurate with the current policy interest in alleviating the impact of downtown gentrification on this group.

6.2 Mechanisms

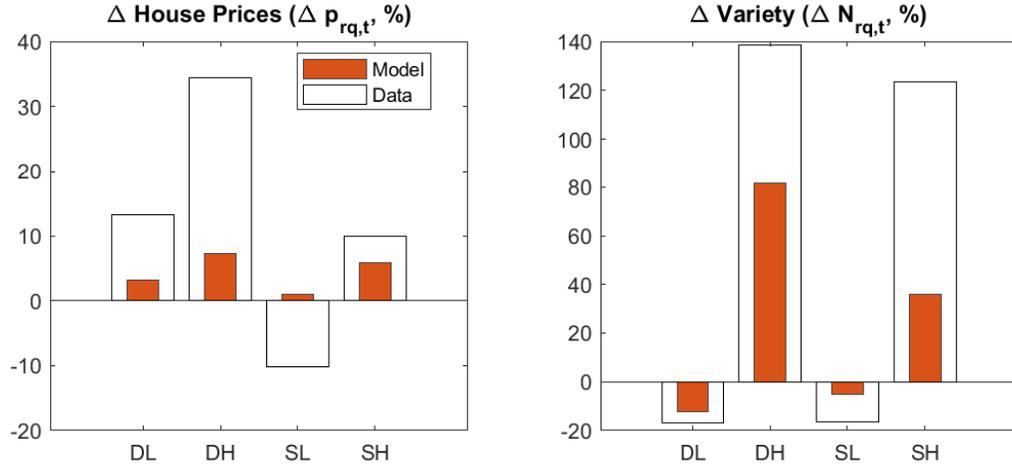
Two main mechanisms drive our sorting and welfare results. The first is the price mechanism that operates through land markets. As the rich get richer, they move downtown to live in high-quality neighborhoods, and to enjoy consumption amenities there. This influx of rich households puts upward pressure on downtown housing prices not only in high-quality neighborhoods, but also in low-quality ones. The left-hand panel of Figure 9 compares these house prices changes, for different neighborhood types, to those observed in the Zillow data. The model predicts that the shift in the income distribution alone generates a 6 percent increase in house prices in high-quality downtown neighborhoods and a 3 percent increase in house prices in low-quality downtown neighborhoods. These predicted increases in downtown house prices amount to about 20 percent of the actual increases observed in the data, which again suggest that other factors (like general CBSA population growth) contribute to house price growth. Housing supply is more elastic in the suburbs than downtown, so the model predicts that house prices increase more in low-quality areas downtown than in low-quality areas in the suburbs (3% vs 1%). This matches the data qualitatively where house price growth between 1990 and 2014 was higher in low-quality downtown neighborhoods than in low-quality suburban neighborhoods. House price growth in low-quality neighborhoods downtown contributes importantly to the welfare losses of the poor renters who remain downtown.

The second key mechanism behind our results is endogenous supply responses and neighborhood change. As the rich move downtown and demand for high-quality neighborhoods increases, developers supply more high-quality neighborhoods. Some of this entry is at the cost of exit of lower quality neighborhoods, so that gentrification takes place. The right-hand panel of Figure 9 shows the growth in supply of neighborhoods in each area and quality level. The model predicts a large proportion of the downtown neighborhood change observed in the data (measured as changes in the number of constant geography Census tracts classified as low and high quality, respectively). Given love of variety preferences, the additional entry of high-quality neighborhoods downtown makes high-income households better off.

Overall, the contraction in the number of low-quality neighborhoods – the gentrification that we also observe in the data – makes low-income households worse off. Finally, we note that the

²⁸In the appendix, we explore counterfactuals where we also allow the population to evolve as is does in the data. In these counterfactuals, not only is the share of richer individuals increasing but the absolute level of richer households are allowed to increase due to population growth. As a result of population growth, our gentrification results are amplified given that there are even more high-income households who want to move downtown in 2014.

Figure 9: Mechanisms

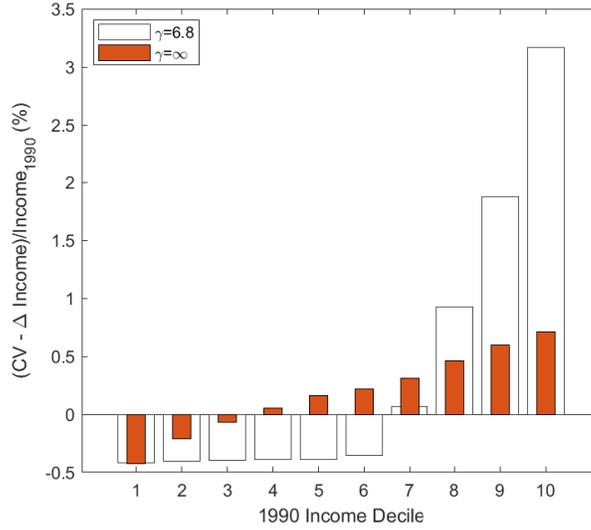


Notes: This figure shows the change in housing costs (on the left) and neighborhood variety (on the right) in each of the four neighborhood types between 1990 and 2014. Red bars are model and wide clear bars are data. Consistent with the model, the change in housing costs are measured in the data for the subset of tracts that maintain the same quality assignment between 1990 and 2014. We measure the change in house prices as the change in the median Zillow 2-bedroom house price index for each neighborhood type between 1996 and 2014.

predicted supply of high-quality neighborhoods also expands in the suburbs, but at a much smaller rate than downtown.

We can use the model to separately quantify the contribution of the price and amenity mechanisms to our welfare results. To that end, we compute a counterfactual that shuts down love-of-variety effects across neighborhoods by setting the between-neighborhood substitution elasticity, γ , to infinity. In this counterfactual, prices respond to changes in the income distribution, but the sorting and welfare effects of these price responses are not amplified by responses in neighborhood (or associated consumption amenity) variety. Welfare results are shown in Figure 10 (solid red bars) and contrasted with our baseline quantification (clear bars). The welfare gap across income groups is mitigated substantially when the love of variety effects are shut down, from 3.6 percentage point in the baseline to 1.1 percentage points without love of variety effects. About two-thirds of the welfare gap in our base results stems from the endogenous private amenities response. The absolute welfare losses for the bottom decile are unaffected, as price increases are not compensated by the gains in consumption amenities that accompany the influx of the rich. Shutting down love of variety effects almost completely eliminates the welfare gains of the rich but does not stem the welfare losses experienced by the poor.

Figure 10: Shutting Down Amplification



Notes: This figure shows the percent welfare growth that households in each income decile are predicted to receive, above and beyond income growth, between the 1990 and 2014 model equilibria. The left-hand panel shows results averaged across homeowners and renters; the right-hand panel focuses on renters only. The clear bars show the results from the baseline calibration and counterfactual where the neighborhood and amenity love of variety parameters (γ and σ) are equal 6.5. The red bars show the welfare results from an alternative calibration and counterfactual where love of variety is shut down by setting these parameters to both equal infinity.

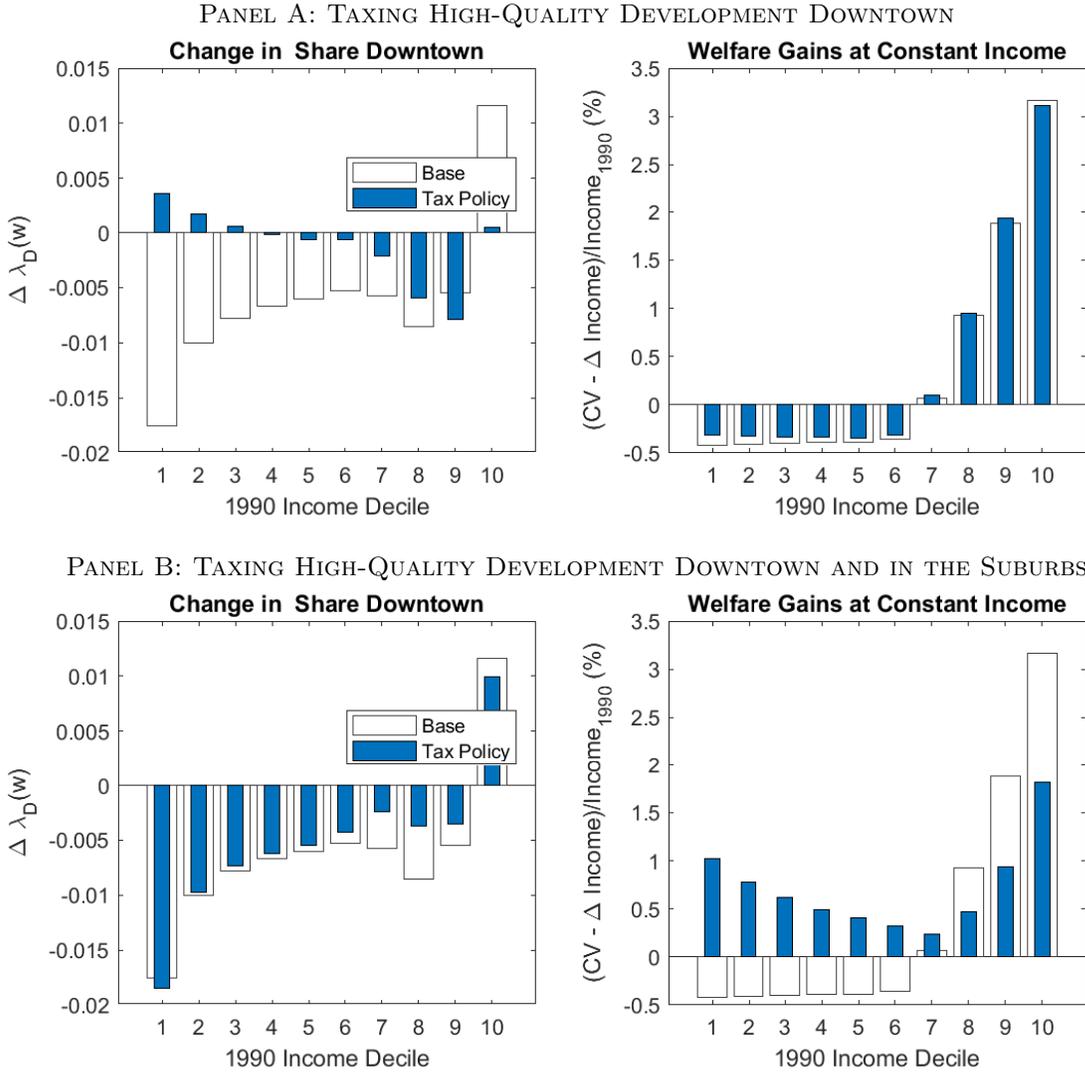
6.3 Gentrification Curbing Policies

Changes in spatial sorting in large US cities have led to a new policy debate on gentrification and housing affordability. In this subsection, we use our model to analyze the potential impact of policies that aim to shape the spatial sorting of heterogeneous households within the city.

Taxing developments We first model a stylized “anti-gentrification” policy, which systematically taxes high-quality housing downtown, and uses the proceeds to subsidize rents in low-quality neighborhoods downtown (the policy is budget neutral). It aims to limit the development of high-quality neighborhoods downtown while helping poorer households to remain located downtown. We compute the counterfactual 2014 spatial equilibrium with a tax on high-quality housing downtown of $t = 5\%$.

Panel A in Figure 11 reports the results and contrasts them with our baseline 1990-2014 counterfactual, in order to evaluate how much such a policy would have curbed the gentrification triggered by changes in the income distribution. The left panel shows that the policy stems part of the gentrification of downtown neighborhoods: the inflow of high-income households downtown is curbed (solid bars), as is the outflow of low-income families, compared to baseline (clear bars). The policy is also effective at stemming part of the land price increase downtown, and limiting quality changes. To the extent that governments intrinsically value social diversity within their downtowns, this simulation suggests that such an anti-gentrification policy can help maintain that target.

Figure 11: Location Choices and Well-Being under “Anti-Gentrification” Policy



Notes: This figure shows the percent change in the propensity to live downtown (on the left) and change welfare (on the right) that result from the change in the income distribution by income decile. The clear bars show the results from the baseline counterfactual. The blue bars show the results from the alternative counterfactual in which development of high-quality neighborhoods and the proceeds of the tax are redistributed to subsidize housing costs in low-quality neighborhoods in the same location. In Panel A, only downtown high-quality neighborhoods are taxed (and low-quality rents subsidized). In Panel B, high-quality neighborhoods are taxed (and low-quality rents subsidized) both downtown and in the suburbs.

The well-being effects of this policy, shown in the right hand plot of Panel A in Figure 11 are, however, much more muted. The policy hardly changes at all either the welfare of high-income households or the welfare losses of low-income households. The policy fails to significantly reduce the losses of low-income households because taxing high-quality development downtown shifts gentrification – i.e., neighborhood quality and price growth – from downtown to the suburbs. As a result, low-income households living in the suburbs experience greater welfare losses relative to baseline. On net, the welfare losses are simply transferred from residents of low-quality downtown

neighborhoods to residents of low-quality suburban neighborhoods.²⁹

Panel B of Figure 11 below shows the impact of an alternative policy that taxes high-quality neighborhood development (and subsidizes rents) both downtown and in the suburbs. This policy has a much stronger progressive effect than the downtown-specific development tax and rent subsidy, largely because a larger share of the population resides in the suburbs. Interestingly, this policy does not limit changes in sorting much, and does not stem gentrification. This is because the tax on high-quality development and subsidy on low-quality housing costs is implemented both downtown and in the suburbs. Therefore, intuitively, changes in relative housing costs (and the amount of high-quality development) are qualitatively similar across the two locations. However, the policy does mitigate inequality. This is not surprising given the policy - by design - is taxing high incomes and distributing the proceeds to low-income households. The endogenous change in amenities stemming from the changing spatial sorting response still makes the rich better off despite them being taxed more. The poor are made better off through the redistribution which is sufficient to compensate them for their increased rental payments in downtown neighborhoods.

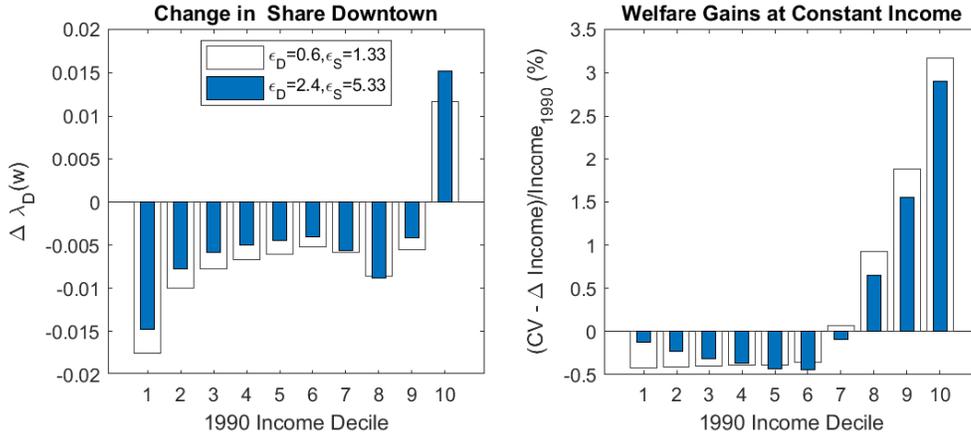
Regulatory constraints on housing supply Finally, we shed light on a policy that has been widely proposed by economists to address the regressive welfare impacts of rising housing costs: relieving regulatory housing supply constraints. Housing regulations do not feature directly into our model, they are instead indirectly captured by the housing supply elasticities that we use in calibration. We now report the effect of quadrupling the elasticity of housing supply both downtown and in the suburbs. Figure 12 shows that doing so does little to stem neighborhood change downtown (in the left panel) but it mitigates the associated welfare losses on the poor (in the right panel). House price growth is reduced by approximately 2 percentage points in all neighborhoods, and effectively shut down in the low-quality neighborhoods. The benefits of this slowed house price growth mostly accrue to the poor – the lowest income decile’s welfare loss is essentially eliminated. Welfare inequality continues to grow, however, because the rich still gain from increased neighborhood variety that persists even with the increased elasticity of housing supply.

7 Robustness

We conclude by performing a series of additional quantitative exercises to explore the robustness of our results.

²⁹In the appendix we show that the results of this policy are very similar to those we obtain when directly modeling zoning regulations; i.e., a policy that imposes a constant relative number of high- to low-quality neighborhoods. The impact on social mixing downtown is significant, but the welfare effects are again very small, as price and quality growth are pushed to the suburbs.

Figure 12: Effects of Increasing Supply Elasticity



Notes: This figure shows the percent change in the propensity to live downtown (on the left) and change welfare (on the right) that result from the change in the income distribution by income decile. The clear bars show the results from the baseline counterfactual. The blue bars show the results from the alternative counterfactual where the elasticity of housing supply is double its baseline level in both the suburbs and downtown.

7.1 Robustness to Key Elasticities

Below our baseline results, Table 3 first shows the sensitivity of these results to ρ , the parameter that governs the extent of sorting by income. For our robustness exercise, we set $\rho = 1.5$ and $\rho = 4.5$, which is roughly a two-standard deviation band around our baseline ρ estimate. As individuals get richer, they are more likely to move downtown when ρ is higher. Additionally, the poor are more likely to migrate out in response to the price increase associated with rich moving downtown as ρ is higher. In other words, gentrification forces increase as ρ increases. Therefore, higher values of ρ amplify our welfare results. However, it is interesting to note that, even when $\rho = 1.5$, accounting for spatial sorting responses increases the inequality between the top and bottom income deciles by 2.6 percentage points.

Next in Table 3, we show the sensitivity of our results to different values of γ . For lower values of γ , endogenous amplification of amenities downtown is stronger. As the endogenous amplification of amenities increases, more high-income individuals move downtown putting further upward pressure on downtown land prices in both high- and low-quality neighborhoods. This increases the welfare differences between individuals in the top and bottom income deciles primarily by increasing the well-being of the rich through higher love-of-variety effects.

