Review of Economic Studies (2024) **91**, 858–898 https://doi.org/10.1093/restud/rdad048

Advance access publication 25 April 2023

Income Growth and the Distributional Effects of Urban Spatial Sorting

Victor Couture
University of British Columbia
Cecile Gaubert
University of California, Berkeley
Jessie Handbury
University of Pennsylvania
and
Erik Hurst
University of Chicago

First version received January 2020*; Editorial decision December* 2021*; Accepted April* 2023 *(Eds.)*

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 neighbourhoods with endogenous amenity quality. As the rich get richer, demand increases for the high-quality amenities available in downtown neighbourhoods. Rising demand drives up house prices and spurs the development of higher quality neighbourhoods 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 neighbourhood 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.

Key words: Spatial sorting, Gentrification, Inequality

JEL codes: D11, R12, R13, R23

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 neighbourhoods 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

The editor in charge of this paper was Kurt Mitman.

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 neighbourhoods with these luxury urban amenities rises and more of them choose to reside downtown, triggering the redevelopment of previously low-income neighbourhoods.

In this paper, we develop a model to formalise this mechanism. We use the model to quantify how top income growth contributes to neighbourhood 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 neighbourhood development: the quality of amenities in city neighbourhoods responds endogenously 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 neighbourhoods within cities.

Figure 1 motivates our analysis by summarising 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 centre 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.

We develop a model of a stylised 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 neighbourhoods that offer different qualities of amenities and housing. Households trade off neighbourhood attractiveness against cost of living. This cost depends on local housing prices and on the cost of commuting to work. Preferences for neighbourhoods are non-homothetic: households with higher incomes are more willing to pay the higher cost of living in desirable neighbourhoods. On the supply side, neighbourhoods are built by private developers who compete for land within each area of the city. As top incomes grow, demand for high-quality neighbourhoods downtown rises leading to an increase in prices throughout downtown, including in low-quality neighbourhoods where the poor live. An increase in the supply of high-quality neighbourhoods, and an associated decrease in the

1. For example, Aguiar and Bils (2015) estimate that restaurant meals and non-durable entertainment are among the goods with the highest income elasticities. Couture and Handbury (2020) document that downtown areas of major cities have a higher density of such amenities.

2. CBSAs 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).

3. See online 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.

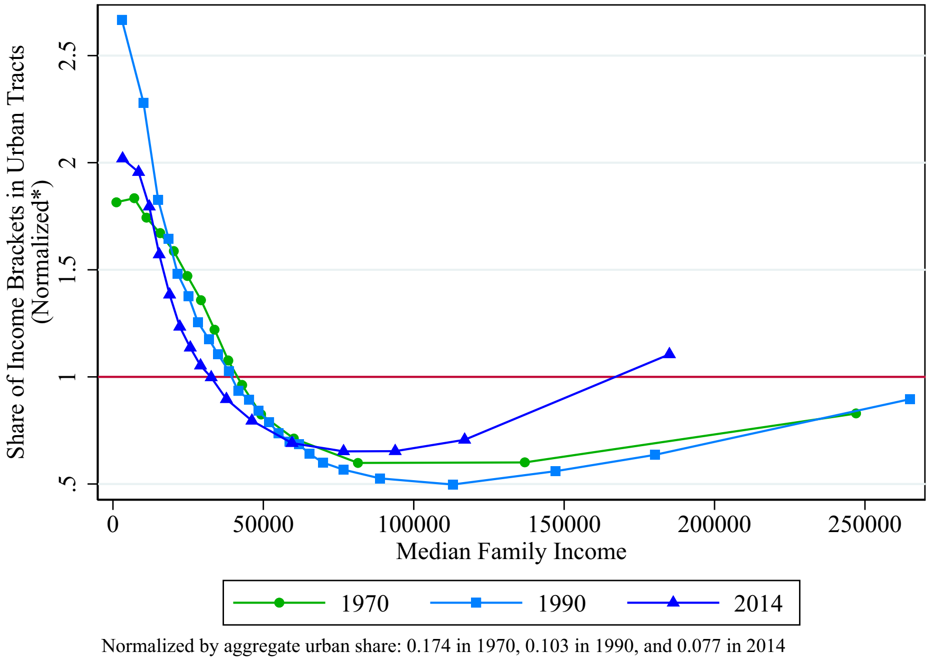

FIGURE 1

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—normalised 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 centre 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 (Ruggles *et al.* 2018) for the corresponding year in the 100 largest CBSAs. The IPUMS microdata are adjusted for top-coding using the generalised Pareto method, as described in online Appendix C.

variety of low-quality neighbourhoods, 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 labour 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 minimise the costs of housing and commuting by residing in low-quality neighbourhoods downtown. At the same time, higher income households are attracted downtown by the density of high-quality, high-amenity luxury neighbourhoods 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 neighbourhoods to change—and (ii) this mechanism can explain roughly sixty percent of the urbanisation of the rich (top income decile) and roughly forty percent of the suburbanisation of the poor (bottom income decile).[4] These findings highlight that, while other forces outside of the model are also arguably contributing to neighbourhood 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 centres.

To further validate the model, we present additional counterfactual analyses. First, we show that a similar procedure applied to both the 1950–70 and the 1970–90 changes in the income distribution lead to spatial sorting responses that are qualitatively different from 1990 to 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 urbanisation of the higher income households. Second, we show that the model also performs quite well at explaining cross-CBSA variation in spatial sorting during the 1990–2014 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 neighbourhood 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 neighbourhood 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 between 2.5 and 4% points (on a base of 19% points) during the past three decades. We further estimate that the welfare of low-income renters was actually reduced by between 0.5% and 1.0% 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.

4. Throughout the paper, we often use "neighbourhood change" of low-income neighbourhoods and "gentrification" interchangeably. We realise 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 neighbourhood 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.

Finally, we use the model to study the effect of stylised "anti-gentrification" policies, for instance one that taxes housing in high-quality neighbourhoods downtown to subsidise housing in low-quality neighbourhoods 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 neighbourhood change on the poor, but does not curb the influx of the rich into downtowns.

Related Literature. This paper contributes to four main Literature studies. 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 endogenous supply responses, especially in the context of international trade.[5] We apply this logic to the endogenous development of neighbourhoods 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 stylised model of a representative city that allows us to model the extensive margin: the number and quality of neighbourhoods in a city 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 neighbourhood 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 behaviour 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 neighbourhoods within

5. Faber and Fally (2022) 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.

6. Our 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, Eckert and Peters (2022) 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.

7. Gaigne *et al.* (2017) theoretically analyse 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.

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 (2022) emphasises the role of rising value of time for high-skilled workers. Our focus on urban amenities as an important dimension of neighbourhood 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 studies 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 neighbourhood 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, 2020; Su, 2022). Couture and Handbury (2020), 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 and changing taste for commutes alone are unlikely to rationalise 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 specifications.

2. A MODEL OF THE RESIDENTIAL CHOICE OF HETEROGENEOUS HOUSEHOLDS

2.1. *Benchmark model*

We propose a model of residential location choice 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 neighbourhoods of heterogeneous qualities. Derivations and proofs are given in online Appendix B.

2.1.1. Demand for neighbourhoods. The city is populated by households who have heterogeneous income, with a continuum of households at each level of income w. The distribution of incomes has cumulative distribution $F(\cdot)$ 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.

8. See, for example, Guerrieri *et al.* (2013), Edlund *et al.* (2022), Ellen *et al.* (2019), Berkes and Gaetani (2018), Vigdor *et al.* (2002), Freeman (2005), McKinnish *et al.* (2010), Ellen and O'Regan (2010), Ding *et al.* (2016), Brummet and Reed (2019), Meltzer and Ghorbani (2017), Lester and Hartley (2014), and Autor *et al.* (2017).

Each household ω makes a discrete choice of a neighbourhood n to reside in. Neighbourhoods are grouped into four broad types. They are characterised, first, by their geographic region r. Specifically, neighbourhoods 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, neighbourhoods can be of High ($q=H$) or Low ($q=L$) quality. Within each of the four types (r,q), there are N_{rq} neighbourhoods to choose from. They are assumed for simplicity to be symmetric. Importantly, the number of neighbourhoods of each type, N_{rq}, is an endogenous variable in the model, as supply of neighbourhoods responds to demand.

