Housing Demand, Inequality, and Spatial Sorting *

John Finlay † Trevor C. Williams ‡

July 21, 2023

Abstract

Skilled workers' incomes have pulled away from those of unskilled workers in recent decades, reflecting increasing skill bias in production. How has this change reshaped the spatial distribution of skill? We show nonhomothetic housing demand connects the aggregate income distribution to spatial sorting. A household's skill level determines its income, and therefore its housing expenditure share, sensitivity to housing costs, and preferences over locations. The result is spatial sorting by skill. Moreover, diverging incomes cause diverging location choices. Using consumption microdata, we estimate that housing is a necessity. Increasing total expenditure by 10% reduces housing expenditure shares by 2.5%. Skilled workers therefore sort into expensive cities, and by raising their relative incomes, increases in aggregate skill bias intensify sorting. Embedding our estimated preferences in a quantitative spatial model, we find that without rising aggregate skill bias, spatial sorting would have grown one quarter less since 1980.

*We thank Costas Arkolakis, Lorenzo Caliendo, Fabian Eckert, Brian Greaney, Sam Kortum, Giuseppe Moscarini, Michael Peters, Vítor Possebom, Jaehee Song, Conor Walsh, and seminar participants at Yale, Imperial College London, the Young Economists' Symposium, and the Urban Economics Association for their helpful comments, feedback, and suggestions. We thank Rachel Ngai for helpful comments and Mark Colas for spotting an error in an earlier draft. We finally thank the Yale Department of Economics for financial support in obtaining restricted-access PSID data. Declarations of interest: none.

†Columbia University. rjf2164@columbia.edu

‡Yale University. trevor.c.williams@yale.edu

Introduction

Income inequality, and in particular the gap between skilled and unskilled workers, has risen sharply since 1980. Between 1980 and 2010, wages for workers with a four year college degree grew 31% faster than those of all other workers, reflecting an intensifying skill bias in production. Hand in hand with this divergence in incomes, the location choices of these two groups have also diverged, with skilled workers increasingly clustering in productive, expensive cities like Boston and San Francisco and unskilled workers moving in the opposite direction. In this paper we investigate the effects of an increase in aggregate skill bias on the spatial distribution of skilled and unskilled workers, focusing on how a spatially neutral increase in the returns to skill can nevertheless have heterogenous effects across locations. In particular, we show that nonhomothetic housing demand links aggregate changes in the income distribution to changes in spatial sorting.

To see why nonhomothetic housing demand should matter here, consider two workers. One is highly skilled and can expect to earn a high income wherever she lives. The other is less skilled and earns a lower income. The key decision facing each worker is where to live. When housing is a necessity, it accounts for a large share of the unskilled worker's expenditure and he finds locations with high housing costs unattractive. The skilled worker, by contrast, spends relatively little on housing and is happy to tolerate high prices in exchange for access to desirable amenities or employment opportunities. The result is spatial sorting: workers with different skills, at different points in the income distribution, make systematically different choices about where to live. Crucially, as inequality rises and the incomes of these workers diverge, their choices about where to live also diverge. Skilled workers with rising incomes cluster in cities with expensive housing and unskilled workers cluster in cities with cheap housing. We formalize this insight in a parsimonious model, quantify its importance by providing new estimates of nonhomothetic housing demand, and show that rising aggregate skill bias has caused a substantial increase in spatial sorting by skill, amounting to roughly one quarter of the increase observed in the data between 1980 and 2010.

We begin with a simple spatial equilibrium model. We introduce heterogeneity across workers by assuming that some are skilled and some are unskilled. These differences in skill give rise to differences in income; these income differences are in turn the key driver of our mechanism. Workers trade off wages and housing costs in choosing where to live. They are freely mobile across space and hence must be compensated for living in a high cost city with higher wages, creating a relationship between housing costs and wages. We show that the slope of this relationship is summarized by a worker's housing expenditure share. It immediately follows that the extent to which higher costs are offset by higher wages varies across workers if and only if their preferences are nonhomothetic. When housing is a necessity skilled workers, with their high incomes and low housing shares, are more willing to tolerate high housing costs. Under general nonhomothetic preferences, we show that the equilibrium features spatial sorting by skill, with skilled workers over-represented in expensive cities. We then introduce rising aggregate skill bias in the form of a proportional increase in skilled productivity in every location and show that as the incomes of skilled workers pull away from those of unskilled workers, this sorting intensifies.

We bring our theory to the data by assuming that workers have nonhomothetic constant elasticity of substitution (NHCES) preferences and estimating these using consumption microdata. A key aspect of our estimation strategy is to control for local housing costs, precisely because workers' sorting decisions introduce a positive correlation between prices and incomes at the city level. We find housing is a necessity, as the housing expenditure share declines with income. For a household in the middle of the expenditure distribution, a 10% increase in total expenditure causes a 2.5% decrease in the housing expenditure share. Our estimation strategy allows us to rule out competing explanations, including adjustment costs in housing and time-invariant heterogeneity in housing expenditure shares. We also reject two alternative specifications that have been widely used used in the literature, Cobb-Douglas and a unit housing requirement.[1]

Finally, to isolate the effects of rising aggregate skill bias, we embed our estimated nonhomothetic preferences a quantitative general equilibrium model with rich heterogeneity in productivities, amenities, and housing costs across locations. Our main counterfactual exercise shuts down the observed increase in skill bias between 1980 and 2010. In line with our theory, in the absence of this force, spatial sorting by skill grows more slowly than it did in the data. Quantitatively, we find removing it would have reduced growth in spatial sorting by skill by 27.4%. We also show that Cobb-Douglas preferences shut down any connection between skill bias and sorting, while a unit housing requirement overstates the strength of this relationship.

We contribute to a literature that studies the causes of spatial sorting by skill, reviewed in Diamond and Gaubert (2022). One strand of this literature focuses on growing urban bias in skill premia (Eckert 2019; Giannone 2022; Eckert, Ganapati, and Walsh 2022; Rubinton 2022), and Diamond (2016) has emphasized how endogenous changes in amenities can amplify the effects of local productivity shocks when these amenities are differentially valued by skilled workers. Our paper differs from this literature in two important ways. First, rather than studying a shock that is inherently urban-biased, we focus on a spatially neutral shock which raises the skill premium proportionally in all locations. Both components, spatially neutral and urban-biased, appear to be important in the data. While the skill premium in a city at the 75th percentile of the city size distribution grew 6 percentage points faster than the skill premium in a city at the 25th percentile between 1980 and 2010, the skill premium still rose by 19 percentage points in the smaller city, indicating that rising returns to skill have been a broad based phenomenon affecting all locations.[2] Second, in all the papers above the key reason for rising spatial sorting is a shock to the production side of the model. Our mechanism is very different: sorting is the result of different incomes across skill groups, rather than their different roles in production. While the particular shock we study is a change in the relative productivity of skilled workers, it only alters sorting through its effect on relative incomes, and other shocks (e.g. changes in the progressivity of income taxation) would have similar effects on sorting insofar as they increase income inequality.

[1]"Unit housing requirement" refers to a model in which each household must purchase one unit of housing, so demand is perfectly price and income inelastic.

[2]These calculations are based on Census data. See Appendix B for full details of our data.

By incorporating nonhomothetic preferences into a spatial model (see also Schmidheiny (2006), Eckert and Peters (2023), and Handbury (2021)), we highlight sorting driven by prices. Two other papers have studied the change in spatial sorting in recent decades in models featuring nonhomothetic housing demand. Ganong and Shoag (2017) connect changes in housing supply regulations to slowing regional income convergence. Aside from this difference in focus, a more fundamental difference between that paper and ours is the mechanism at work. Ganong and Shoag (2017)'s results are driven by location-specific shocks to housing supply regulations, whereas we explore the consequences of a shock which is inherently neutral across locations — an increase in the relative productivity of skilled workers — but which, as we show, nevertheless has sharply different consequences across locations. Gyourko, Mayer, and Sinai (2013) study a similar shock to our paper, linking shifts in the income distribution to diverging housing prices across cities. While we allow housing costs to evolve endogenously in our quantitative model, such changes are not central to our mechanism. Our focus, instead, is on how the distribution of skill across cities shifts in response to diverging incomes across skill groups.

Our paper also relates to a recent literature that connects changes in the income distribution to changes in sorting across neighborhoods within a single city. Fogli and Guerrieri (2019) find that residential segregation amplifies increases in income inequality via human capital spillovers. Couture et al. (2023) show that rising income inequality can explain the revitalization of inner cities observed in the US in recent decades. Relative to these papers, we make two key contributions. First, we show that spatial sorting driven by nonhomothetic housing demand is not only an important consideration within cities; instead, the same force also shapes the distribution of skill across cities. Second, instead of assuming a unit housing requirement, we estimate flexible nonhomothetic preferences using consumption microdata, and show that the strength of the relationship between aggregate skill bias and spatial sorting is closely tied to the parameters we estimate.

At the level of cities a common assumption, even in models with heterogeneous households, is that preferences are Cobb-Douglas and therefore homothetic (see, e.g., Eeckhout, Pinheiro, and Schmidheiny (2014), Diamond (2016), Fajgelbaum and Gaubert (2020)). The Cobb-Douglas assumption is often justified by the fact that housing expenditure shares vary little across cities with very different income levels (Davis and Ortalo-Magné 2011). We offer an alternative explanation for the similarity of housing expenditure shares across cities: offsetting price and income effects, a view shared by Albouy, Ehrlich, and Liu (2016). Our demand elasticities are broadly similar to those in Albouy, Ehrlich, and Liu (2016), though we estimate housing demand using consumption microdata whereas they rely on city-level variation in incomes, prices and rental expenditures. Our estimation strategy thus avoids any assumptions about aggregating preferences within a city or about the relationship between income and expenditure. More broadly, we contribute to the literature that links inequality and housing, surveyed by Ioannides and Ngai (2023), by providing new estimates of nonhomothetic housing demand and tracing out their consequences for the connection between inequality and spatial sorting.

The rest of the paper proceeds as follows. Section 1 connects aggregate skill bias to spatial sorting in a simple model of location choice with general nonhomothetic preferences. Section 2 estimates nonhomothetic housing demand. Section 3 embeds these preferences in a quantitative spatial model, and Section 4 uses the calibrated model to quantify the effect of the rising aggregate skill bias on sorting. Section 5 concludes.

1 A Theory of Inequality and Spatial Sorting

We present a model linking changes in aggregate skill bias to changes in the spatial distribution of workers of different skill levels. Nonhomothetic housing demand is the force that connects these two objects. We prove our main results under general preferences, then use a particular parameterization of utility to derive an intuitive expression for the relationship between aggregate skill bias and spatial sorting. Proofs of all the propositions in this section are in Appendix A.

1.1 Environment

There are two types of household, skilled and unskilled, with types denoted by $i = s, u$.[3] The total mass of each type of household is equal to one. There are $n = 1, \ldots, N$ locations. In location n households of type i earn wages w_{in} and derive utility from the consumption of a tradable consumption good, whose price is normalized to one everywhere, and housing, whose price is location-specific, exogenous, and denoted by p_n.

The indirect utility function of a household facing wages w_{in} and housing prices p_n is $v(w_{in}, p_n)$. Free mobility implies that the following spatial indifference condition holds in equilibrium

$$v(w_{in}, p_n) = v(w_{im}, p_m) = v_i \quad \forall n, m \quad \text{and} \quad i = s, u \tag{1}$$

where v_i is the endogenous common utility level for all households of type i. Wages in each location n are determined by the following downward-sloping labor demand curves

$$w_{sn} = A z_n \ell_{sn}^{\alpha-1}, \tag{2}$$
$$w_{un} = z_n \ell_{un}^{\alpha-1}. \tag{3}$$

where $A > 1$ shifts the relative productivity of skilled households in all locations, z_n is a location-specific productivity term, ℓ_{in} is the mass of household of type i in location n, and $\alpha < 1$ governs

[3]Our focus on a binary definition of skill yields a simple model which can be taken to the data in a straightforward way, but it does prevent us from confronting trends in the labor market beyond the widening gap between skilled and unskilled workers. Acemoglu and Autor (2011) point out that labor market polarization is an important trend that cannot be captured in a binary framework, and Davis, Mengus, and Michalski (2021) embed it in a spatial model. Extending our analysis to more than two skill groups is beyond the scope of our paper, but we conjecture a richer model with more than two types would yield similar results. Polarization implies falling relative incomes for middle-skilled workers. Our model would predict that their location choices would move towards those of low-skilled workers and diverge from high-skilled workers. If we then bucket low- and middle-skilled workers into a broad 'unskilled' category – which would be consistent with our construction of these categories in the data, see Appendix B – our model would predict rising spatial sorting.

returns to scale in production. The wage system (2) and (3) is a straightforward way of introducing congestion in local labor demand while maintaining a productivity shifter unique to skilled workers. A large literature in economics has explored the causes and consequences of skill-biased technical change, through which the relative demand for skilled workers has grown over the past forty years. In this spirit, we will refer to the productivity shifter A as aggregate skill bias. An equilibrium of this economy is a vector of employment levels ℓ_{in} (summing to one for each type) and wages w_{in} satisfying (1) – (3).

Our focus is spatial sorting by skill, by which we mean the extent to which skilled households and unskilled households choose to live in *different* locations. We define sorting in terms of the log skill ratio in each location, denoted by μ_n,

$$\mu_n \equiv \log\left(\frac{\ell_{sn}}{\ell_{un}}\right).$$

This object is analytically convenient not only in our model, but also in more general quantitative spatial settings with log-linear labor demand and supply. Our proposed measure of sorting, $\mathcal{M}$, is (proportional to) the variance of μ_n,[4]

$$\mathcal{M} = \frac{1}{2}\text{Var}\left(\mu_n\right). \tag{4}$$

$\mathcal{M}$ is zero when skilled workers are distributed in proportion to unskilled workers across space, and rises as each group clusters in different cities. $\mathcal{M}$ is invariant to proportional increases in the number of skilled workers in all locations. This invariance property will be useful when we turn to the data because the share of skilled workers in the US has grown over time.

1.2 Sorting and Aggregate Skill Bias

The spatial indifference condition (1) implies that labor supply in location n is perfectly elastic with respect to the utility level offered there, so that, as in the canonical models of Rosen (1979) and Roback (1982), in equilibrium higher prices p_n must be offset by higher wages w_{in} to leave utility in n unchanged. We express this relationship using the Hicksian demand for housing. Denoting the housing expenditure share of a household of type i in location n by η_{in}, we can write

$$\eta_{in} = \eta(p_n, v_i). \tag{5}$$

Then totally differentiating (1) and applying Roy's identity shows that perfectly elastic labor supply imposes the following relationship between local prices and the wages of type i,

$$\frac{d\log w_{in}}{d\log p_n} = \eta(p_n, v_i). \tag{6}$$

[4]The division by two is for ease of notation in subsequent expressions and has no effect on any of our results.

Intuitively, if housing is a large share of a household's budget, then the increase in wages needed to compensate them for a given increase in housing costs must be large. When housing demand is nonhomothetic, this elasticity will generally vary across types, and will also depend on aggregate skill bias A.

Differencing (6) between skilled and unskilled households shows that local prices shift the relative labor supply curve

$$\frac{d\log\left(\frac{w_{sn}}{w_{un}}\right)}{d\log p_n} = \eta(p_n, v_s) - \eta(p_n, v_u). \tag{7}$$

Wages consistent with free mobility are pinned down by the local housing price. Equation 7 clarifies the role of nonhomotheticity. If housing is a necessity — which will turn out to be the empirically relevant case — then the housing expenditure share falls as utility rises. In equilibrium skilled workers have higher wages, higher utility, and therefore lower housing expenditure shares. Then $\eta(p_n, v_s) < \eta(p_n, v_u)$ and the left-hand side of (7) is negative. This result is familiar from Black, Kolesnikova, and Taylor (2009), who showed that the skilled wage premium is lower in expensive locations when housing is a necessity. The logic is the same here: unskilled households are relatively more exposed to high housing prices, and so must be compensated with relatively higher nominal wages in expensive cities.

Finally, taking the ratio (2) and (3) yields a relative labor demand curve that is downward sloping in relative quantities,

$$\log\left(\frac{w_{sn}}{w_{un}}\right) = \log A + (\alpha - 1)\,\mu_n \tag{8}$$

Combining relative labor demand with relative labor supply (7) yields the log skill ratio as a function of local prices p_n and utilities v_s and v_u.

$$\frac{d\mu_n}{d\log p_n} = \left(\frac{1}{1-\alpha}\right)(\eta(p_n, v_u) - \eta(p_n, v_s)). \tag{9}$$

Notice that in (9), the *only* reason for different location choices between skilled and unskilled workers is differences in their utilities v_i. This makes the role of skill in our model precise; the high relative productivity of skilled workers leads them to enjoy higher utility in every location, and this is what drives sorting. Proposition 1 formally spells out the implications of (9).

Proposition 1 *Suppose housing is a necessity. Then μ_n, the log skill ratio, is a strictly increasing function of housing prices p_n, and sorting $\mathcal{M} > 0$. If instead housing demand is homothetic, skilled and unskilled workers are distributed in proportion to one another in every location and $\mathcal{M} = 0$.*

So far, we have focused on the level of sorting. Now we turn to changes in sorting caused by economywide changes in skill bias. Proposition 2 states our main theoretical result.