Finally, Table 3 shows that land supply elasticities downtown and in the suburbs are a crucial determinant of the welfare losses to poor renters. This is not surprising given the policy counterfactual experiments highlighted in the prior section. Much of the welfare effect on the poor stems from them paying higher rents downtown as the rich move in. The more inelastic the downtown housing supply (in both absolute terms and relative to the suburbs), the more house prices move, generating modest additional growth in the welfare gap between the poor and the rich. The growth in the welfare gap masks heterogeneity between owners and renters. Additional price growth mitigates

Table 3: Robustness of Welfare Estimates to Key Parameters

Decile:	$(\Delta CV - \Delta Inc)/Inc_{1990}$						Δ Urban Share			
	All Households			Renters Only			Predicted (p.p.)		Share of Actual	
	Top	Bottom	Diff.	Top	Bottom	Diff.	Top	Bottom	Top	Bottom
Base Specification	3.17	-0.42	3.59	1.81	-0.74	2.55	1.16	-1.75	57%	41%
Elasticity of Substitution between Neighborhood Types (base: $\rho = 3$)										
$\rho = 1.5$	2.26	-0.30	2.57	1.16	-0.60	1.75	1.36	-1.96	67%	46%
$\rho = 4.5$	3.73	-0.46	4.19	2.29	-0.76	3.05	1.40	-2.09	69%	49%
Elasticity of Substitution between Same-Type Neighborhoods (base: $\gamma = 6.8$)										
$\gamma = 4$	5.34	-0.42	5.75	3.36	-0.55	3.91	2.38	-3.62	117%	85%
$\gamma = 8$	2.75	-0.42	3.17	1.52	-0.78	2.30	1.07	-1.59	53%	37%
$\gamma = \infty$	0.71	-0.42	1.13	0.09	-1.02	1.11	0.82	-1.06	41%	25%
Housing/Land Supply Elasticities (base: $\epsilon_D = 0.6, \epsilon_S = 1.33$)										
$\epsilon_D = 0.3, \epsilon_S = 1.33$	3.19	-0.43	3.62	1.81	-0.77	2.57	1.05	-2.03	52%	48%
$\epsilon_D = \epsilon_S = 1.33$	3.12	-0.39	3.52	1.83	-0.68	2.51	1.38	-1.20	68%	28%
$\epsilon_D = 2.4, \epsilon_S = 5.33$	2.90	-0.12	3.02	2.04	-0.09	2.13	1.51	-1.47	75%	35%

Notes: This table summarizes the sensitivity of our welfare results and changing location choice predictions to alternate parameter values. The first three columns report the sensitivity of the absolute change in welfare of the top and bottom decile of the income distribution (columns 1 and 2) and the relative change in welfare between these deciles (column 3) to values of the key parameters, while feeding in the same income shock. The next three columns show the same welfare statistics for renters (i.e., households not receiving any share of the house price appreciation mutual fund). The final columns summarize the model predictions for the urbanization of top income decile households and the suburbanization of bottom income deciles, first in absolute percentage point terms and then as a share of the respective 2.3 percentage point inflow and 4 percentage point outflow observed in the data.

welfare losses for poor owners downtown, but exacerbates losses to poor renters.³⁰

Overall, this variation in our welfare and spatial sorting estimates to different parameter values is useful for understanding the forces driving our results. But we note that over reasonable parameter ranges, our welfare results are fairly similar. Our main qualitative results are not reversed by any of these perturbations: poor households (particularly renters) are worse off in both absolute terms and relative to the wealthy from the spatial sorting response to top income growth between 1990 and 2014.

7.2 Targeting an Additional Moment

In this subsection, we explore how the calibrated model matches housing spending share. In the model, the unit housing requirement means that, within neighborhood type, all households spend the same amount on housing regardless of income, so the income share of housing expenditure is mechanically downward sloping in income. This slope is mitigated by non-homothetic sorting across neighborhoods: higher incomes sort into the more expensive neighborhood types so their income share of housing does not fall proportionally with income. With only four neighborhood

³⁰We also explored the robustness of our results to the public finance parameters (T_D , T_S , and Ω). The choice of these parameters had little influence on either our welfare or spatial sorting results.

types in the quantified model, this sorting goes a long way in replicating the income share of housing in the data. However, as we discuss further below, the implied housing “Engel” curve from the quantitative model is still larger than what is found in the data.

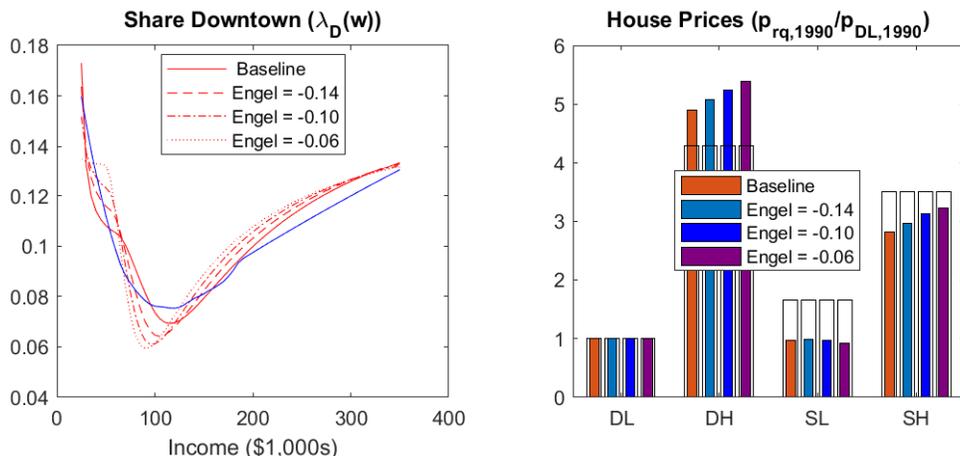
To create a data analog, we use reported spending patterns on “housing” by income deciles reported in public release tables from the Consumer Expenditure Survey (CES). For our empirical measure of housing expenditures, we use the CES’s combined reported expenditure on “Shelter” and “Utilities”.³¹ Both in the data and the model we regress the housing share of total expenditure on log total income. Using the model, we get a housing share Engel curve semi-elasticity of -0.19; this implies that a 10 percent increase in income is associated with the housing spending share falling by 1.9 percentage points. When we run the same regression on the CEX generated data, we get a housing share Engel curve semi-elasticity of -0.11. Moreover, in the CEX data, we can run the housing expenditure share on log expenditure (as opposed to log income); in our static model, a household’s log income equals their log expenditure. When we run this latter regression in the CEX data, we get a housing share Engel curve semi-elasticity of -0.06. In other word, our model is generating a stronger relationship between the housing share of expenditure and total income (expenditure) relative to the data.

In our baseline specification, we target the relative housing prices across neighborhoods and the U-shape downtown sorting patterns when calibrating our model. However, given that our model is off in matching how the housing spending share varies with income, as a robustness exercise we use the empirical housing Engel curve slope from the CEX as an additional model target. These results are shown in Figure 13 and Table 4. Specifically, we target housing Engel curve slopes of -0.14, -0.10, and -0.06. These values span the estimates from the CEX micro data. In all cases, when we target these values, our model is able to match the Engel curve slopes exactly. However, as seen in Figure 13, in order to target a flatter housing Engel curve, the model needs to generate a higher housing price in downtown high-quality neighborhoods. In particular, if richer households are going to spend more on housing in our model with a unit housing requirement, the model needs to make the housing predominantly bought by higher income households more expensive. It is also interesting to note that our fit matching the U-shape sorting pattern is essentially invariant to the targeted Engel curve slope.

Table 4 shows how our welfare and gentrification results change as we target different Engel curve slopes. The higher price of downtown high-quality neighborhoods implies that fewer household can afford to live in those neighborhoods even with income growth. As a result, the model’s implied gentrification gets slightly smaller as we target a flatter housing Engel curve. As the Engel curve gets flatter, our welfare results also get mitigated slightly. However, even when we target a housing Engel curve of -0.06, we still find that the welfare of the top income deciles grows by 3.1 percentage

³¹See <https://www.bls.gov/cex/tables/calendar-year/mean-item-share-average-standard-error.htm#cu-income> for the Consumer Expenditure Survey (CES) public release tables. We omit the “Other Lodging” component of the “Shelter” category when making our empirical measure. The “Other Lodging” component includes the household’s spending on hotels, vacation homes, and college dorm fees. We discuss the details of the mapping of CES measures of housing to our model analogs in the online appendix.

Figure 13: Sensitivity of Calibrated 1990 Urban Shares and Neighborhood Prices to Targeting Housing Share Moment



Notes: These figures show the fit of the calibrated model to the two main targeted moments, when targeted alone and then along with a moment targeting an Engel slope. The left-hand plot shows the share of households in each \$5,000 income bracket that reside downtown in 1990. The solid blue line shows the data, while the solid red line shows the prediction of the base calibrated model. The dashed lines then show the prediction of the model calibrations that target different Engel curve slopes. The clear bars in the right-hand plot shows the average Census home value in tracts of each location-quality type, normalized by the average index in low-quality tracts downtown, in 1990. The solid red bars show the predicted relative housing costs predicted by the baseline calibrated model, while the solid bars next to the baseline show the predictions of the model when calibrated to different Engel curve slopes as well as these two moments.

points relative to the bottom decile in response to the observed income growth between 1990 and 2014.

7.3 Additional Potential Mitigating Forces

A limitation of the benchmark model is the assumption that increased variety of neighborhoods of a given rq type only benefits inhabitants of that type of neighborhood. In reality, the gentrification of downtown neighborhoods can benefit all inhabitants of the city, to the extent those inhabitants can travel to consume urban amenities there. In an additional robustness specification, we modify the model to allow for individuals to consume amenities in other neighborhoods types. We discipline this model extension using (1) expenditure data from the Consumer Expenditure Survey on spending on amenities like restaurants and entertainment venues and (2) proprietary cell-phone data which maps the extent to which individuals travel to restaurants and entertainment options outside of the neighborhood where they live. Our conjecture was that such a model extension would mitigate the welfare differences between high and low income decile residents stemming from the changing spatial sorting response to the shift in the income distribution between 1990 and 2014. As the influx of the rich created more high-quality downtown neighborhoods, lower income households would get additional utility from consuming the amenities of those neighborhoods. While our conjecture was qualitatively correct, allowing for this channel had a very small quantitative effect on our welfare results. The reason for the small adjustment to our welfare results stemmed from the fact that

Table 4: Robustness of Welfare Estimates to Targeting Housing Share Moment

Engel Slope	$(\Delta CV - \Delta Inc)/Inc_{1990}$						Δ Urban Share			
	All Households			Renters Only			Predicted (p.p.)		Share of Actual	
	Top	Bottom	Diff.	Top	Bottom	Diff.	Top	Bottom	Top	Bottom
Base Specification										
-0.19	3.17	-0.42	3.59	1.81	-0.74	2.55	1.16	-1.75	57%	41%
Targeting Engel Slope										
-0.14	3.20	-0.37	3.57	1.85	-0.60	2.45	1.05	-1.59	52%	38%
-0.10	3.08	-0.31	3.38	1.80	-0.45	2.25	0.94	-1.40	46%	33%
-0.06	2.84	-0.22	3.06	1.69	-0.27	1.97	0.85	-1.21	42%	29%

Notes: This table summarizes the sensitivity of our welfare results and changing location choice predictions to adding a moment targeting an Engel slope to the calibration. The structure of the table replicates that of Table 3.

empirically, the expenditure share on urban amenities was relatively small and the cell-phone data highlighted that lower income households rarely consume amenities in high-quality neighborhoods. Given the small quantitative results, we omitted the details of this extension from the current version of the paper. However, the full details can be found in the NBER working paper version of our paper (Couture et al., 2019).³²

8 Conclusion

We set out to explore the link between rising incomes at the top of the income distribution and changes in the urban landscape of U.S. cities in the past few decades: high-income households have been moving into downtowns, where housing prices have gone up and neighborhoods have been changing dramatically. These changes have led to anti-gentrification protests and a renewed interest in policy circles for maintaining social diversity in urban neighborhoods. To study this phenomenon, we develop a spatial model of a city with heterogeneous agents, neighborhoods of different qualities, and non-homothetic preferences. We quantify the model and use it to tease out how much of the change in spatial sorting patterns by income over time can be plausibly traced back to changes in the income distribution, tilted towards higher incomes.

Our estimates suggest that rising incomes at the top of the distribution were a substantive contributor to increased urban neighborhood change during the last 25 years within the U.S. The analysis also suggests that neighborhood change resulting from the increased incomes of the rich did make poorer residents worse off. Accounting for the spatial sorting response resulting from the change in income distribution between 1990 and 2014 exacerbates the growing inequality between the top and bottom income deciles by an additional 3.6 percentage points.

We explore possible policy responses to mitigate these distributional impacts of neighborhood

³²In the working paper, we also allowed for commuting costs to differ between the suburbs and downtown. We used data from the National Household Transportation Survey to discipline the differential commuting costs. Allowing for commuting differences between the suburbs and downtown did not alter our quantitative results in any way.

change. We find that policies that contain gentrification seem to only lead to a very modest mitigation of the increase in well-being inequality, which could arguably be targeted more efficiently by direct redistribution. On the other hand, these policies are effective at maintaining social diversity in urban neighborhoods, arguably one of the goals of such policies. However, policies that relax land supply constraints can mitigate welfare losses to the poor.

In this paper, we have focused on the within-city consequences of a rise in top incomes. By doing so, we have highlighted one mechanism that has contributed to shape neighborhood change in the past twenty five years: the rising incomes of the rich coupled with non-homothetic preferences for location across income groups. In order to conduct this analysis, we have developed a model that is stylized in some dimensions, but is very flexible. It is in particular amenable to study other sources of changes in within-city spatial sorting that are potentially empirically relevant. Using our framework to study other potential causes of neighborhood change is a natural avenue for future research.

References

- Aguiar, M. and M. Bilal (2015). Has consumption inequality mirrored income inequality? *American Economic Review* 105(9), 2725–56.
- 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 (2015). Optimal city structure. *Yale University, mimeograph*.
- Almagro, M. and T. Dominguez-Iino (2019). Location sorting and endogenous amenities: Evidence from Amsterdam.
- Armour, P., R. V. Burkhauser, and J. Larrimore (2016). Using the Pareto distribution to improve estimates of topcoded earnings. *Economic Inquiry* 54(2), 1263–1273.
- Athey, S., D. Blei, R. Donnelly, F. Ruiz, and T. Schmidt (2018). Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data. *AEA Papers and Proceedings* 108, 64–67.
- Autor, D. H., C. J. Palmer, and P. A. Pathak (2017). Gentrification and the amenity value of crime reductions: Evidence from rent deregulation. (23914).
- Baum-Snow, N. and L. Han (2019). The microgeography of housing supply.
- Baum-Snow, N. and D. Hartley (2020). Accounting for central neighborhood change, 1980–2010. *Journal of Urban Economics* 117, 103228.
- Baum-Snow, N. and J. Marion (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics* 93(5-6), 654–666.
- Berkes, E. and R. Gaetani (2018). Income segregation and rise of the knowledge economy. *Ohio State University Mimeo*.
- Borusyak, K., P. Hull, and X. Jaravel (2021). Quasi-experimental shift-share research designs. *The Review of Economic Studies*. rdab030.
- 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.
- Brummet, Q. and D. Reed (2019). The effects of gentrification on original neighborhood residents: Evidence from longitudinal Census microdata. Working Paper 19-30, Federal Reserve Bank of Philadelphia.
- Chen, M. K. and R. Rohla (2018). The effect of partisanship and political advertising on close family ties. *Science* 360(6392), 1020–1024.

- Couture, V., J. Dingel, A. Green, J. Handbury, and K. Williams (2021). Measuring movement and social contact with smartphone data: a real-time application to Covid-19. *Journal of Urban Economics*.
- Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2019). Income growth and the distributional effects of urban spatial sorting. Technical report, National Bureau of Economic Research.
- Couture, V. and J. Handbury (2017). Urban revival in America, 2000 to 2010. (24084).
- Davis, D. R. and J. I. Dingel (2020). The comparative advantage of cities. *Journal of International Economics* 123(C).
- Davis, D. R., J. I. Dingel, J. Monras, and E. Morales (2019). How segregated is urban consumption? *Journal of Political Economy* 127(4), 1684–1738.
- Davis, M. A. and F. Ortalo-Magne (2011). Household expenditures, wages, rents. *Review of Economic Dynamics* 14(2), 248–261.
- Diamond, R. (2016). The determinants and welfare implications of US workers’ diverging location choices by skill: 1980-2000. *American Economic Review* 106(3), 479–524.
- Diamond, R. and T. McQuade (2019). Who wants affordable housing in their backyard? an equilibrium analysis of low-income property development. *Journal of Political Economy* 127(3), 1063–1117.
- Diamond, R., T. McQuade, and F. Qian (2019, September). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco. *American Economic Review* 109(9), 3365–94.
- Ding, L., J. Hwang, and E. Divringi (2016). Gentrification and residential mobility in Philadelphia. *Regional Science and Urban Economics* 61, 38–51.
- Dingel, J. I. (2016). The determinants of quality specialization. *The Review of Economic Studies* 84(4), 1551–1582.
- Edlund, L., C. Machado, and M. Sviatchi (2019). Bright minds, big rent: Gentrification and the rising returns to skill. (21729).
- Ellen, I. G., K. M. Horn, and D. Reed (2019). Has falling crime invited gentrification? *Journal of Housing Economics* 46, 101636.
- Ellen, I. G. and K. O’Regan (2010). Crime and urban flight revisited: The effect of the 1990s drop in crime on cities. *Journal of Urban Economics* 68(3), 247–259.
- Eriksen, M. D. and S. S. Rosenthal (2010). Crowd out effects of place-based subsidized rental housing: New evidence from the LIHTC program. *Journal of Public Economics* 94(11-12), 953–966.

- Faber, B. (2014). Trade liberalization, the price of quality, and inequality: Evidence from Mexican store prices. *UC-Berkeley Working Paper*.
- Faber, B. and T. Fally (2021). Firm heterogeneity in consumption baskets: Evidence from home and store scanner data. *The Review of Economic Studies*.
- Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011). Income distribution, product quality, and international trade. *Journal of Political Economy* 119(4), 721–765.
- Fajgelbaum, P. D. and C. Gaubert (2020). Optimal spatial policies, geography, and sorting*. *The Quarterly Journal of Economics* 135(2), 959–1036.
- Fajgelbaum, P. D. and A. K. Khandelwal (2016). Measuring the unequal gains from trade. *The Quarterly Journal of Economics* 131(3), 1113–1180.
- Fajgelbaum, P. D., E. Morales, J. C. Suárez Serrato, and O. Zidar (2018). State taxes and spatial misallocation. *The Review of Economic Studies* 86(1), 333–376.
- Fogli, A. and V. Guerrieri (2019). The end of the American dream? Inequality and segregation in US cities. *University of Chicago Booth School of Business*.
- Gaigne, C., H. Koster, F. Moizeau, and J.-F. Thisse (2017). Amenities and the social structure of cities. HSE Working paper WP BRP 162/EC/2017, National Research University Higher School of Economics.
- Gaubert, C. and F. Robert-Nicoud (2021). Sorting to expensive cities. Technical report.
- Glaeser, E., J. Kolko, and A. Saiz (2001). Consumer city. *Journal of Economic Geography* 1(1), 27–50.
- Guerrieri, V., D. Hartley, and E. Hurst (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics* 100, 45–60.
- Gyourko, J., C. Mayer, and T. Sinai (2013). Superstar cities. *American Economic Journal: Economic Policy* 5(4), 167–199.
- Hoelzlein, M. (2019). Two-sided sorting and spatial inequality in cities.
- Holian, M. J. and M. E. Kahn (2012). The impact of center city economic and cultural vibrancy on greenhouse gas emissions from transportation. MTI Report 11-13, Mineta Transportation Institute.
- Hsieh, C.-T. and E. Moretti (2019, April). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics* 11(2), 1–39.
- Jaravel, X. (2018). The unequal gains from product innovations: Evidence from the US retail sector. *The Quarterly Journal of Economics* 134(2), 715–783.