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

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

In this expression, w^{ω} is the income of household ω, p_n the cost of living in neighbourhood n, B_n summarises the quality of amenities in neighbourhood n, and b_n^{ω} is household ω's idiosyncratic preference for neighbourhood 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 neighbourhoods 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 neighbourhoods, which have different house prices.

Households have idiosyncratic preferences for neighbourhoods denoted b_n^{ω}. They are drawn iid from a generalised extreme value (GEV) distribution,

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

with $\rho \leq \gamma$, where $\mathcal{R}_{rq}$ is the set of neighbourhoods of type (r,q). The parameter ρ captures the substitutability of neighbourhoods of different quality and location types, while γ governs the higher substitutability across neighbourhoods 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 modelling. Households first make a choice over quality and location of neighbourhoods (the "upper nest"), then over horizontally differentiated neighbourhoods within this category (the "lower nest").[10]

9. We make the implicit assumption that, in a first step that is not modelled, 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 neighbourhoods are encompassed in expression (1), which can be derived from the more general specification $v_n(\omega) = \left((1-\tau_n)w^{\omega} - \tilde{p}_n\right)\tilde{B}_n b_n^{\omega}$, where $\tilde{p}_n$ is the cost of housing in neighbourhood 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).

10. See *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 Frechet distribution when $\gamma = \rho$, in which case the independence of irrelevant alternatives applies and all neighbourhoods are equal 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.

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

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

In this expression,

$$V_{rq}(w)=N_{rq}^{\frac{1}{\gamma}}B_{rq}\left(w-p_{rq}\right) \tag{4}$$

is the inclusive value of neighbourhoods of type (r,q) for a household with income w, summing across all the possible choices of neighbourhood r of type (r,q). Importantly, this inclusive value increases with the number of neighbourhoods N_{rq} to choose from. This is because of a love-of-variety effect: given the idiosyncratic preferences (2), having more neighbourhoods to choose from leads to better matches between households and neighbourhoods on average, yielding higher utility. Conveniently, this love of variety for neighbourhoods 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 neighbourhood type (r,q). This amenity composite is an endogenous variable in the model, because the provision of neighbourhoods 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 neighbourhood depends positively on the quality of its amenities, positively on the variety of neighbourhoods of that type the household can choose from, and negatively on its housing cost.

2.1.2. Supply of neighbourhoods. Neighbourhoods are supplied by private developers, who choose the quality and location (r,q) of the neighbourhood they develop. To develop a neighbourhood 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:

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

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

The pricing and entry behaviour of developers is summarised here and described in further detail in online Appendix B.2. Given unit housing demand and monopolistic competition, equations (3) and (4) lead to the following housing pricing formula:

$$p_{rq} = \frac{\gamma}{\gamma+1} k_{rq} R_r + \frac{1}{\gamma+1} \bar{\mathcal{W}}_{rq} \tag{6}$$

where $k_{rq}R_r$ is the unit cost of housing in (r,q) and $\bar{\mathcal{W}}_{rq}$ is a measure of demand for a neighbourhood 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 neighbourhoods are within a type, and the higher the market power of developers.

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

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

This equation determines how the supply of neighbourhoods 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 neighbourhoods N_{rq} such that (i) households maximise their utility; (ii) developers and landowners maximise 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 neighbourhoods.*

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 neighbourhoods.

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

12. The 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.

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 neighbourhood types by income can be written as

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

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 neighbourhoods, 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 stylised 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,

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

Recalling that, in the model, there are four neighbourhood 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:*

$$p_{DL} < p_{SL} < p_{SH} < p_{DH}. \tag{9}$$

The formal proof of Lemma 2 establishes that the share of households living in the most expensive neighbourhoods in the city necessarily increases with income, while the share of households living in the least costly neighbourhoods of the city necessarily decreases with income. When both the most expensive and the least expensive neigbhorhoods 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 neighbourhoods, while low-income households are attracted downtown by a relatively low cost of living in low-quality neighbourhoods.

13. In 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.

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 neighbourhood characteristics: an "amenity composite" ($N_{rq}^{\frac{1}{\gamma}} B_{rq}$), which is endogenous through N_{rq}—the supply of neighbourhoods 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 $\imath$, $F_\imath(w)$, ordered in the sense of FOSD, that is,

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

Starting at income distribution $F_\imath(w)$, we consider an infinitesimal increase in $\imath$. 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 $\imath$ in the relative number of high-income households, downtown land prices increase: $\frac{\partial R_D}{\partial \imath} \geq 0$.*

Through the market for urban land, a demand shock for housing in high-quality neighbourhoods pushes up downtown land prices, transmitting the shock to the entire urban area, including in low-quality downtown neighbourhoods. 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 neighbourhood 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 $\imath$ in the relative number of high-income households, the perceived quality of the most expensive neighbourhood option increases ($\frac{\partial N_{DH}}{\partial \imath} > 1$), while the perceived quality of the least expensive neighbourhood option decreases ($\frac{\partial N_{DL}}{\partial \imath} < 1$).*

This result is quite intuitive: as the number of high-income households increases in the city, increased demand for the luxury option *DH* results in an increased supply of such high-quality neighbourhoods downtown. Importantly, this endogenous increase in the supply of *DH* neighbourhood makes the *DH* option even more attractive, because preferences for neighbourhoods 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: neighbourhoods endogenously become higher quality—a phenomenon that looks like urban gentrification. Our framework therefore formalises a model for the gentrification of poorer neighbourhoods in a city, as a result of demand and supply shifts for these neighbourhoods.

2.3. *Extensions and empirical implementation*

Our benchmark model is stylised 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 neighbourhoods is driven by public amenities such as parks, public schools, or policing, financed by local governments. They respond to taxes according to

$$B_{rq} = B^{o}_{rq} \left(G_r\right)^{\Omega}, \tag{10}$$

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

$$G_r = \int \lambda_r(w) T_r(w)\, dF(w). \tag{11}$$

As the tax base downtown increases, government revenues G_r increase, which raises amenities for all households downtown, irrespective of the quality of their neighbourhood.

2.3.2. Homeownership. The benchmark model is static, hence it cannot capture the important notion that, as neighbourhoods downtown gentrify, homeowners in gentrifying neighbourhoods 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 labour 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

$$w + \chi\,(w) - T_r(w). \tag{12}$$

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 neighbourhoods 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 neighbourhoods. 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 (2022) investigates job urbanisation by skill from 1994 to 2010 and finds no difference in job location trends between high- and low-skilled jobs: they both suburbanise somewhat over that time period, at the same rate. Second, Baum-Snow and Hartley (2020) decompose the role of residential demand versus labour demand in explaining 2000–10 downtown gentrification. They conclude that labour demand plays little role. Third, Couture and Handbury (2020) conclude that high wage jobs—in the upper third of the wage distribution—do not urbanise fast enough from 2002 to 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 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 neighbourhood supply amplifies 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 households' location decisions and housing costs mapped into the spatial units employed in the model. Herein, we describe this mapping and summarise our main data sources. The full exposition of the data work is detailed in online Appendix C.

3.1.1. Spatial units and classifications. We equate the notion of a city to a Core-Based Statistical Area (CBSA), and that of a neighbourhood to a census tract. All data are interpolated to constant 2010-boundary tracts using U.S. Census Tract Boundaries from Esri (2022) 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 centre of the CBSA that accounted for 10% 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). Figure 2 shows that downtown includes the lower half of Manhattan and much of Brooklyn in New York and a semi-circle from Montrose Beach to Hyde Park in Chicago. Online appendix A.1 shows which tracts in New York, Chicago, Philadelphia, San Francisco, Boston, and Las Vegas are classified as downtown and suburban.[16]

14. Moreover, data from the U.S. Census show 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 centres to reverse commute—*i.e.* to live downtown and work in the suburbs.