Proposition 2 *Suppose housing is a necessity. Consider an increase in aggregate skill bias $\Delta \log A > 0$. Then, the change in the skill ratio $\Delta\mu_n$ is a strictly increasing function of prices p_n. Sorting rises, $\Delta\mathcal{M} > 0$. If instead housing demand is homothetic, then $\Delta\mu_n = 0$ for all n and $\Delta\mathcal{M} = 0$.*

As aggregate skill bias rises, so does skilled utility v_s . Assuming housing is a necessity, the Hicksian demand (5) implies that skilled housing expenditure shares fall. Skilled households are now more willing to tolerate high prices, and so move towards expensive cities. This process continues until skilled wages in those cities have fallen enough that they once again satisfy (6). In the new equilibrium, skilled workers are more clustered in expensive cities and spatial sorting is higher. Thus, increases in aggregate skill bias can potentially account for at least some of the increase in spatial sorting by skill observed since 1980. At the same time, our model makes the counterfactual prediction that the relative wages of skilled workers should have fallen in expensive cities, highlighting the importance of other forces – for example, faster skill-biased technical change in urban areas. We return to this point in the context of our quantitative model in Section 4.

1.3 Nonhomothetic Constant Elasticity of Substitution Preferences

Our results so far do not hinge on a particular choice of utility function. As long as housing is a necessity, higher skill bias implies higher spatial sorting. In order to quantify the importance of this mechanism, however, we will have to parameterize the utility function. In our context nonhomothetic constant elasticity of substitution (NHCES) preferences, recently highlighted by Comin, Lashkari, and Mestieri (2021), are a convenient choice.

We drop location and type subscripts for the moment. The utility u of a household consuming h units of housing and c units of the consumption good is implicitly defined by

$$u^{\frac{\sigma-1}{\sigma}} = \Omega^{\frac{1}{\sigma}} h^{\frac{\sigma-1}{\sigma}} u^{\frac{\epsilon}{\sigma}} + c^{\frac{\sigma-1}{\sigma}}, \tag{10}$$

where $0 < \sigma < 1$, $\epsilon \geq \sigma - 1$, and $\Omega > 0$ are parameters.[5] The household maximizes u subject to the budget constraint,[6]

$$ph + c \leq w. \tag{11}$$

NHCES preferences admit a straightforward Hicksian demand function $\eta(p, v)$, where v is the common utility level attained by households of a given type. Minimizing expenditure subject to (10) yields

$$\log\left(\frac{\eta(p,v)}{1-\eta(p,v)}\right) = \log\Omega + (1-\sigma)\log p + \epsilon \log v. \tag{12}$$

We can see σ determines the sensitivity of housing expenditure to prices, and ϵ determines how

[5]The restriction $\sigma < 1$ implies housing demand is price-inelastic, which turns out to be the empirically relevant case. We impose $\sigma < 1$ purely for ease of exposition. NHCES preferences in general do allow $\sigma > 1$.

[6]Relative to a fully general formulation, (10) normalizes an ϵ-parameter for the numéraire consumption good to zero. Comin, Lashkari, and Mestieri (2021) show that in a single-location model this normalization is without loss of generality. It is also without loss of generality in the multi-location model we develop in Section 1, because we assume an isoelastic spatial labor supply function. See Appendix A.4 for a proof.

housing expenditure varies with utility. In particular $\epsilon < 0$ implies the housing expenditure share falls with utility, whereas $\epsilon > 0$ implies the opposite. Because utility is monotonically increasing in total expenditure, ϵ determines the sign of the income elasticity of housing expenditure. If $\epsilon < 0$, housing is a necessity and its expenditure share falls with income, whereas if $\epsilon > 0$, it is a luxury and its share rises with income.

Conveniently, NHCES preferences nest the two specifications commonly used in the spatial literature. Cobb-Douglas preferences are obtained by taking $\epsilon = 0$ and $\sigma \to 1$ in (12). In this case, the expenditure share is constant and equal to

$$\eta = \frac{\Omega}{\Omega + 1}. \tag{13}$$

The opposite case, a unit housing requirement, is obtained by taking $\epsilon = -1$ and $\sigma \to 0$. Each household consumes Ω units of housing. In this case, the expenditure share is

$$\eta = \Omega \left(\frac{p}{w} \right). \tag{14}$$

With these preferences in hand, we return to our model of spatial sorting. Propositions 1 and 2 continue to apply, but NHCES preferences allow us to make the relationship between skill bias A and spatial sorting $\mathcal{M}$ more concrete. In particular, we have the following proposition.

Proposition 3 *Assume all households have identical NHCES preferences, and consider a small increase in the aggregate skill bias term $\Delta \log A > 0$. Then the change in sorting, $\Delta \mathcal{M}$, is, to a first-order approximation,*

$$\Delta \mathcal{M} = \left(\frac{-\epsilon}{1-\sigma} \right) \times \left(\frac{1}{1-\alpha} \right) \times \text{Cov}\left(\eta_{sn}, \mu_n \right) \times \chi_s \times \Delta \log A \tag{15}$$

where $\chi_s \equiv \left(1 + \frac{\epsilon}{1-\sigma} \bar{\eta}_s \right)^{-1} > 0$ is the elasticity of skilled utility with respect to A and $\bar{\eta}_s$ is the average skilled housing expenditure share.

Equation 15 clarifies the forces at work in our model. The first term in parentheses captures the role of preferences. To gain intuition for this term, consider the ideal price index of a skilled household in location n,

$$P_{sn} \equiv \frac{w_{sn}}{v_{sn}} = \left(1 + \Omega p_n^{1-\sigma} v_s^{\epsilon} \right)^{\frac{1}{1-\sigma}}. \tag{16}$$

The NHCES ideal price index closely resembles the standard homothetic CES price index, except that the weight on housing prices decreases with utility when $\epsilon < 0$. An increase in skill bias causes skilled utility to rise. Then (16) gives us the corresponding change in the ideal price index,

$$\frac{\Delta \log P_{sn}}{\Delta \log v_s} = \left(\frac{\epsilon}{1-\sigma} \right) \eta_{sn}. \tag{17}$$

We can see that the ideal price index in n declines whenever $\epsilon < 0$, and this decline is steeper whenever ϵ is more negative. The elasticity of substitution σ controls the curvature of P_{sn} and so determines how strongly the increase in v_s is reflected in the ideal price index. Therefore ϵ and σ jointly determine how much ideal price indices fall in response to an increase in A, and thus how much sorting this change induces.

The second term in parentheses captures congestion forces that limit sorting. In our simple model, the only source of congestion is decreasing returns to scale in production. When α is small, decreasing returns to scale set in quickly and congestion is severe. As a result the relationship between sorting and skill bias is weaker. The quantitative model of Section 3 adds a richer set of congestion forces, driven by imperfect labor mobility and imperfect substitutability between skilled and unskilled labor in production.

Finally, the third term shows the role of the pre-existing distribution of skilled and unskilled workers, as captured by the covariance of housing expenditure shares and the log skill ratio. From (17), we can see that an increase in skilled utility lowers the ideal price index of skilled workers more in high η_{sn} locations. This pushes skilled workers towards expensive locations. When those locations are initially skill intensive, and this covariance is positive, spatial sorting rises. In our simple model in which prices are ultimately the only source of sorting, Proposition 1 guarantees that this covariance is positive. In the more general model of Section 3, sorting will be driven by a rich set of location-specific fundamentals and the magnitude and sign of this covariance will come from the data.

Comparison to Type-Specific Cobb-Douglas

Above we have emphasized that when $\epsilon = 0$ and preferences are homothetic, households do not sort on prices and their sorting decisions do not diverge as the skill premium rises. It is reasonable to ask whether the same mechanism might be captured by allowing for exogenous, skill-specific differences in the housing expenditure share while retaining a Cobb-Douglas specification within each skill group. This has appeared in the literature as a tractable stand-in for nonhomothetic preferences (Diamond 2016; Notowidigdo 2020; Colas and Hutchinson 2021; Diamond and Gaubert 2022). Cobb-Douglas preferences imply the utility of type i in location n is

$$v_{in} = w_{in} p_n^{-\kappa_i}$$

where κ_i is the housing expenditure share of type i. With these preferences, the skill ratio in each location is

$$\mu_n = \frac{\kappa_u - \kappa_s}{1-\alpha} \log p_n + \varkappa \tag{18}$$

where $\varkappa$ is constant across locations. Equation (18) shows that when $\kappa_u > \kappa_s$, skilled households will sort into high price locations, just as in the model above. However, unlike in our explicitly nonhomothetic model, changes in aggregate skill bias do not cause changes in spatial sorting. We conclude that in order to capture the mechanism we focus on, it is not enough to impose different expenditure shares by type. Instead, incomes must alter the weight each skill group places on housing costs.

2 Estimating Nonhomothetic Housing Demand

The simple model above highlights the role of preferences, and in particular the NHCES parameters ϵ and σ, in translating changes in aggregate skill bias into changes in spatial sorting. We now estimate those preference parameters. In doing so we confront three main challenges. First, we require expenditure data because the key parameters of the model are the elasticities of the housing expenditure share with respect to total expenditure and local prices.[7] Second, OLS estimates are biased by measurement error in expenditure, so we require an instrument. Finally, and most importantly, the price of housing varies widely across space, and is correlated with household income. Therefore, we need to control for variation in housing prices. As we show below, failing to do so would strongly bias our results toward homotheticity.

Combining consumption microdata with detailed geographical information also allows us to advance the literature (reviewed in Subsection 2.6) in two ways: we avoid any assumptions about aggregating demand across agents within a location; and we investigate the potential role of permanent, unobserved heterogeneity across agents in driving the relationship between total expenditure and housing expenditure shares.

2.1 Data

We use the restricted-access Panel Study of Income Dynamics (PSID), which identifies households' county of residence (University of Michigan Institute for Social Research 2021). Since 2005, the PSID has collected information on essentially all consumption covered by the Consumer Expenditure Survey (CEX) (Andreski et al. 2014). We use the 2005-2017 biennial surveys. Our baseline sample is restricted to renting households because they have a clear measure of housing consumption, but we also find similar results using homeowners in Appendix C.

The PSID has two advantages relative to the CEX. One, we can link price data to about 90% of households in the PSID. By contrast, the CEX has geographic identifiers only for households in 24 large cities, which is less than half the CEX sample. Two, the PSID follows the same households over time, so we can study how housing expenditure responds to changes in total expenditure within the same household.

In the model of Section 1, housing consumption h is a scalar. In reality, housing consumption is determined by a bundle of characteristics like square footage and number of bedrooms. We therefore estimate the price of housing for each Metropolitan Statistical Area (MSA) with a hedonic regression as in Albouy (2016), so that local prices $\log p_n$ correspond to the fixed effects in a regression of log household rent on observed housing unit characteristics, both taken from the American Community Survey (ACS). In Appendix C.1 we clarify the rationale underlying this approach by modeling housing consumption as an aggregate over characteristics, and show that this approach yields price indices consistent with those obtained from our hedonic regressions.

[7] At the risk of ambiguity, we use the familiar term "income elasticity" as shorthand for "expenditure elasticity" throughout the rest of the paper.

For more details of our data, sample selection, and price indices, see Appendix B.

2.2 Estimation

As in Section 1, we assume households have NHCES preferences over housing and a numéraire consumption good. We do not use any of the other structure of the model in this section. We also do not aggregate households into two skill groups, as in the model, but instead use the whole income distribution, although we will show in Table 2 that this choice has no impact on our results. Finally, we denote total expenditure by e rather than w, to emphasize that in reality expenditure and income are distinct concepts.

To take the Hicksian demand function (5) to the data, we substitute out unobservable utility v.[8] This yields an expression that implicitly defines η as function of expenditure, prices, and parameters:

$$\eta = \Omega e^{\epsilon} p^{1-\sigma} (1-\eta)^{1+\frac{\epsilon}{1-\sigma}} . \tag{19}$$

Equation 19 is written entirely in observables and unknown parameters, and is the relationship we will take to the data.

We consider households indexed by i in years t. Households reside in MSAs indexed by n. Housing prices vary by MSA and year and are denoted by p_{nt}, whereas the price of the consumption good is assumed not to vary across space and is normalized to one. We assume a common housing market within each MSA, so that prices p_{nt} do not vary within a city. We interpret Ω as an idiosyncratic shock to an individual household's taste for housing, so that (19) becomes

$$\eta_{it} = \Omega_{it} e_{it}^{\epsilon} p_{nt}^{1-\sigma} (1-\eta_{it})^{1+\frac{\epsilon}{1-\sigma}} . \tag{20}$$

To build intuition for our estimation strategy, we log-linearize (20) around the median housing share $\bar{\eta}$ to obtain

$$\hat{\eta}_{it} = \left(\frac{1-\bar{\eta}}{1-\bar{\eta}+(\frac{\epsilon}{1-\sigma}+1)\bar{\eta}} \right) \left(\hat{\Omega}_{it} + \epsilon \hat{e}_{it} + (1-\sigma)\hat{p}_{nt} \right) , \tag{21}$$

where $\hat{x}$ denotes the log deviation of a variable x from its median. Equation (21) reads as

$$\hat{\eta}_{it} = \omega_{it} + \beta \hat{e}_{it} + \psi \hat{p}_{nt}, \tag{22}$$

where $\omega_{it} \equiv \left(\frac{1-\bar{\eta}}{1-\bar{\eta}+(\frac{\epsilon}{1-\sigma}+1)\bar{\eta}} \right) \hat{\Omega}_{it}$ and β and ψ are defined analogously. Under the null of homothetic preferences, $\epsilon = \beta = 0$. We bring (22) to the data by modeling the demand shifter ω_{it} as a function of observable demographic characteristics, year fixed effects, and an additive error. Formally,

$$\hat{\eta}_{it} = \omega_t + \omega' X_{it} + \beta \hat{e}_{it} + \psi \hat{p}_{nt} + \xi_{it}, \tag{23}$$

where X_{it} is a vector with the age, gender, and race of the household head, household size, and the

[8] From the Hicksian demand for the consumption good, $v = (1-\eta)^{\frac{1}{1-\sigma}} e$.

number of earners in the household. We observe total expenditure e_{it}, the housing expenditure share η_{it}, and prices p_{nt}. The error term ξ_{it} represents measurement error in expenditure and random shocks to housing demand.[9]

2.3 Main Results

Table 1, columns (1) - (4), show the estimates of (23). Note that because columns (1) and (2) do not attempt to estimate the coefficient on prices, they cannot recover the structural parameters ϵ and σ. Column (1) estimates (23) by OLS without controlling for price $\hat{p}_{nt}$. The point estimate indicates significant nonhomotheticity, but two sources of bias are evident. First, measurement error in expenditure is likely to bias $\hat{\beta}$ downwards.[10] Second, a positive correlation between prices and expenditure, reflecting the sorting of high-income households into high-price MSAs, will bias $\hat{\beta}$ upwards.

Column (2) addresses measurement error by instrumenting for log expenditure using log income, following Lewbel (1996), Davis and Ortalo-Magné (2011), and Aguiar and Bils (2015). As expected, $\hat{\beta}$ rises toward zero. The exclusion restriction here is that income is unrelated to the housing share, conditional on the true level of expenditure. One threat to identification is that if housing expenditure is subject to some adjustment costs, it may react to income changes more slowly than overall expenditure. This would bias our estimates downwards. Another threat is that there may be permanent, unobservable differences in housing demand across households which are correlated with income. We address both these concerns with alternative specifications in Table 2.

Column (3) of Table 1 returns to OLS but addresses omitted variable bias by controlling for prices. Relative to column (1) the coefficient on log expenditure falls, implying very income-inelastic housing demand. This result is consistent with high-income households sorting into high-price MSAs — exactly the pattern the model of Section 1 predicted. Using (21) we also back out estimates of the structural parameters ϵ and σ.

Together columns (2) and (3) show that failing to instrument for expenditure and control for prices introduces offsetting biases in the coefficient on expenditure. Column (4) corrects for both biases simultaneously by instrumenting for expenditure using income and controlling for prices.

Prices are potentially endogenous because they are a function of housing demand. For example, a city-level shock to housing demand might increase expenditure shares and, consequently, prices. We therefore instrument for prices in column (4) using Saiz (2010)'s measures of regulatory and geographical constraints on housing construction. These instruments are relevant if tight constraints force up local housing costs. They satisfy the exclusion restriction if, conditional on prices and total expenditure, they don't have an effect on housing expenditure. Thus, for example, the

[9]In Appendix E.4 we show that the log linear specification (23) is exactly the demand curve obtained under the assumption of Price-Independent Generalized Linear (PIGL) preferences. PIGL preferences are a workhorse choice for nonhomothetic preferences in the structural change (Boppart 2014) and spatial (Eckert and Peters 2023) literatures.

[10]Because expenditure appears in the denominator of $\hat{\eta}$, the bias in $\hat{\beta}$ is not standard classical measurement error. See Appendix C for a short proof.

existence of a correlation between supply restrictions and local productivity or amenity shocks — a possibility highlighted by Davidoff (2016) — would not violate this assumption as long as these shocks do not alter the relative attractiveness of housing versus nonhousing consumption (i.e Ω_{it}), given prices and total expenditure.[11]

The results in column (4) imply that housing is a necessity, and is moderately price inelastic. A 10% increase in total expenditure causes a 2.5% decrease in the housing expenditure share. Finally, column (5) shows our preferred specification. Here, we estimate the nonlinear equation (19) directly by GMM, rather than estimating a linearized version as in columns (1)-(4).[12] Similarly to column (4), we instrument for expenditure and prices, and allow ω_{it} to vary with demographic characteristics and year. The estimated ϵ and σ are close to their values in column (4).