- Lance Freeman (2005). Displacement or succession?: Residential mobility in gentrifying neighborhoods. *Urban Affairs Review* 40(4), 463–491.
- Lee, S. (2010). Ability sorting and consumer city. *Journal of Urban Economics* 68(1), 20–33.
- Lester, T. W. and D. A. Hartley (2014). The long term employment impacts of gentrification in the 1990s. *Regional Science and Urban Economics* 45(1), 80–89.
- Levitt, S. D. (2004). Understanding why crime fell in the 1990s: Four factors that explain the decline and six that do not. *Journal of Economic Perspectives* 18(1), 163–190.
- McKinnish, T., R. Walsh, and T. Kirk White (2010). Who gentrifies low-income neighborhoods? *Journal of Urban Economics* 67(2), 180–193.
- Meltzer, R. and P. Ghorbani (2017). Does gentrification increase employment opportunities in low-income neighborhoods? *Regional Science and Urban Economics* 66(C), 52–73.
- Neumark, D., J. Zhang, and B. Wall (2007). Employment dynamics and business relocation: New evidence from the National Establishment Time Series. In *Aspects of worker well-being*, pp. 39–83. Emerald Group Publishing Limited.
- Peters, M., F. Eckert, et al. (2018). Spatial structural change. In *2018 Meeting Papers*, Number 98. Society for Economic Dynamics.
- Piketty, T., E. Saez, and G. Zucman (2017, 10). Distributional national accounts: Methods and estimates for the United States*. *The Quarterly Journal of Economics* 133(2), 553–609.
- Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative spatial economics. *Annual Review of Economics* 9(1), 21–58.
- Redding, S. J. and D. M. Sturm (2016). Estimating neighborhood effects: evidence from war-time destruction in London. *Princeton University, mimeograph*.
- Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas, and M. Sobek (2018). IPUMS-USA.
- Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics* 125(3), 1253–1296.
- Su, Y. (2018). The rising value of time and the origin of urban gentrification.
- Tsivanidis, N. (2019). The aggregate and distributional effects of urban transit infrastructure: Evidence from Bogotá’s TransMilenio. Working paper, University of Chicago Booth School of Business.
- Vigdor, J. L., D. S. Massey, and A. M. Rivlin (2002). Does gentrification harm the poor? *Brookings-Wharton Papers on Urban Affairs*, 133–182.

Online Appendix:

“Income Growth and the Distributional Effects of Urban Spatial Sorting”

Appendix A U-Shaped Sorting Patterns

In this section, we detail the construction of the U-Shape urbanization patterns by income documented in Figure 1 in the main text. We also highlight the robustness of the U-Shape patterns to alternate downtown definitions and within detailed demographic groups.

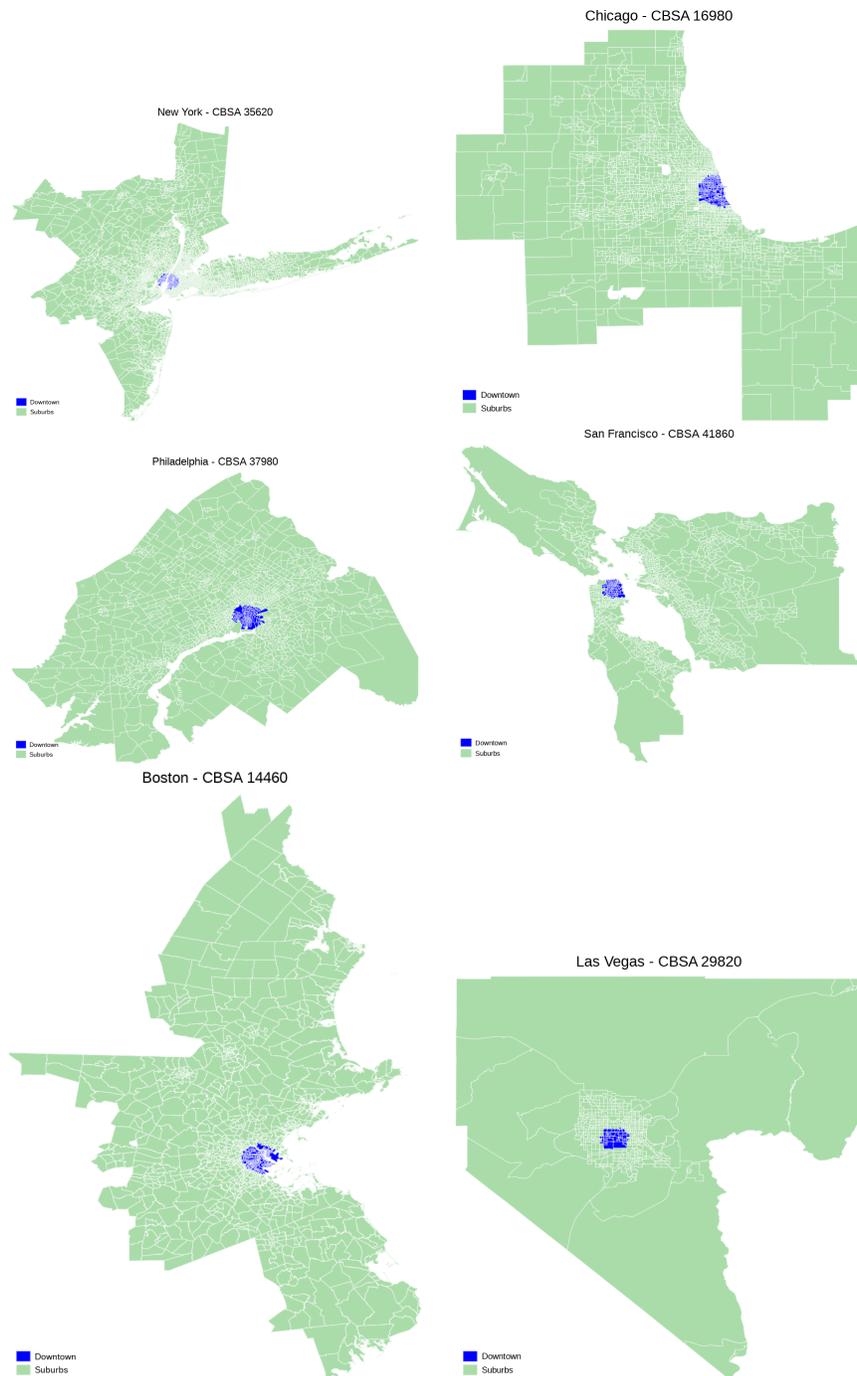
Appendix A.1 Main U-Shape Figure

Figure 1 of the main text summarizes the Engel curve for residing downtown. It shows the relative propensity of families to reside downtown by income, and its evolution over time. The stylized facts in this figure are based on data from the 1970, 1990, and 2000 U.S. Censuses, as well as from the 2012-2016 American Community Surveys (ACS). We refer to the 2012-2016 pooled ACS data as the 2014 ACS. We use census tract level data published by the National Historical Geographic Information System (NHGIS). All data are interpolated to constant 2010-boundary tracts and 2014-boundary CBSAs using the Longitudinal Tract Data Base (LTBD). We complement Census tables with microdata from the Integrated Public Use Micro-data Series (Ruggles et al., 2018), adjusted for top-coding using the generalized Pareto method. We use the 1% IPUMS sample in 1970, and the 5% IPUMS samples in 1990, 2000, and 2012-2016. In what follows, all income measures are CPI-adjusted to 1999 dollars. We provide a detailed discussion of this data (including our top-coding procedure) below in Appendix C.

With this data, we measure the location choice of households with differing levels of income. As described in the main text, we classify as downtown the set of tracts closest to the city center that accounted for 10 percent of the CBSA’s population in 2000. This defines a spatial boundary of downtown, which we keep constant across all years. Figure A.1 shows the downtown and suburban tracts for the CBSAs of New York, Chicago, Philadelphia, San Francisco, Boston, and Las Vegas. Figure A.2 shows income growth in downtown and selected suburban tracts within the central county of each of these CBSAs.

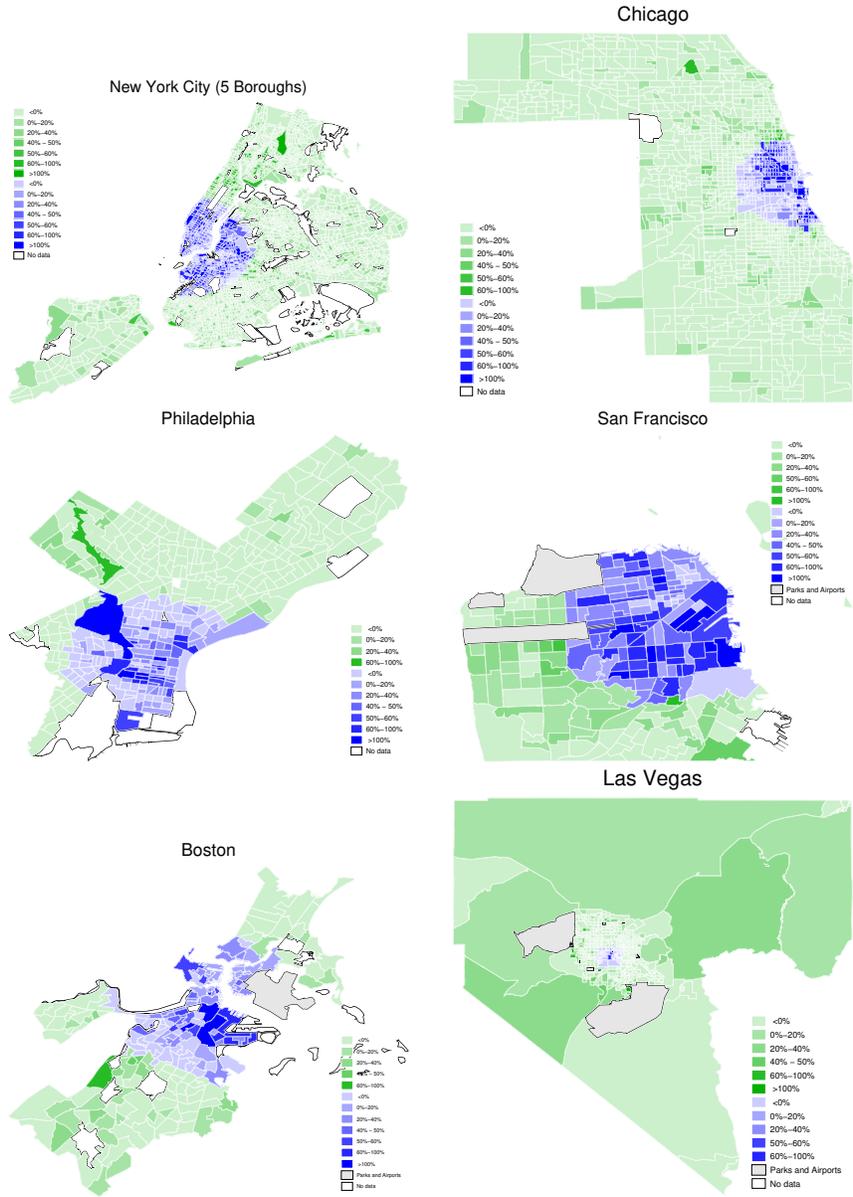
Each point in Figure 1 represents the share of families, in a given Census income bracket, who reside downtown in a given year – normalized by the share of all families who reside downtown that year. This normalization allows us to abstract from the suburbanization of the population as a whole over this period due to general population growth. The share of families at all income levels that live downtown was 0.1 in 2000 (by construction) but was 0.17 in 1970 and 0.08 in 2014. The x-axis features the median family income for that bracket in the same year, in 1999 dollars, computed using IPUMS micro data. The number of points on the graph is limited by the number of income brackets reported by the Census for tract-level information.

Figure A.1: Downtown and Suburban Tracts in Selected CBSAs.



Note: Downtown tracts in dark blue consists of all tracts closest to the city center and accounting for 10% of total CBSA population in 2000.

Figure A.2: Income Growth in Tracts within Central County of Selected CBSA



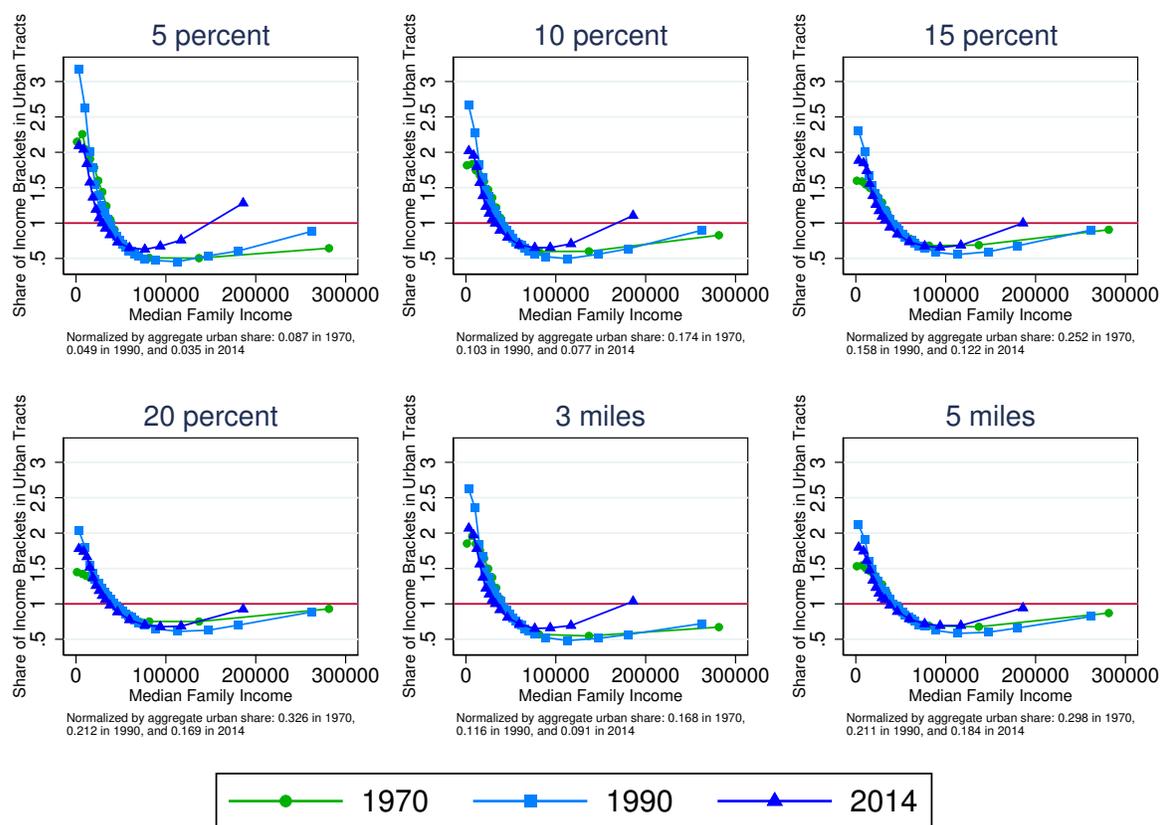
Note: Each map shows the central county of a given CBSA, except for New York which shows the five counties (boroughs) of New York City. Downtown tracts in blue consist of all tracts closest to the city center and accounting for 10% of total CBSA population in 2000. The shading of each tract shows its percent growth in median household income between 1990 and 2014.

Finally, we note that the patterns in Figure 1 robust to CBSA sub-samples, the choice of price deflator, and the use of household income rather than family income. In the next subsections, we explore robustness to downtown definitions, and within detailed demographic groups.

Appendix A.2 U-Shape for Different Downtown Definitions

We now verify that the U-shape patterns of urbanization by income are not qualitatively sensitive to reasonable variation in downtown definitions. Figure A.3 shows a replication of Figure 1 in the main text for 5, 15, and 20% downtown population cut-offs, in addition to our baseline 10% cut-off. The figure also shows alternative downtown definitions - similar to that in Baum-Snow and Hartley (2020) - that include all tracts whose centroid are within 3 or 5 miles from the city center. All these definitions similarly confirm that U-shape urbanization patterns by income become more pronounced from 1990 to 2014. Unsurprisingly, this uptick becomes less pronounced as the definition of downtown becomes geographically larger.

Figure A.3: U-shape of different downtown definitions



Note: This figure uses Census data on family income for the 100 largest CBSAs in 1970, 1990, and 2014. Urban tracts consists of all tracts closest to the city center that account for 5%, 10%, 15%, or 20% of the CBSA's population in 2000, or all tracts with centroid within 3 miles or 5 miles of a CBSA's city center. Each dot in the figure corresponds, on the x-axis, to the median family income within each Census bracket. We compute this median using IPUMS microdata for the corresponding year in the 100 largest CBSAs. All incomes are in real 1999 dollars.

Appendix A.3 U-Shape with Demographic Controls

One may think that these U-shape patterns reflect demographic characteristics that are correlated with income, and/or that the changes in the U-shape pattern over time simply reflect demographic shifts that are correlated with income and that took place between 1990 and 2014. In this section, we replicate Figure 1 showing normalized urban shares by income bracket, but controlling for demographic characteristics.

Unlike Figure 1 that uses Census tables from the 100 largest CBSAs, here we use our 27 CBSAs with constant downtown geography, which allows us to control for demographic characteristics of households. We create demographic control dummies for race, age, family type, nationality of birth, and dual income status.¹ For race, we use the IPUMS definitions.² For age, we construct 5-year age buckets. For family type, we define four categories: Unmarried - No Children, Married - No Children, Youngest Child < 5, and Youngest Child > 5. For nationality of birth, we define two categories: native born and foreign born. We define dual income households as having two adults working at least 30 hours per week. We include interactions of dual income status and family type (after merging married and unmarried with no children into a "no children" family type), which allows us to capture dual income households with no kids (DINK).