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

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

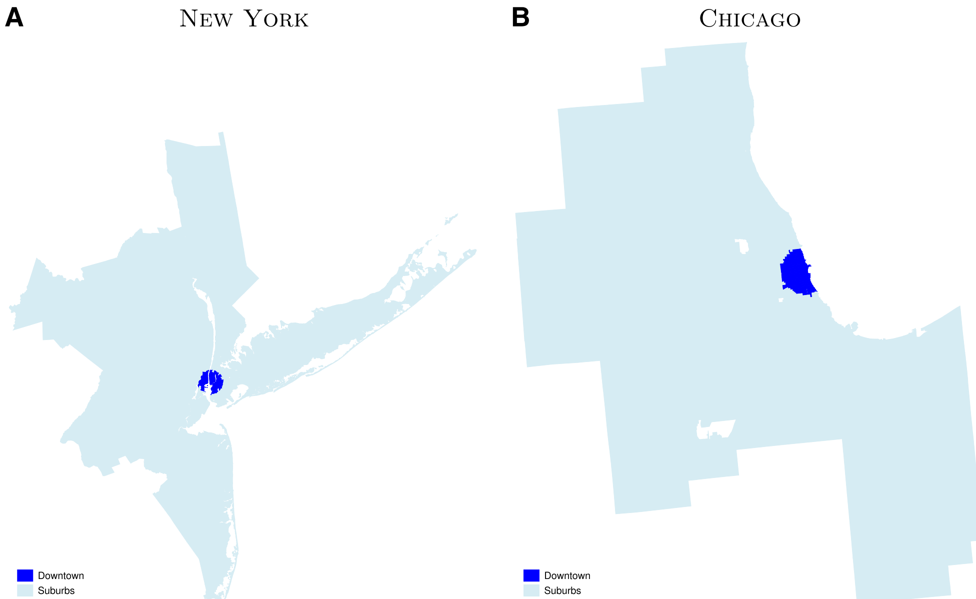

FIGURE 2
Downtown definitions

Note: Panel A shows the five central counties (boroughs) of New York City and Panel B shows Cook County, Chicago. Downtown tracts, shaded in dark blue, consist of all tracts closest to the city centre and accounting for 10% of total CBSA population in 2000.

We define neighbourhood 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 neighbourhood to be high quality if at least 40% of residents between the ages of 25 and 65 have at least a bachelor's degree. Under this definition, 15, 22, and 32% of census tracts in the top 100 CBSAs are, respectively, classified as high quality in 1990, 2000, and 2014.

3.1.2. Income and house prices. Our data on location choice by household income are from the 1990 Census and the 2012–2016 American Community Surveys (henceforth referred to as the "2014 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 summarise 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.

17. We 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 top-coding in these data can be found in online Appendix C.

We use the National Historical Geographic Information System (NHGIS) 1990 Census tables from Manson *et al.* (2022) 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 neighbourhood 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* from Zillow (2019) when measuring changes in housing prices over time within an area–quality pair (r, q). The Zillow neighbourhood house price indices are specifically designed to capture house price changes over time for a constant set of housing units within a given neighbourhood. 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 1996–8.[18] We pool prices for years 2012–6 for our end period (2014). Finally, the Zillow indices measure neighbourhood house price changes at the level of zip codes which are larger than census tracts. We use crosswalks from the Department of Housing and Urban Development (HUD , 2022) to map zip codes to census tracts. We then aggregate tracts to area–quality pairs.[19] In the CBSA-level analysis herein, for example, $p^h_{rq,ct}$ 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% in 1990 and 4.6% in 2014 from the Lincoln Institute of Land Policy.

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 ("urbanise") relative to the average household, while poorer households move out ("suburbanise"). 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:

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

18. House prices were relatively flat over the 1990–5 period suggesting, that this measurement issue is unlikely to bias our results in any meaningful way.

19. Given 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.

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. $\Delta Income_c$ 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 3 plots point estimates and 95% confidence bands for the β^m coefficients from equation (13). The left-hand plot shows that richer households urbanised more, and poorer households less, in CBSAs that experienced more average income growth. The right-hand plot shows even stronger patterns of urbanisation of rich households and suburbanisation of poor households in CBSAs with more top income growth. A 10% increase in the 95/50 ratio of incomes is associated with a 20% increase in the propensity of households earning more than \$150,000 to live downtown relative to the average household, and a 10% 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 urbanisation of the rich and suburbanisation of the poor.

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 neighbourhoods. 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,

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

Normalising brackets m and m' such that $w_m > w_{m'}$, the dependent variable is positive if the richer income bracket urbanises 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

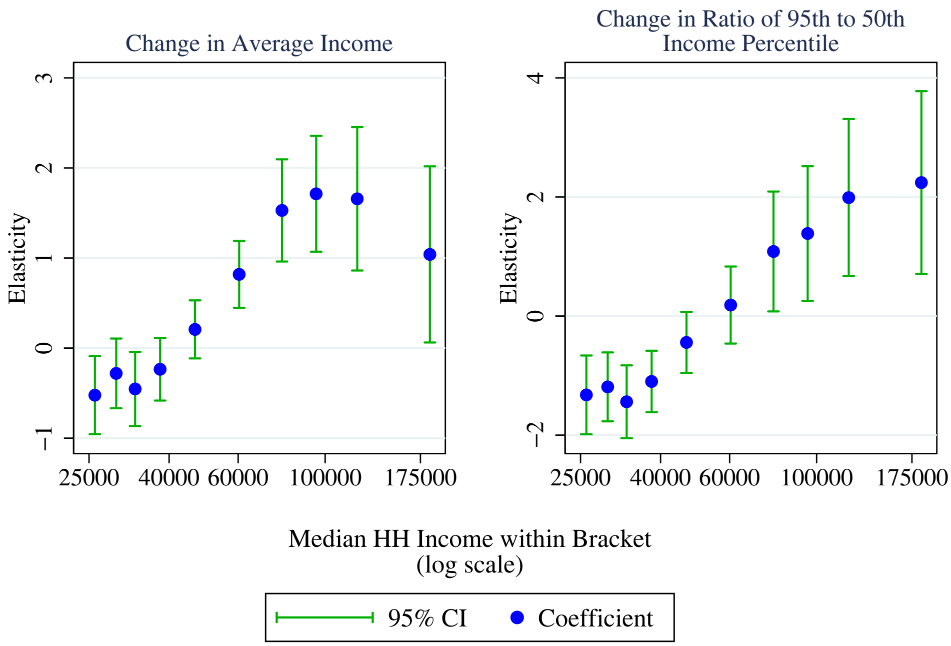

FIGURE 3

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% 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 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.

urbanise 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 affect 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

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

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 neighbourhoods of a given quality tier in a CBSA will not compromise identification. These shocks are captured by the common utility shifter $(\sum_{r',q'} V^{\rho}_{r'q'}(w))$ 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 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 neighbourhood 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 neighbourhood 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.* (2022): 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

21. Note that taking a difference across two timeperiods (1990–2014) 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.

are most likely to work.[22] This isolates wage growth in suburbanised 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 disproportionately 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 centres.

Finally, we conduct pre-trend and balance tests similar to those proposed by Borusyak *et al.* (2022) and find no evidence of pre-trends or balance violation. In addition, we re-estimated ρ from a regression in which the dependent variable (change in urbanisation by income) is from the pre-period of 1970–90, but everything else is as in equation (14). We find a small $\rho = 0.64$ that is not significantly different from zero. A full discussion of these results can be found in online 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 $\Delta \ln(p^h_{Dq,c}/p^h_{Sq,c})$ (on the y-axis) in the left panel of Figure 4. 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:

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

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 4 reports estimates from equation (15), along with their 95% confidence intervals. We find that a CBSA income shock causes differential spatial sorting responses from the rich versus 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% 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 summarise, Figure 4 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.

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

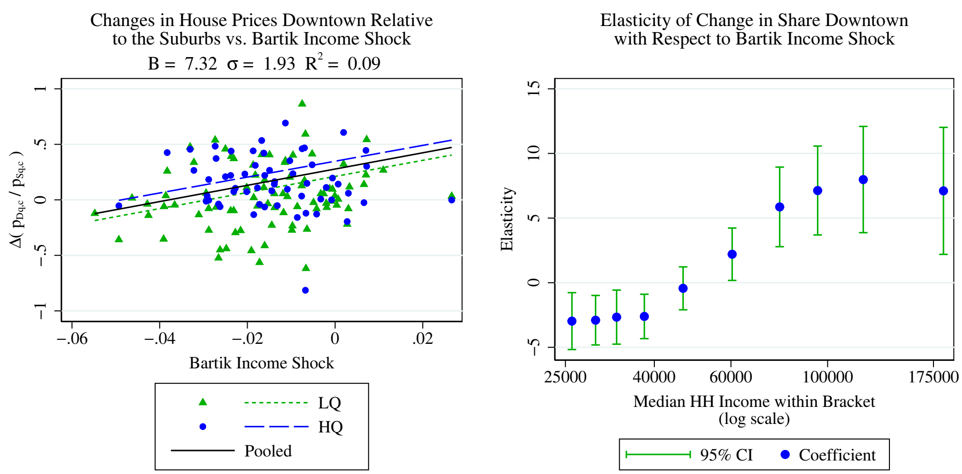

FIGURE 4

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 are from the Zillow 2 Bedroom Index in 1996–8 and 2012–6. We drop the top and bottom 1% of $\Delta(p_{Dq,c}/p_{Dq,c})$ from the plot. On the right, we show income bracket-specific coefficients, along with 95% confidence intervals, from equation (15) on the *y*-axis (regression of Bartik income shock on changes in normalised 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.