We now compare the preferences estimated in Table 1 to two benchmarks from the literature:

Table 1: Preference Estimates

Dependent variable: Log housing share

	(1) OLS	(2) 2SLS	(3) OLS	(4) 2SLS	(5) GMM
ϵ			-0.436	-0.291	-0.306
			(0.018)	(0.037)	(0.036)
σ			0.436	0.542	0.522
			(0.039)	(0.079)	(0.075)
Log expenditure	-0.298	-0.162	-0.393	-0.248	
	(0.028)	(0.039)	(0.021)	(0.035)	
Log price			0.508	0.390	
			(0.026)	(0.057)	
Demographic controls	✓	✓	✓	✓	✓
R^2	0.12		0.20		
First-stage F-stat.		1,264.4		107.3	
N	12,351	12,351	12,351	10,678	10,678
No. of clusters	484	484	484	217	217

Source: PSID, Census, and Saiz (2010)

Note: Columns (1)-(4) estimate the linear equation 23 and column (5) estimates the nonlinear equation 19. Renters only. Instruments are log household income (columns (2), (4), and (5)) and housing supply constraints (column (5)). Demographic controls are bins for family size, number of earners, and sex, race, and age of household head. All regressions include year fixed effects. Standard errors clustered at MSA level. See Appendix B for further details of sample construction.

[11] One may still be concerned that supply constraints directly influence Ω_{it}. Column (2) of Table 2 addresses this concern by repeating the estimation with MSA-fixed effects. The relationship between total expenditure and the housing share is virtually unchanged relative to column (4) of Table 1.

[12] In stating our preferences, we imposed $\epsilon > \sigma - 1$ and $0 < \sigma < 1$. We do not impose these restrictions in our estimation procedure, but they are satisfied by the values obtained in column (5).

Figure 1: Housing Expenditure Shares

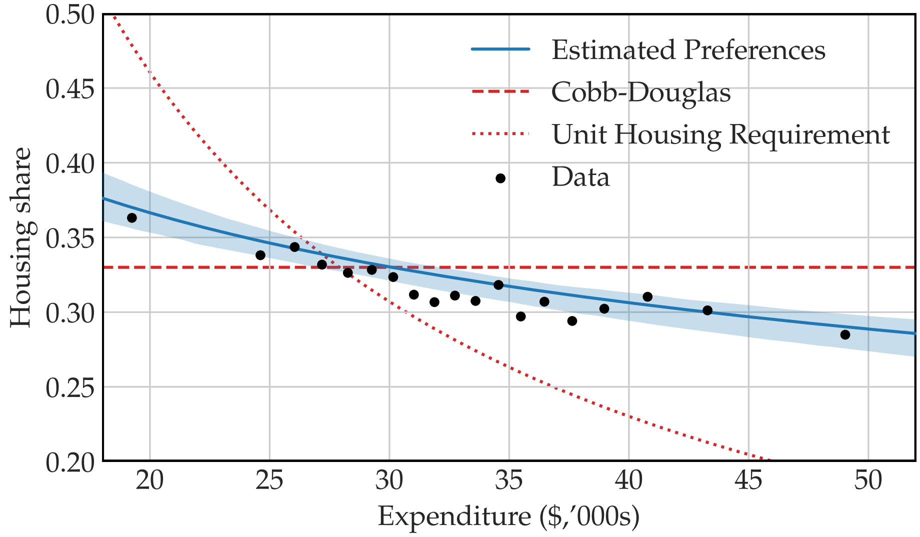

Notes: 'Estimated Preferences' plots (19) at the parameter values obtained in Table 1, column (5). The shaded area represents a 95% confidence interval.'Cobb-Douglas' and 'Unit Housing Requirement' plot the preferences described by (13) and (14), respectively, with the scale parameter Ω chosen to match an expenditure share of 0.33 at the median level of total expenditure. 'Data' plots the average housing share in twenty evenly sized bins defined by predicted total expenditure, whose construction is described in the text.

Cobb-Douglas preferences and a unit housing requirement. We begin with formal statistical tests. The NHCES preferences estimated above nest both of these special cases. The null hypothesis of Cobb-Douglas preferences, corresponding to $\epsilon = 0$ and $\sigma = 1$, can be rejected at the 1% level. A unit housing requirement corresponds to $\epsilon = -1$ and $\sigma = 0$ — again, column (5) allows us to reject this null hypothesis at the 1% level. Although our NHCES specification is more flexible than these special cases, it still imposes a particular functional form on the relationship between total expenditure and housing expenditure. To assess the validity of this assumption we construct a binned scatterplot of expenditure against housing shares.[13] The results are shown in Figure 1 alongside our estimated preferences (the solid line). Our estimated preferences appear to fit the data well. For comparison we also plot Cobb-Douglas preferences and a unit housing requirement, given by the dashed and dotted lines respectively. Neither alternative comes close to matching the data.

2.4 Alternative Specifications

Table 2 shows alternative specifications. We discuss each in detail below.

[13]We do not use expenditure directly, since as discussed above measurement error contaminates the relationship between expenditure and the housing share. Instead, we predict total expenditure for each household using the instruments and covariates in column (5) of Table 1, then split households into twenty bins of predicted expenditure and calculate the average housing share in each bin, partialling out covariates.

Household Fixed Effects

We consider the possibility of permanent, unobservable differences in housing demand across households. We parameterize the demand shifter Ω_{it} as follows

$$\log \Omega_{it} = \omega_i + \omega_t + \omega' \tilde{X}_{it} + \xi_{it}$$

where ω_i is a household fixed effect, $\tilde{X}_{it}$ is the subset of demographic controls which are time-varying and ξ_{it} is an idiosyncratic error term. Taking logs of (19) then yields

$$\log \eta_{it} = \omega_i + \omega_t + \omega' \tilde{X}_{it} + \epsilon \log e_{it} + (1-\sigma) \log p_{nt} + \left(1 + \frac{\epsilon}{1-\sigma}\right) \log (1-\eta_{it}) + \xi_{it} \qquad (24)$$

Equation (24) allows for permanent, unobservable differences in housing demand across households, captured by ω_i. If ω_i happens to be negatively correlated with income, this specification could generate the negative relationship between expenditure and η found in Table 1 even when $\epsilon = 0$ and preferences are homothetic. Such permanent differences in housing demand are sometimes used in the literature as a tractable alternative to explicitly nonhomothetic preferences (Diamond 2016; Notowidigdo 2020; Colas and Hutchinson 2021; Diamond and Gaubert 2022). As we have shown in Section 1, however, distinguishing between such demand shifters and explicitly nonhomothetic preferences is critical for the mechanism we focus on in this paper.

We demean (24) at the household level so that ω_i drops out.[14] We estimate the demeaned equation by GMM, using the same instruments as in column (5) of Table 1. Since the instruments for p_{nt} do not vary over time, σ is identified only by households who face different prices because they move between MSAs. The results are reported in column (1) of Table 2. The point estimate for ϵ falls relative to our baseline, indicating somewhat stronger nonhomotheticity, but the two estimates are not significantly different. The price elasticity σ is very close to its baseline value. We are still able to reject both Cobb-Douglas preferences and a unit housing requirement. We conclude that permanent, unobservable differences in housing demand across households are not driving our baseline results: even within a single household, an increase in total expenditure decreases the housing expenditure share.

Alternative Instruments

A natural concern is that housing expenditure is relatively insensitive to total expenditure because housing expenditure can only be adjusted slowly while total expenditure may fluctuate with transitory income shocks. Column (2) addresses this concern by instrumenting for expenditure using the household's education level. Since differences in education across households are permanent,[15] slow adjustment of housing expenditure to transitory shocks is irrelevant in this

[14] In Appendix Table C.3, column (7), we pursue an alternative estimation strategy by log-linearizing (24) and using 2SLS with household fixed effects. We find almost identical point estimates.

[15] For 90% of households education level does not change while they are in the sample.

specification. The point estimates in column (2) are similar to those in our baseline specification and again indicate that housing is a necessity.

Within-MSA Results

The specifications in Table 1 identify ϵ and σ using variation both within and across MSAs. One might therefore be concerned about the role of sorting across MSAs in driving our results. While we have controlled for price differences in columns (3) - (5), differences in an unobservable shock to the taste for housing across MSAs, captured in the error term ξ_{it}, might be playing a role. For example, suppose individuals with a strong taste for housing (conditional on total expenditure and demographics) sort into low price MSAs. Then we would see relatively high housing expenditure shares in low price cities and would underestimate the sensitivity of expenditure shares to prices.

To investigate this possibility, we return to the linearized specification (23) but replace the

Table 2: Preferences, Alternative Specifications
Dependent variable: Log housing share

	(1) GMM	(2) GMM	(3) 2SLS	(4) GMM	(5) GMM
ϵ	-0.465	-0.271		-0.300	-0.322
	(0.109)	(0.065)		(0.041)	(0.085)
σ	0.511	0.532		0.389	0.515
	(0.198)	(0.077)		(0.090)	(0.040)
Log expenditure			-0.261		
			(0.032)		
Household FE	✓				
MSA FE			✓		
Non-housing prices				✓	
Demographic controls	✓	✓	✓	✓	
IV	Income	Education	Income	Income	Income
Aggregation	Household	Household	Household	Household	City × Education
First-stage F-stat.			1,054.6		
N	8,670	10,271	12,257	8,183	1,912
No. of clusters	197	216	390	208	217

Source: PSID, Census, and Saiz (2010)

Note: Renters only. Columns (1), (3), (4), and (5) use housing supply constraints as instruments. Demographic controls are bins for family size, number of earners, and sex, race, and age of household head. Column (4) includes only time-varying demographic controls. All regressions include year fixed effects. Standard errors clustered at MSA level. See Appendix B for further details of sample construction.

prices p_{nt} with MSA fixed effects.[16] This specification therefore exploits only within-MSA variation The results are shown in column (3) of Table 2. The estimated coefficient on log expenditure is -0.261, very close to the value of -0.248 that we estimated in column (4) of Table 1. The similarity of the two coefficients is reassuring. It implies that differences in unobservables across MSAs are not driving the estimated relationship between total expenditure and the housing expenditure share.

Now, the specification estimated in column (3) of Table 2 does not allow us to directly infer the preference parameters ϵ and σ. Instead, it gives us the composite parameter β defined in (23). For a given value of σ, however, we can use this estimate to back out an implied value of ϵ. Varying σ between 0.25 and 0.75 (recall from column (5) of Table 1 that our central estimate for σ is 0.522), we obtain values of ϵ between -0.249 and -0.318, not too far from our preferred estimate of -0.306. We conclude that sorting across MSAs based on unobservables is not playing an important role in our estimation of ϵ. As we will see in Section 4, this will turn out to be the crucial parameter in relating changes in the income distribution to changes in spatial sorting.

Nonhousing Prices

Housing is not the only good whose price varies across space, and variation in other prices might in principle bias our estimates of ϵ and σ. However, a quick glance at the data suggests any potential misspecification is quantitatively small. Using the Bureau of Economic Analysis (BEA) Metropolitan Regional Price Parities (Bureau of Economic Analysis 2020) for 2008-2017, we calculate the standard deviations of rental prices, goods prices, and service prices across MSAs. The vast majority of spatial variation in cost of living comes from rents — the standard deviations of goods prices and service prices are roughly one eighth and one fifth as large as the standard deviation of rental prices, respectively — suggesting that omitting other prices from our main estimation is not likely to have a large impact on ϵ and σ. To verify this intuition, we incorporate nonhousing prices, denoted by q_{nt}, into our NHCES preferences. Equation (20) becomes

$$\eta_{it} = \Omega_{it} \left(\frac{e_{it}}{q_{nt}} \right)^{\epsilon} \left(\frac{p_{nt}}{q_{nt}} \right)^{1-\sigma} (1 - \eta_{it})^{1+\frac{\epsilon}{1-\sigma}} . \tag{25}$$

The nonhousing price index q_{nt} is the price of a Cobb-Douglas aggregate of goods and nonhousing services constructed from the Regional Price Parities. The weight on goods is 0.51 and on nonhousing services 0.49, in line with the weights used by the BEA in constructing the price indices. The results of estimating (25) by GMM are shown in column (4) of Table 2. The point estimate of ϵ is virtually unchanged relative to its value in column (5) in Table 1, while the estimate of σ is somewhat smaller — but not significantly so.

[16] Note that, unlike the prices p_{nt}, the MSA fixed effects do not vary with time. We have experimented with MSA-by-year fixed effects, and have found they do not change our results.

Abstracting from Household Heterogeneity

A potential concern is that the model in Section 1 and the estimation here operate at different levels of aggregation. The income distribution in the model has only two points in each city, corresponding to skilled and unskilled workers, whereas our estimates so far use variation from the entire income distribution. We address this inconsistency by aggregating households in the data into two groups, namely those with and without a four year college degree, which we map to skilled and unskilled workers in the model. We then compute the mean housing share, expenditure, and income by city, skill level, and year. Column (5) of Table 2 reports the GMM estimates of this aggregated specification using the same instruments as in column (5) of Table 1. They are virtually identical to the baseline estimates.

Fixed Costs in Housing

We have modeled the choice of housing as a smooth, unconstrained problem. In practice, housing costs may have a fixed component — one can only have so few square feet and so many roommates, which puts a lower bound on housing expenditure. In Appendix C.3 we consider an alternative specification of housing demand in which underlying preferences are homothetic but the household faces fixed and variable costs of housing. We show that this system is isomorphic to one with Stone-Geary preferences and no fixed costs. We then estimate those preferences and again find that housing is a necessity. The main results of the paper do not depend on the underlying interpretation of preferences; what is important, instead, is how housing expenditure shares decline with income.

Appendix Specifications

Finally, we explore a number of alternative specifications and data sources in Appendix C. We assess the robustness of our results to controlling for liquid wealth; to removing demographic controls; to using prices directly rather than instrumenting for them; to using alternative data sources and geographies for prices; to splitting the sample into movers and non-movers; and to using alternative instruments for expenditure. We continue to find that housing is a necessity. We replicate our results using the CEX, and then extend them to include homeowners (Bureau of Labor Statistics 2020a). The estimated parameters look very similar when we include homeowners.

2.5 Housing Expenditure Shares Across Space

Cobb-Douglas preferences have been a popular choice in quantitative spatial models, because prior work (Davis and Ortalo-Magné 2011) has documented that housing expenditure shares vary relatively little across cities with widely different levels of average income.[17] The dots in panel (a) of Figure 2 show a scatterplot of median housing expenditure shares against median income

[17] Behrens, Duranton, and Robert-Nicoud (2014) and Albouy, Ehrlich, and Liu (2016) both note that the offsetting price and income effects we discuss here can generate constant expenditure share across space.

Figure 2: Housing Shares Across Space

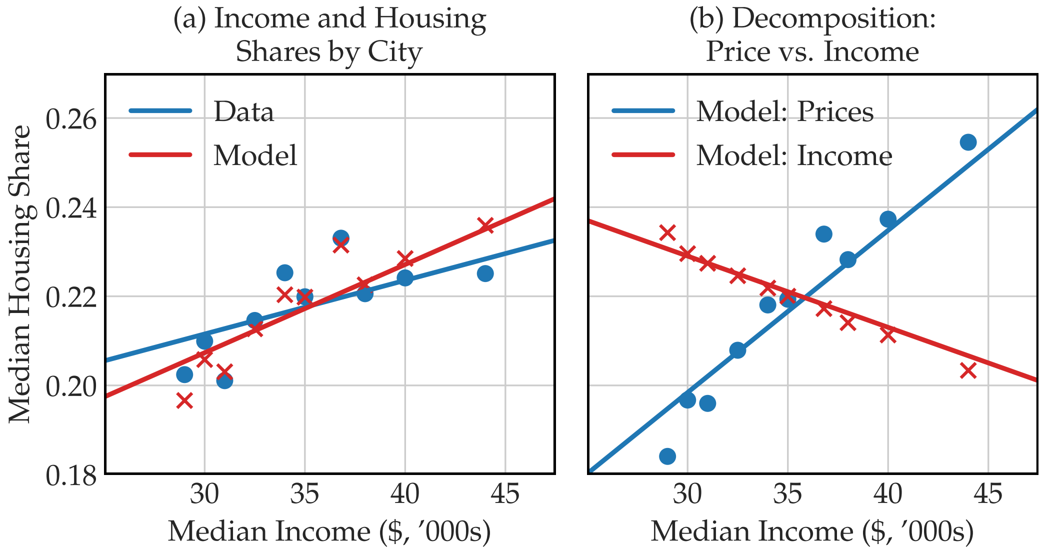

Source: 2000 Census data on rental expenditure and household income, renters only. *Notes*: Panel (a): 'Data' plots median income against median housing share. Housing share calculated as the ratio of rental expenditure to income, as in Davis and Ortalo-Magné (2011). Dots show averages in ten bins and line shows a fitted regression, weighted by 2000 employment. 'Model' shows the same objects generated by the model in column (5) of Table 1. Note that the parameter Ω is chosen so that the average expenditure share produced by the model matches the average in the data. Panel (b): 'Model: Prices' shows expenditure shares predicted by the model with prices varying as in the data, but with incomes held constant at their average across cities. 'Model: Incomes' instead holds prices constant and varies incomes.

at the MSA level; the solid line shows a fitted regression. In constructing this plot we closely follow Davis and Ortalo-Magné (2011) and calculate housing expenditure shares as the ratio of rental expenditure to household income using Census data from 2000. We can see that in the data there is a modest positive relationship between the two variables. High income MSAs have relatively high housing expenditure shares. The crosses in panel (a) show predicted expenditure shares produced by our model with parameters taken from column (5) of Table 1. The model produces a moderately positive relationship, matching the data well. Panel (b) of Figure 2 shows why. Here we again plot the expenditure shares predicted by our model, but now we separate out the roles of incomes and prices. Specifically, the dots are the expenditure shares that result from allowing incomes to vary as in the data, but holding prices constant; and the crosses vary prices but hold incomes constant. The price and income effects work in opposite directions, and offset one another to produce the mildly positive relationship seen in the left panel.