To compute urban shares within each income bracket without demographic controls, we estimate the following equation, separately in 1990 and in 2014:

$$\text{UrbanWeight}_i = c + \sum_{k \in K} \beta_k \text{IncomeDummy}_{ki}, \tag{A.1}$$

where UrbanWeight_i is the urban weight of household i , which equals 1 if the household is assigned entirely to the urban area of its CBSA.³ IncomeDummy_{ki} is a dummy equal to 1 if household i is in income bracket k .⁴ The fitted values from this regression are urban shares within each income bracket. To normalize these shares relative to the average household, we divide the fitted value for each income bracket, $\hat{c} + \hat{\beta}_k$, by a weighted average of all fitted values, where each fitted value is weighted by the total number of households in that income bracket. Plotting these normalized fitted values against median income within each income bracket replicates Figure 1 in the paper, but using IPUMS data for our 27 constant geography CBSAs instead of Census tables.

To compute urban shares that control for demographic characteristics, first denote each group

¹The age, race, and nationality at birth are that of the head of household. The youngest age for a head of household with nonzero income is 15.

²We have to merge three categories in 2014 so the definitions are consistent across both periods. These three categories are 'Other race' 'Two major races' and 'Three or more major races'. We match all of these to the 1990 definition 'Other race'. Hispanic is a separate variable in IPUMS. For this analysis, we do not distinguish whether a person is hispanic or not.

³We use a weight instead of a 0/1 dummy because we only know household location at the PUMA level, and some PUMAs span both the urban and suburban areas.

⁴We assign each household into 100 evenly log-spaced household income brackets. We adopt this methodology so brackets are directly comparable between 1990 and 2014, and to ensure large enough population counts in higher income brackets. We merge the bottom 61 brackets with income below \$10,000, and then we drop all income between \$200,000 and \$400,000 that is heavily impacted by topcoding in IPUMS. See Appendix C for further discussion. Our results are robust to different methods of adjusting for topcodes.

of controls (race, age, family type, race, nationality, dual income) by g , and each category within a group by d (e.g., 30-34 year olds). The estimating equation becomes:

$$\text{UrbanWeight}_i = c + \sum_{k \in K} \beta_k \text{IncomeDummy}_{ki} + \sum_{g \in G} \sum_{d \in D} \gamma_{gd} \text{DemoDummy}_{igd}, \quad (\text{A.2})$$

where DemoDummy_{igd} is equal to 1 if household i is in category d within group of controls g . To obtain urban shares within each income bracket k that control for demographic characteristics, we compute fitted values of equation A.2 under the assumption that demographic shares within each income brackets are exactly representative of the demographic shares within the total population. Under this assumption, fitted urban shares are equal to:

$$c + \widehat{\beta}_k + \sum_{g \in G} \sum_{d \in D} \text{SharePop}_{gd} \times \widehat{\gamma}_{gd},$$

where SharePop_{gd} is the share of total population in each category d (e.g., share of 30-34 year olds).

Figure A.4 shows normalized urban shares for each income bracket in 1990 and 2014. The left-hand plots shows estimates from equation (A.1) (without demographic controls) and the right-hand plot shows estimates from equation (A.2) (with demographic controls.)

Our key finding is that controlling for demographics makes the U-shape even more pronounced at the top of the income distribution, in both 1990 and 2014. The regression results from equation A.2 show what drives this finding. In Table A.1, column 1 and 2 show the coefficient on each demographic group dummy in 1990 and 2014, column 3 and 4 show the correlation of each demographic group dummy with household income in 1990 and 2014, and column 5 and 6 show the share of the population within each demographic group. The table shows that there is almost an exact correspondence between the demographic groups that are most suburbanized, wealthiest, and largest. This explains why the urban share of high-income households is larger once we control for demographics; high-income households would be even more urbanized if they weren't also white, middle-aged, and with older children, all of which are suburbanized demographic categories. These first order correlations hold in both 1990 and 2014, so the uptick in the U-shape from 1990 to 2014 largely persists after adding demographic controls.

To further assess whether the U-shape patterns that we document are specific to certain demographic categories, in Figure A.5 we plot normalized urban shares separately by demographic category within each group. To get large enough samples, we further aggregate age (25-34, 35-44, 45-64, 65+), and race (we keep "white", "Black", and an "other" category comprised mostly of Asian, Indigeneous, or multiethnic households.) Remarkably, we find a U-shape pattern, and an urbanization of the richest households from 1990 to 2014 in each category within each group of demographic characteristics.⁵

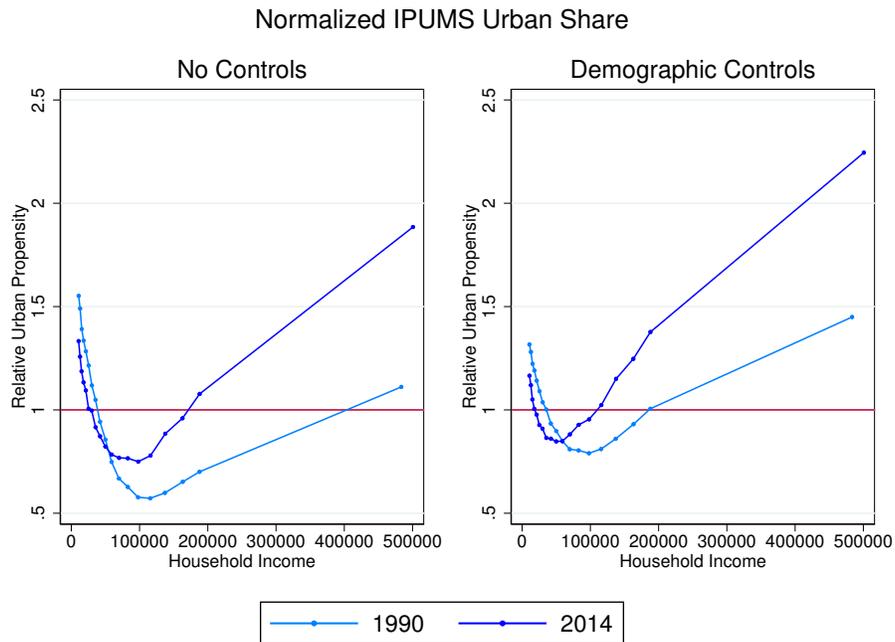
⁵The figure only show dual income with no kids (DINK) and non-DINK households, but we also find U-shape patterns for all dual income households and all non-dual income households.

Table A.1: Coefficients on Demographic Control Dummies in 1990 and 2014

Variable	Coefficient		Correlation with HH Income		Share of Pop Within Group	
	1990	2014	1990	2014	1990	2014
Age < 24 (omitted)	.	.	-0.114	-0.091	0.047	0.032
Age 25-29	0.000	0.021	-0.051	-0.051	0.101	0.071
Age 30-34	-0.006	0.004	0.006	-0.005	0.123	0.091
Age 35-39	-0.007	-0.025	0.049	0.030	0.118	0.092
Age 40-44	-0.014	-0.044	0.098	0.051	0.109	0.098
Age 45-49	-0.016	-0.057	0.118	0.063	0.088	0.102
Age 50-54	-0.018	-0.068	0.099	0.062	0.071	0.106
Age 55-59	-0.018	-0.072	0.067	0.049	0.066	0.100
Age 60-64	-0.019	-0.073	0.010	0.014	0.068	0.087
Age 65-69	-0.030	-0.080	-0.063	-0.014	0.067	0.072
Age 70-74	-0.034	-0.083	-0.096	-0.044	0.055	0.052
Age 75-79	-0.036	-0.087	-0.109	-0.061	0.044	0.038
Age 80-84	-0.040	-0.091	-0.098	-0.066	0.027	0.029
Age 85+	-0.037	-0.098	-0.087	-0.081	0.017	0.030
Native Born (omitted)	.	.	0.041	0.048	0.847	0.752
Foreign Born	0.039	0.028	-0.041	-0.048	0.149	0.248
White (omitted)	.	.	0.150	0.113	0.780	0.690
Black	0.145	0.058	-0.148	-0.132	0.141	0.158
Native American	0.047	0.037	-0.014	-0.017	0.004	0.004
Chinese	0.088	0.042	0.011	0.023	0.010	0.021
Japanese	0.033	0.034	0.016	0.010	0.004	0.003
Other Asian	0.013	-0.000	0.019	0.053	0.020	0.049
Other Race (including mixed raced)	0.120	0.038	-0.073	-0.070	0.041	0.074
Unmarried-No children (omitted)	.	.	-0.288	-0.244	0.343	0.380
Married-No children	-0.052	-0.054	0.111	0.128	0.209	0.200
Youngest Child ≤ 5	-0.040	-0.101	0.012	0.031	0.138	0.105
Youngest Child > 5	-0.029	-0.056	0.189	0.125	0.310	0.315
Non-dual Income (omitted)	.	.	-0.384	-0.299	0.522	0.607
Dual Income	-0.010	-0.009	0.384	0.299	0.478	0.393
Dual Income*No Children (omitted)	.	.	0.201	0.186	0.126	0.103
Dual Income*Youngest Child ≤ 5	-0.034	0.001	0.061	0.068	0.116	0.077
Dual Income*Youngest Child > 5	-0.039	-0.024	0.249	0.175	0.237	0.212

Note: Columns 1 and 2 report the coefficients from equation (A.2) for years 1990 and 2014 with all demographic controls included. The standard errors are small and not shown. Columns 3 and 4 show the pairwise correlation of each demographic control dummy and household income. Columns 5 and 6 report the total share of population falling into each demographic category.

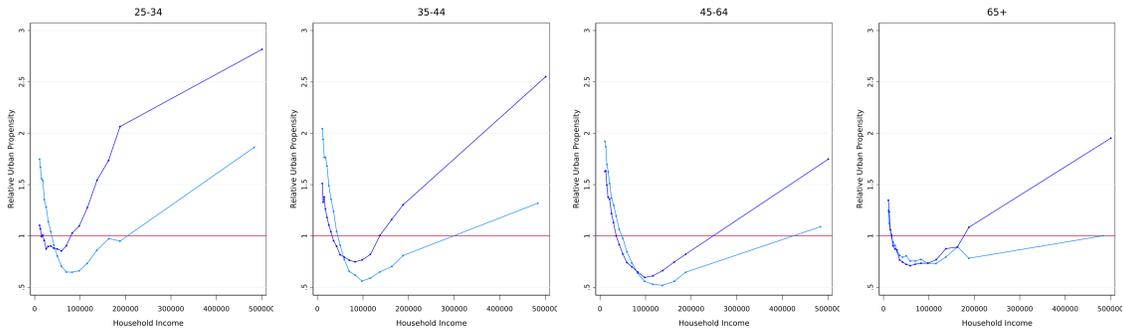
Figure A.4: Impact of Demographic Controls on Relative Urbanization by Income



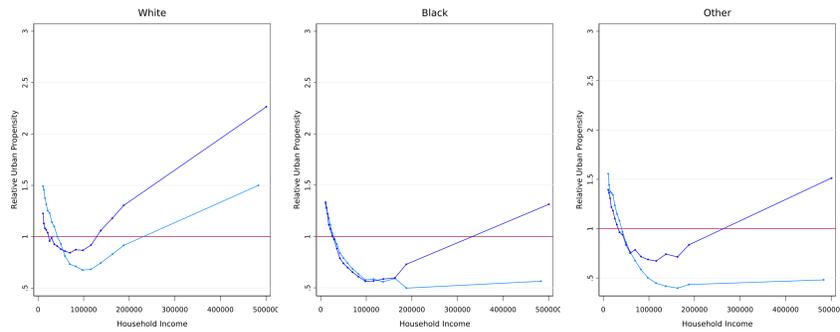
Note: This figure shows urban shares normalized by the aggregate urban share in each year with and without demographic controls. The left plot shows the coefficients from equation (A.1). The right plot shows the coefficients from equation (A.2). We drop household with income below \$10,000 and between \$200,000 and \$400,000. IPUMS data from 1990 and 2014 in 27 CBSAs with constant downtown areas.

Figure A.5: Normalized Urban Shares by Demographic Categories in 1990 and 2014

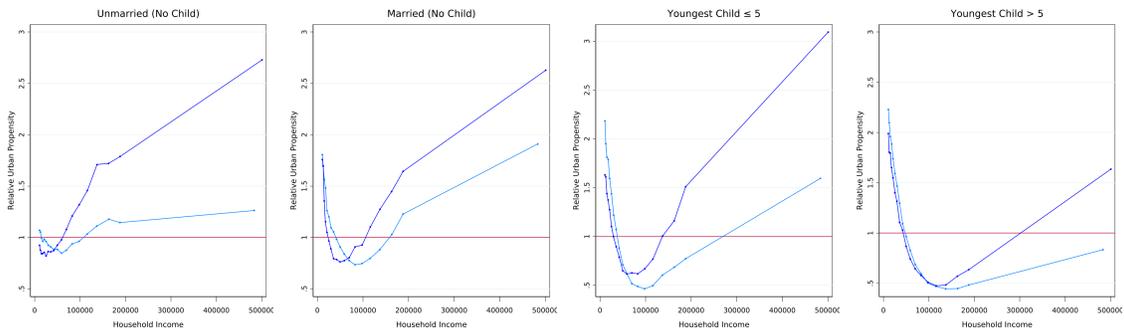
PANEL A: AGE



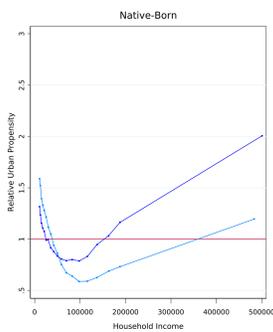
PANEL B: RACE



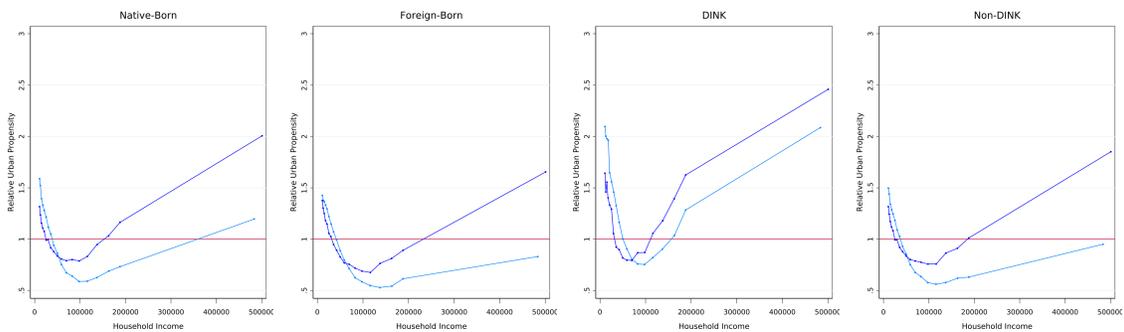
PANEL C: FAMILY TYPE



Panel D: Foreign Status



Panel E: Dual Income with No Kids (DINK)



—●— 1990 —●— 2014

Note: This figure shows normalized urban shares from equation (A.1), plotted separately for each demographic category. The right plot shows the coefficients from equation (A.2). We drop household with income below \$10,000 and between \$200,000 and \$400,000. IPUMS data from 1990 and 2014 in 27 CBSAs with constant urban areas.

Appendix B Model Appendix

Appendix B.1 Income Elasticity of Housing consumption

Neighborhoods differ in the size and quality of their housing units, driven by the quality shifter k_{rq} that in turn impacts housing prices p_{rq} (see equation (6)), but housing is homogeneous within a neighborhood. Therefore, in the model, variation in housing spending patterns across income groups arises solely from variation in the fraction of households who choose neighborhoods of various types. Despite this simple setup, we show now that the model is able to capture salient features of housing consumption in the data. Denoting with $\bar{p}(w) \equiv \sum_{r,q} \lambda_{rq}(w) p_{rq}$ the average expenditure on housing for households of income w , the income elasticity of housing consumption in the model is:

$$\frac{\partial \log \bar{p}(w)}{\partial \log w} = \rho \frac{w}{\bar{p}(w)} \sum_{rq} (p_{rq} - \bar{p}(w)) \lambda_{rq}(w) x_{rq}(w), \quad (\text{B.3})$$

where $x_{rq}(w) = (w - p_{rq})^{-1}$. This income elasticity $\frac{\partial \log \bar{p}(w)}{\partial \log w}$ is strictly positive as soon as the city has more than one type of neighborhoods to choose from (and would be trivially 0 otherwise).

Proof. If there is only one type of neighborhood, one can factorize $x(w)$ and get that:

$$\frac{\partial \log \bar{p}(w)}{\partial \log w} = \rho x(w) \frac{w}{\bar{p}(w)} \sum_{rq} (p_{rq}^h - \bar{p}(w)) \lambda_{rq}(w) = 0,$$

by definition of \bar{p} . If there are several types of neighborhoods, note that $x_{rq}(w)$ increases with p_{rq} . Therefore, $\text{cov}(p_{rq}^h \lambda_{rq}(w) - \bar{p}(w) \lambda_{rq}(w), x_{rq}(w)) > 0$ for any w . The result follows. \square

Appendix B.2 Closing the Model: Neighborhood Development

A developer of neighborhood n of type rq faces total operating cost $\int_w \lambda_r(w) k_{rq} R_r dF(w)$ to serve its demand, for revenues $\int_w \lambda_n(w) p_{rq} dF(w)$, so that its operating profits are:

$$\pi_n = \left[\int_w \lambda_n(w) dF(w) \right] (p_n - k_{rq} R_r). \quad (\text{B.4})$$

Among households with income w , the share that locates in a particular neighborhood n of type (r, q) is $\lambda_n(w) = \lambda_{rq}(w) \lambda_{n|rq}(w)$, where the notation $\lambda_{n|rq}$ indicates the share of workers who choose neighborhood n conditional on choosing a neighborhood of quality q in location r . Given the structure of the idiosyncratic preference shocks, the conditional probability of choosing n among other (r, q) choices is:⁶

$$\lambda_{n|rq}(w) = \frac{V_n(w)^\gamma}{\sum_{n' \in \mathcal{R}(rq)} V_{n'}(w)^\gamma} = \frac{V_n(w)^\gamma}{V_{rq}(w)^\gamma}, \quad (\text{B.5})$$

⁶In equilibrium, all neighborhoods are symmetric within type, so that $\lambda_{n|rq}(w) = \frac{1}{N_{rq}}$.

where $V_{rq}(w)$ is defined in (4) and $V_n(w)$ is the inclusive value of neighborhood n :

$$V_n(w) = (w - p_n)B_{rq(n)}. \quad (\text{B.6})$$

The probability $\lambda_{rq}(w)$ that the neighborhood chosen is of type (r, q) is given by (3). Plugging in the expression for λ_n as a function of price and taking the developer's first order condition for profit maximization leads to the following pricing formula:

$$p_n = \frac{\gamma}{\gamma + 1} k_{rq} R_r + \frac{1}{\gamma + 1} \mathcal{W}_{rq}(p_n), \quad (\text{B.7})$$

where $\mathcal{W}_{rq}(p) = \frac{\int_w (w-p)^{-1} \delta(w) w dF(w)}{\int_w (w-p)^{-1} \delta(w) dF(w)}$ with $\delta(w) = 1\{w - p > 0\}$.