TABLE 1

Estimation of elasticity ρ

	(1)	(2)	(3)	(4)	(5)
$\hat{\rho}$	2.48	3.05	2.70	2.26	3.21
	(0.34)	(0.70)	(0.76)	(0.77)	(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 with neighbourhood quality defined from education mix of residents. Each observation is weighted by the number of households in the income bracket. Columns 3–5 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 columns (1)–(5).

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 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 versus 3.05), 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 urbanised, 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 labour demand shocks

TABLE 2
Key model parameters

Parameter	Description	Value	Source
	Non-homotheticity		
ρ	Between-type neighbourhood substitution elasticity	3.0	Estimation
	Land Price Responses		
ϵ_D	Downtown land supply elasticity	0.6	Calibrated to Saiz (2010)
ϵ_S	Suburban land supply elasticity	1.3	Calibrated to Saiz (2010)
	Amplification		
γ	Within-type neighbourhood substitution elasticity	6.8	Literature
	Other		
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

that are concentrated in urban centres of CBSAs and that disproportionately affect high-income households. In online 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 summarise, 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 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 online Appendix E.

4. MODEL QUANTIFICATION

Having established the empirical relevance of the key model mechanism and providing 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 parametrisation*

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.

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

4.1.2. Key amplification parameter (γ). Income sorting (governed by ρ) is amplified in the model by endogenous neighbourhood development. The intensity of this effect is governed by the shape parameter γ, which controls gains from variety in residential neighbourhood choice. Ahlfeldt *et al.* (2015) estimate a within-type neighbourhood 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 neighbourhoods of a given quality estimated in Germany many not map exactly to the preferences of Americans, we show the sensitivity of our results to alternate values of γ.

4.1.3. 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 annualised 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).

4.1.4. 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 homeownership 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 labour income level capital gains corresponding to their average real estate portfolio. This transfer equals the average house price growth in the neighbourhoods 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 homeownership increases systematically with labour 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 neighbourhood type $N_{rq}^{\frac{1}{\gamma}} B_{rq}$, and (ii) the price of housing in each neighbourhood 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 summarise 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 neighbourhood 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 \$5000 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 neighbourhoods, by a revealed preference approach applied to our non-homothetic demand function: the same level of price and quality of a neighbourhood generates different demand patterns at different levels of income. Concretely, the identification relies on the equation

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

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 normalisation level, whose value does not impact the counterfactuals done in the following section.

4.2.1. Moment fit.

The moment fit is presented in Figure 5. 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 neighbourhoods. In 1990, the model and Census data house prices in low-quality downtown neighbourhoods and in the suburbs are close. In both the model and data, high-quality suburban neighbourhoods have housing prices about three times higher than low-quality suburban neighbourhoods. Importantly, in both the model and the data, prices in downtown high-quality neighbourhoods are between four and five times higher than prices in low-quality downtown neighbourhoods.

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 rationalise the observed changes in spatial sorting within the city. We start by analysing the 1990–2014 period before turning to 1970–90 and 1950–70.

23. The IPUMS data identify the locations of respondents at the PUMA (Public Use Microdata Area), each of which contains approximately 100,000 individuals, relative to the 4000 contained in each Census tract. To replicate the urban share for each fine income bracket, 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 online Appendix C for more details.

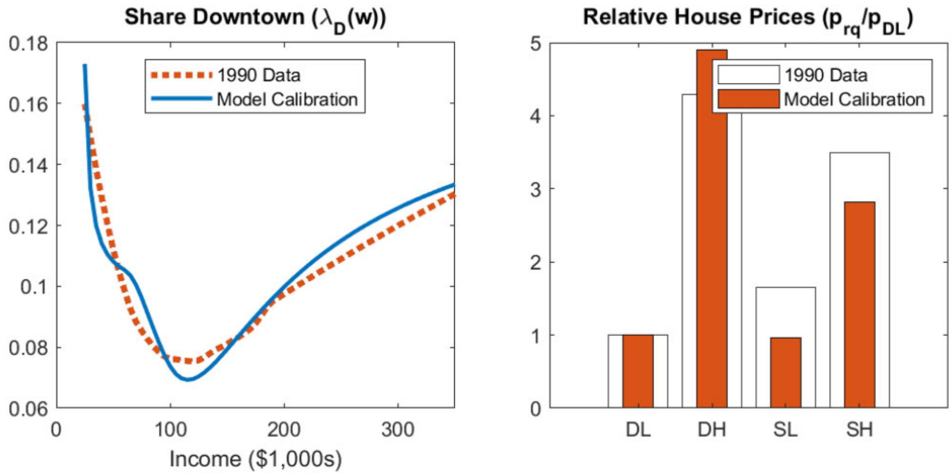

FIGURE 5

Calibration to 1990 urban shares and neighbourhood 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 $5000 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 online 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, normalised by the average index in low-quality tracts downtown, in 1996. The solid red bars show the relative housing costs predicted by the calibrated model.

Given the structure of the model, a counterfactual equilibrium can be computed with the following parameters on hand, as detailed in online Appendix G: (i) the model elasticities $\{\rho, \gamma, \epsilon_r, \Omega\}$, and (ii) the initial equilibrium values for population in each neighbourhood 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 6 summarises 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%, while for the top decile, income increased by about 20%. Overall, the 90–10 income gap widened by 21% 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 parameter 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 6, 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, summarising 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.

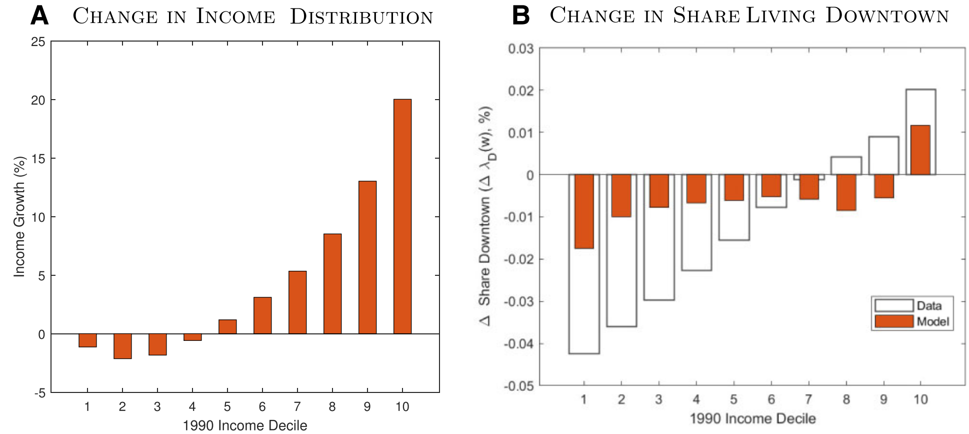

FIGURE 6

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 summarises 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.

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% of the suburbanisation of the bottom decile of the income distribution, and about 60% of the urbanisation 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 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 1950–70 and 1970–90 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 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. In particular, we do not recalibrate our model; instead, we

24. Note that the predicted urbanisation of the highest income decile reflects both a shift along the calibrated U-shape of Figure 5 as well as an endogenous uptick in the U-shape, generated by the change in the income distribution. In the online Appendix we also show that the model also explains a significant portion of the observed uptick itself.

use our baseline paramaterisation calibrated to 1990 data and feed in the aggregate 1950 and 1970 income distributions for the United States as opposed to the 2014 income distribution as in Section 5.1.[25] Incomes unambiguously grew in the 1950–70 and 1970–90 periods, as the top panel of Figure 7 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 7, meanwhile, show that, through the lens of the model, this income growth generated different changes in sorting from 1950–70 and 1970–90 than that predicted in our main counterfactual for 1990–2014. For example, our model predicts lower income individuals *urbanised* and higher income households *suburbanised* during the 1950–70 period. Additionally, the model predicts little change in spatial sorting patterns by income during the 1970–90 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 to 1970, \$100,000–125,000 from 1970 to 1990, and then \$125,000–150,000 from 1990 to 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 suburbanisation 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 urbanisation 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 urbanised 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 CBSAs. For each CBSA, we calibrate the model by targeting the 1990 distribution of location choice by income within the CBSA

25. The 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 generalised Pareto method, as described in online Appendix C.