2.6 Connections to prior work

A wide range of estimates of the income elasticity of housing expenditure shares exist in the literature. We summarize these estimates in Appendix C.6, and also note the extent to which each paper addresses the three challenges we highlighted at the start of this section. Our preferred estimate of about -0.25 lies around the middle of the estimates we survey.

Closest to our approach is Albouy, Ehrlich, and Liu (2016), who allow for nonhomotheticity when estimating preferences over housing and nonhousing consumption and also find that housing expenditure shares decline with income. That paper aggregates to the MSA level and uses data on income rather than expenditure, while we take individual households as our unit of analysis and use expenditure data. We view our results as complementary to Albouy, Ehrlich, and Liu (2016)'s, but note that our approach avoids some assumptions which are inherent in theirs. Our estimation procedure does not assume that demands can be aggregated across households of different income levels. Furthermore, by directly using data on expenditure we avoid assumptions on the relationship between expenditure and income. Finally, using variation within a household allows us to reject the hypothesis that the observed negative relationship between the housing share and total expenditure is driven by permanent, unobservable household characteristics. This is not possible when the data are aggregated to the MSA level.

3 Quantitative Model and Calibration

We now turn to quantifying the implications of our estimates of ϵ and σ for the relationship between the aggregate skill bias of labor demand and spatial sorting. We start by enriching the simple model of Section 1 on a number of dimensions: imperfect labor mobility, constant elasticity of substitution (CES) production, progressive income taxation, and inelastic housing supply.

We also add a full set of type- and location-specific productivity and amenity shifters.[18] We do this for two reasons. First, as we show below, this allows the model to exactly match any observed equilibrium. Recall from equation 15 in Section 1 that the quantitative importance of the mechanism we are interested in depends on the covariance of housing expenditure shares with the skill ratio, capturing the extent to which skilled workers sort into expensive location. Therefore exactly matching the distribution of skilled and unskilled workers across space allows us to accurately measure the contribution of changes in aggregate skill bias to changes in spatial sorting. Second, saturating the model with shifters in this way allows us to perform counterfactual experiments in a transparent way by using the 'hat algebra' technique that is standard in quantitative spatial models (Redding and Rossi-Hansberg 2017).

3.1 Quantitative Model

Preferences and Location Choice

As in the simple model of Section 1, household types are indexed by i, and households are either skilled ($i = s$) or unskilled ($i = u$). They live in locations indexed by n, where they have expenditure levels e_{in} and face local housing prices p_n. We continue to assume NHCES preferences,

[18] We allow all of these shifters to vary over time, but for clarity suppress time subscripts while laying out the model.

so that indirect utility v_{in} is implicitly defined by

$$v_{in} = e_{in} \left(1 + \Omega p_n^{1-\sigma} v_{in}^{\epsilon}\right)^{\frac{-1}{1-\sigma}}. \tag{26}$$

Just as in (19), these preferences imply a housing expenditure share η_{in} for each type i and location n

$$\eta_{in} = \Omega e_{in}^{\epsilon} p_n^{1-\sigma} \left(1 - \eta_{in}\right)^{1+\frac{\epsilon}{1-\sigma}}. \tag{27}$$

The simple model made the stark assumption that households are freely mobile across space. In reality moving can involve large costs (Kennan and Walker 2011), and employment is less than perfectly elastic with respect to the utility offered by a location. Capturing this phenomenon is especially important given that Proposition 3 highlighted the importance of congestion forces (broadly defined) in mediating the relationship between aggregate skill bias and spatial sorting. We therefore allow for imperfect mobility by assuming households have heterogeneous preferences over locations as in Redding (2016). Each household of type i draws an vector of idiosyncratic shocks b and chooses a location n to solve

$$n = \operatorname{argmax}_{m=1,\dots N} \{b_m v_{im}\}.$$

The preference shock b_m is drawn from a Fréchet distribution with cdf G_{im},

$$G_{im}(b) = \exp\left(-B_{im} b^{\theta^{-1}}\right).$$

All locations have a common shape parameter θ. However, we allow the distribution of draws in each location to differ in its scale parameter, which we denote by B_{im}. A high B_{im} corresponds to a high average amenity for type i in location m.[19] The probability a household of type i chooses location n is then

$$\ell_{in} = \frac{v_{in}^{\theta^{-1}} B_{in}}{\sum_m v_{im}^{\theta^{-1}} B_{im}}, \tag{28}$$

and with a continuum of households this choice probability is exactly location n's employment share among households of type i. We impose that total population sums to one, and denote the share of skilled workers by $\phi \in [0,1]$. Skilled employment in location n, denoted L_{sn}, is therefore equal to $\phi \ell_{sn}$, and similarly unskilled employment is $L_{un} = (1-\phi)\,\ell_{un}$.

Production and Taxation

The labor demand curves (2) and (3) in the simple model have two counterfactual implications: wages for type i in location n are independent of labor supply of type $j \neq i$ in the same location; and relative wages in location n depend only on relative labor supply. In reality, skilled wages

[19] Note that amenities B_{im} do not enter the problem (26) which defines v_{im} and η_{im} and so do not directly affect housing demand. Of course, amenities may still influence housing demand through their effect on endogenous wages and prices, but this poses no threat to the identification strategy pursued in Section 2.

respond to unskilled labor supply and vice versa, and locations with high skill ratios typically have high skill premia, suggesting a role for relative demand shifters. To address these shortcomings, we enrich the model with a constant elasticity of substitution (CES) production function with location-specific skill bias. Given skilled employment L_{sn} and unskilled employment L_{un}, output in location n is

$$y_n = F(L_{sn}, L_{un}) = Zz_n \left((Aa_n L_{sn})^{\frac{\rho-1}{\rho}} + L_{un}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}, \tag{29}$$

with $\rho > 1$ the elasticity of substitution between skilled and unskilled labor. Implied wages are

$$w_{un} = (Zz_n)^{\frac{\rho-1}{\rho}} L_{un}^{\frac{-1}{\rho}} y_n^{\frac{1}{\rho}}, \tag{30}$$

$$w_{sn} = (Aa_n)^{\frac{\rho-1}{\rho}} (Zz_n)^{\frac{\rho-1}{\rho}} L_{sn}^{\frac{-1}{\rho}} y_n^{\frac{1}{\rho}}. \tag{31}$$

As in the simple model, the quantitative model contains a location-specific productivity shock z_n and an aggregate skill bias term A. We additionally allow for location-specific skill bias using the shifter a_n, so that skilled households may have a comparative advantage in working in, say, San Francisco relative to Detroit. The result is that the skill premium may be higher in expensive, skill-intensive cities, in line with the data. The economy-wide productivity shifter Z is for notational convenience when we conduct counterfactuals.[20]

Finally we relate wages w_{in} to expenditure e_{in}. There are two differences between income and expenditure. The first is that the relevant quantity for expenditure is permanent income, but in the data we only observe current income. However, aggregating to the level of a skill group averages away any transitory income shocks, making this less of a concern. Second, taxes create a wedge between income and expenditure. We incorporate this wedge into our model following Heathcote, Storesletten, and Violante (2017),

$$e_{in} = T w_{in}^{1-\tau}. \tag{32}$$

where we impose that expenditure is equal to after-tax income. τ determines the progressivity of the tax system and T is an endogenous variable chosen so that the government budget balances.

Housing Costs

In the simple model, the price of housing p_n was exogenous. In reality, increases in A push skilled households toward expensive cities, putting upward pressure on housing costs and crowding out unskilled households. The quantitative model captures this feedback to house prices by including inelastic housing supply as in Hsieh and Moretti (2019). The price of housing in location n is given by

$$p_n = \Pi_n (HD_n)^{\gamma_n}, \tag{33}$$

[20]The levels of a_n and z_n are not separately identified from A and Z. We therefore assume that the geometric means of a_n and z_n, weighted by 1980 employment shares, are both equal to one.

where HD_n is housing demand in n and Π_n is an exogenous price shifter. The parameter γ_n, which governs the elasticity of housing supply, is allowed to vary by location to reflect different physical or regulatory constraints on building. Housing demand is the sum of housing expenditure by both types of households, equal to[21]

$$HD_n = \sum_i \eta_{in} e_{in} L_{in}. \tag{34}$$

Equilibrium

Given parameters $(\epsilon, \sigma, \Omega, \theta, \rho, \tau, \{\gamma_n\})$, location-specific fundamentals $(a_n, z_n, B_{un}, B_{sn}, \Pi_n)$ for all n, aggregate fundamentals (Z, A), and the aggregate skill share ϕ, an equilibrium is a vector of indirect utilities v_{in}, expenditure shares η_{in}, employment shares ℓ_{in}, wages w_{in}, total expenditures e_{in}, housing demands HD_n, and prices p_n satisfying equations (26) - (34).

A Neutrality Result

We conclude this subsection by extending part of Proposition 2 to the quantitative model. Crucially, although our quantitative model accommodates rich patterns of sorting based on location-specific skill biases a_n and amenities B_{in}, the following proposition shows that homotheticity shuts down any relationship between aggregate skill bias A and sorting.

Proposition 4 *Suppose $\epsilon = 0$ so that preferences are homothetic. Then, $\mathcal{M}$, the level of sorting, does not depend on aggregate skill bias A.*

See Appendix A.5 for a proof. The intuition for this result is straightforward. When $\epsilon = 0$, local prices p_n do not drive spatial sorting by skill. Instead, the skill ratio in n is can be written

$$\mu_n = \vartheta_0 + \vartheta_1 \log a_n + \vartheta_2 \log \left(\frac{B_{sn}}{B_{un}} \right),$$

where ϑ_0 does not vary across locations and ϑ_1 and ϑ_2 depend only on parameters. The striking feature of this expression is that μ_n is entirely determined by location-specific fundamentals a_n and B_{in}. Changes in A have no impact on the distribution of skill ratios, and thus no impact on sorting, when preferences are homothetic. Proposition 4 is useful because it implies any changes in sorting in our quantitative model following changes in A, or indeed any aggregate fundamental, are ultimately the result of nonhomothetic housing demand.

[21] Alternatively, we could specify the housing price as a function of the quantity of housing demanded, H_n, rather than expenditure $HD_n = H_n \times p_n$. These formulations are equivalent. To see this, let $p_n = \widetilde{\Pi}_n H_n^{\widetilde{\gamma}_n}$; after some algebra, this reduces to $p_n = \widetilde{\Pi}^{\frac{1}{\widetilde{\gamma}_n+1}} (HD_n)^{\frac{\widetilde{\gamma}_n}{\widetilde{\gamma}_n+1}}$, which is our setup with $\gamma_n = \frac{\widetilde{\gamma}_n}{\widetilde{\gamma}_n+1}$.

3.2 Calibration

Data

Location-level information on wages, rents and employment are from IPUMS (Ruggles et al. 2020). We use the 5% population samples of the 1980, 1990, and 2000 decennial censuses and the 3% population sample from the 2009-2011 ACS. Our census sample consists of prime-age adults who report strong labor-force attachment. Workers with a four-year college degree or higher level of education are classified as "skilled", while remaining workers are classified as "unskilled." We have a balanced panel of 269 locations: 219 MSAs and the 50 non-metropolitan portions of states. Wages and rents are deflated by the Consumer Price Index (CPI) excluding shelter (Bureau of Labor Statistics 2020b). Location-level price indices are constructed each year from a hedonic rents regression as in Section 2. See Appendix B for more details on the data and definitions.

Hat Algebra and Fundamentals

Before we can quantify the implications of our theory, we need to calibrate two sets of objects: the parameters $(\epsilon, \sigma, \Omega, \theta, \rho, \tau, \{\gamma_n\})$ and the fundamentals: the location-specific productivity, amenity and housing-supply shifters $\left(a_n^t, z_n^t, B_{un}^t, B_{sn}^t, \Pi_n^t\right)$, the aggregate productivity parameters A^t and Z^t, and the aggregate skill share ϕ_t. Note we have added a time superscript because we allow all fundamentals to vary by year.

As is standard in quantitative spatial models (Redding and Rossi-Hansberg 2017), the model's fundamentals can be chosen so that the model rationalizes any observed equilibrium.[22] The next proposition shows that recovering these fundamentals explicitly is not even necessary for solving the model relative to some observed observed equilibrium because it can be written in the familiar 'hat algebra' form.

Proposition 5 *Suppose we have observations on wages $\bar{w}_{in}$ and labor supplies $\bar{L}_{in}$ for types $i = s, u$ and locations $n = 1, ..., N$ and local housing costs $\bar{p}_n$ for locations $n = 1, ...N$, generated by an equilibrium with parameters $(\epsilon, \sigma, \Omega, \theta, \rho, \tau, \{\gamma_n\})$ and some fundamentals. Consider another equilibrium generated by a different set of fundamentals, and denote by $\hat{x}$ the ratio of any variable in this new equilibrium relative*

[22]See Appendix A for a formal statement and proof.

to the original one. Then the new equilibrium solves

$$\hat{y}_n = \hat{Z}\hat{z}_n \left(\bar{W}_n \left(\hat{A}\hat{a}_n \hat{L}_{sn} \right)^{\frac{\rho-1}{\rho}} + (1-\bar{W}_n)\hat{L}_{un}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \tag{35}$$

$$\hat{w}_{sn} = \left(\hat{A}\hat{a}_n \hat{Z}\hat{z}_n\right)^{\frac{\rho-1}{\rho}} \left(\frac{\hat{y}_n}{\hat{L}_{sn}}\right)^{\frac{1}{\rho}}, \quad \hat{w}_{un} = \left(\hat{Z}\hat{z}_n\right)^{\frac{\rho-1}{\rho}} \left(\frac{\hat{y}_n}{\hat{L}_{un}}\right)^{\frac{1}{\rho}} \tag{36}$$

$$\hat{e}_{in} = \hat{T}\hat{w}_{in}^{1-\tau} \tag{37}$$

$$\hat{T} = \left(\frac{\sum_i \sum_n \bar{L}_{in}\bar{w}_{in}\hat{L}_{in}\hat{w}_{in}}{\sum_i \sum_n \bar{L}_{in}\bar{w}_{in}^{1-\tau}\hat{L}_{in}\hat{w}_{in}^{1-\tau}} \right) \tag{38}$$

$$\hat{\eta}_{in} = \hat{e}_{in}^{\epsilon}\hat{p}_n^{1-\sigma} \left(\frac{1-\hat{\eta}_{in}\bar{\eta}_{in}}{1-\bar{\eta}_{in}} \right)^{1+\frac{\epsilon}{1-\sigma}} \tag{39}$$

$$\hat{p}_n = \hat{\Pi}_n \left(\bar{H}_n \hat{\eta}_{sn}\hat{e}_{sn}\hat{L}_{sn} + (1-\bar{H}_n)\hat{\eta}_{un}\hat{e}_{un}\hat{L}_{un} \right)^{\gamma_n} \tag{40}$$

$$\hat{v}_{in} = \hat{e}_{in} \left(\frac{1-\hat{\eta}_{in}\bar{\eta}_{in}}{1-\bar{\eta}_{in}} \right)^{\frac{1}{1-\sigma}} \tag{41}$$

$$\hat{\ell}_{in} = \hat{B}_{in}\hat{v}_{in}^{\frac{1}{\theta}} \left(\sum_{m=1}^{N} \bar{\ell}_{im}\hat{B}_{im}\hat{v}_{im}^{\frac{1}{\theta}} \right)^{-1} \tag{42}$$

$$\hat{L}_{sn} = \hat{\ell}_{sn}\hat{\phi}, \quad \hat{L}_{un} = \hat{\ell}_{un} \left(\frac{1-\bar{\phi}\hat{\phi}}{1-\bar{\phi}} \right) \tag{43}$$

where $\bar{W}_n$ is the share of skilled workers in the wage bill of location n and $\bar{H}_n$ is the share of skilled workers in total housing expenditure in location n.

See Appendix A for a proof. Proposition 5 clarifies the data we need to solve our model for a given set of parameters: the distribution of skilled and unskilled workers across locations, their wages in each location, and their housing expenditure shares in each location. Labor supplies and wages can be taken directly from the Census for each decade 1980-2010. We then solve for the expenditure shares for each location and type implied by our estimated price indices.

Parameters

The key parameters in our model are the NHCES preference parameters ϵ and σ. Following column (5) of Table 1, we set $\epsilon = -0.306$ and $\sigma = 0.522$. We choose the scale parameter Ω in each year such that the model matches the average housing expenditure share from the CEX in each year.[23] The average expenditure share rose from 0.32 in 1980 to 0.41 in 2010.

[23] Our hedonic regressions do not recover the level of price indices in each year, so we normalize prices to have mean one in each year.