By symmetry, all neighborhoods of type (r, q) have the same price in equilibrium, which we denote as p_{rq} , given in (6). Finally, under free entry, the number of developers entering location r at quality q becomes:

$$N_{rq} = \frac{1}{f_{rq}} \left[\int_w \lambda_{rq}(w) (p_{rq} - k_{rq} R_r) dF(w) \right]. \quad (\text{B.8})$$

Appendix C Data Sources and Sample Descriptions

This section details the data sources used in our various empirical analyses. We also discuss how we adjust our income data for topcoding in the IPUMS data.

Appendix C.1 Census Data and ACS Data

Census Tract Data For our work at the neighborhood level, we assemble a database of constant 2010 geography census tracts using the Longitudinal Tract Data Base (LTDB) and data from the National Historical Geographic Information System (NHGIS) for the 1970-2000 censuses and the 2012-2016 ACS. In each of the censuses from 1970 to 2000, some tracts are split or consolidated and their boundaries change to reflect population change over the last decade. The LTDB provides a crosswalk to transform a tract level variable from 1970 to 2000 censuses into 2010 tract geography. This reweighting relies on population and area data at the census block level, which is small enough to ensure a high degree of accuracy. We combine these reweighted data with the 2012-2016 ACS data ('2014 ACS'), which already uses 2010 tract boundaries.

The Census house price data in 1990 and 2014 is the median house value for all owner-occupied housing units within each census tract. The Census family income data is the count of all families by income brackets in 1970, 1990, and 2014 within each census tract. The Census household income data is the count of all households in each income bracket in 1970, 1990, and 2014 within each census tract.⁷

CBSA Definitions Core Based Statistical Areas (CBSAs) refer collectively to metropolitan and micropolitan statistical areas. CBSAs consist of a core area with substantial population, together with adjacent communities that have a high degree of economic and social integration with the core area. We assign 2010 census tracts to CBSAs based on 2014 CBSA definitions. Our model estimation sample consists of the 100 metropolitan area CBSAs with the largest population in 1990s.

IPUMS Data PUMA geography is also not constant from 1990 to 2014, so we use a crosswalk between PUMAs (Public-Use Microdata Areas) and CBSAs in each year to link each PUMA to a CBSA. To construct constant downtowns from PUMAs across years, we follow the methodology in (Couture and Handbury, 2017). We first intersect PUMA geographies in 1990 and 2014 with our constant downtown geography described in the main text, defined out of tracts closest to the city center accounting for 10 percent of a CBSA's population in 2000. PUMAs generally intersect with both the urban and suburban region of a CBSA, so we assign an urban weight to each PUMA equal to the percentage of that PUMA's population falling within the urban region (i.e., downtown) of

⁷A family is a group of two or more people related by birth, marriage or adoption and residing together. The main difference between households and families in the Census is that families exclude persons living alone, or groups of unrelated people living together. Figure 1 uses family income instead of household income, because household income is top-coded at a much lower value in 1970, at only \$25,000. The family income data is top-coded at \$50,000 in 1970, \$150,001 in 1990, and \$200,001 in 2014.

that CBSA. We compute the urban and suburban population of each PUMA using the population of all census blocks whose centroid falls in a given region.

In most of the 100 CBSAs, PUMAs are too large to accurately represent downtowns. We therefore enforce an inclusion criteria where we only keep CBSAs for which 60% of the urban population lives in PUMAs whose population is at least 60% urban. Under this restriction, we find a set of 27 CBSAs for which we can define urban areas in 1990 and 2014.

Topcoding in IPUMS Data IPUMS data is topcoded by income component. Household and family income reported in the IPUMS data is sum of total individual income for all members of the household or family. Total individual income is the sum of income components where each component has a unique topcode. Table C.2 shows each of the income components that contribute to total individual income and their respective topcodes for 1990 and 2014.

Table C.2: Topcoded Income Components Contributing to Total Individual Income

1990			2000+		
Variable	Description	Topcode (nominal)	Variable	Description	Topcode (nominal)
INCWAGE	Pre-tax wage and salary income	\$140,000	INCWAGE	Pre-tax wage and salary income	99.5th Percentile in State
INCBUS	Non-farm business and/or professional practice income	\$90,000	INCBUS00*	Business and farm income and/or professional practice income	99.5th Percentile in State
INCFARM	Farm	\$54,000	INCSS	Social security and disability	Not Topcoded
INCSS	Social security and disability	\$17,000	INCWELFR	Other government assistance	Not Topcoded
INCWELFR**	Other government assistance	10,000	INCSUPP	Supplementary Security Income	Not Topcoded
INCINVST	Rents, interests, dividends, etc.	\$40,000	INCINVST	Rents, interests, dividends, etc.	99.5th Percentile in State
INCRETIR	Retirement income other than social security	\$30,000	INCRETIR	Retirement income other than social security	99.5th Percentile in State
INCOTHER	Income not included above	\$20,000	INCOTHER	Income not included above	99.5th Percentile in State

* 1990 equivalent is INCBUS + INCFARM

** 2014 equivalent is INCWELFR + INCSUPP

In 1990, component topcodes are the same across all states. Table C.3 shows the percent of all units impacted by topcoding for each component for individuals, households, and families. For households and families, we assume that any household or family whose reported component level income is above the person-level topcode is subject to topcoding for that component. The last row of the table shows the percent of total aggregate income impacted where we apply the individual-level topcode for wages.

In the 2012-2016 ACS, component topcodes vary both across states and year. State-specific topcodes for wages range from \$105,000 to \$280,000 in 1999 dollars. Because of this high variance, we allow each state to retain a state-specific topcode: the minimum topcode in 1999 dollars across

Table C.3: Percent of Income Impacted by Topcoded Components in 1990

Variable	Person	Household	Family
incwage	0.60%	1.39%	1.66%
incbus	3.70%	4.25%	4.62%
incfarm	2.79%	3.28%	3.48%
incss	0.73%	3.67%	5.60%
incwelfr	3.18%	5.80%	6.83%
incinvest	2.11%	2.98%	3.18%
incretir	3.20%	4.15%	5.08%
incother	2.28%	2.51%	2.43%
TOTAL	0.62%	1.78%	2.28%

Note: This table shows the percent of income at or above the topcode value in 1990 among the set of observations where income is non-missing and greater than \$0.

Table C.4: Percent of Income Impacted by Topcoded Components in 2014

Variable	Person	Household	Family
incwage	1.2%	3.5%	4.5%
incbus	3.7%	4.4%	5.1%
incinvest	3.6%	4.6%	5.2%
incretir	3.5%	5.6%	7.2%
incother	3.4%	4.0%	4.0%
TOTAL	1.5%	4.1%	5.8%

Note: This table shows the percent of income at or above the topcode value in 2014 among the set of observations where income is non-missing and greater than \$0. The topcode value is set at the minimum topcode in each state across the five years of ACS (2012-2016).

the 5 years of ACS. Table C.4 shows the percent of income impacted by topcodes for different components and units of observation. The last row of the table shows the percent of total aggregate income impacted where we apply the state specific individual-level topcode for wages.

To interpolate the income distribution above the topcoded values, we turn to the Piketty et al. (2017) methodology to construct We use the R package *gpinter* that Piketty et al. (2017) developed to estimate the generalized Pareto curves for each state s , region r , and period t .^{8, 9,10} The gener-

⁸Since we are only able to define urban cores for the set of 27 CBSAs with sufficiently small PUMAs in 1990 and 2014, we estimate the distribution only for the portion of state s that is covered by a CBSA in that sample.

⁹Generalized Pareto curves allow the pareto coefficient to vary with income. In the U.S. context, the relationship between the Pareto parameter and income has become increasingly U-shaped over time. This would suggest that a simple Pareto, allowing for a single Pareto coefficient, would underestimate the fatness of the right tail of the income distributions, especially in 2014 relative to 1990.

¹⁰The *gpinter* package approximates the income distribution using a set of income percentiles and the average income between each percentile. For each state, area, and period we use the same set for the first 6 percentiles: [10,30,45,60,75,85]. Then based on where the topcode falls for that particular distribution we allow the set of top percentiles to vary. If the topcode percentile p_t falls between the 85th and 92nd percentile, we include no additional moments between the 85th and p_t . If p_t falls between the 92nd and 93rd percentile, our top two percentiles are [

alized Pareto methodology requires an unbiased estimate of average income for some top quantile of income. We approximate the average income above the topcode using a type I (simple) Pareto distribution as in Armour et al. (2016).¹¹ We also use the Armour et al. (2016) simple Pareto to interpolate the income distribution above the topcode for all regions and periods in any state for which the generalized Pareto estimation routine converged to a degenerate distribution for any region-by-period.

We combine the income distribution observed in the IPUMS data below the topcode with the approximated distribution above the topcode. To do this, we first construct a kernel-smoothed CDF using the IPUMS data below the topcode. We then join the below-topcode CDF with the above-topcode CDF from the generalized Pareto distribution. To avoid any kinks around the join point we first adjust the above-topcode CDF such that it matches the CDF at the topcode for the below-topcode. We use numerical differentiation of this CDF to derive the full PDF adjusting for topcoding. To further avoid any kinks around the topcode, we cut incomes within \$1,500 of the topcode and interpolate through the PDF. Using the total population in region r , state s , and period t we use this smoothed PDF to get a population estimate at each \$5,000 interval. Finally, we can aggregate across states to get the urban, suburban, and total distribution for each of the 27 CBSAs in our calibration sample in 1990 and 2014 and across samples to get the urban, suburban, and total distribution for our pooled “representative city” sample.

Our main 1990 calibration and 2014 counterfactual uses the distribution of household income. Only family income is available in earlier decades (1950 and 1970) so, to study backward-looking counterfactuals, we re-calibrate the model using family income, interpolating the distributions for each state and area above the respective topcodes in 1990 and 2014 using the method outlined above, from which we construct the aggregate income distribution and U-shape moment describing the downtown share at each income level. For 1950 and 1970, years in which PUMAs are too large to permit reasonably accurate depictions of the downtown regions of most of the CBSAs we use for our counterfactuals, we instead interpolate a single aggregate income distribution above the topcodes for each year (\$10,000 in 1950 and \$50,000 in 1970). This aggregate income distribution

$85 + \frac{p_t - 85}{2}, p_t]$. If p_t falls between 93rd and 96th percentile are top 3 percentiles are $[85 + \frac{(p_t - 85)}{3}, 85 + \frac{2(p_t - 85)}{3}, p_t]$. If p_t falls above the 96th percentile the top 4 percentiles are $[85 + \frac{(p_t - 85)}{4}, 85 + \frac{2(p_t - 85)}{4}, 85 + \frac{3(p_t - 85)}{4}, p_t]$.

¹¹Armour et al. (2016) apply a type I Pareto distribution to estimate the tail of topcoded income in survey data. Equation C.9 shows their formula for estimating the pareto shape parameter α_{tsn} for time period t , state s , and area n .

$$\widehat{\alpha}_{tsn} = \frac{M_{tsn}}{T_{tsn} \ln(X_{Ts}) + \sum_{x_{m_{tsn}} \leq x_i < x_{T_{tsn}}} \ln(x_i) - (M_{tsn} + T_{tsn}) \ln(x_{m_{tsn}})} \quad (\text{C.9})$$

M_{tsn} is the number of households or families with earnings between the lower cutoff $x_{m_{tsn}}$ and the censoring point $x_{T_{tsn}}$. In 1990, we assign the censoring point as the topcode for a single-wage household adjusted to 1999 dollars (\$188,160). In 2014, for state s we choose censoring point as the topcode for a single-wage household for the year with the lowest topcode of the 5 years surveyed. T_{tsn} is the number of households with income at or above censoring point $x_{T_{tsn}}$. We choose the lower cutoff $x_{m_{tsn}}$ as the 95% income in state s for period t and area n . This is consistent with Armour et al. (2016). We add the additional restriction that at least 1.5% of the total income distribution falls between $x_{m_{tsn}}$ and $x_{m_{tsn}}$ to ensure we have a sufficient data to estimate the shape parameter. If less than 1.5% of the total income distribution falls between those two points we lower the percentile cutoff by 1% percentage point until that condition is met.

is sufficient to feed into the model to conduct the counterfactual prediction for the share downtown in each year. We aggregate this counterfactual prediction to broad income brackets so that it can be compared to the share that we observe in bracketed tract-level data from the Census.

Zillow House Price Indexes Our 2 bedroom index is the Zillow House Value Index (ZHVI) for all two-bedroom homes (i.e., single family, condominium, and cooperative), which is available monthly for 8,030 zip codes in 1996, 8,031 zip codes in 2000, 8,575 zip codes in 2012, and 8,898 zip codes in 2016.¹² In robustness checks, we use the per square foot Zillow House Value Index for All Homes, which is available monthly for 14,417 zip codes in 1996, 14,421 zip codes in 2000, and 15,500 zip codes in 2014. The Zillow indexes are median price estimates for a fixed (over time) set of homes within each zip code. As a result, changes in house prices in the Zillow data can be interpreted as appreciation of the typical home.

For each zip code in the Zillow data, we compute a yearly index by averaging over all months of the year. We map zip codes to tracts with a crosswalk from HUD. We compute the tract-level index as the weighted average of the home value index across all zip codes overlapping with the tract, using as weights the share of residential addresses in the tract falling into each zip code. For tracts falling partly into missing zip codes, we normalize the residential share in zip codes with available data to one. The final data set contains the 2 bedroom index for around 35,000 tracts in each year, and the all home index for around 53,000 tracts in each year.

Appendix C.2 Variable Definitions

This subsection details the computation of variables used in Section 3 onwards in the paper.

CBSA Level Wage Bartik shock We use a Bartik wage shock to predict CBSA-wide average income growth between 1990 and 2014. We determine industry growth using 3-digit Census industry codes in 1990. The Census Bureau provides crosswalks between 2012 and 1990 industry codes. Examples of 3-digit industry categories includes “Aluminum production and processing”, “Shoe Stores”, “Retail Florists”, and “Real Estate”.

To calculate national wage growth for each industry between 1990 and 2014, we use person-level IPUMS data in 1990 and 2014. We keep the sample of people between 21 and 55 years who work at least 35 hours a week in a non-farm profession. We use annual pre-tax wage and salary income for individual earners. As is standard we compute a CBSA-leave out growth for each CBSA.

As in Diamond (2016), we compute wage growth in each industry as the (leave out) difference in average log wage across years. Our Bartik income shock is then wage growth weighted by initial 3-digit 1990 industry shares in each CBSA.

For our robustness specifications, we compute Bartik shocks leaving out two major industry

¹²We downloaded the data in February 2019. The index and methodology are available at: <http://www.zillow.com/research/data/>.

Table C.5: Most and Least Urbanized Industries

Code	Industry	Urban Share
872	Museums, art galleries, historical sites, and similar institutions	24.8%
762	Traveler accommodation	20.6%
151	Cut and sew apparel manufacturing	20.5%
800	Motion pictures and video industries	20.2%
750	Car Washes	19.5%
900	Executive offices and legislative bodies	19.2%
761	Private households	18.9%
721	Advertising and related services	18.4%
951	U. S. Coast Guard	18.4%
542	Apparel, fabrics, and notions wholesalers	17.6%
...
352	Aircraft and Parts	4.0%
622	Other motor vehicle dealers	3.8%
311	Agricultural implement manufacturing	3.6%
950	U. S. Marines	3.4%
561	Farm supplies wholesalers	3.2%
41	Coal Mining	2.8%
362	Guided Missles, Space Vehicles, and Parts	2.7%
821	Office of chiropractors	2.5%
590	Miscellaneous retail stores	2.3%
11	Animal production	1.2%

Notes: This table shows the 3-digit industries with the highest share of workers who live in urban areas in row 1 to 10, and industries with the lowest share in row 11 to 20. IPUMS data from the 27 CBSAs with constant geography urban area in 1990 and 2014. These urban areas contain 10 percent of each CBSA's population in 2000.

categories: Finance, Real Estate and Insurance (1990 industry codes 700-712), and Technology.¹³. In another robustness specification, we drop the most urbanized industries from our Bartik instrument. Table C.5 shows the 10 most urbanized and 10 least urbanized industries, along with the share of urban workers in that industry. The table highlights that even for the most urbanized industry (Museum, Art Galleries, Historical Sites, and Similar Institutions) the share of urban workers is only 24 percent, so that most of the Bartik variation comes from the suburbs. This is because our urban areas, by construction, are small relative to the suburbs.

Median Income within Census Table Brackets The U-shape plot in Figure 1 shows median income within each family income brackets from the NHGIS Census tables. To find the median income within each census bracket, we use the distribution of family income within the 100 largest

¹³181 = Pharmaceuticals; 342 = Electronic component and product manufacturing ; 352 = Aircraft and Parts ; 362 = Aerospace products and parts manufacturing ; 891 = Scientific research and development services ; 732 = Computer systems design and related services + Software Publishing + Data processing, hosting, and related services; 882 = Architectural, engineering, and related services.

CBSAs in the IPUMS microdata in the corresponding year. To adjust for topcoding in IPUMS, we estimate the shape of the IPUMS income distribution above the 95th percentile assuming a Pareto distribution.

The estimation of ρ also requires median income within each census bracket. In this case, however, the estimation requires constant brackets over time. To do this, we assume that households are uniformly distributed within each bracket, except for the top bracket. We can then map each CPI-adjusted census brackets in 1990 onto 2014 bracket definitions, setting median income w as the mid-point of these constant brackets. For the top bracket (above \$140,600 in 1999 dollars), we determine median income using 2000 IPUMS microdata.

Yearly User Cost of Housing ($p_{rq,c}$). We first compute a population weighted-median house price over all tracts in a given area quality pair in a given CBSA. To obtain $p_{rq,c}$ that we use in estimation and calibration, we multiply this median house value by a user cost of housing equal to 5.0 percent of house value in 1996, 4.7 percent in 2000 and 4.6 percent in 2014. These rent-price ratios come from the Lincoln Institute of Land Policy.¹⁴

Property Taxes as a Share of $p_{rq,c}$ Using CPI-adjusted tract-level ACS and Census estimates of the median property taxes for owner-occupied units, we find the population-weighted median amount paid in property taxes in 1990, 2000, and 2014 for each area quality pair. We then divide this amount by $p_{rq,c}$.