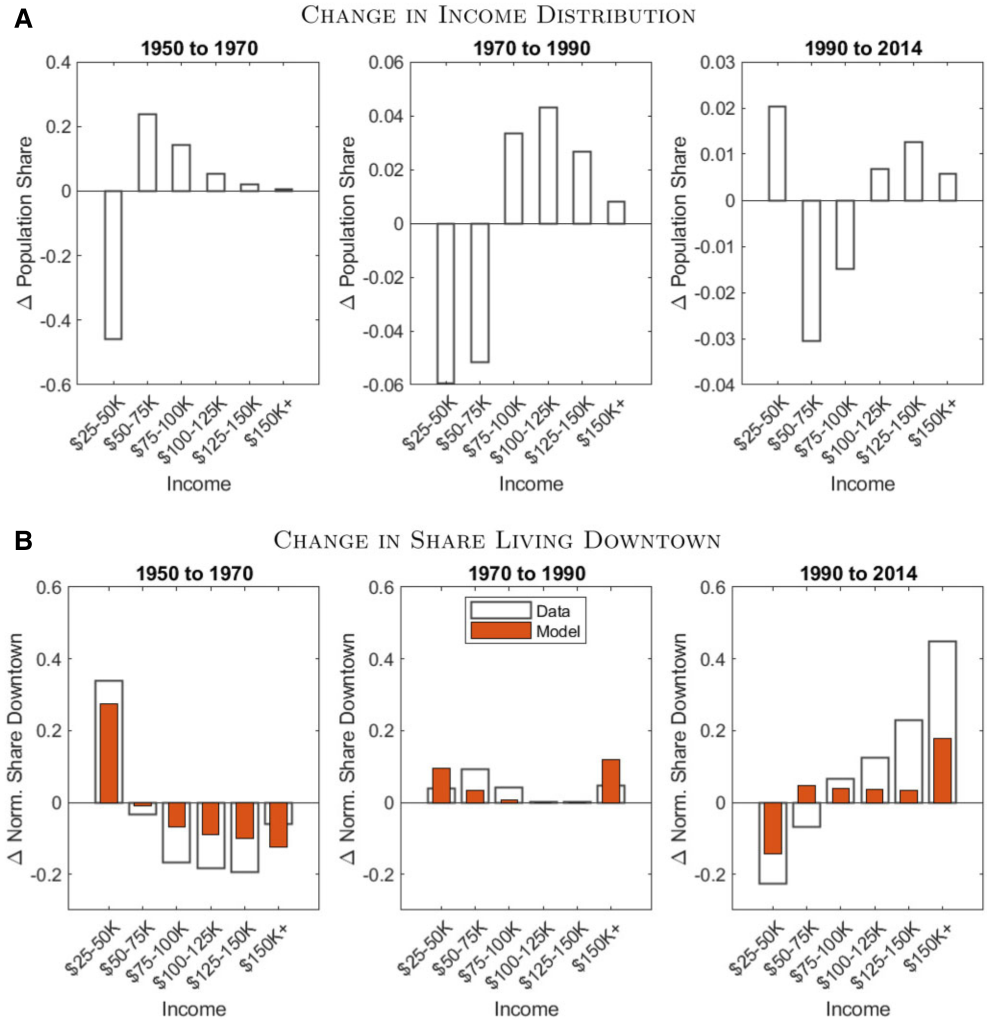

FIGURE 7

1950 and 1970 counterfactuals

Notes: Panel A shows the change in the share of the population in each income bracket in the 27 CBSAs modelled 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.

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 8 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

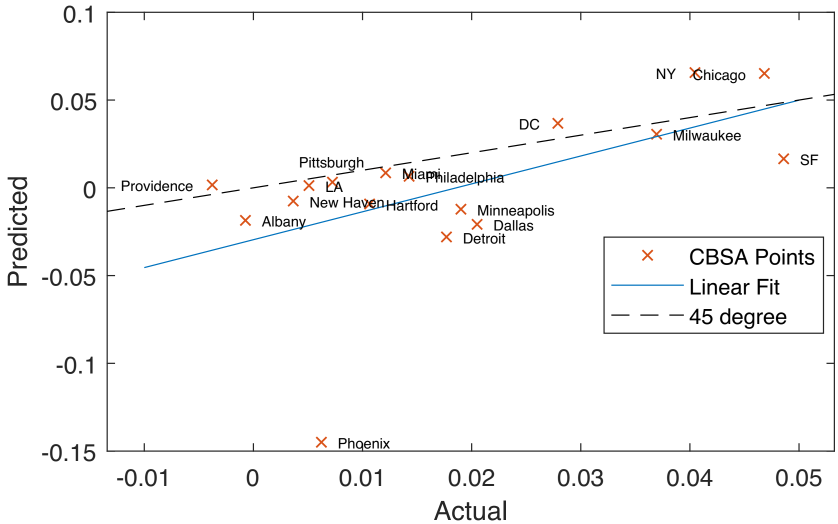

FIGURE 8

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 versus as observed in the data over the same period.

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.

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 neighbourhoods 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 analyse 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 neighbourhood change and spatial re-sorting triggered by top income growth between 1990 and 2014. 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.

26. For this analysis, we can only use 17 of 27 CBSAs. These are the CBSAs for which there is sufficient coverage above the state-specific IPUMS income top-codes to implement the generalised Pareto interpolation procedure that we use to measure CBSA-level income distributions and U-shapes.

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:

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

where V_{rq} is the inclusive value of all neighbourhoods 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:

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

where $m_t(i)$ is the 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 9 reports our headline welfare results for each income decile. It shows the welfare gains from the 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% points. As high earners move downtown, the supply of high-quality neighbourhoods that they value rises endogenously, making them better off. House prices increase as well, but for high earners the amenity benefit of neighbourhood 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 neighbourhood variety benefits. As a result, well-being changes are even more negative than income changes in the bottom decile, by an additional 0.4% points. Overall, the well-being gap between the top and bottom deciles of the income distribution grew by an additional 3.6% points, compared to the 19% point growth in the nominal income gap. That is, within-city spatial sorting amplifies the growing welfare gap between the rich and the poor from rising income inequality by more than 15% (3.6/21).

Comparing the welfare results for all households in the left panel of Figure 9 with those for renters only in the right panel, we see that capital gains from house price appreciation benefits homeowners at all income levels, to some extent. For example, about 30% 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% 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 neighbourhood 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% point growth in welfare.

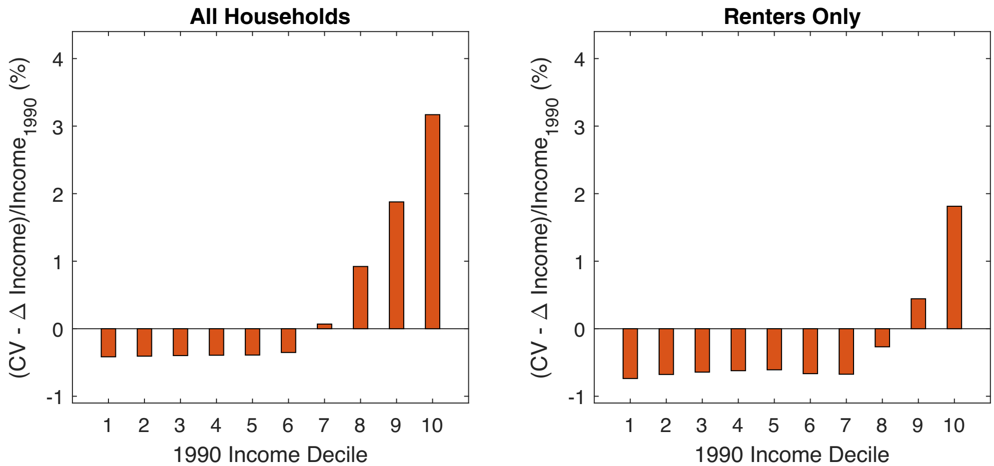

FIGURE 9

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.