Table 3: Calibrated Parameters

Parameter	Value	Role	Source
ϵ	−0.306	Income elasticity	Estimated in Column (5), Table 1 using PSID
σ	0.522	Price elasticity	Estimated in Column (5), Table 1 using PSID
ρ	3.850	Production	Card (2009)
τ	0.174	Taxation	PSID following Heathcote, Storesletten, and Violante (2017)
$\{\gamma_n\}$	0.630[a]	Housing supply	Census following Saiz (2010)
θ	0.285	Migration elasticity	Targeting Hornbeck and Moretti (2022)

[a] Employment-weighted mean

The remaining parameters $(\theta, \rho, \tau, \{\gamma_n\})$ are standard, and we describe their calibration briefly here and in more detail in Appendix D. We set the elasticity of substitution in production $\rho = 3.85$ using evidence from Card (2009), who estimates this parameter at the MSA level using using immigration as an instrument for labor-supply changes.[24] We calibrate the tax progressivity parameter following Heathcote, Storesletten, and Violante (2017) and obtain $\tau = 0.174$. We follow Saiz (2010) and model the housing supply elasticities γ_n as linear functions of geographical and regulatory constraints. We estimate the parameters of this model using Bartik shocks to labor demand as an instrument for housing demand as in Diamond (2016), and obtain values for γ_n very similar to those in Saiz (2010).[25] We set the (inverse) migration elasticity θ to 0.285 to match the elasticity of local employment to nominal wages reported by Hornbeck and Moretti (2022), who instrument for wages using shocks to manufacturing TFP. Table 3 summarizes our calibrated parameters.

4 Aggregate Skill Bias and Spatial Sorting

Aggregate skill bias rose sharply between 1980 and 2010. How did this change spatial sorting by skill? Subsection 4.1 begins by discussing trends in the data in aggregate skill bias and spatial sorting by skill. Subsection 4.2 then uses our calibrated model to quantitatively explore the relationship between the two. We find that rising skill bias caused just over one quarter of the observed increase in sorting. Subsection 4.3 unpacks the mechanisms behind this number, and finally subsection 4.4 discusses extensions and robustness checks.

4.1 Trends in Skill Bias and Sorting

To identify trends in aggregate skill bias A_t we rely on the aggregate skill premium, which we define as

$$\omega_t = \sum \lambda_n \log \left(\frac{w_{snt}}{w_{unt}} \right) \tag{44}$$

[24] See Table 5, column (7), in Card (2009) for the negative inverse elasticity of $-\frac{1}{\rho} = -0.26$.

[25] Implicitly, our model has only one sector while Bartik shocks leverage variation in sectoral employment shares. In Appendix D, we show that, with some parameter restrictions, our one sector model is isomorphic to a multisector model.

with λ_n the employment share of location n in 1980. Column (1) of Table 4 shows the evolution of this measure between 1980 and 2010. The skill premium has grown rapidly, from 0.362 in 1980 to 0.591 in 2010.[26]

The CES production function (29) suggests a simple framework for understanding this trend. In particular, by taking the ratio of skilled and unskilled wages in each location and aggregating over locations, we obtain

$$\omega_t = \left(\frac{\rho - 1}{\rho}\right) \log A_t - \left(\frac{1}{\rho}\right) \log \left(\frac{\phi_t}{1 - \phi_t}\right) - \left(\frac{1}{\rho}\right) \sum \lambda_n \log \left(\frac{\ell_{snt}}{\ell_{unt}}\right). \tag{45}$$

Equation (45) shows that three forces drive changes in the skill premium. First, increases in aggregate skill bias A_t mechanically raise the skill premium. Second, the CES structure of production implies that increases in the share of skilled workers ϕ_t push the skill premium down. Column (2) of Table 4 shows that this share has risen substantially over time, from 23% in 1980 to 36% in 2010. The third term captures changes in the skill premium driven by reallocation across space; this term turns out to be quantitatively small, so our focus in what follows will be on the the aggregate terms A_t and ϕ_t. Given observations on employment shares across space and our calibrated value for ρ, equation (45) allows us to back out A_t. Column (3) of Table 4 reports the implied A_t between 1980 and 2010. In line with the literature documenting skill-biased technical change (Acemoglu and Autor 2011) we find that A_t has risen sharply and in 2010 was 73% above its 1980 level.

As in Section 1 we measure sorting using the variance of the log skill ratio across cities, but in taking this measure to the data we now weight by 1980 employment

$$\mathcal{M}_t = \frac{1}{2} \sum \left(\lambda_n \left(\mu_{nt} - \bar{\mu}_t\right)^2\right) \tag{46}$$

where μ_{nt} is the log skill ratio in location n in year t. Column (4) of Table 4 reports $\mathcal{M}_t$ relative to its value in 1980. Consistent with a literature documenting divergence across space (Moretti 2012; Diamond 2016), $\mathcal{M}_t$ has risen by about one third since 1980, indicating that skilled and unskilled workers have been making increasingly different choices about where to live.[27]

4.2 Main results

What were the consequences of the increase in A_t reported in Table 4 for spatial sorting by skill? The theory introduced in Section 1 suggests that it caused sorting to intensify but is silent on its quantitative importance. To answer this question, we use the calibrated model to perform the following experiment. For each census year 1980-2010, we solve for the counterfactual allocations

[26] One can imagine alternative measures of the skill premium. Ours abstracts from differences in the distribution of skilled and unskilled workers across space by weighting by 1980 employment shares. The log difference in average wages is slightly higher and has grown slightly faster: it was 0.384 in 1980 and 0.651 in 2010, an increase of 0.267 log points. This is almost identical to the figure reported in Acemoglu and Autor (2011), who additionally control for changes in the demographic composition of each skill group.

[27] $\mathcal{M}_t$ is of course not the only possible measure of spatial sorting by skill. In Appendix E.3 we show that other measures have followed a similar trend.

Table 4: Aggregate Trends

	(1) ω_t Skill premium	(2) ϕ_t Skill share	(3) $\log A_t$ Skill bias	(4) $\mathcal{M}_t$ Sorting
1980	0.362	0.225	0.042	1.000
1990	0.483	0.265	0.288	1.203
2000	0.550	0.302	0.430	1.313
2010	0.591	0.357	0.588	1.328
Change, 1980-2010	0.229	0.132	0.546	0.328

Note: Column (4) reports values of $\mathcal{M}_t$ relative to its value in 1980. The final row of each column reports the difference between each variable in 1980 and 2010.

that would have been observed had A_t remained at its 1980 level and all other fundamentals, both aggregate and local, had evolved as they did in the data. Formally we solve the hat algebra system described in Proposition 5 for each year t, setting

$$\hat{A}_t = \frac{A_{1980}}{A_t}$$

where the A_t are the values reported in Column (4) of Table 4.

Column (1) of Table 5 reports our main results. The first row shows how spatial sorting changes in the counterfactual between 1980 and 2010, measuring how much sorting would have risen *without* the observed increase in A_t. Our results imply that sorting would have risen by roughly 24%, relative to just under 33% in the data. To give a sense of how the income distribution changes in the counterfactual, the second row reports how the skill premium evolves. With A_t constant but ϕ_t rising as it did in the data, the skill premium falls by roughly 17%. The final row puts the effect of rising skill bias in context by reporting the percentage difference between the observed increase in sorting and the counterfactual increase, capturing the share of the observed increase accounted for by this force. We find that this force explains just over one quarter of the observed increase in $\mathcal{M}_t$. The remaining columns of Table 5 study related counterfactuals.

Column (2) in investigates an alternative way of performing the same experiment. Here we hold all fundamentals other than A_t constant, and feed the model the increase in A_t reported in Column 4 of Table 4. Now the *only* change in fundamentals is the observed increase in A_t. Sorting rises by about 10%. As in column (1) we report the change in the aggregate skill premium, which naturally is larger than the observed rise because it is not offset by the increase in ϕ_t reported in column (2) of Table 4. The final row reports the share of the observed increase in $\mathcal{M}_t$ explained by the counterfactual. We find that this counterfactual explains just over one quarter of the observed increase in $\mathcal{M}_t$, similar to the result in Column 1.[28]

[28]The shocks in Columns (1) and (2) of Table 5 are the same, but they hit different economies with different sets of fundamentals, and thus different distributions of skilled and unskilled workers across space, and so in principle

The decomposition (45) emphasizes the role of two aggregate forces: changes in aggregate skill bias A_t and changes in the skill share ϕ_t. In column (3) we ask how sorting would have evolved if both of these forces were shut down. Formally, in addition to the $\hat{A}_t$ sequence we considered in Column (1), we also set

$$\hat{\phi}_t = \frac{\phi_{1980}}{\phi_t}$$

so that the skill share remains constant at 0.225. Sorting rises by roughly 26%. Consistent with (45), the skill premium stays approximately constant at its 1980 level. The final row shows that this counterfactual explains about one fifth of the observed increase in spatial sorting. The fact that column (3) implies a smaller share explained than Column (1) is natural. When we shut down changes in both A_t and ϕ_t, the skill premium stays constant rather than falling. Our theory emphasizes that spatial sorting is, at least in part, driven by differences in incomes across skill groups, so when these incomes diverge by less sorting also rises by less.

Our main result in Column (1) changes only the aggregate skill bias term A_t, but this implies changes in both the level of income as well as how it is distributed across skilled and unskilled workers. Indeed, average wages in the counterfactual in 2010 are 22% lower than in the data. Although the theory developed in Section 1 emphasizes the importance of inequality as a driver of spatial sorting, with nonhomothetic preferences the level of income may also matter. In column (4) we present one way of exploring this possibility. In addition to the $\hat{A}_t$ we considered in column (1), we feed the model a sequence of shocks to aggregate productivity $\hat{Z}_t$ such that average wages evolve as they did in the data. The results in column (4) are very similar to those in column (1), suggesting that changes in the level of income are not the main driver of our counterfactual results. Instead, they mainly reflect changes in inequality between skilled and unskilled workers.

Above we have focused on the evolution of relative quantities (i.e the skill ratio) across space. A literature (Baum-Snow and Pavan 2012; 2013; Baum-Snow, Freedman, and Pavan 2018) has also

Table 5: Counterfactual Results

	(1)	(2)	(3)	(4)
$\Delta\mathcal{M}_t$, 1980-2010	0.237	0.095	0.255	0.228
$\Delta\omega_t$, 1980-2010	-0.169	0.398	-0.001	-0.170
% of observed $\Delta\mathcal{M}_t$ accounted for	27.4%	29.2%	21.6%	30.1%

Note: Column (1) reports results from solving the model when A_t remains constant at its 1980 level and all other fundamentals evolve as they did in the data. Column (2) holds all fundamentals but A_t constant at their 1980 level and allows A_t to evolve as in Column (4) of Table 4. Column (3) repeats the exercise from (1) but additional holds ϕ_t, the skill share, constant at its 1980 level. Column (4) repeats the exercise from (1) but varies aggregate productivity Z_t so that GDP per capita evolves as in the data. In Columns (1), (3), and (4), the final row reports the difference between counterfactual $\Delta\mathcal{M}_t$ and its value in the data as percentage of its value in the data. In Column (2) the final row reports the counterfactual $\Delta\mathcal{M}_t$ as a percentage of its value in the data.

might have very different effects. The fact that the results are so similar suggests the mechanism we focus on is not too sensitive to the changes in fundamentals that have occurred 1980-2010.

studied trends in the relative prices (i.e the skill premium) across space. Notably, these two objects have become increasingly tightly linked over time. Appendix Table E.1 documents that in 1980 the correlation of (log) skill ratios and (log) skill premia across locations was 0.208, whereas by 2010 this correlation was 0.618. This was mainly the result of faster growth in the skill premium in initially skill intensive locations; the correlation of 1980 skill ratios and 2010 skill premia is 0.559, only slightly lower than the contemporaneous correlation.

How did the increase in aggregate skill bias contribute to this trend? Appendix Table E.1 reports the correlations generated by the counterfactual experiment corresponding to Column (1) of Table 5. We can see that the model explains essentially none of the rise in this correlations since 1980. Moreover, if we consider the correlation of 1980 skill ratios and 2010 skill premia, the model explains a slightly negative share (-5.35%) of the trend in the data. The reason is simple. Increases in A_t push skilled workers towards initially skill intensive locations. From (31), the skill premium in these locations then falls. We conclude that while our mechanism is informative about the evolution of relative quantities across space, it is less useful in understanding relative prices. Instead, the model accounts for these with exogenous shocks to the local skill bias terms a_n. These might represent, for example, falling communication costs (Eckert 2019) or the introduction of new skill-biased technologies (Rubinton 2022).

4.3 Understanding the magnitudes

Our main result, from column (1) of Table 5, is that the increase in aggregate skill bias accounts for roughly one quarter of the observed increase in spatial sorting by skill. What are the forces behind this number? Proposition 3 from Section 1 offers a simple decomposition: the strength of the relationship between A_t and $\mathcal{M}_t$ depends on (i) preferences, (ii) congestion forces, and (iii) the spatial distribution of skill ratios. In our quantitative model this proposition no longer holds exactly, but still provides a useful framework around which to organize our results.

Preferences

Panels (a) and (b) of Figure 3 show how our counterfactual results change as we vary ϵ and σ, the two key parameters of our NHCES preferences.[29] The y-axis plots the effect of the increase in A_t on $\mathcal{M}_t$, expressed as a percentage of the observed increase in sorting as in the final row of Table 5. The vertical lines in each plot show our baseline estimates of $\epsilon = -0.306$ and $\sigma = 0.522$. We can see that ϵ plays a key role, whereas σ is less important. When $\epsilon = 0$ and preferences are homothetic, in line with Proposition 4 changes in A_t have no effect on $\mathcal{M}_t$. As ϵ becomes more negative, the strength of the relationship between A_t and $\mathcal{M}_t$ rises, with the share explained reaching 43% at $\epsilon = -0.478$.

[29] Our plots do not consider the entire range of possible ϵ and σ, because we restrict them to satisfy $\epsilon \geq \sigma - 1$. For each choice of ϵ and σ we recalibrate the shifter Ω so that the average housing expenditure share matches CEX data, just as in our main counterfactual exercise.

We can also consider what our results would look like under two extreme specifications common in the spatial literature: Cobb-Douglas, corresponding to $\epsilon = 0$ and $\sigma = 1$, and a unit housing requirement, corresponding to $\epsilon = -1$ and $\sigma = 0$. For Cobb-Douglas preferences, Proposition 4 immediately implies that the share explained must zero. Nonhomotheticity is needed to generate a relationship between aggregate skill bias and spatial sorting. When we instead assume a unit housing requirement, we obtain a share explained of 35.3%, substantially above our baseline estimate of 27.4%. The implication is that carefully estimating nonhomothetic preferences, rather than assuming a convenient specification, is crucial for gauging the quantitative importance of the mechanism we study.

Congestion forces

Panels (c) and (d) consider the role of two key congestion forces: preference heterogeneity, as captured by θ; and imperfect substitutability between skilled and unskilled labor in production, as captured by ρ.[30]

We begin by varying θ between 0 and 0.83 in panel (c) of Figure 3, with the dashed line showing our benchmark value of 0.285. A value of $\theta = 0$ corresponds to perfect mobility, while $\theta = 0.83$ represents a relatively low level of mobility.[31] Unsurprisingly, a higher θ implies a given change in the A_t has a smaller effect on $\mathcal{M}_t$. When θ is large, preferences are very dispersed and the shock to A_t causes relatively few households to relocate.

The elasticity of substitution ρ plays two roles in our counterfactual. First, from (45), it determines the A_t sequence we infer from observed path of the aggregate skill premium, and thus the magnitude of the shock we feed the model. In general a smaller ρ will correspond to a larger shock. Second, ρ also determines how fast the relative wage of skilled workers falls in a location as its skill ratio rises, and therefore naturally tends to dampen increases in sorting. The solid line in panel (d) shows how the results of our counterfactual vary as ρ^{-1} moves between 0, corresponding to perfect substitutability in production, and $\rho^{-1} = 1/1.60$, the value for this parameter typically found in time series studies (Katz and Murphy 1992).[32] Overall a higher ρ corresponds to a larger change in spatial sorting, but the relationship is relatively shallow and nonmonotonic because of the two offsetting forces highlighted above.

[30]Inelastic housing supply also acts as a congestion force in our model, but actually tends to amplify the effects of increase in A_t. This is because as skilled households move towards expensive cities, they bid up housing costs and force out unskilled households, further raising the skill ratio in those locations. This effect turns out to be quantitatively modest: even when housing supply is perfectly elastic, the share explained by the counterfactual in column (1) of Table 5 only falls by four percentage points.

[31]In particular, 0.83 is the largest value (corresponding to the lowest degree of mobility) among the estimates of this parameter in Suárez Serrato and Zidar (2016).

[32]As we discuss in Appendix D, studies which use spatial variation — more appropriate in the context of our spatial model — typically estimate values for ρ between 3 and 5, and our baseline value of 3.85 lies around the middle of such estimates.