Appendix D Robustness of Estimation of ρ

In this appendix, we show the robustness of our estimates of ρ to many different specifications, including different time periods, different house price measures, different definitions of downtown area, and different quality cut-offs. We also conduct pre-trends and balance tests for our Bartik instrument.

Table D.6 replicates the ρ estimation shown in Table 1 of the main text, but replacing the Zillow 2 Bedroom house price index by Census house prices.

Table D.7 shows variants of our preferred OLS and IV estimates of ρ from column 1 and 2 of Table 1. In column 1 and 2 of Table D.7, we change the time period from 1990-2014 to 2000-2014.¹⁵ In column 3 and 4, we change the house price index from Zillow 2 Bedroom to Zillow All Home, and in column 5 and 6 we change the house price index to Census house prices that we hedonically adjust for variation in number of rooms across areas.¹⁶ In column 7 and 8, we change the downtown

¹⁴Data collected in October 2018 from <https://datatoolkits.lincolninst.edu/subcenters/land-values/rent-price-ratio.asp>.

¹⁵When computing the Bartik shock from 2000 to 2014, we keep the initial industry shares fixed to 1990.

¹⁶To compute the hedonically adjusted Census house prices in 1990 or 2014, we regress house prices on number of rooms at the tract-level, and compute a fitted value of house prices in each tract at the national average number of rooms in that year. Note that in 1990, the only available data on rooms is for all housing units, so our computations impute the number of owner-occupied units in 1990 using data from 2000 on the share of all housing units that are owner-occupied, for each number of rooms.

Table D.6: Census Price ρ Estimation

	(1)	(2)	(3)	(4)	(5)
$\hat{\rho}$	1.49 (0.42)	4.06 (2.29)	4.93 (3.35)	5.15 (3.78)	3.19 (1.85)
Instrument	None	Base	Omit Top Urban Industries	Omit FIRE Industries	Omit Hi-Tech Industries
R ²	0.061				
KP F-Stat		2.04	1.63	1.63	1.74
Obs	7,226	7,226	7,226	7,226	7,226

Notes: This table shows estimates from equation (14). Data from 100 largest CBSAs from 1990 and 2014. Each observation is weighted by the number of households in the income bracket with the fewest households amongst the four brackets in each independent variable. Column 3 to 5 also control for share of omitted industries. Standard errors clustered at the CBSA-quality level are in parentheses. KP F-Stat = Kleibergen-Paap Wald F statistic.

definition from 10% to 15% of the population living downtown in 2000, and in column 9 and 10 we change the downtown definition to all tracts with centroids within 5 miles of the CBSA center.

Table D.8 shows robustness to changes in our high quality cut-off from at least 40 percent college share, to at least 30, 50, or 60 percent college share. Across all three robustness tables D.6, D.7 and D.8, estimates range from 1.43 to 5.15. This range roughly overlaps with 1.64 to 4.44, which is a two standard deviation range around our preferred estimate of 3.04 from column 2 of Table 1.

Finally, we conduct pre-trend and balance tests for our Bartik shock instrument similar to those proposed by Borusyak et al. (2021). To test for pre-trends, we regress the pre-period dependent variable in our ρ estimation (change in urbanization by income, from equation 14, computed for the pre-period of 1970 to 1990), on our Bartik shock (computed from 1990 to 2014.) Table D.9 shows the result of this regression, in row 1 of Panel A. The coefficient of the Bartik shock on the pre-period change in urbanization by income is not significant. We note that our definition of a high-quality neighborhoods - at least 40 percent college share - may be too restrictive for 1970, when college shares were lower. As an additional test in row 2, we reduce the college share quality cut-off to 30 percent, and also find no evidence of pre-trends.

The balance tests are shown in Panel B of Table D.9. We estimate the correlation between initial 1990 CBSA characteristics and the Bartik shock from 1990 to 2014. Row 1 shows the Bartik shock correlation with initial level of house prices in downtown relative to the suburbs, row 2 shows the Bartik shock correlation with the initial share of households in high-quality tracts in downtown relative to the suburbs, and row 3 shows the Bartik shock correlation with the initial share of high-income households (earning more than \$100,000) in downtown relative to the suburbs. All of these correlations are smaller than 0.14. Coefficients from the bivariate regressions corresponding to these correlations are similarly insignificant.

Table D.7: Robustness Exercise for ρ Estimation

	2000-2014		Zillow All Home		Hedonic Census		15 pct Downtown		5 miles Downtown	
	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV	(7) OLS	(8) IV	(9) OLS	(10) IV
$\hat{\rho}$	1.70 (0.33)	3.19 (1.11)	2.45 (0.36)	2.91 (0.96)	1.43 (0.31)	1.78 (0.59)	2.21 (0.50)	2.80 (0.84)	2.73 (0.33)	2.66 (0.91)
Instrument	None	Base	None	Base	None	Base	None	Base	None	Base
R ²	0.12		0.19		0.10		0.16		0.24	
KP F-Stat		8.82		5.40		4.11		25.1		26.8
Obs	6111	6111	6879	6879	7156	7156	6426	6426	6496	6496

Notes: This table shows estimates from equation (14). Data from 100 largest CBSAs, from 2000 and 2014 in column 1 and 2 and from 1990 to 2014 in all other columns. Downtown of each CBSA consists of all tracts closest to the center of each CBSA that account for 10% of that CBSA's population in 2000 in column 1 to column 6, 15% of the CBSA population in column 7 and 8, and all tracts with centroid within 5 miles of the city center in column 9 and 10. The house price index used in estimation is the Zillow All Home Index in column 3 and 4, Census median house price hedonically adjusted for number of rooms in column 5 and 6, and the Zillow 2 Bedroom Index in all other columns. Each observation is weighted by the number of households in the income bracket with the fewest households amongst the four brackets in each independent variable. Standard errors clustered at the CBSA-quality level are in parentheses. KP F-Stat = Kleibergen-Paap Wald F statistic.

Table D.8: ρ College Cutoff Robustness

	30		40		50		60	
	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV	(7) OLS	(8) IV
$\hat{\rho}$	1.84 (0.47)	2.02 (0.53)	2.48 (0.34)	3.04 (0.70)	2.51 (0.35)	4.48 (1.04)	2.33 (0.37)	4.34 (1.40)
Instrument	None	Base	None	Base	None	Base	None	Base
R ²	0.15		0.21		0.22		0.18	
KP F-Stat		23.9		26.3		17.0		14.5
Obs	6,235	6,235	5,878	5,878	5,452	5,452	4,734	4,734

Notes: This table shows estimates from equation (14). Data from 100 largest CBSAs in 1990 and 2014. Columns 1 and 2 define high-quality tracts as those that contain 30% of residents with at least a bachelor's degree, column 3 and 4 as at least 40%, column 5 and 6 as at least 50% and column 7 and 8 as at least 60%. Each observation is weighted by the number of households in the income bracket with the fewest households amongst the four brackets in each independent variable. Standard errors clustered at the CBSA-quality level are in parentheses. KP F-Stat = Kleibergen-Paap Wald F statistic.

Table D.9: Pre-Trend and Balance Tests of Bartik Instrument

PANEL A: PRE-TREND				
Dependent variable from ρ estimation in pre-period (1970 to 1990)	Coefficient of 90-14 Bartik	Std. Er.	Obs.	
Main Specification	0.28	(0.92)	5,833	
College Share Cut-off of 30%	2.23	(1.42)	7,023	

PANEL B: BALANCE TESTS				
Variable in initial level (1990)	Coefficient of 90-14 Bartik	Std Er.	Correlation with 90-14 Bartik	Obs.
Relative house prices (p_{cD}/p_{cS})	0.0012	(0.0027)	0.14	1,376
Relative share of households in high-quality tracts ($\lambda_{cDH}/\lambda_{cD}$)/($\lambda_{cSH}/\lambda_{cS}$)	0.000016	(0.000029)	0.10	1,449
Relative share of high-income households ($> \$100,000$), ($\lambda_{cD,100k}/\lambda_{cD}$)/($\lambda_{cS,100k}/\lambda_{cS}$)	0.0035	(0.0039)	0.10	1,600

Notes: Panel A of the table shows estimates from a regression of the dependent variable from the ρ estimation in equation 14 computed in the pre-period of 1970 to 1990, on the 1990 to 2014 Bartik income shock. Row 1 shows estimates using our baseline definition of high-quality neighborhoods as having a college share above 40%, and row 2 shows estimate using a 30% college share. Standard errors are clustered at the CBSA-quality level. Each observation is weighted by the number of households in the income bracket with the fewest households amongst the four brackets in each independent variable. Panel B shows coefficients from CBSA-level regressions with the 1990 to 2014 Bartik shock as an independent variable and initial 1990 levels, as shown in the table, as dependent variables. The downtown of each CBSA consists of all tracts closest to the center of each CBSA that account for 10% of that CBSA's population in 2000. Data from the 100 largest CBSA.

Appendix E Reduced Form Evidence of Amenity Response to CBSA Income Growth

In the model, changes in demand for locations are amplified by endogenous neighborhood change. As high-income households move downtown, neighborhoods with low-quality amenities are replaced by neighborhoods with high-quality amenities. In this appendix, we provide some reduced form empirical evidence for this model mechanism. Specifically, we plot the growth in restaurant quality downtown minus that for the suburbs against the Bartik income shock. To measure tract-level restaurant quality, we first leverage novel smartphone movement data (Couture et al., 2021), described below, to measure the quality of the 100 largest restaurant chains in the smartphone data. We define chain quality as the propensity of high-income individual to visit a given chain, relative to the propensity of the average individual, controlling for their relative proximity to establishments. Then, we measure tract-level restaurant quality using the National Establishment Time-Series (NETS) geocoded census of local restaurants in 2000 and 2012.

We now describe the smartphone and NETS data, before explaining in detail how we compute tract-level restaurant quality.

Smartphone Movement Data The smartphone movement data is from October 2016 to August 2018. Our data provider aggregates data from multiple apps' location services.¹⁷ Each visit comes from raw movement data intersected with a basemap of polygons (usually buildings). Each visit receives a unique location, device, and time stamp.

We define the permanent home location of each device as in Couture et al. (2021), using 90 billion visits to residential establishments. We first identify a individual's weekly home location as the residential location where it spends most night hours, conditional on visiting that location at least three different nights that week. We then assign permanent home location to any device that has the same weekly home location for three out of four consecutive weeks. We are able to identify permanent homes for 87 million devices between 2016 and 2018. We refer to this location as the device's home location.

We have 600 million visits to chain restaurants in our smartphone sample. In our estimation, we restrict attention to 60 million of these visits that we can identify as starting from home using the time stamp and duration of each visit. We define a visit as starting from home if the previous visit was to home and ended less than 60 minutes earlier.¹⁸ We refer to Couture et al. (2021) for additional details on that data.

¹⁷Athey et al. (2018) and Chen and Rohla (2018) use similar smartphone data from a different provider. We refer to Couture et al. (2021) for evidence that the spatial distribution of smartphone devices provides a balanced representation of the U.S. population along a number of dimensions (CBSA, income, race, education); the distance traveled to different destinations implied by the smartphone data resembles that from the NHTS travel survey; and the mapping of commercial establishments visited by smartphone users to the business registry is relatively complete.

¹⁸We do not observe all travel by individuals, so visit duration is a lower bound and missing in some cases. This explains why we are only able to ascertain 10 percent of trips as starting from home, whereas for instance about 30 percent of trips to restaurants in the NHTS start from home.

NETS Data The 2012 National Establishment Time-Series (NETS) Database includes 52.4 million establishments with time-series information about their location, industries, performance and headquarters from 1990-2012. The NETS dataset comes from annual snapshots of U.S. establishments by Duns and Bradstreet (D&B). D&B collects information on each establishment through multiple sources such as phone surveys, Yellow Pages, credit inquiries, business registrations, public records, media, etc. Walls & Associates converts D&B’s yearly data into the NETS time-series. The NETS data records the exact address for about 75 percent of establishments. In the remaining cases, we observe the establishments zip code and assign it’s location to the zip code centroid.

Neumark et al. (2007) assess the NETS reliability by comparing it to other establishment datasets (i.e., QCEW, CES, SOB and BED data). Their conclusions support our use of the NETS data to compute a long 12-year difference from 2000 to 2012. They report that NETS has better coverage than other data sources for very small establishments (1-4 persons), which is often the size of consumption amenity establishments.

Table E.10 shows the number of establishment in the smartphone data basemap for the ten largest restaurant chains, compared with recent estimates of the actual number that we found online and the 2012 NETS data. This comparison shows that the smartphone basemap is nearly complete, with one exception, Starbucks, where almost half of the establishments are missing from the smartphone basemap. Additionally Table E.10 suggests that the NETS database is less complete than the smarphone basemap, but some of this difference is due to the earlier count. The NETS contains a majority of establishments for nine of the ten largest chains, with the exceptions of Subway where the NETS misses more than half the actual number of establishments.

Estimating Chain Quality We define quality for the 100 largest restaurant chains, with the most establishments in the smartphone data basemap. We index block groups by i , venues by j , and chains by c . We denote by N_{ic} the total number of visits by individuals living in block i that start from home and end in venues in chain c within its CBSA. Restricting our sample of chain visits to those that start from a person’s home isolates the choice of visiting a chain from other considerations of travelers (e.g., eating during lunch at work).

We further control for proximity to venues within that chain to isolate chains that high-income people like from chains that simply co-locate with them. Our main specification has two controls for proximity of block i to venues in chain c : first the normalized straightline distance between the centroid of block i and the closest venue j in chain c , denoted by $dist_{ic(closest)}$, and second the normalized number of establishments in chain c within 5 miles of block i , denoted by $num5mil_{ic}$.¹⁹ Our estimation sample consists of 2.3 million block*chain pairs with at least one within-CBSA visit from home. In a first step, we purge the number of visits from the impact of proximity to chains

¹⁹We normalize $dist_{ic(closest)}$ to equal 1 at the median distance of the closest venue for that chain, computed across all blocks with at least one visit to that chain. The variable $dist_{ic(closest)}$ is then in multiples of that median distance. We do this to ensure that our distance-adjusted number of visits remains unchanged for a block at median distance from chain c .

Table E.10: Ten Largest Restaurant Chains in NETS vs Smartphone Data

Chain	NETS 2012 Rank	Smartphone 2016 Rank	NETS 2012 Count	Smartphone 2016 Count	Most Recent Actual Count
Subway	1	1	10,946	25,889	24,000+
McDonalds	2	2	9,889	14,914	14,000+
Starbucks	3	3	6,581	7,636	14,000+
Pizza Hut	4	6	5,754	6,695	7500+
Burger King	5	5	5,660	7,011	6500+
Wendys	6	8	4,127	5,683	5000+
Dunkin Donuts	7	4	4,030	7,418	8500+
KFC	8	14	3,997	3,157	4000+
Taco Bell	9	7	3,544	6,102	6000+
Dairy Queen	10	10	3,380	4,199	3500+

Notes: The data source from the most recent actual count obtained on 19 December 2018 from the following websites:

Subway: <https://www.subway.com/en-US/exploreourworld>

McDonald: <https://news.mcdonalds.com/our-company/restaurant-map>

Starbucks: <https://www.loxcel.com/sbux-faq.html>

Pizza Hut: <https://locations.pizzahut.com/>

Burger King: <https://locations.bk.com/index.html>

Wendys: <https://locations.wendys.com/united-states>

Dunkin Donuts: <https://www.dunkindonuts.com/en/about/about-us>

KFC <http://www.yum.com/company/our-brands/kfc/>

Taco Bell: <http://www.yum.com/company/our-brands/taco-bell/>

Dairy Queen: <https://www.qsrmagazine.com/content/23-biggest-fast-food-chains-america>

by running:

$$\ln N_{ic} = \beta_1 + \beta_2 \ln(\text{dist}_{ic(\text{closest})}) + \beta_3 \ln(\text{num5mil}_{ic}) + \epsilon_{ic}.$$

We then compute a number of visits purged of proximity as:

$$\widehat{N}_{ic} = \exp\left(\ln N_{ic} - \widehat{\beta}_2 \ln(\text{dist}_{ic(\text{closest})}) - \widehat{\beta}_3 \ln(\text{num5mil}_{ic})\right)$$

In the next step, we compute the relative propensity of high-income devices to visit each chain, relative to the average device. We assign income at the block group level, and define as high-income block groups that had median income of \$100,000 per year in 1999 dollars in the 2014 ACS. The share of visits to chain c out of total visits to the 100 largest chains, among individuals living in high-income block groups, is:

$$S_c^{High} = \frac{\sum_{i \in I_c} \widehat{N}_{ic}^{High}}{\sum_{c=1}^{100} \sum_{i \in I_c} \widehat{N}_{ic}^{High}},$$

where I_c is the set of block groups with a positive number of visits to chain c . We can then define the quality of chain c as the propensity of individuals in high-income block groups to visit chain c relative to that of individuals in the average block:

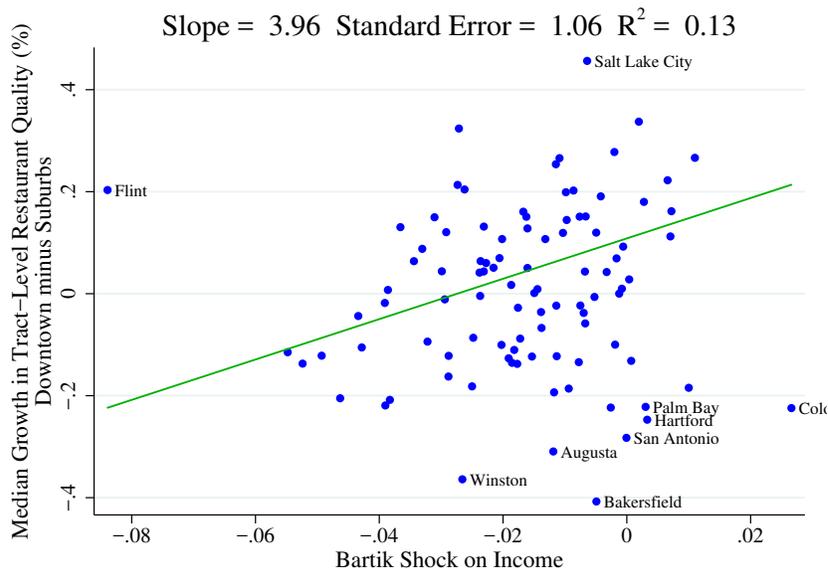
$$Quality_c = \frac{S_c^{High}}{S_c},$$

where $Quality_c = 1$ means that high-income individuals are as likely to visit chain c as the average device, controlling for differences in proximity to venues in chains c . Among the restaurant chains with the highest quality are smaller gourmet chains like Shake Shack (1st), Zoës Kitchen (2nd) and California Pizza Kitchen (3rd), as well as large national chains like Chipotle (6th), Panera Bread (7th), and Starbucks (14th).