Before analysing the main mechanisms at play behind these findings, it is worth commenting on the magnitudes of these welfare effects. Figure 9 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% 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 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, the 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 neighbourhoods and to enjoy consumption amenities there. This influx of rich households puts upward pressure on downtown housing prices not only in high-quality neighbourhoods, but

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

28. In the online Appendix, we explore counterfactuals where we also 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.

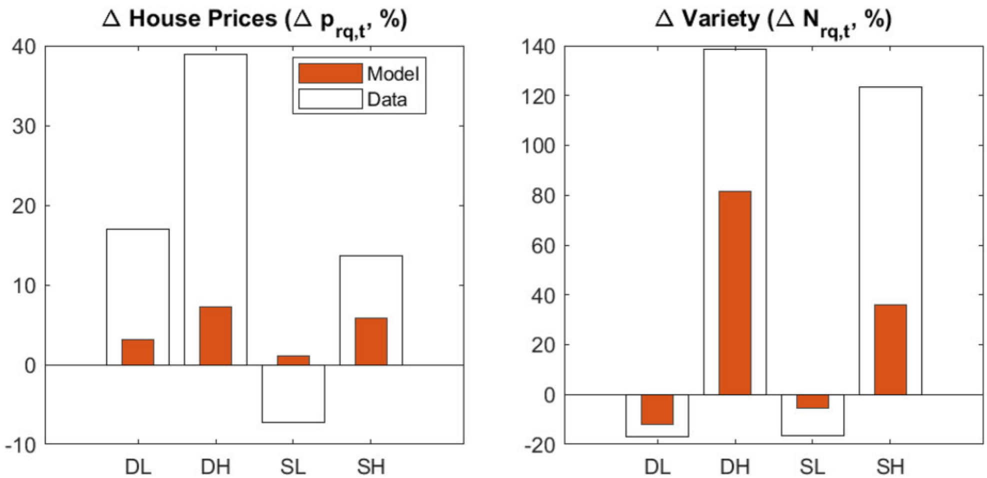

FIGURE 10
Mechanisms

Notes: This figure shows the change in housing costs (on the left) and neighbourhood variety (on the right) in each of the four neighbourhood types between 1990 and 2014. Red bars are model and wide clear bars are data. Consistent with the model, the change in housing costs is 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 neighbourhood type between 1996 and 2014.

also in low-quality ones. The left-hand panel of Figure 10 compares these house prices changes, for different neighbourhood types, to those observed in the Zillow data. The model predicts that the shift in the income distribution alone generates a 6% increase in house prices in high-quality downtown neighbourhoods and a 3% increase in house prices in low-quality downtown neighbourhoods. These predicted increases in downtown house prices amount to about 20% 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% versus 1%). This matches the data qualitatively where house price growth between 1990 and 2014 was higher in low-quality downtown neighbourhoods than in low-quality suburban neighbourhoods. House price growth in low-quality neighbourhoods 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 neighbourhood change. As the rich move downtown and demand for high-quality neighbourhoods increases, developers supply more high-quality neighbourhoods. Some of this entry is at the cost of exit of lower quality neighbourhoods, so that gentrification takes place. The right-hand panel of Figure 10 shows the growth in the supply of neighbourhoods in each area and quality level. The model predicts a large proportion of the downtown neighbourhood 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 neighbourhoods downtown makes high-income households better off.

Overall, the contraction in the number of low-quality neighbourhoods—the gentrification that we also observe in the data—makes low-income households worse off. Finally, we note that the predicted supply of high-quality neighbourhoods also expands in the suburbs, but at a much smaller rate than downtown.

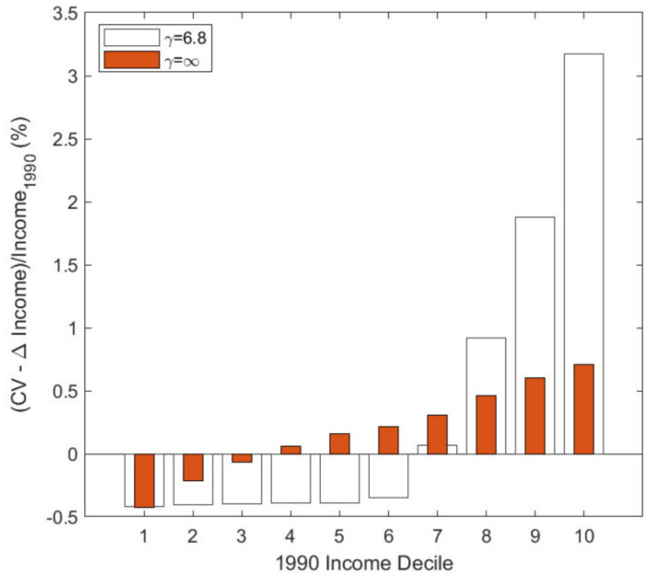

FIGURE 11

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 neighbourhood 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.

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 neighbourhoods by setting the between-neighbourhood 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 neighbourhood (or associated consumption amenity) variety. Welfare results are shown in Figure 11 (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% point in the baseline to 1.1% points without love-of-variety effects. About two-thirds of the welfare gap in our base results stems from the endogenous private amenity 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.

6.3. *Gentrification curbing policies*

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

6.3.1. Taxing developments. We first model a stylised "anti-gentrification" policy, which systematically taxes high-quality housing downtown and uses the proceeds to subsidise rents in low-quality neighbourhoods downtown (the policy is budget neutral). It aims to limit the development of high-quality neighbourhoods 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 12 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 neighbourhoods: 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.

The well-being effects of this policy, shown in the right hand plot of Panel A in Figure 12 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.* neighbourhood 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 neighbourhoods to residents of low-quality suburban neighbourhoods.[29]

Panel B of Figure 12 shows the impact of an alternative policy that taxes high-quality neighbourhood development (and subsidises 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 by much and does not stem urban 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 neighbourhoods.

6.3.2. 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 13 shows that doing so does little to stem neighbourhood change downtown (in the left panel) but it mitigates the associated welfare losses

29. In the online Appendix we show that the results of this policy are very similar to those we obtain when directly modelling zoning regulations; *i.e.* a policy that imposes a constant relative number of high- to low-quality neighbourhoods. 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.

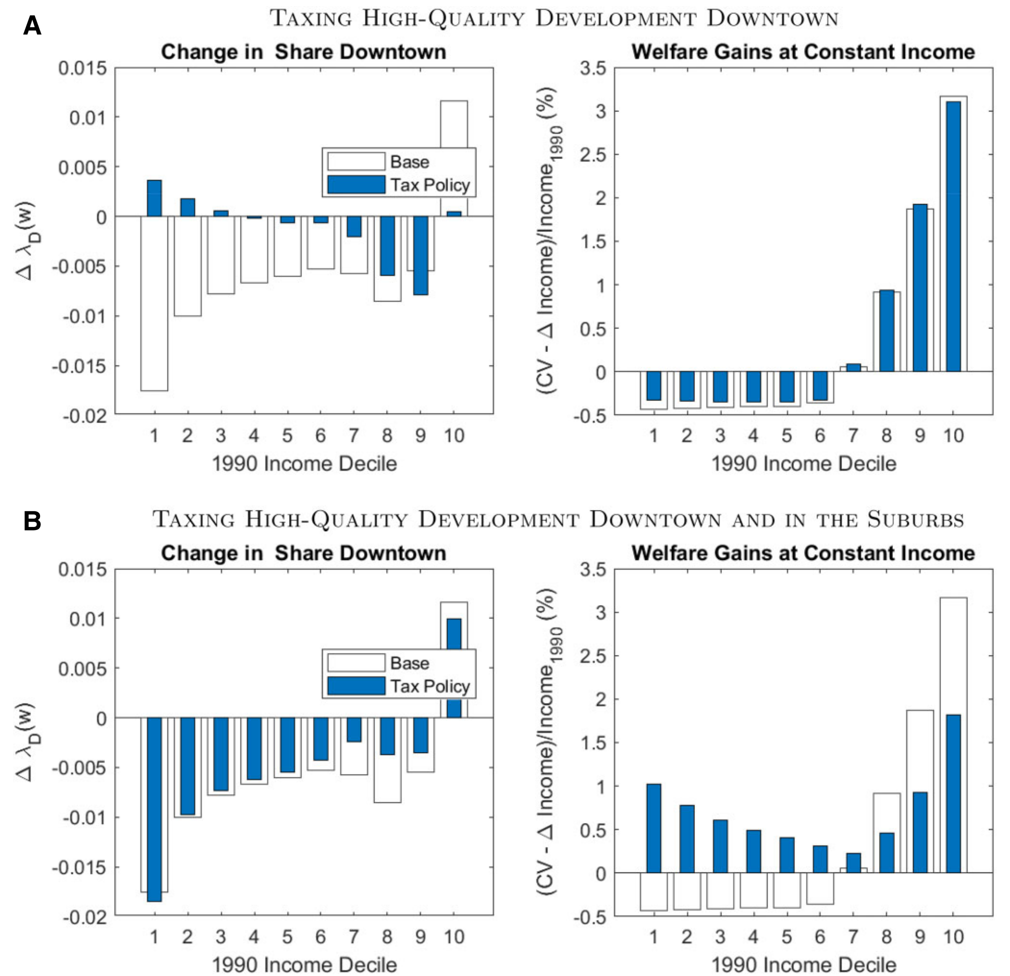

FIGURE 12

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 the development of high-quality neighbourhoods is taxed and the proceeds of the tax are redistributed to subsidise housing costs in low-quality neighbourhoods in the same location. In Panel A, only downtown high-quality neighbourhoods are taxed (and low-quality rents subsidised). In Panel B, high-quality neighbourhoods are taxed (and low-quality rents subsidised) both downtown and in the suburbs.