Figure 3: Understanding the magnitudes

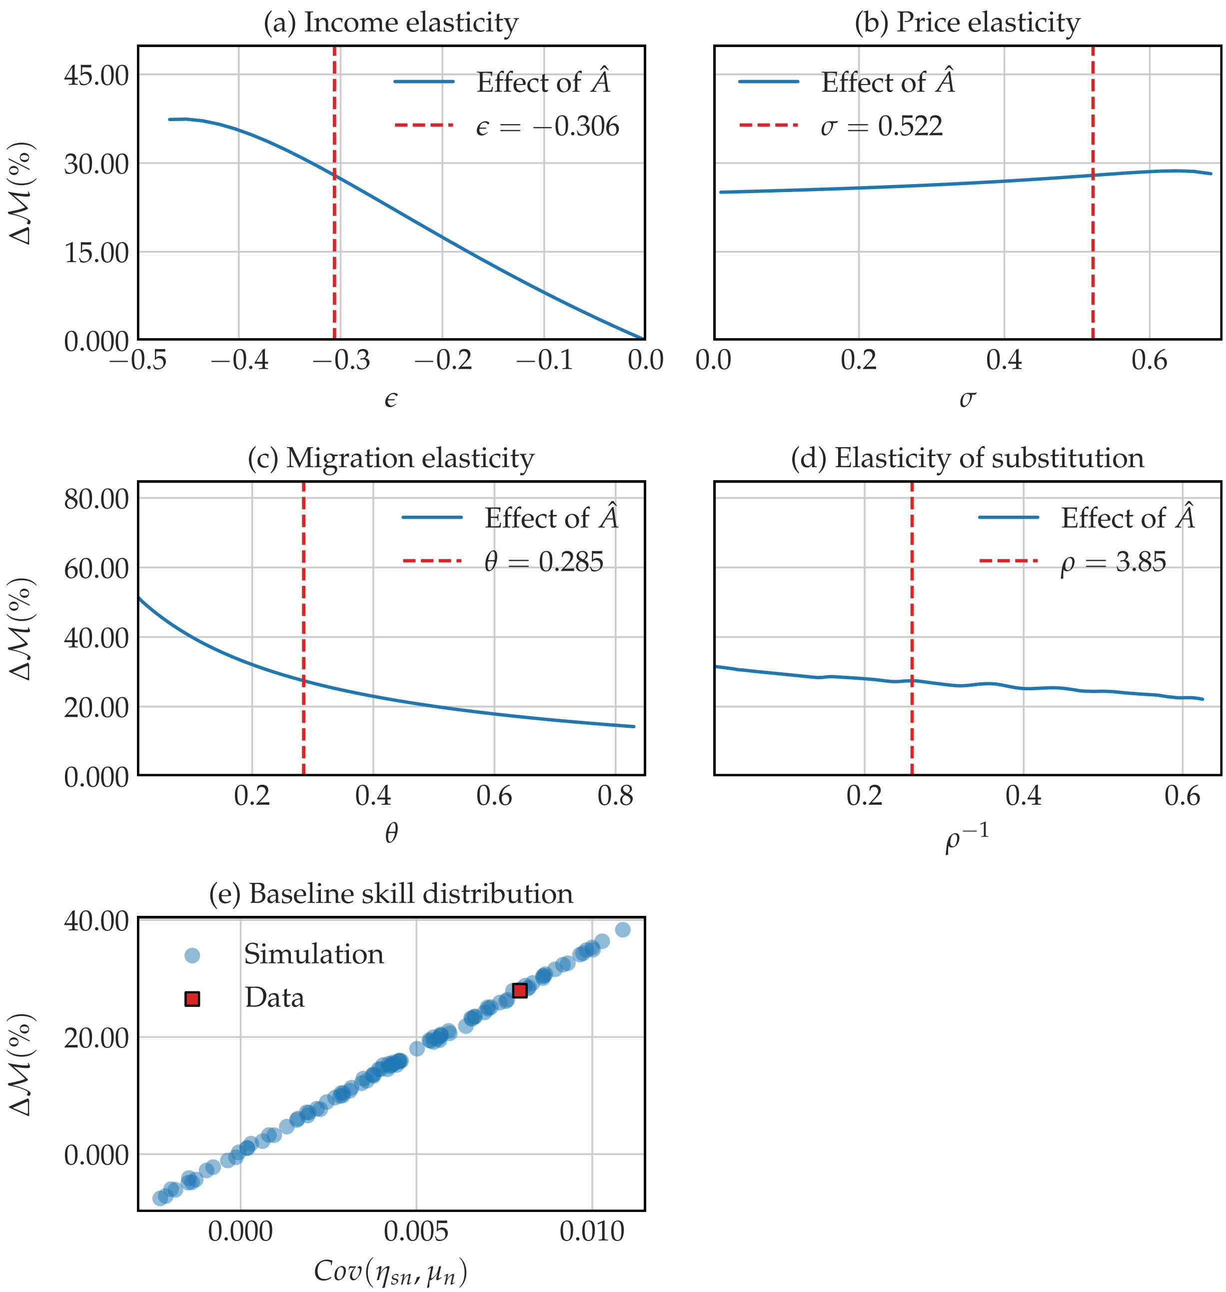

Source: Panel (a) varies ϵ and for each value of this parameter plots the increase in $\mathcal{M}_t$ caused for by A_t, expressed as a percentage of the change in $\mathcal{M}_t$ observed in the data. Panel (b) performs the same exercise but varies σ. In both cases Ω varies so that the average housing expenditure shares matches CEX data. Panel (c) varies θ between 0 — corresponding to perfect mobility — and 0.83 and plots the same outcome as in Panel (a). The solid line in Panel (d) does the same but for ρ^{-1}, with ρ^{-1} corresponding to perfect substitutability in production. Panel (e) show the results of the main counterfactual exercise performed around simulated baseline allocations with varying degrees of covariance between expenditure shares and skill ratios, plotted on the x-axis. The square marker in Panel (e) corresponds to the baseline allocation taken from the data.

Baseline distribution of skill

In general, the relationship between aggregate skill bias and spatial sorting will depend not only on the model's parameters, but also on the baseline allocation around which we perform counterfactuals. In the simple model of Section 1, Proposition 3 tells us that the baseline allocation matters through just one moment: the covariance between log skill ratios μ_n and the expenditure shares of skilled workers η_{sn}. This makes intuitive sense. When housing is a necessity, increases in aggregate skill bias A_t push skilled workers towards relatively expensive locations with high η_{sn}. This pattern results in an increase in spatial sorting if those locations are also relatively skill intensive, i.e. have high μ_n.

Panel (e) of Figure 3 investigates how this covariance shapes our results in the quantitative model. We do this by creating many simulated datasets which match the real data in every respect apart from the log skill ratios μ_n. We draw new log skill ratios from a normal distribution, varying how tightly they covary with expenditure shares η_{sn} to explore how this covariance matters for our results.[33] For each simulated dataset, we then re-run our main counterfactual exercise, corresponding to column (1) of Table 5. The x-axis of panel (e) shows the covariances produced by repeating this procedure 100 times and the y-axis shows the effect of A_t on $\mathcal{M}_t$ in 2010, expressed as a share of the increase in sorting observed between 1980 and 2010.

In line with Proposition 3, there is a positive relationship between the covariance between η_{sn} and μ_n and the sensitivity of sorting to aggregate skill bias. Furthermore, this relationship is very tight. Although the different dots in (e) reflect different random draws of μ_n, and thus different spatial distributions of skill, we can accurately predict the results of the counterfactual using only the covariance between η_{sn} and μ_n. Therefore, just as in the simple model, the role of the baseline allocation in the counterfactual results can (almost) be summarized by a single moment. Finally, the red square in (e) shows the covariance and counterfactual result corresponding to the real data. We can see that is in very much in line with the results from simulated data. Overall, panel (e) shows that this covariance, reflecting the sorting of skilled workers into expensive locations, is the crucial moment in the data which, alongside the parameters discussed above, determines the magnitude of our results. By matching this moment our model is able to accurately quantify the relationship between aggregate skill bias and spatial sorting.

4.4 Extensions and Robustness

In Appendix E we run a number of robustness checks. We consider alternative measure of sorting in E.3 and find similar results to those obtained using our baseline measure. We experiment with an alternative specification of nonhomothetic preferences in E.4 by re-calibrating the model to Price-Independent Generalized Linear (PIGL) preferences. This change has little effect on our main results.

In Appendix E.5, we extend our model to incorporate endogenous amenities following Dia-

[33] Appendix E.2 describes how we generate the simulated datasets in detail.

mond (2016). We first show the neutrality result in Proposition 4 continues to apply. If housing demand is homothetic, changes in aggregate skill bias to have no effect on sorting. We then explore the quantitative importance of this force by estimating a model of labor supply that incorporates endogenous amenities. Our estimation procedure follows Diamond (2016) closely, except that we use NHCES rather than Cobb-Douglas preferences, and we obtain broadly similar results. Plugging these estimates into our model and repeating our main counterfactual, we find that increasing aggregate skill bias accounts for 53.5% of the observed increase in spatial sorting, up from 27.4% in our baseline model. Endogenous amenities amplify the effects of an increase in A_t because as a location becomes more skill intensive, its amenities become relatively more attractive to skilled workers, encouraging further in-migration by skilled workers. We conclude that endogenous amenities are potentially quantitatively important in amplifying the effects of changes in aggregate skill bias on sorting, but we also emphasize that they do not create an independent link between these two objects. Nonhomothetic preferences remain central to our results.

5 Conclusion

Housing is a necessity, implying housing expenditure shares fall with income. Skilled households have high incomes, low housing shares, and are insensitive to high house prices. The opposite is true for unskilled households. Growing aggregate skill bias since 1980 has amplified the cost-of-living wedge, causing skilled households to move toward expensive cities and unskilled households to move toward cheap ones. It has thus increased spatial sorting by skill. Our quantitative model finds that the increase caused by rising skill bias amounts to just over one quarter of the increase in sorting observed since 1980.

We have made these points in a simple environment, deliberately abstracting from spillovers in production or consumption in order to focus on the link between skill bias, nonhomothetic housing demand, and spatial sorting. Incorporating such spillovers into our model would create interesting feedbacks from the spatial distribution of skill to the aggregate income distribution, as well as raising the question of how a social planner might optimally respond to the intensification of spatial sorting that we have studied. Such extensions are an exciting avenue for future research.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, tasks and technologies: Implications for employment and earnings". In *Handbook of Labor Economics*, 4:1043–1171.

Aguiar, Mark, and Mark Bils. 2015. "Has consumption inequality mirrored income inequality?" *American Economic Review* 105 (9): 2725–56.

Albouy, David. 2016. "What are cities worth? Land rents, local productivity, and the total value of amenities". *Review of Economics and Statistics* 98 (3): 477–487.

Albouy, David, Gabriel Ehrlich, and Yingyi Liu. 2016. "Housing demand, cost-of-living inequality, and the affordability crisis". Working Paper.

Andreski, Patricia, Geng Li, Mehmet Zahid Samancioglu, and Robert Schoeni. 2014. "Estimates of annual consumption expenditures and its major components in the PSID in comparison to the CE". *American Economic Review* 104 (5): 132–35.

Attanasio, Orazio P, Renata Bottazzi, Hamish W Low, Lars Nesheim, and Matthew Wakefield. 2012. "Modelling the demand for housing over the life cycle". *Review of Economic Dynamics* 15 (1): 1–18.

Attanasio, Orazio P., and Luigi Pistaferri. 2016. "Consumption inequality". *Journal of Economic Perspectives* 30 (2): 3–28.

Baum-Snow, Nathaniel, Matthew Freedman, and Ronni Pavan. 2018. "Why has urban inequality increased?" *American Economic Journal: Applied Economics* 10 (4): 1–42.

Baum-Snow, Nathaniel, and Ronni Pavan. 2013. "Inequality and city size". *Review of Economics and Statistics* 95 (5): 1535–1548.

— . 2012. "Understanding the city size wage gap". *The Review of economic studies* 79 (1): 88–127.

Beaudry, Paul, Mark Doms, and Ethan Lewis. 2010. "Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas". *Journal of Political Economy* 118 (5): 988–1036.

Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud. 2014. "Productive cities: Sorting, selection, and agglomeration". *Journal of Political Economy* 122 (3): 507–553.

Bilal, Adrien, and Esteban Rossi-Hansberg. 2021. "Location as an Asset". *Econometrica* 89 (5): 2459–2495.

Black, Dan, Natalia Kolesnikova, and Lowell Taylor. 2009. "Earnings functions when wages and prices vary by location". *Journal of Labor Economics* 27 (1): 21–47.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption inequality and family labor supply". *American Economic Review* 106 (2): 387–435.

Boppart, Timo. 2014. "Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences". *Econometrica* 82 (6): 2167–2196.

Bound, John, Jeffrey Groen, Gabor Kezdi, and Sarah Turner. 2004. "Trade in university training: cross-state variation in the production and stock of college-educated labor". *Journal of Econometrics* 121 (1-2): 143–173.

Bureau of Economic Analysis. 2020. *Metropolitan Area Regional Price Parities*. Tech. rep. https://apps.bea.gov/iTable/iTable.cfm?reqid=70&step=30&isuri=1&major_area=5&area=xx&year=-1&tableid=104&category=8101&area_type=0&year_end=-1&classification=non-industry&state=5&statistic=1&yearbegin=-1&unit_of_measure=levels, last accessed 06/01/21. Washington, DC.

Bureau of Labor Statistics. 2020a. *Consumer Expenditure Survey Public Use Microdata*. Tech. rep. bls.gov/cex/pumd.htm, last accessed 12/01/20. Washington, DC.

— . 2020b. *Consumer Price Index for All Urban Consumers: All Items Less Shelter in U.S. City Average*. Tech. rep. Retrieved from FRED, Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/CUSR0000SA0L2, last accessed 06/01/21.

Card, David. 2009. "Immigration and inequality". *American Economic Review* 99 (2): 1–21.

Colas, Mark, and Kevin Hutchinson. 2021. "Heterogeneous workers and federal income taxes in a spatial equilibrium". *American Economic Journal: Economic Policy* 13 (2): 100–134.

Comin, Diego A, Danial Lashkari, and Martí Mestieri. 2021. "Structural change with long-run income and price effects". *Econometrica* 89 (1): 311–374.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst. 2023. "Income Growth and the Distributional Effects of Urban Spatial Sorting". *The Review of Economic Studies* (). ISSN: 0034-6527. doi:10.1093/restud/rdad048.

Davidoff, Thomas. 2016. "Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors". *Critical Finance Review* 6.

Davis, Donald R., Eric Mengus, and Tomasz Michalski. 2021. "Labor Market Polarization and The Great Urban Divergence". Working Paper.

Davis, Morris A, and François Ortalo-Magné. 2011. "Household expenditures, wages, rents". *Review of Economic Dynamics* 14 (2): 248–261.

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

Diamond, Rebecca, and Cecile Gaubert. 2022. "Spatial sorting and inequality". *Annual Review of Economics* 14:795–819.

Eckert, Fabian. 2019. "Growing apart: Tradable services and the fragmentation of the US economy". Working Paper.

Eckert, Fabian, Sharat Ganapati, and Conor Walsh. 2022. "Urban-biased growth: a macroeconomic analysis". Working Paper.

Eckert, Fabian, and Michael Peters. 2023. "Spatial structural change". Working Paper.

Eeckhout, Jan, Roberto Pinheiro, and Kurt Schmidheiny. 2014. "Spatial sorting". *Journal of Political Economy* 122 (3): 554–620.

Fajgelbaum, Pablo D, and Cecile Gaubert. 2020. "Optimal spatial policies, geography, and sorting". *The Quarterly Journal of Economics* 135 (2): 959–1036.

Fogli, Alessandra, and Veronica Guerrieri. 2019. "The end of the American Dream? Inequality and segregation in US cities". Working Paper.

Ganong, Peter, and Daniel Shoag. 2017. "Why has regional income convergence in the US declined?" *Journal of Urban Economics* 102:76–90.

Giannone, Elisa. 2022. "Skilled-biased technical change and regional convergence". Working Paper.

Greaney, Brian. 2023. "The distributional effects of uneven regional growth". Working Paper.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai. 2013. "Superstar cities". *American Economic Journal: Economic Policy* 5 (4): 167–99.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008. "A new measure of the local regulatory environment for housing markets: the Wharton Residential Land Use Regulatory Index". *Urban Studies* 45 (3): 693–729.

Handbury, Jessie. 2021. "Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities". *Econometrica* 89 (6): 2679–2715.

Hansen, Julia L, John P Formby, W James Smith, et al. 1996. "The income elasticity of demand for housing: evidence from concentration curves". *Journal of Urban Economics* 39 (2): 173–192.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2017. "Optimal tax progressivity: an analytical framework". *The Quarterly Journal of Economics* 132 (4): 1693–1754.

Hornbeck, Richard, and Enrico Moretti. 2022. "Estimating who benefits from productivity growth: local and distant effects of city productivity growth on wages, rents, and inequality". *The Review of Economics and Statistics*: 1–49.

Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing constraints and spatial misallocation". *American Economic Journal: Macroeconomics* 11 (2): 1–39.

Ioannides, Yannis, and Rachel Ngai. 2023. "Housing and Inequality". Working Paper.

Ioannides, Yannis M, Jeffrey E Zabel, et al. 2008. "Interactions, neighborhood selection and housing demand". *Journal of Urban Economics* 63 (1): 229.

Katz, Lawrence F, and Kevin M Murphy. 1992. "Changes in relative wages, 1963–1987: supply and demand factors". *The Quarterly Journal of Economics* 107 (1): 35–78.

Kennan, John, and James R Walker. 2011. "The effect of expected income on individual migration decisions". *Econometrica* 79 (1): 211–251.

Larsen, Erling Røed. 2014. "The Engel curve of owner-occupied housing consumption". *Journal of Applied Economics* 17 (2): 325–352.

Lewbel, Arthur. 1996. "Demand estimation with expenditure measurement errors on the left and right hand side". *The Review of Economics and Statistics*: 718–725.

Lewbel, Arthur, and Krishna Pendakur. 2009. "Tricks with Hicks: The EASI demand system". *American Economic Review* 99 (3): 827–63.

Moretti, Enrico. 2012. *The New Geography of Jobs*. Boston: Houghton Mifflin Harcourt.

Notowidigdo, Matthew J. 2020. "The incidence of local labor demand shocks". *Journal of Labor Economics* 38 (3): 687–725.