We perform a number of robustness checks. First, we note that excluding block*chain pairs with zero visits from home is likely to bias our quality index against chains that locate far from high-income residents. We experiment with including all block*chain pairs with zero visits in our regression and index computation, and obtain an index with a correlation of 0.94 with our preferred index.²⁰ We also experiment with different income cut-off and find that an index defining high-income blocks as having median income above \$75,000 has a correlation of 0.93 with our preferred index. Finally, we experiment with adding controls for number of chains farther away than 5 miles, and for demographic similarity between block i and the block in which the closest venue in chain c is located (median income difference, age difference, share college difference, EDD measure of racial dissimilarity in Davis et al. 2019). The correlation of these indices with our preferred chain quality index is above 0.98.

²⁰In that case, $N_{ic} = 0$ gets adjusted upward if the closest venue to block i is farther than median distance, and therefore included as a positive number of visits in the index computation, possibly creating the opposite bias as in our preferred specification. We use the invert hyperbolic sine transform to allow for log of zeros.

Figure E.6: Difference in Amenity Quality Growth between Downtown and the Suburbs Across CBSAs with Different Aggregate Bartik Income Shocks



Note: This figure plots the difference in the change in the amenity quality of tracts downtown between 1990 and 2014 and the amenity quality of tracts in the suburbs over the same period across CBSAs. This quality growth differential is plotted against CBSA-level Bartik income shocks over the same period for the largest 100 CBSAs in 1990. Tract amenity quality growth in an area is measured as the percent change in the population-weighted median restaurant chain quality index across tracts in the area. The line through the scatterplot shows the CBSA-population weighted linear fit.

From Chain Quality to Tract-level Restaurant Quality In the NETS data, we can find all of the 100 largest chains in the smartphone data in 2012, accounting for 64,000 establishments, and 96 chains in 2000, accounting for 49,000 establishments.²¹ We compute restaurant quality at the tract level as the average quality of all chains within the tract. If a tract contains fewer than 3 chains, we take the average over all tracts with centroid within 0.25 mile from the tract, and so on in further 0.25 mile increment until there are at least 3 chains. We set a limit of 1.5 miles in urban areas, and 3 miles in suburban areas, below which we set quality to missing if there are still fewer than 3 chains within that limit. This procedure generates 4 percent missing tracts in urban areas, and 15 percent in suburban areas.²²

Results Figure E.6 shows the change in population-weighted median tract restaurant quality downtown minus that for the suburbs against the CBSA Bartik income shock (defined in the main paper). The figure shows that downtowns restaurant quality became relatively higher than that in the suburbs in response to CBSA income growth between 1990 and 2014. The association between

²¹The earliest NETS data is in 1992, but we cannot reliably define tract quality so far back in the past, because too many of the largest chains in our 2016-2018 smartphone data only experienced national growth after 1992.

²²For urban tract, there are at least three chains within tract for 15 percent of tracts, within 0.5 miles for 29 percent of tracts, and within 1 mile for 73 percent of tracts. For suburban tracts, there are at least three chains within tracts for 20 percent of tracts, within 0.5 miles for 25 percent of tracts, within 1 mile for 55 percent of tracts, and within 2 miles for 86 percent of tracts.

large income growth and outsized amenity quality growth downtown relative to the suburbs is driven, for the most part, by absolute amenity quality growth downtown. This finding is consistent with the prediction of our model’s key amplification mechanism. As high-income households move downtown, the amenity mix of downtown neighborhoods change towards those consumed by high-income individuals.

Appendix F Empirically Measuring Housing Expenditure Engel Curves

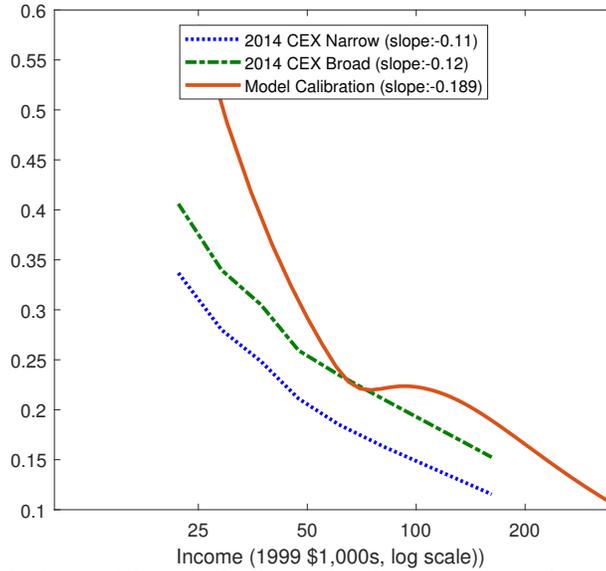
To create empirical measures of how the share of total expenditure spent on housing varies with income we use “Housing” spending by income deciles as reported in public release tables from the Consumer Expenditure Survey (CEX).²³ The CEX “Housing” category includes the following sub-categories: *Shelter*, *Utilities*, *Household Operations*, *Housekeeping Supplies*, and *Household Furnishings*. *Shelter* expenditures include rent paid by renters, mortgage and interest charges paid by homeowners, property taxes paid, and maintenance expenditure. *Shelter* also includes a sub-category of “Other Lodging” which aggregates spending on hotels, vacation rentals and child dorm expenditures. *Household Operations* and *Household Supplies* both include spending done to clean and maintain the household. The former measures spending done to pay others to maintain the household and includes, for example, spending on housekeeping services, gardening and lawn care services, dry cleaning services and home security services. The latter measures spending done if the individual was going to maintain the house themselves and includes, for example, spending on cleaning supplies, spending on laundry supplies, and spending on gardening items. Finally, *Household Furnishings* include spending on furniture, home appliances, and other household furnishings (like lamps, rugs, closet organizers, and home computers).

In our model, the notion of housing expenditure is any spending needed to own and maintain a house in a given neighborhood type. There is some inherent judgment on how to match the model notion of housing expenditure to its empirical analog. We explore two empirical notions of housing expenditure from the CE. First, we use a broad measure of housing expenditures defined as all CEX *Housing* expenditures excluding expenditures on *Other Lodging*. This includes spending on *Shelter*, *Utilities*, *Household Operations*, *Housekeeping Supplies* and *Household Furnishings* but excludes spending on hotels, vacation homes and child dorm expenditures. Second, we also explore a narrower measure of spending that focuses on just *Shelter* (less *Other Lodging*) and *Utility* expenditures. This narrower measure excludes the other components of the *Housing* expenditure category like furniture and spending associated with gardening and cleaning. Figure F.7 below shows the share of housing spending relative to income in the model relative to the data. For the data, we show the patterns using both the broad and narrow spending measures.

As seen from the figure, we compare the slopes of the housing spending “Engel Curves” for

²³See <https://www.bls.gov/cex/tables/calendar-year/mean-item-share-average-standard-error.htm#cu-income>.

Figure F.7: Housing Expenditure Share by Income



Notes: This figure shows both the model implied relationship between share of income allocated to housing (solid line) and the corresponding data from the Consumer Expenditure Survey in 2014 (the dashed line uses the broad housing expenditure measure, the dotted line uses the narrow housing expenditure measure). In published tables, the CEX reports average household income and expenditure for each household income decile. The earliest year that the CEX reported average expenditures by income decile was 2014. The CEX line includes only eight data points, one for each decile where decile income exceeded \$20,000 in 1999 dollars.

the model and for the two separate empirical measures. To do so, we simply regress each housing expenditure share measure on log income. As discussed in the paper paper, the model yields a semi-elasticity of -0.19; a 10 percent in income is associated with the housing spending share falling by -1.9 percentage points. The semi-elasticities associated with the broad and narrow expenditure measures are -0.12 and -0.11, respectively. Notice, while the level of the spending share varies with how we define spending, the slope of the relationship is roughly similar regardless of whether we use the broad or narrow CEX expenditure measure.

There is also a question of whether we should use an empirical measure of housing expenditure shares relative to income or relative to expenditure. In our static model, total income and total expenditure are the same. Empirically, an individual’s current income can differ from their expenditure due to borrowing/savings or due to differential measurement error in the reporting of income and expenditure. As a result, we also explore the empirical relationship of our broad and narrow housing expenditure measures relative to total CEX expenditures. For our broad housing expenditure measure relative to total expenditures, we get a semi-elasticity estimate of -0.067. For our narrow housing expenditure measure relative to total expenditures, we get a semi-elasticity estimate of -0.073. The empirical estimates of the housing Engel curves using housing shares relative to expenditure are even flatter than what we find relative to income. Given these results, in the paper, we explore the robustness of our results when we target various empirical measures of the relationship between housing expenditure and income (or expenditure).

Appendix G Quantification Appendix

This appendix provides more details on how we select the parameters used in the baseline calibration and the method of moments procedure to calibrate the model to the baseline equilibrium for 1990.

Appendix G.1 Parametrization of Land Market Transmission Mechanism (ϵ_S and ϵ_D)

In the model, the area-specific elasticity of land supply ϵ_r is equivalent to an elasticity of housing supply. This elasticity determines the strength of an important welfare transmission mechanism through land markets. When housing supply is inelastic, an influx of rich households in high-quality neighborhoods downtown raises rents for poor incumbent households in low-quality neighborhoods. Saiz (2010) provides housing supply elasticity estimates ϵ_c for 95 large Metropolitan Statistical Areas, based on geographical constraints and housing regulations. We match 83 of these MSAs to our CBSA sample. Unfortunately, these are not estimated separately for downtown and suburban areas. To calibrate ϵ_D and ϵ_S , we posit that housing supply elasticities vary systematically, in equilibrium, with average household density ($density_c$), and estimate the following log-linear regression of ϵ_c on $density_c$:

$$\ln(\epsilon_c) = 1.97 - \frac{0.30}{(0.07)} \ln(density_c) + \xi_c^e, R^2 = 0.21 \quad (\text{G.10})$$

We rely on cross-CBSA variation to estimate this equation. We then define $\hat{\epsilon}_D$ and $\hat{\epsilon}_S$ as the fitted values from equation (G.10) computed at typical density of D and S neighborhoods in the 100 largest CBSAs. We find $\hat{\epsilon}_D = 0.60$ and $\hat{\epsilon}_S = 1.33$.²⁴ We use these values in our baseline calibration and test the sensitivity of our results to alternative parameter values.

Appendix G.2 Parameterization of Public Amenities and Homeownership

A household earning labor income w , receives a transfer of $\chi(w) = OS(w)\lambda_{1999,rq}(w) \sum_{rq} (p_{2014,rq}^h - p_{1990,rq}^h)$, where $OS(w)$ is the share of households with income w who reported owning homes in the 2000 IPUMS data (see Table G.11). This allows us to forgo taking a stance on the initial level of $\chi(w)$ and instead only focus on the changes in $\chi(w)$ over time that results from house price growth due to the income inequality shock that we study.

²⁴In our downtowns, the average CBSA population-weighted household density is 4,300 households per square mile, versus 300 in the suburbs. The highest density CBSA, New York, has 850 households per square mile, so the average density in D is out-of-sample. However, $\hat{\epsilon}_D = 0.60$ turns out to equal the elasticity of housing supply in Miami, which is the metropolitan area with the most inelastic housing supply in Saiz (2010).

Table G.11: Income and Homeownership Rates by Income Decile

Decile	Income (\$1,000s)			2000 Homeowner Share	
	1990	2014	% Growth	Downtown	Suburbs
1	28.3	27.9	-1.14	32%	49%
2	34.7	33.9	-2.14	35%	53%
3	41.2	40.5	-1.83	39%	57%
4	48.2	47.9	-0.59	43%	62%
5	55.8	56.4	1.18	47%	68%
6	64.4	66.4	3.10	51%	73%
7	74.5	78.5	5.34	55%	77%
8	87.8	95.2	8.52	60%	82%
9	108.6	122.7	13.02	65%	86%
10	164.2	197.0	20.02	71%	91%

Notes: This table shows the average income in 1990 and 2014 and 2000 homeownership rate of households by income decile in 2000, using data from IPUMS. See Appendix C for details on this data source.

Appendix G.3 Calibration Details

The calibration procedure minimizes the following objective function over parameters θ :

$$\hat{\theta} = \arg \min_{\theta = \left\{ \frac{p_{rq}}{p_{DL}}, \frac{B_{rq}}{B_{DL}} \right\}_{rq \in \{DH, SL, SH\}}} \hat{g}(\theta)' \mathbf{W} \hat{g}(\theta)$$

where $g(\mathbf{X}; \theta) = [g^1(\mathbf{X}; \theta), g^2(\mathbf{X}; \theta)]$ consists of two vectors of moments. The first set of moments matches the propensity to reside downtown by income w as predicted by the model ($\lambda_D(w; \theta) = \lambda_{DL}(w; \theta) + \lambda_{DH}(w; \theta)$ for $\lambda_{rq}(w; \theta)$ defined in 3) to that observed in the data ($\hat{\lambda}_D(w)$):

$$g_1(\theta) = \lambda_D(w; \theta) - \hat{\lambda}_D(w)$$

for incomes w in \$5,000 increments from \$25,000 to \$500,000 (in 1999 US dollars). The second set of moments matches the predicted price in each of $rq \in \{DH, SL, SH\}$ relative to that in the DL (using equation 6 to the relative price observed in the data (\hat{p}_{rq}):

$$g_2(\theta) = p_{rq}(\theta) - \hat{p}_{rq}$$

The weighting matrix used in the baseline calibration attributes a weight of 2,000,000/96 to each element of the first (U-shape) moment and 50/3 to each element of the second (relative price) moment.

Appendix H Counterfactual and Welfare Appendix

In this appendix we provide further details on the counterfactual simulation procedure and welfare calculations from Section 5 of the paper. We then present robustness on these results.

Appendix H.1 Computing Counterfactuals

In this section, we derive the system of equations that we use to solve for a counterfactual equilibrium for a different income distribution $F'(w)$, conditional on (i) an initial calibration corresponding to the income distribution $F(w)$, and (ii) the model elasticities $\{\rho, \gamma, \epsilon_r, \Omega\}$. To compute this equilibrium, we require the calibrated values at the initial equilibrium for $\{\lambda_{rq}(w), p_{rq}\}$, where L_{rq} is the total population living in neighborhoods of type $\{r, q\}$ in the initial equilibrium, i.e.:

$$L_{rq} = \int F(w) \lambda_{rq}(w) dw.$$

We write a counterfactual equilibrium in changes relative to the initial equilibrium, denoting by $\hat{x} = \frac{x'}{x}$ the relative change of the variable x between the two equilibria. The counterfactual equilibrium is the solution to the following set of equations for $\{(p_{rq})', \lambda'_{rq}(w), L'_{rq}\}$ –or, equivalently, their “hat” values.

First, changes in housing costs are given by:

$$\hat{R}_r = \left(\sum_q \hat{R}_r \hat{L}_{rq} \right)^{\frac{1}{1+\epsilon_r}}. \quad (\text{H.11})$$

Note that $\hat{L}_{rq} = \frac{\int \lambda'_{rq}(w) dF'(w)}{\int \lambda_{rq}(w) dF(w)}$ where $\lambda'_{rq}(w)$ is unknown and a solution of the system of equations described here and the counterfactual distribution of income $F'(w)$ is taken as given.

Second, housing prices in the new equilibrium are defined by²⁵:

$$(p_{rq})' = \frac{\gamma}{\gamma+1} k_{rq} R_r \hat{R}_r + \frac{1}{\gamma+1} \mathcal{W}'_{rq}(p'_{rq}), \quad (\text{H.12})$$

where the function $\mathcal{W}'_{rq}(p)$ is defined by:

$$\mathcal{W}'_{rq}(p) = \frac{\int_w \Lambda'_{rq}(p, w) [w + \chi(w)'] F'(w) dw}{\int_w \Lambda'_{rq}(p, w) F'(w) dw}, \quad (\text{H.13})$$

with $\Lambda'_{rq}(p, w) = \frac{\lambda'_{rq,r}(w)}{[w + \chi(w) - p]}$. The $\chi(w)$ are assumed constant between the two equilibria.

Third, define overall neighborhood attractiveness $\tilde{B}_{rq} = N_{rq}^{\frac{1}{\gamma}} B_{rq} G_r^\Omega$. The change in neighborhood attractiveness is:

²⁵Note that $k_{rq} R_r$ is known in the initial equilibrium using equation B.7 and the known variables $p_{rq}, \lambda_{rq,r}(w), F(w)$

$$\widehat{B}_{rq} = \widehat{N}_{rq}^{\frac{1}{\gamma}} \widehat{G}_r^{\Omega} \quad (\text{H.14})$$

In this expression, the change in the number of neighborhoods is given by:

$$\widehat{N}_{rq} = \widehat{L}_{rq} \frac{(p_{rq})' - k_{rq} R_r'}{p_{rq} - k_{rq} R_r},$$

Change in tax levied is:

$$\widehat{G}_r = \frac{\int F'(w) \left(\sum_q \lambda'_{rq}(w) \right) T_r'(w) dw}{\int F(w) \left(\sum_q \lambda_{rq}(w) \right) T_r(w) dw.}$$

where $T_r'(w)$ is the tax scheme in the counterfactual equilibrium, possibly different (exogeneously so) from the one in the reference equilibrium.

Finally, the counterfactual location choice of workers can be simply expressed as a function of initial location choices λ_{rq} , changes in neighborhood quality and prices defined above, and changes in income, which we take as an exogenous input to the counterfactual. Specifically, changes in location choices are given by:

$$\widehat{\lambda}_{rq}(w) = \frac{\widehat{B}_{rq}^{\rho}}{\widehat{V}^{\rho}(w)} \frac{[w + \chi'(w) - p'_n]^{\rho}}{[w + \chi(w) - p_n]^{\rho}}, \quad (\text{H.15})$$

In parallel, we get the change in welfare given by:

$$\widehat{V}^{\rho}(w) = \sum_{rq} \widehat{B}_{rq}^{\rho} \frac{[w + \chi'(w) - p'_n]^{\rho}}{[w + \chi(w) - p_n]^{\rho}} \lambda_{rq}(w), \quad (\text{H.16})$$

Values for $\{p'_{rq}, \lambda'_{rq}(w), L'_{rq}, R'_{rq}\}$ are the solutions of equations (H.11)-(H.15) that define a counterfactual equilibrium of the economy corresponding to an alternative distribution of income $F'(w)$ in the city.