on the poor (in the right panel). House price growth is reduced by approximately 2% points in all neighbourhoods, and effectively shut down in the low-quality neighbourhoods. 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 neighbourhood 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.

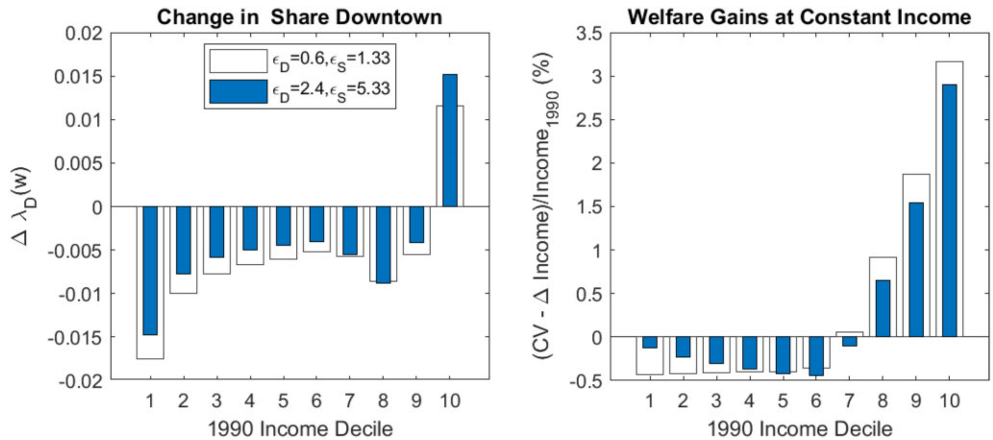

FIGURE 13
Effects of increasing supply elasticity

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 the 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% points.

Next in Table 3, we show the sensitivity of our results to different values of γ. For lower values of γ, the 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 neighbourhoods. 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 welfare losses for poor owners downtown, but exacerbates losses to poor renters.[30]

30. We 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.

TABLE 3
Robustness of welfare estimates to key parameters

	$(\Delta CV - \Delta Inc)/Inc_{1990}$						Δ Urban Share			
	All Households			Renters Only			Predicted (p.p.)		Share of Actual	
Decile:	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 Neighbourhood 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 Neighbourhoods (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 summarises 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 summarise the model predictions for the urbanisation of top income decile households and the suburbanisation of bottom income deciles, first in absolute percentage point terms and then as a share of the respective 2.3% point inflow and 4% point outflow observed in the data.

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 shares by income. In the model, the unit housing requirement means that, within a neighbourhood 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 neighbourhoods: higher incomes sort into the more expensive neighbourhood types so their income share of housing does not fall proportionally with income. With only four neighbourhood 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 analogue, we use reported spending patterns on "housing" by income deciles reported in public release tables from the Consumer Expenditure Survey (CEX). For our empirical measure of housing expenditures, we use the CEX'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% increase in income is associated with the housing spending share falling by 1.9% 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 words, 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 neighbourhoods 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 14 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 14, in order to target a flatter housing Engel curve, the model needs to generate a higher housing price in downtown high-quality neighbourhoods. 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 neighbourhoods implies that fewer household can afford to live in those neighbourhoods 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% points relative to the bottom decile in response to the observed income growth between 1990 and 2014.

7.3. *Additional potential mitigating forces*

Alimitation of the benchmark model is the assumption that an increased variety of neighbourhoods of a given *rq* type only benefits inhabitants of that type of neighbourhood. In reality, the gentrification of downtown neighbourhoods 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 neighbourhoods 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

31. See https://www.bls.gov/cex/tables/calendar-year/mean-item-share-average-standard-error.htm#cu-income for the Consumer Expenditure Survey (CEX) 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 CEX measures of housing to our model analogues in the online Appendix.

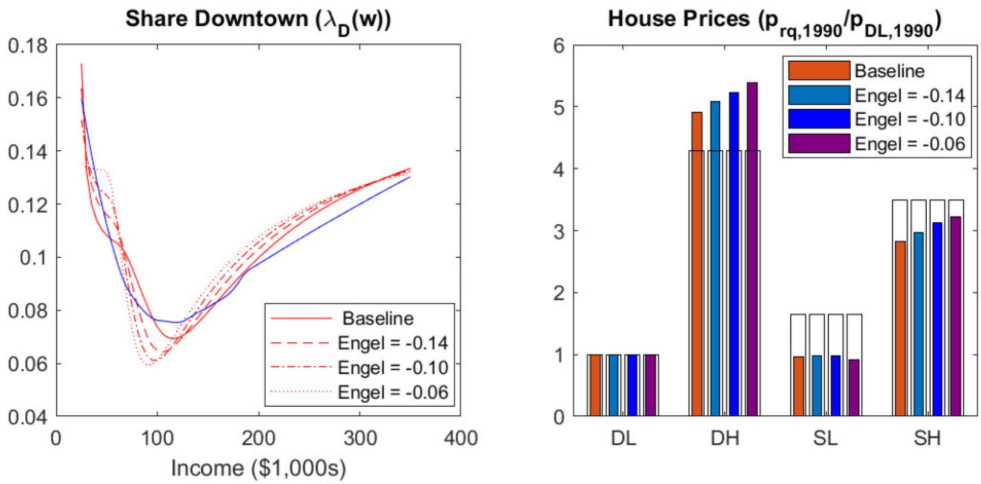

FIGURE 14

Sensitivity of calibrated 1990 urban shares and neighbourhood 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 \$5000 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, normalised 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.

TABLE 4

Robustness of welfare estimates to targeting housing share moment

	$(\Delta CV - \Delta Inc)/Inc_{1990}$						Δ Urban Share			
	All Households			Renters Only			Predicted (p.p.)		Share of Actual	
Engel Slope	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 summarises 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.

to restaurants and entertainment options outside of the neighbourhood 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 neighbourhoods, lower income households would get additional utility from consuming the amenities of those neighbourhoods. 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 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 neighbourhoods. 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 neighbourhoods 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 neighbourhoods. To study this phenomenon, we develop a spatial model of a city with heterogeneous agents, neighbourhoods 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 neighbourhood change during the last 25 years within the United States. The analysis also suggests that neighbourhood 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 2.5–4% points across a range of parameter values.

We explore possible policy responses to mitigate these distributional impacts of neighbourhood 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 neighbourhoods, 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 shaping neighbourhood 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 stylised 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 neighbourhood change is a natural avenue for future research.

Acknowledgments. 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.

32. In 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.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data underlying this article are available on Zenodo at https://doi.org/10.5281/zenodo.7779687

REFERENCES

AGUIAR, M. and BILS, M. (2015), "Has Consumption Inequality Mirrored Income Inequality?", *American Economic Review*, **105**, 2725–2756.

AHLFELDT, G. M., REDDING, S. J. and STURM, D. M., *et al.* (2015), "The Economics of Density: Evidence from the Berlin Wall", *Econometrica*, **83**, 2127–2189.

ALLEN, T., ARKOLAKIS, C. and LI, X. (2015), *Optimal City Structure* (New Haven: Yale University, Mimeograph).

ALMAGRO, M. and DOMINGUEZ-IINO, T. (2019), Location Sorting and Endogenous Amenities: Evidence from Amsterdam.