Poterba, James, and Todd Sinai. 2008. "Tax expenditures for owner-occupied housing: Deductions for property taxes and mortgage interest and the exclusion of imputed rental income". *American Economic Review* 98 (2): 84–89.

Redding, Stephen J. 2016. "Goods trade, factor mobility and welfare". *Journal of International Economics* 101:148–167.

Redding, Stephen J, and Esteban Rossi-Hansberg. 2017. "Quantitative spatial economics". *Annual Review of Economics* 9:21–58.

Roback, Jennifer. 1982. "Wages, rents, and the quality of life". *Journal of political Economy* 90 (6): 1257–1278.

Rosen, Sherwin. 1979. "Wage-based indexes of urban quality of life". *Current issues in urban economics*: 74–104.

Rosenthal, Stuart S. 2014. "Are private markets and filtering a viable source of low-income housing? Estimates from a "repeat income" model". *American Economic Review* 104 (2): 687–706.

Rubinton, Hannah. 2022. "The geography of business dynamism and skill biased technical change". *FRB St. Louis Working Paper*, no. 020.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. 2020. *IPUMS USA: Version 10.0*. usa.ipums.org/usa, last accessed 12/01/20, Minneapolis, MN.

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

Schmidheiny, Kurt. 2006. "Income segregation from local income taxation when households differ in both preferences and incomes". *Regional Science and Urban Economics* 36 (2): 270–299.

Straub, Ludwig. 2019. "Consumption, savings, and the distribution of permanent income". Working Paper.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. "Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms". *American Economic Review* 106 (9): 2582–2624.

University of Michigan Institute for Social Research. 2021. *Panel Study of Income Dynamics*. Tech. rep. psidonline.isr.umich.edu, last accessed 05/01/21. Ann Arbor, MI.

Zabel, Jeffrey E. 2004. "The demand for housing services". *Journal of Housing Economics* 13 (1): 16–35.

Zillow. 2017. *Zillow Rent Index (ZRI), Metro/County, Per Sq. Ft., All Homes*. Tech. rep. data.world/zillow-data, last accessed 08/14/20. Seattle, WA.

A Theory

A.1 Proof of Proposition 1

First we establish that $A > 1$ implies $v_s > v_u$. From the relative labor demand curve

$$\log(w_{sn}) - \log(w_{un}) = \log A + (1-\alpha)(\log \ell_{un} - \log \ell_{sn}).$$

Since ℓ_{sn} and ℓ_{un} both sum to one, there must be some location n for which $\ell_{sn} \leq \ell_{un}$. Then the labor demand curve in this location implies $w_{sn} > w_{un}$, because $A > 1$. Then from (1), $v_s = v(w_{sn}, p_n) > v(w_{un}, p_n) = v_u$.

Given $v_s > v_u$, the result obtains trivially. If housing is a necessity then the housing expenditure share $\eta(p_n, v)$ is decreasing in v, and so $\eta(p_n, v_s) < \eta(p_n, v_u)$ for every p_n. Then (9) implies the derivative of μ_n with respect to p_n is always strictly positive and the result follows. By contrast, if housing demand is homothetic then $\eta(p_n, v_s) = \eta(p_n, v_u)$ and μ_n is a constant function of p_n. Then μ_n does not vary across locations and its variance, $\mathcal{M}$, is zero.

A.2 Proof of Proposition 2

First we establish that $\Delta \log A > 0$ implies $\Delta v_s > 0$. For every location n, differencing the log of (2) yields

$$\Delta \log w_{sn} = \Delta \log A - (1-\alpha)\Delta \log \ell_{sn}.$$

Since $\Delta \ell_{sn}$ sums to zero, there must be some location for which $\Delta \ell_{sn} \leq 0$. Then $\Delta \log \ell_{sn} \leq 0$ and from the labor demand curve $\Delta w_{sn} > 0$. Since skilled wages in n rise and prices by assumption do not change, we must have $\Delta v(w_{sn}, p_n) > 0$. Then by (1), $\Delta v_s > 0$. At the same time, unskilled wages do not change, so $\Delta v_u = 0$.

Now suppose housing is a necessity. Then

$$\Delta \eta(v_s, p_n) < 0, \quad \Delta \eta(v_u, p_n) = 0$$

which from (9) implies that the derivative of μ_n with respect to p_n grows at every point. It immediately follows that $\Delta \mu_n$ is a strictly increasing function of p_n. To show that $\Delta \mathcal{M} > 0$, use the definition (4),

$$\Delta \mathcal{M} = \frac{1}{2}\text{Var}(\mu_n + \Delta \mu_n) - \mathcal{M} = \frac{1}{2}\text{Var}(\mu_n) + \frac{1}{2}\text{Var}(\Delta \mu_n) + \text{Cov}(\Delta \mu_n, \mu_n) - \mathcal{M} > \text{Cov}(\Delta \mu, \mu)$$

where the last inequality follows because $\text{Var}(\Delta \mu_n)$ is always positive. Now, μ_n and $\Delta \mu_n$ are both strictly increasing functions of p_n. Therefore

$$\Delta \mathcal{M} > \text{Cov}(\Delta \mu_n, \mu_n) > 0,$$

as claimed.

Finally, suppose housing demand is homothetic. From (9), μ_n is a constant function of p_n for any value of A. It immediately follows that $\Delta\mu_n = 0$ and $\Delta\mathcal{M} = 0$.

A.3 Proof of Proposition 3

We start by deriving an expression for the elasticity of skilled utility with respect to A, which we denote by

$$\chi_s \equiv \frac{d\log v_s}{d\log A}.$$

First, combining the spatial indifference condition (1) with the definition of the NHCES price index (16), skilled utility in any location n can be written

$$v_s = w_s\left(1 + p_n^{1-\sigma}v_s^{\epsilon}\right)^{\frac{1}{\sigma-1}}.$$

Then use (2) to substitute out wages and rearrange

$$\ell_{sn} = \left(v_s^{-1}Az_n\left(1 + p_n^{1-\sigma}v_s^{\epsilon}\right)^{\frac{1}{\sigma-1}}\right)^{\frac{1}{1-\alpha}}.$$

Now exploit the fact that the ℓ_{sn} must sum to one to get an expression that implicitly defines v_s

$$v_s = A\left(\sum_n\left(z_n\left(1 + p_n^{1-\sigma}v_s^{\epsilon}\right)^{\frac{1}{\sigma-1}}\right)^{\frac{1}{1-\alpha}}\right)^{1-\alpha}.$$

Taking logs and differentiating with respect to $\log A$ yields

$$\chi_s = \left(1 + \frac{\epsilon}{1-\sigma}\sum_n \ell_{sn}\eta_{sn}\right)^{-1} \equiv \left(1 + \frac{\epsilon}{1-\sigma}\bar{\eta}_s\right)^{-1},$$

as claimed. To see the implications for labor supply, log-differentiate the expression for ℓ_{sn} above to obtain

$$\frac{d\log\ell_{sn}}{d\log A} = \left(\frac{1}{1-\alpha}\right)\left(1 - \chi_s - \left(\frac{\epsilon}{1-\sigma}\right)\chi_s\eta_{sn}\right)$$

and note that the only term here that depends on n is η_{sn}. Now by definition $\mu_n = \log\ell_{sn} - \log\ell_{un}$, and since changes in A do not affect unskilled workers, we obtain

$$\frac{d\mu_n}{d\log A} = \frac{d\log\ell_{sn}}{d\log A}.$$

So, in line with Proposition 2, the change in μ_n caused by an increase in A is an increasing function of η_{sn}, and this p_n, as long as $\epsilon < 0$. To see the implications for sorting $\mathcal{M}$, note that

$$\frac{d\mathcal{M}}{d\log A} = \text{Cov}\left(\frac{d\mu_n}{d\log A}, \mu_n\right).$$

Plugging in the expression for the derivative of μ_n above, we obtain

$$\frac{d\mathcal{M}}{d\log A} = \left(\frac{1}{1-\alpha}\right)\left(\frac{\epsilon}{1-\sigma}\right)\chi_s \text{Cov}(\eta_{sn}, \mu_n)$$

where constant terms have dropped out of the covariance. It follows that the expression in Proposition 3 approximates the effect of a small change $\Delta \log A$, to first order.

A.4 Irrelevance of income elasticity normalization

In Section 1, we introduced NHCES preferences as

$$U^{\frac{\sigma-1}{\sigma}} = \Omega^{\frac{1}{\sigma}} h^{\frac{\sigma-1}{\sigma}} U^{\frac{\epsilon}{\sigma}} + c^{\frac{\sigma-1}{\sigma}}. \tag{47}$$

Here, we show that this is equivalent to the more general formulation

$$\tilde{U}^{\frac{\sigma-1}{\sigma}} = \Omega_h^{\frac{1}{\sigma}} h^{\frac{\sigma-1}{\sigma}} \tilde{U}^{\frac{\epsilon_h}{\sigma}} + \Omega_c^{\frac{1}{\sigma}} c^{\frac{\sigma-1}{\sigma}} \tilde{U}^{\frac{\epsilon_c}{\sigma}} \tag{48}$$

where ϵ_h, ϵ_c, Ω_h and Ω_c are parameters.

First, observe that (47) is a special case of (48) with $\epsilon_c = 0$, $\Omega_c = 1$, $\Omega_h = \Omega$, and $\epsilon_h = \epsilon$. Second, let us take ϵ_c and $\Omega_c > 0$ as given. It is straightforward to show that by choosing Ω_h and ϵ_h correctly, we can produce preferences which yield identical housing demand functions as in (47). To see this, divide both sides of (48) by $\Omega_c^{\frac{1}{\sigma}} \tilde{U}^{\frac{\epsilon_c}{\sigma}}$ to obtain

$$\Omega_c^{-\frac{1}{\sigma}} \tilde{U}^{\frac{\sigma-1}{\sigma} - \frac{\epsilon_c}{\sigma}} = \Omega_h^{\frac{1}{\sigma}} \Omega_c^{-\frac{1}{\sigma}} h^{\frac{\sigma-1}{\sigma}} \tilde{U}^{\frac{\epsilon_h - \epsilon_c}{\sigma}} + c^{\frac{\sigma-1}{\sigma}}. \tag{49}$$

Now set

$$\Omega_h = \Omega_c^{\epsilon(1-\sigma)+1}\Omega, \quad \epsilon_h = \epsilon + \epsilon_c\left(1 - \frac{\epsilon}{\sigma - 1}\right).$$

Inserting these expressions into (49), we obtain

$$\left(\Omega_c^{1-\sigma} \tilde{U}^{1-\frac{\epsilon_c}{\sigma-1}}\right)^{\frac{\sigma-1}{\sigma}} = \Omega^{\frac{1}{\sigma}} h^{\frac{\sigma-1}{\sigma}} \left(\Omega_c^{1-\sigma} \tilde{U}^{1-\frac{\epsilon_c}{\sigma-1}}\right)^{\frac{\epsilon}{\sigma}} + c^{\frac{\sigma-1}{\sigma}}. \tag{50}$$

By comparing this with (47), we can see that

$$\tilde{U} = \Omega_c^{\frac{1}{\sigma-1-\epsilon_c}} U^{\frac{\sigma-1}{\sigma-1-\epsilon_c}}. \tag{51}$$

That is, $\tilde{U}$ is a monotonically increasing transformation of U and so represents the same preferences over housing and non-housing consumption.

Finally, it could in principle be the case that (48), when incorporated into the quantitative spatial model developed in Section 3, might lead to different preferences over locations than (47). But for our isolastic model of labor supply, this is not the case. To see this, consider the location choice equation (28) using the preferences defined in (48). We obtain

$$\ell_{in} = \frac{\tilde{v}_{in}^{\theta^{-1}} B_{in}}{\sum_m \tilde{v}_{im}^{\theta^{-1}} B_{im}}, \tag{52}$$

where

$$\tilde{v}_{in} = \Omega_c^{\frac{1}{\sigma-1-\epsilon_c}} v_{in}^{\frac{\sigma-1}{\sigma-1-\epsilon_c}}. \tag{53}$$

Combining these two expressions, we obtain

$$l_{in} = \frac{v_{in}^{\tilde{\theta}^{-1}} B_{in}}{\sum_m v_{im}^{\tilde{\theta}^{-1}} B_{im}}, \tag{54}$$

where

$$\tilde{\theta} = \theta \left(\frac{\sigma - 1 - \epsilon_c}{\sigma - 1} \right). \tag{55}$$

That is, choosing $\epsilon_c > 0$ just proportionally rescales the migration elasticity θ. Following the calibration strategy outlined in Section 3 with $\epsilon_c > 0$, we would just estimate a rescaled version of θ, and all of the model's predictions would be unchanged. Therefore, assuming $\epsilon_c = 0$ and $\Omega_c = 1$ is without loss of generality.

A.5 Proof of Proposition 4

Suppose $\epsilon = 0$, so that preferences are homothetic. Then for each type of household, utility v_{in} can be written

$$v_{in} = e_{in} P_n^{-1} = T w_{in}^{1-\tau} P_n^{-1},$$

where P_n is the price index common to all households in location n. Taking the ratio of v_{sn} to v_{un} in any location n yields

$$\frac{v_{sn}}{v_{un}} = (a_n A)^{\frac{\rho-1}{\rho}(1-\tau)} \left(\frac{\ell_{sn}}{\ell_{un}} \right)^{\frac{1-\tau}{-\rho}}.$$

Now, (28) implies

$$\mu_n = \frac{1}{\theta} \log \left(\frac{v_{sn}}{v_{un}} \right) + \log \left(\frac{B_{sn}}{B_{un}} \right) + \varpi$$

where ϖ is constant across locations. Plugging in the expression for the ratio of indirect utilities and rearranging yields

$$\mu_n = \vartheta_0 + \vartheta_1 a_n + \vartheta_2 \log \left(\frac{B_{sn}}{B_{un}} \right)$$

where the ϑ's are all constant across locations and ϑ_1 and ϑ_2 depend only on parameters. Notice that variation across locations in μ_n is entirely determined by local amenity and productivity shifters, and aggregate skill bias A plays no role. Proposition 4 follows immediately.

A.6 Description of quantitative model inversion

For each year, the quantitative model of Section 3 has $5N$ local fundamentals: amenities B_{sn} and B_{un}; productivities a_n and z_n; and housing supply shifters Π_n. The model also has two aggregate fundamentals A and Z. Below we show that given data on employment shares ℓ_{in}, wages w_{in} and local prices p_n, these fundamentals are uniquely identified, up to a scaling factor in the case of the amenities B_{sn} and B_{un}. We proceed in six steps.

1. Taking the ratio of (30) and (31) in location n yields an expression for Aa_n in terms of relative wages and labor supplies in n. Then the assumption that $\sum \mu_n \log a_n = 0$ separately identifies A and a_n.

2. From the fact that the production function (29) has constant returns to scale, we can identify total output y_n by summing skilled and unskilled wage bills. Then plugging Aa_n and labor supplies into (29) identifies Zz_n. Again the assumption that $\sum \mu_n \log z_n = 0$ separately identifies Z and z_n.

3. Given data on w_{in} and ℓ_{in}, we can construct expenditures e_{in} using (32).

4. Expenditures e_{in} and prices p_n can be used to calculate indirect utilities v_{in}.

5. For each type i, we normalize the $B_{i1} = 1$. Then we identify the amenity in every other location using $B_{in} = \left(\ell_{in}/\ell_{i1}\right)\left(v_{i1}/v_{in}\right)^{\theta^{-1}}$.

6. Finally, e_{in} and p_n can also be used to calculate housing expenditure shares η_{in}. Then (34) allows us to calculate housing demand HD_n. Then data on prices p_n identifies the housing supply shifters Π_n.

A.7 Proof of Proposition 5

Here we describe how to obtain the 'hat' system of Proposition 5, and in doing so prove that the solving this system is the same as solving the original equilibrium system. (36) follows from taking the ratios of (30) and (31), respectively, between the observed and counterfactual equilibria. (37) follows from taking the ratio of (32) between the observed and counterfactual equilibria. (39) follows from taking the ratio of (27) between the observed and counterfactual equilibria. (40) follows from substituting (34) into (33) and taking the ratio between observed and counterfactual equilibria. (41) follows from taking the ratio of (26) between the observed and counterfactual equilibria and using the definition of η in (27). Finally (42) follows from taking the ratio of (28) between the observed and counterfactual equilibria.

B Data

B.1 PSID

The primary consumption microdata come from the Panel Study of Income Dynamics (PSID). The PSID is administered biannually, with about 9,000 households in each wave. It included a consumption module starting in 1999 and added several categories in 2005. The survey now covers about 70% of spending in the national accounts (Blundell, Pistaferri, and Saporta-Eksten 2016). Total expenditure is computed as the sum of all reported consumption categories: rent, food, utilities, telephone and internet, automobile expenses (including car loans, down payments, lease payments, insurance, repairs, gas, and parking), other transportation expenses, education, childcare, healthcare, home repairs, furniture, computers (2017 only), clothing, travel, and recreation. The PSID imputes a small number of observations to handle invalid responses. To match the definition in IPUMS, housing expenditure is equal to rent plus utilities. Homeowners were not asked to estimate the rental value of their home until 2017, so we restrict attention to renters and analyze homeowners with the CEX.

We use the 2005-2017 waves of the PSID and select our sample according to the following criteria. We drop respondents in the top and bottom 1% of the pre-tax income distribution in each year. We then select households in which the head is prime-age (25-55, inclusive) and attached to the labor force (head or spouse reports usually working at least 35 hours per week). The controls included in the regressions are dummies for family size bins, number of earners, age bins, sex of household head, race of household head, and year. Education is defined as years of schooling of the highest-earning household member. We use the PSID sample weights in all regressions.