Given that the model is over-identified, the baseline model matches the 1990 data imperfectly. We treat the log-differences between data and model as measurement error, and hold it constant across periods when we conduct counterfactuals. Formally, for our main counterfactual, let Θ^* denote the estimated parameters coming from the minimum distance procedure, let I^{1990} denote the income distribution in 1990, and let ε^{1990} denote a vector of measurement error defined as the difference between model-based (log) variables and data (log) variables in year 1990:

$$\log x^{1990} = \log \bar{x}^{1990} (\Theta^*, I^{1990}) + \varepsilon^{1990}. \quad (\text{H.17})$$

Similarly, we have:

$$\log x^{2014} = \log \bar{x}^{2014} (\Theta^*, I^{2014}) + \varepsilon^{2014}. \quad (\text{H.18})$$

We are reporting such counterfactual changes of the form:

$$\log \bar{x}^{2014} (\Theta^*, I^{2014}) - \log \bar{x}^{1990} (\Theta^*, I^{1990}),$$

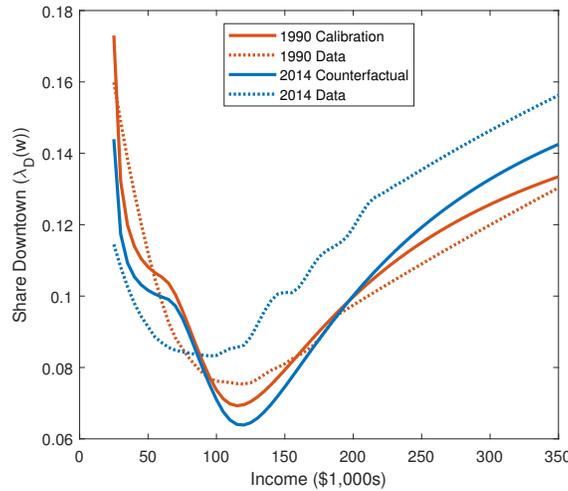
where model-based changes can be computed using the "exact hat" procedure explicated above. When reporting counterfactuals, we keep measurement errors unchanged at their 1990 level.

Appendix H.2 Additional Counterfactual and Welfare Results

Appendix H.2.1 Alternative Representation of Main Counterfactual Result

Figure H.8 shows the predicted change in sorting patterns in our main counterfactual between 1990 (solid orange) and 2014 (solid blue), compared with the actual change (the corresponding dashed lines) at each household income level between \$25,000 and \$350,000.

Figure H.8: Counterfactual impact of shift in income distribution on the U-shape



Notes: This figure shows the change in the propensity to live downtown between 1990 (orange) and 2014 (blue) for households in different \$5,000 income brackets between \$25,000 and \$350,000. The solid lines compare the predicted share downtown by income in the model calibrated to the actual data in 1990 to the predicted share downtown that results from the change in the income distribution between 1990 and 2014. The dashed lines compare the share downtown in the data for 1990 and 2014.

Appendix H.2.2 Importance of Endogenous Public Amenities in Base Specification

As downtown gets richer, taxes collected are higher and public amenities respond for all households downtown. This effect makes both richer and poorer households better off downtown, but housing prices respond to this amenity increase which tends to hurt poorer households. To quantify the net effect, we compute a 1990-2014 counterfactual with no endogeneous response in public amenities.

Figure H.9 reports the results (red bars), and compares it to the baseline (clear bars). We see that endogenous public amenities tend to mitigate the welfare differential between richer and poorer households somewhat, but are far from strong enough to overturn the general tendency of spatial sorting responses to increase well-being inequality.

Figure H.9: Shutting Down Public Amenities

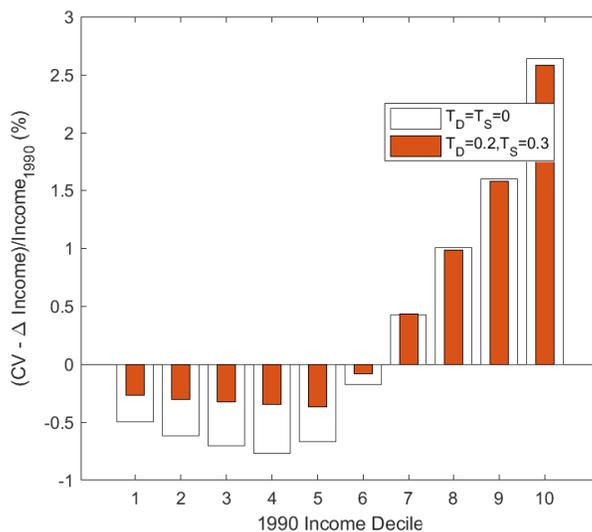


Table H.12 shows the robustness of our results to alternative values of the parameters that govern the public amenity response (the property tax rates downtown T_D and in the suburbs T_S and the elasticity of public amenities Ω). The response of our welfare estimates to the elasticity of the endogenous component of public amenities confirms that low-income households benefit from the increases in local tax revenues that accompany gentrification. However, as we highlighted in our discussion of Figure H.9 the effects of changing the parameters governing endogenous public amenities on welfare is quantitatively small.

Appendix H.2.3 Welfare Impact of Alternative Changes to Income and Population

In the analysis above, we have studied the effects of changes in the observed income distribution holding everything else, including population, constant. We complement this analysis by studying the implication of total population change itself. Further, to tease out what characteristic of the 1990-2014 income shock are important in driving our result, we explore alternative changes in the income distribution.

Table H.13 reports the results. Row 1 re-displays our baseline results. In the second row, we feed in both the actual population change and the change in the income distribution between 1990 and 2014. The third row isolates the effects of population growth separately from income growth, by feeding in only the observed change in population, holding the underlying income distribution constant. Accounting for growth in population results in a larger increase in welfare inequality

Table H.12: Robustness of Welfare Estimates to Alternate Public Amenity Parameters

Decile:	$(\Delta CV - \Delta Inc)/Inc_{1990}$						Δ Urban Share			
	All Households			Renters Only			Predicted (p.p.)		Share of Actual	
	Top	Bottom	Diff.	Top	Bottom	Diff.	Top	Bottom	Top	Bottom
Base Specification	3.17	-0.42	3.59	1.81	-0.74	2.55	1.16	-1.75	57%	41%
Property Tax Rates (base: $T_D = 0.2$ and $T_S = 0.3$)										
$T_D = 0.15; T_S = 0.25$	3.30	-0.43	3.73	1.91	-0.79	2.70	1.16	-1.73	57%	41%
$T_D = 0.25; T_S = 0.25$	3.04	-0.41	3.46	1.79	-0.71	2.51	1.13	-1.68	56%	40%
Public Amenity Elasticity (base: $\Omega = 0.05$)										
$\Omega = 0$	2.64	-0.47	3.11	1.56	-0.83	2.40	1.31	-1.90	65%	45%
$\Omega = 0.03$	2.64	-0.47	3.11	1.56	-0.83	2.40	1.31	-1.90	65%	45%
$\Omega = 0.08$	2.64	-0.47	3.11	1.56	-0.83	2.40	1.31	-1.90	65%	45%

compared to our baseline. The larger increase stems from two forces. First, population growth amplifies the love of variety effects described above. Second, the increase in population drives up rents everywhere but more so in the downtown areas where land is more constrained. Given our unit housing assumption, this impacts poorer households disproportionately. Changing both population and income increases the well-being gap between high and low income residents by over 5.7 percentage points (on a base of 19 percentage points). Additionally, poorer renters are made worse off in absolute terms by an amount equal to 3.3 percent of their income.

In the fourth row of the table, we return to holding population fixed, and we now assume that all households experience the same income growth equal to the 1990-2014 per capita average. Interestingly, under this alternative income change, the poor are much more worse off in absolute terms relative to our base specification. This happens because a broad based increase in income generates a stronger spatial sorting response, with many middle class individuals in the suburbs moving up their residential Engel curves. This rising demand for downtown living puts more upward pressure on house prices than in our baseline counterfactual, where incomes rise for only a few households at the top of the distribution. As a result, the increase in welfare inequality due to spatial responses is higher with broad based income growth than in our baseline case, at about 2.7 percent (instead of 1.7 percent).

In the final rows (5 through 7) of the table we explore crude predictions about the potential future welfare impact of neighborhood change. Specifically, we hold population growth fixed and ask what happens through the lens of our model when income growth increases by an additional 10, 20, and 30 percent for everyone, starting from the actual 2014 income distribution. These counterfactuals shed some light on the potential effects of future economic growth on the spatial distribution of residents within cities. Holding population fixed, the quantified model suggests that the spatial sorting response from an additional 10 percent income growth for all individuals (which does not impact income inequality) further increases well-being inequality. The mechanisms are the same as what we highlight above. Our model predicts that if income growth in the U.S.

Table H.13: Welfare Estimates under Different Counterfactual Income Distributions

	$(\Delta CV - \Delta Inc)/Inc_{1990}$						
	Decile:	All Households			Renters Only		
		Top	Bottom	Diff.	Top	Bottom	Diff.
Base Specification (aggregate population fixed)							
[1]	3.17	-0.42	3.59	1.81	-0.74	2.55	
Alternative Driving Forces (1990-2014)							
[2] Allowing for population growth	8.37	-1.16	9.53	3.91	-3.58	7.49	
[3] Only population growth, no change to income distribution	4.30	-0.70	5.00	1.65	-2.76	4.41	
[4] No population growth, income distribution shifts rightwards	2.76	-1.10	3.86	2.05	-2.34	4.39	
Projected Further Welfare Changes from Further Income Growth from 2014 Onward							
[5] $HHInc_i = HHInc_{i,2014} \times 110\%$	3.28	-1.27	4.55	2.28	-2.59	4.87	
[6] $HHInc_i = HHInc_{i,2014} \times 120\%$	10.41	-3.51	13.92	6.59	-6.42	13.01	
[7] $HHInc_i = HHInc_{i,2014} \times 130\%$	14.77	-5.33	20.10	8.01	-10.20	18.21	

continues, even without further increase in income inequality, additional gentrification and within city neighborhood change will be an enduring feature of the urban landscape. This suggests that it is not income inequality per se that drives our results, but instead an increase in the absolute number of high-income households regardless of what is happening to the rest of the income distribution.

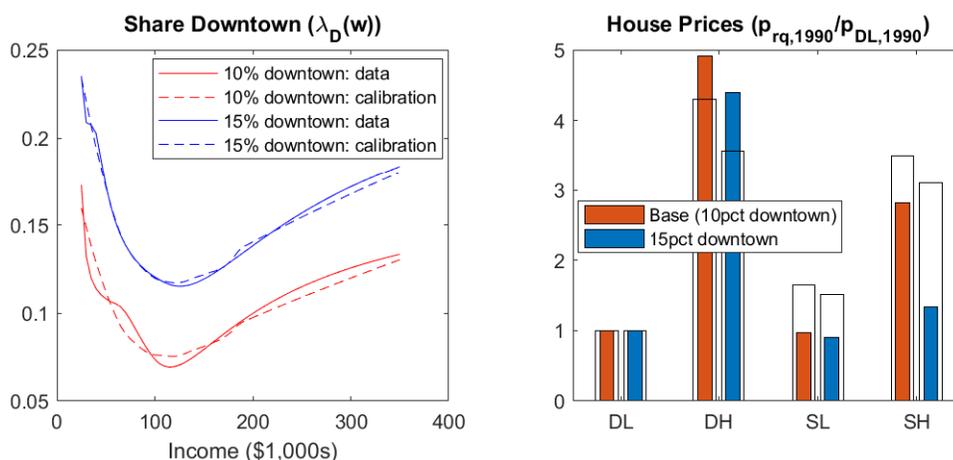
Before turning to counterfactual policy analysis, we now analyze two counterfactuals that provide additional model validation.

Appendix H.2.4 Robustness to Alternate Definition of Downtown

Figure H.10 compares the fit of our model calibrated to data based on our baseline downtown definition, where “downtown” includes all Census tracts closest to the city center that include 10% of the population in 2000, with a larger alternative definition that contains 15% of the population.²⁶ Table H.14 replicates the welfare results for our main counterfactual exercise using these two downtown definitions. The welfare effects using the larger downtown definition are qualitatively similar, though less pronounced (consistent with the evidence in Couture and Handbury (2017) that neighborhood change was highly localized over this period).

²⁶Our calibration requires Census microdata that are only available at the geography of Public Use Microdata Areas (PUMA). This geography is too large to construct U-shape that we calibrate to for a smaller 5% downtown definition. For example, in 1990 we found a median share of 17% across the 100 largest CBSAs.

Figure H.10: Robustness of Calibration to Alternative Downtown Definition



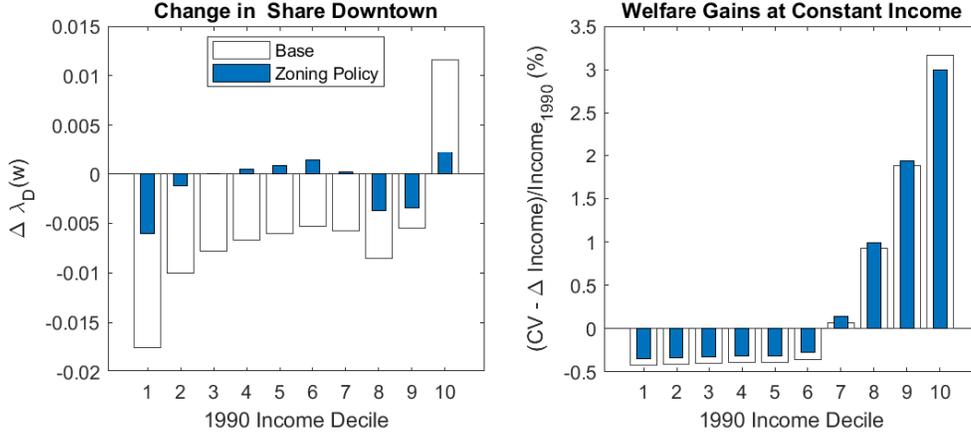
Notes: These figures compare the fit of the calibrated model to the two targeted moments, under two different definitions of the downtown area. The left-hand plot shows the share of households in each \$5,000 income bracket that reside downtown in 1990. The dashed line shows the data, while the solid line shows the prediction of the calibrated model. The data are constructed from micro IPUMS data and reflect the propensity to reside downtown by income in the 27 CBSAs in which PUMAs (the finest spatial unit the IPUMS data) are small enough relative to the downtown definition to make useful inference here. The curve is interpolated to address top-coding in the IPUMS data. See Appendix C for more details. The clear bars in the right-hand plot shows the average Zillow 2 Bedroom Home Value Index in tracts of each location-quality type, normalized by the average index in low-quality tracts downtown, in 1996. The solid red bars show the predicted relative housing costs predicted by the calibrated model.

Table H.14: Robustness of Welfare Estimates to Alternative Downtown Definition

	$(\Delta CV - \Delta Inc) / Inc_{1990}$						
	All Households			Renters Only			
	Decile:	Top	Bottom	Diff.	Top	Bottom	Diff.
Base Specification		3.17	-0.42	3.59	1.81	-0.74	2.55
Alternative (15pct) Downtown Definition		1.35	0.23	1.12	1.09	0.02	1.07

Appendix H.2.5 Alternative Policy Counterfactual

Figure H.11: Location Choices and Well-Being under Zoning Policy



Notes: This figure shows the change in the propensity to live downtown (on the left) and change welfare (on the right) that result from the change in the income distribution by income decile. The clear bars show the results from the baseline counterfactual. The blue bars show the results from the alternative counterfactual with zoning that preserves the neighborhood mix (i.e., the share of neighborhoods of each rq type) at its level from the 1990 equilibrium.

Appendix H.2.6 Additional Potential Mitigating Force

A limitation of the benchmark model is the assumption that increased variety of neighborhoods of a given rq type only benefits inhabitants of that type of neighborhood. In reality, the gentrification of downtown neighborhoods can benefit all inhabitants of the city, to the extent they can travel to consume urban amenities there - as they do in the data. We embed this amenity consumption into the model, by assuming that non-housing expenditure is spent on a Cobb-Douglas aggregate of private urban amenities a and the traded good c , and the demand for urban amenities for a household living in r is a CES aggregate of amenity options in all neighborhoods. That is, the utility of household ω that lives in neighborhood $n \in \mathcal{R}_{rq}$ is:

$$U_n(\omega) = B_n \left(\frac{a_n}{\alpha} \right)^\alpha \left(\frac{c_n}{1-\alpha} \right)^{1-\alpha} b_n^\omega \quad (\text{H.19})$$

$$a_n = \left(\sum_{n'} (\beta_{j(n)j'(n')})^{\frac{1}{\sigma}} (a_{nn'})^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (\text{H.20})$$

where $a_{nn'}$ is consumption of amenities in n' for a household living in n , $\sigma > 1$ is the elasticity of substitution between amenities in different neighborhoods, and the disutility term $\beta_{j(n)j'(n')}$ depends on the dissimilarity in quality between a household's own neighborhood and the destination neighborhood.²⁷ Commuting to amenities is costly, with a cost that increases with distance. The

²⁷We normalize $\beta_{qq} = 1$. Typically $\beta_{qq'} \leq 1$ if $j \neq j'$, so that households value horizontal differentiation within quality level, but to a lesser degree outside of their preferred quality type.

cost of consuming amenities in neighborhood n' for a household living in n is $d_{nn'}^\delta p_{n'}^a$, where δ governs how the cost of commuting to consume amenities in n' varies with the distance $d_{nn'}$ between n and n' , and $p_{n'}^a$ is the price of amenities in n' . Given our symmetry assumptions, neighborhoods in \mathcal{R}_{rq} offer the same price index for amenity consumption, P_{rq}^a .

These amenities are provided by the private developers, who build both housing units and retail space when they develop neighborhoods. Beyond renting out housing units, they operate non-tradable services like restaurants and entertainment options that are marketed to households living in the neighborhood as well as to those living in other parts of the city. Developers use K^h and K^a units of land to build H_{rq}^h housing units and H_{rq}^a retail areas of quality j in location n . Land markets clear, accounting for demand coming from both housing units and retail space. Developers in the retail amenity sector are monopolistically competitive, so they set amenity prices, denoted p_{rq}^a , at a constant markup over marginal costs, i.e.:

$$p_{rq}^a = \frac{\sigma}{\sigma - 1} k_{rq}^a R_r. \quad (\text{H.21})$$

The equilibrium price for housing is still pinned down by (6). The number N_{rq} of neighborhoods of type (r, q) adjusts so that $\pi_{rq} + \pi_{rq}^a - f_{rq} = 0$, where π_{rq}^i for $i = h, a$ is the operating profit of a developer of type (r, q) in activity i .

Through this channel, households benefit from an increased supply of urban amenities outside of their own type of neighborhood. More details on this channel and its quantification can be found in the working paper version (Couture et al., 2019).