AUTOR, D. H., PALMER, C. J. and PATHAK, P. A. (2017), "Gentrification and the Amenity Value of Crime Reductions: Evidence from Rent Deregulation" (Working Paper 23914, National Bureau of Economic Research).

BAUM-SNOW, N. and HAN, L. (2019), The Microgeography of Housing Supply.

BAUM-SNOW, N. and HARTLEY, D. (2020), "Accounting for Central Neighborhood Change, 1980–2010", *Journal of Urban Economics*, **117**, 103228

BAUM-SNOW, N. and MARION, J. (2009), "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods", *Journal of Public Economics*, **93**, 654–666.

BERKES, E. and GAETANI, R. (2018), *Income Segregation and Rise of the Knowledge Economy* (Toronto: Ohio State University Mimeo).

BORUSYAK, K., HULL, P. and JARAVEL, X. (2022), "Quasi-Experimental Shift-Share Research Designs", *The Review of Economic Studies*, **89**, 181–213.

BRUECKNER, J. K., THISSE, J.-F. and ZENOU, Y. (1999), "Why is Central Paris Rich and Downtown Detroit Poor?: An Amenity-Based Theory", *European Economic Review*, **43**, 91–107.

BRUMMET, Q. and REED, D. (2019), "The Effects of Gentrification on Original Neighborhood Residents: Evidence from Longitudinal Census Microdata" (Working Paper 19–30, Federal Reserve Bank Of Philadelphia).

COUTURE, V., GAUBERT, C. and HANDBURY, J., *et al.* (2019), "Income Growth and the Distributional Effects of Urban Spatial Sorting" (Working Paper 26142, National Bureau of Economic Research).

COUTURE, V. and HANDBURY, J. (2020), "Urban Revival in America", *Journal of Urban Economics*, **119**, 103267.

DAVIS, D. R. and DINGEL, J. I. (2020), "The Comparative Advantage of Cities", *Journal of International Economics*, **123**, 103291.

DAVIS, M. A. and ORTALO-MAGNE, F. (2011), "Household Expenditures, Wages, Rents", *Review of Economic Dynamics*, **14**, 248–261.

DIAMOND, R. (2016), "The Determinants and Welfare Implications of USs Workers' Diverging Location Choices by Skill: 1980–2000", *American Economic Review*, **106**, 479–524.

DIAMOND, R. and MCQUADE, T. (2019), "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development", *Journal of Political Economy*, **127**, 1063–1117.

DIAMOND, R., MCQUADE, T. and QIAN, F. (2019), "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco", *American Economic Review*, **109**, 3365–94.

DING, L., HWANG, J. and DIVRINGI, E. (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**, 1551–1582.

ECKERT, F. and PETERS, M. (2022), "Spatial Structural Change" (Working Paper 30489, National Bureau of Economic Research).

EDLUND, L., MACHADO, C. and SVIATSCHI, M. (2022), "Gentrification and the Rising Returns to Skill", *Economica*, **89**, 258–292.

ELLEN, I. G., HORN, K. M. and REED, D. (2019), "Has Falling Crime Invited Gentrification?", *Journal of Housing Economics*, **46**, 101636.

ELLEN, I. G. and O'REGAN, K. (2010), "Crime and Urban Flight Revisited: The Effect of the 1990s Drop in Crime on Cities", *Journal of Urban Economics*, **68**, 247–259.

ERIKSEN, M. D. and ROSENTHAL, S. S. (2010), "Crowd out Effects of Place-Based Subsidized Rental Housing: New Evidence from the LIHTC Program", *Journal of Public Economics*, **94**, 953–966.

Esri (2022), "U.S. Census Tract Boundaries". ArcGIS Online.

FABER, B. (2014), "Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices" (Working Paper, UC-Berkeley).

FABER, B. and FALLY, T. (2022), "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data", *The Review of Economic Studies*, **89**, 1420–1459.

FAJGELBAUM, P. D. and GAUBERT, C. (2020), "Optimal Spatial Policies, Geography, and Sorting", *The Quarterly Journal of Economics*, **135**, 959–1036.

FAJGELBAUM, P., GROSSMAN, G. M. and HELPMAN, E. (2011), "Income Distribution, Product Quality, and International Trade", *Journal of Political Economy*, **119**, 721–765.
FAJGELBAUM, P. D. and KHANDELWAL, A. K. (2016), "Measuring the Unequal Gains from Trade", *The Quarterly Journal of Economics*, **131**, 1113–1180.
FAJGELBAUM, P. D., MORALES, E. and SUÁREZ SERRATO, J. C., *et al.* (2018), "State Taxes and Spatial Misallocation", *The Review of Economic Studies*, **86**, 333–376.
FOGLI, A. and GUERRIERI, V. (2019), "The End of the American Dream? Inequality and Segregation in US Cities" (Working Paper 26143, National Bureau of Economic Research).
FREEMAN, L. (2005), "Displacement or Succession?: Residential Mobility in Gentrifying Neighborhoods", *Urban Affairs Review*, **40**, 463–491.
GAIGNE, C., KOSTER, H. and MOIZEAU, F., *et al.* (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 ROBERT-NICOUD, F. (2021), Sorting to Expensive Cities.
GLAESER, E., KOLKO, J. and SAIZ, A. (2001), "Consumer City", *Journal of Economic Geography*, **1**, 27–50.
GUERRIERI, V., HARTLEY, D. and HURST, E. (2013), "Endogenous Gentrification and Housing Price Dynamics", *Journal of Public Economics*, **100**, 45–60.
GYOURKO, J., MAYER, C. and SINAI, T. (2013), "Superstar Cities", *American Economic Journal: Economic Policy*, **5**, 167–199.
HOELZLEIN, M. (2019), Two-Sided Sorting and Spatial Inequality in Cities.
HOLIAN, M. J. and KAHN, M. E. (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 MORETTI, E. (2019), "Housing Constraints and Spatial Misallocation", *American Economic Journal: Macroeconomics*, **11**, 1–39.
HUD (2022), "USPS Zip Code Crosswalk Files" (United States, Department of Housing and Urban Development). available at https://www.huduser.gov/portal/datasets.
JARAVEL, X. (2018), "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector", *The Quarterly Journal of Economics*, **134**, 715–783.
LEE, S. (2010), "Ability Sorting and Consumer City", *Journal of Urban Economics*, **68**, 20–33.
LESTER, T. W. and HARTLEY, D. A. (2014), "The Long Term Employment Impacts of Gentrification in the 1990s", *Regional Science and Urban Economics*, **45**, 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**, 163–190.
MANSON, S., SCHROEDER, J. and RIPER, D. V., *et al.* (2022), *IPUMS National Historical Geographic Information System: Version 17.0* (Minneapolis, MN: IPUMS).
MCKINNISH, T., WALSH, R. and WHITE, T. K. (2010), "Who Gentrifies Low-Income Neighborhoods?", *Journal of Urban Economics*, **67**, 180–193.
MELTZER, R. and GHORBANI, P. (2017), "Does Gentrification Increase Employment Opportunities in Low-Income Neighborhoods?", *Regional Science and Urban Economics*, **66**, 52–73.
REDDING, S. J. and ROSSI-HANSBERG, E. (2017), "Quantitative Spatial Economics", *Annual Review of Economics*, **9**, 21–58.
REDDING, S. J. and STURM, D. M. (2016), *Estimating Neighborhood Effects: Evidence from War-Time Destruction in London* (Princeton: Princeton University, mimeograph).
RUGGLES, S., FLOOD, S. and GOEKEN, R., *et al.* (2018), "Integrated Public Use Microdata Series: Version 8.0 [dataset]", available at http://doi.org/10.18128/D010.V8.0.
SAIZ, A. (2010), "The Geographic Determinants of Housing Supply", *The Quarterly Journal of Economics*, **125**, 1253–1296.
SU, Y. (2022), "The Rising Value of Time and the Origin of Urban Gentrification", *American Economic Journal: Economic Policy*, **14**, 402–439.
TSIVANIDIS, N. (2019), "The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogota's TransMilenio" (Working paper, University of Chicago Booth School of Business).
VIGDOR, J. L., MASSEY, D. S. and RIVLIN, A. M. (2002), Does Gentrification Harm the Poor? *Brookings-Wharton Papers on Urban Affairs* 133–182.
Zillow (2019), "Zillow Housing Data" Zillow Inc. available at http://www.zillow.com/research/data/.