Using the restricted access county identifiers, we can assign local prices to 92% of households in the PSID sample. The remaining households live in rural counties for which we do not construct rental price indices.

B.2 Rental Price Indices

To include prices in the preference estimates, we compute metropolitan area rental price indices from ACS data following Albouy (2016). We estimate a standard hedonic regression model of the form

$$\log \text{rent}_{int} = \log p_{nt} + X'_{int}\delta_t + v_{int} \tag{56}$$

where i denotes households, n denotes cities, and X is a set of observed dwelling characteristics: number of rooms, number of bedrooms, the interaction of the two, building age, number of units in the building, type of kitchen, type of plumbing, plot size, a dummy for whether the unit is a condo, and a dummy for whether the unit is a mobile home. v_{int} is an idiosyncratic error term. The estimate of p_{nt}, an MSA by year fixed effect, is the rental price index. We run the regression separately for each two-year window starting in 2005 and restrict the sample to renting households in the ACS. We follow the same procedure with the 1980/1990/2000/2010 Census/ACS data to construct rental price indices for the quantitative model.

Regressing log $\hat{p}_{nt}$ on MSA average log rent yields a slope coefficient of 0.79 (population-weighted) and an R^2 of 0.90. In a robustness exercise, we use the Metropolitan Regional Price Parities published by the BEA (Bureau of Economic Analysis 2020). The BEA estimates MSA-level price indices for rents, goods and other services. Regressing our rental index on the BEA rental index yields a slope coefficient of 0.84 and an R^2 of 0.98.

B.3 CEX

We append the 2006-2017 Consumer Expenditure Surveys (CEX) together and annualize at the household level. We define rental expenditure as actual rent paid for renters (`rendwe`) and self reported rental-equivalent (`renteqvx`) for owners. As in PSID, we add utilities `util` to be consistent with the data available in the Census. To solicit rental equivalent, homeowners are asked "If someone were to rent your home today, how much do you think it would rent for monthly, unfurnished and without utilities?" We define total consumption expenditure as equal to total reported expenditure `totexp` *less* retirement and pension savings `retpen`, cash contributions `cashco`, miscellaneous outlays `misc` (which includes mortgage principal), and life and personal insurance `lifins`. For homeowners, we subtract `owndwe` and add `renteqvx`. We apply the exact same sample selection criteria and controls in the CEX as in the PSID (see Section B.1). We use CEX sample weights in all regressions.

In 2006, the CEX added more detailed geographic identifiers in the variable `psu`. The primary sampling unit, i.e. the MSA of residence, is available for a subset of households. The CEX identifies twenty-four large MSAs, which cover about 45% of households in the survey.

B.4 Census

We use the 5% public use samples from the 1980, 1990, and 2000 Censuses. For the final period of data, we use the 2009-2011 American Community Survey, a 3% sample. For convenience we refer to this as the "2010 data." IPUMS attempts to concord geographic units across years, although complete concordance is not possible because of data availability and disclosure rules. We classify MSAs according to the variable `metarea`. We produce a balanced panel using the following rule: if an MSA appears in all four years, then it is kept. If an MSA does not appear in all four years, then we assign all individuals in that MSA across all years to a residual state category. For example, Charlottesville, VA appears in 1980, 2000, and 2010, but not in 1990. Therefore we assign all individuals in Charlottesville in every year to "Virginia." This procedure gives us 219 MSAs (including Washington, D.C.) and 50 residual state categories, for a total of 269 regions. The share of national employment which can be assigned to an MSA, rather than a state residual, is 70% in 1980, 72% in 1990, and 75% in 2000 and 2010.

A worker is considered skilled if she or he has completed at least a four year college degree according to the variable `educ`. By this metric, the national fraction of workers who are skilled is

22.5% in 1980, 26.5% in 1990, 30.2% in 2000, and 35.7% in 2010.

We compute average wages and employment for each region, skill group, and year. Wages are from the IPUMS variable `incwage`. To be included in the wage and employment sample, workers must be between 25 and 55 years old, inclusive; not have any business or farm income; work at least 40 weeks per year and 35 hours per week; and earn at least one-half the federal minimum wage. Wages are adjusted to 2000 real values using BLS' Non-Shelter CPI.

Households within a skill group, location, and year are assumed to have expenditure given by the average post-tax wage income of group, $e_{int} = \overline{y}_{int} \equiv T_t \left(\overline{w}_{int}\right)^{1-\tau}$, where T_t is chosen to balance the budget. This assumes that the elasticity of expenditure to permanent post-tax income is unity. Households save in the data, but savings wash out in the aggregate since we focus on permanent income.

To construct Bartik instruments, we use the industry categories in the Census variable `ind1990`. Harmonizing the industries with our own crosswalk yields 208 industries which are consistently defined over all four periods. We drop individuals who cannot be classified into any industry ($\approx$ 0.3% of workers) or who are in the military ($\approx$ 0.9% of workers).

An alternative way to classify workers' skill would be assign workers with 'some college' partially to both the unskilled and skilled categories, as in Katz and Murphy (1992). That paper estimates that workers with some college contribute a weight of 0.69 labor units to the high school labor force and 0.29 units to the college labor force. We find that these two definitions of the skill ratio are highly correlated across space (0.98), the skill premium rises more (0.27) under the Katz-Murphy measure than under our baseline (0.23), and that spatial sorting rises under this measure, but by less than using our baseline measure (11% vs 32%). However, to measure sorting, we prefer a definition of skill which always assigns each worker to one skill group, rather than partially to both skill groups. Decisions about where to live are made at the level of an individual, so imposing that some individuals carry e.g. 0.69 units of unskilled labor and 0.29 units of skilled labor enforces that these units of skilled and unskilled labor move across space together. This mechanically limits the intensity of spatial sorting: even if 'high school or less' and 'some college' workers perfectly segregate themselves from 'college of more' workers, a measure of sorting which partially assigns the 'some college' workers to both groups will always imply some mixing. Given this constraint, we then choose to group 'some college' workers with 'high school or less' because their wages and employment shares tend to look more similar to one another, relative to college graduates. Figure B.1 presents binned scatter plots of employment shares against city size (left) and log average wages against city size (right), for each of the three groups. The employment share and wages of college graduates increase with city size. By contrast, 'some college' and 'high school' employment shares decline with city size, and their wages are relatively flat with respect to city size.

Figure B.1: Employment and wages by education level, 2010

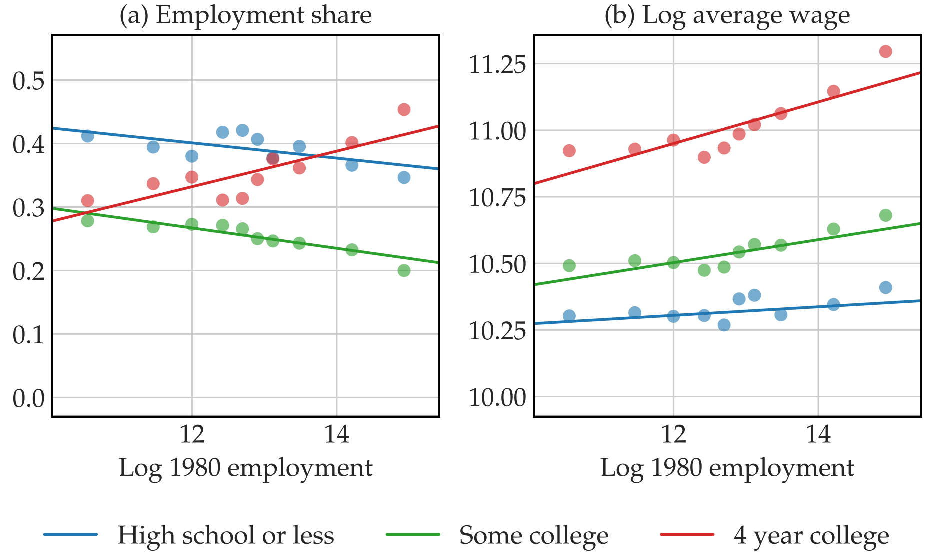

Source: Census. Workers with 1, 2, or 3 years of college are assigned to 'some college' while workers with 4 years of college or a graduate degree are assigned to '4 year college.'

C Estimation

C.1 A Model of Housing Characteristics

In the model we estimate in Section 2, each agent i in location n chooses a scalar h_{in} that represents their housing consumption, facing a single price p_n. This is an incomplete description of housing demand in reality, where houses can differ on a number of dimensions, for example number of bedrooms, age, or proximity to parks and restaurants. Here we introduce an extension to our model which allows agents to consume a bundle of characteristics, and in doing so make an explicit connection between our model of housing demand and the hedonic regression (56) we use to back out local prices p_n.

We assume that a 'house' is a bundle of characteristics denoted by $x \in \mathbb{R}^J$, where J is the number of characteristics. Agent i chooses each element of x, as well as nonhousing consumption c, subject to a budget constraint

$$c + \sum q_{jn} x_{jin} = e_{in} \tag{57}$$

where q_{jn} is the price of characteristic j in location n.

As in Section 2, we assume that agents have NHCES preferences over housing h and nonhousing consumption c, but now housing consumption is an aggregator over the characteristics x. We denote this aggregator by $H(x)$, and the flow of housing consumption the household enjoys is $h = H(x)$. So that we can talk about a single housing cost index p_n in each location, we assume that this aggregator has constant returns to scale. Denoting total housing expenditure by R_{in}, this immediately implies

$$R_{in} \equiv \sum q_{jn} x_{jin} = p_n H(x_{in}). \tag{58}$$

for some price index p_n. Notice that the price index p_n does not vary across households. This is a consequence of constant returns to scale in the aggregator H. We note that there is no tension between this assumption and the assumption that housing demand overall is nonhomothetic; it merely imposes that relative preferences over different characteristics do not vary with overall housing consumption, without restricting how overall housing consumption varies with income. To complete the connection to the NHCES model, notice that the budget constraint becomes

$$c + p_n h = e_{in}$$

just as in (11).

We now consider how this relates to our hedonic regressions. Taking logs of the expression above yields

$$\log R_{in} = \log p_n + \log H(x_{in}). \tag{59}$$

This is close to the hedonic regression (56). To get all the way there, however, we need one more assumption. There are two possibilities. The first is to assume that H is a Cobb-Douglas aggregator with weights β_j. Then denoting $X_{jin} = \log x_{jin}$, we obtain

$$\log R_{in} = \log p_n + \sum \beta_j X_{jin}, \tag{60}$$

replicating (56). The second possibility is to assume that the price of characteristic j relative to characteristic k does not vary across locations. With this assumption, plus the assumption that H has constant returns to scale, one can show that

$$\log R_{in} \simeq \chi + \log p_n + \sum \beta_j X_{jin}, \tag{61}$$

where χ is a constant that does not vary across locations and β_j is the share of housing expenditure devoted to characteristic j, constant across locations by virtue of the assumption that relative prices are constant. Since we never use the level of the prices p_n, the presence of the constant χ is not a problem.

Finally, we comment on the role of two assumptions in our approach. The first we have made explicit: the weights β_j do not vary across locations. To assess the plausibility of this assumption, we re-run our hedonic regression allowing these coefficients to vary with city size. We report the results in Table C.1, and find that the coefficients are quite stable across the city size distribution. Crucially, the price indices recovered from this more flexible specification are highly correlated (0.98) with those from our baseline specification. We conclude that this is not an unreasonable assumption in our context.

The second assumption is implicit: we assume that we observe all the relevant characteristics of a house. For physical characteristics this is not obviously unreasonable, but for others — such as proximity to restaurants — it is harder to justify. It turns out that we can relax this assumption if we are willing to assume that the relative prices of characteristics do not vary across space. To see this, suppose that one of the characteristics, for concreteness indexed k, is not observed. Write:

$$\log R_{in} = \chi + \log p_n + \sum_{j \neq k} \beta_j X_{jin} + \beta_k X_{ikn}. \tag{62}$$

By the assumption that relative prices do not vary across space, we can write

$$p_n q_k x_{ikn} = \beta_k R_{in}$$

or equivalently

$$\log p_n + \log q_k + X_{ikn} = \log \beta_k + \log R_{in}.$$

Substituting this into the expression above, rearranging and absorbing constants into χ yields

$$\log R_{in} = \chi + \log p_n + \sum_{j \neq k} (1 - \beta_k)^{-1} \beta_j X_{jin}. \tag{63}$$

We can see that in this case, even though some characteristics are not observed by the econometrician, the location fixed effect in (56) remains the ideal price index p_n.

C.2 Measurement error

Recall that the log-linearized estimating equation is

$$\hat{\eta}_{int} = \omega_t + \omega' X_{int} + \beta \hat{e}_{int} + \psi \hat{p}_{nt} + \xi_{int}$$

We address measurement error in expenditure in the following way. First, partialling out observable demographics and prices, the reduced-form relationship between expenditure shares and total expenditure is

$$\eta = \beta e + \xi \tag{64}$$

where each variable is residualized, and hats and subscripts are suppressed for notational convenience. Expenditure and rental expenditure are measured with error: $\tilde{e} \equiv e + v^e$, $\tilde{r} \equiv r + v^r$, and $\tilde{\eta} \equiv \tilde{r} - \tilde{e}$. v^e and v^r are assumed to be uncorrelated with e, r, and ξ.

The OLS estimate of β is asymptotically

$$\begin{aligned}\widehat{\beta}_{OLS} &= \frac{\text{Cov}(\tilde{\eta}, \tilde{e})}{\text{Var}(\tilde{e})} \\ &= \frac{\beta\sigma_e^2 + \sigma_{v^r,v^e} - \sigma_{v^e}^2}{\sigma_e^2 + \sigma_{v^e}^2} \\ &= \beta\frac{\sigma_e^2}{\sigma_e^2 + \sigma_{v^e}^2} + \frac{\sigma_{v^r,v^e} - \sigma_{v^e}^2}{\sigma_e^2 + \sigma_{v^e}^2}\end{aligned}$$

The attenuation bias $\sigma_e^2/(\sigma_e^2 + \sigma_{v^e}^2)$ is familiar from classical measurement error. There are two additional sources of bias: (1) measurement error in expenditure appears on both the left- and right-hand sides of (64) and (2) measurement errors in expenditure and rent are mechanically correlated. The direction of the bias is ambiguous, but is likely to be downward if measurement error in expenditure is large and not too highly correlated with measurement error in rent.

Table C.1: Housing prices, hedonic regression

Characteristic	(1) # Rooms	(2) # Bedrooms	(3) Built before 1940	(4) Single-family detached
Panel A: Baseline				
	0.044	0.130	-0.247	0.171
	(0.004)	(0.018)	(0.011)	(0.029)
Panel B: Interactions				
Employment $<$ 50k	0.028	0.146	-0.256	0.180
	(0.008)	(0.015)	(0.022)	(0.040)
50k–100k	0.050	0.150	-0.255	0.209
	(0.005)	(0.007)	(0.018)	(0.027)
100k–250k	0.039	0.133	-0.283	0.174
	(0.004)	(0.009)	(0.011)	(0.020)
250k–500k	0.039	0.143	-0.263	0.174
	(0.005)	(0.009)	(0.014)	(0.022)
500k–1m	0.044	0.150	-0.267	0.153
	(0.005)	(0.010)	(0.016)	(0.021)
$>$ 1m	0.054	0.097	-0.214	0.158
	(0.007)	(0.043)	(0.017)	(0.070)

Source: ACS, 2009-2011.

Note: Renters only, $N = 940,947$. Full set of housing characteristics include number of rooms, bedrooms, the interaction rooms $\times$ bedrooms, and a set of dummies for building age, structure type, and fuel source. Correlation of city price indices from model without interactions (Panel A) and from model with interactions (Panel B) is 0.98. Standard errors in parentheses clustered at MSA level.

C.3 Fixed costs in housing

Throughout the paper we assume that a household in location n pays price p per unit of housing. In practice, there is a lower bound on housing expenditure: one can only add so many roommates or rent so small of a studio. Here, we suppose that households also bear a fixed cost of housing, $f > 0$. In this case, fixed costs will generate nonhomothetic behavior even when underlying preferences are homothetic.

Consider a household with total expenditure e having Cobb-Douglas preferences over housing h and non-housing consumption c. The Cobb-Douglas weight on housing is κ. Letting h^* be the optimal choice of housing, the housing expenditure share is

$$\begin{aligned}\eta^{FC} &= \frac{ph^* + f}{e} \\ &= \kappa + (1-\kappa)\,\frac{f}{e},\end{aligned} \tag{65}$$

which declines with total expenditure; i.e., housing is a necessity.

This setup is isomorphic to an alternative, explicitly nonhomothetic model. Let households have Stone-Geary preferences with weight κ on housing and subsistence housing requirement $\underline{h} > 0$, so that the utility function is $u(h, c) = c^{1-\kappa}(h - \underline{h})^{\kappa}$. By standard arguments, the housing expenditure share is

$$\eta^{SG} = \kappa + (1-\kappa)\,\frac{p\underline{h}}{e}. \tag{66}$$

The expenditure shares in the two models (65) and (66) coincide when $f = p\underline{h}$ — that is, when the fixed component of housing costs is proportional to the variable component. In other words, Cobb-Douglas preferences combined with a particular structure of fixed costs is isomorphic to Stone-Geary preferences.

We estimate (66) in the PSID, without controls, and report the results in Table C.2. To facilitate the interpretation, prices are rescaled to have mean unity. The typical household faces a subsistence housing requirement of about $3,000. The subsistence requirement is about 8% of the typical household's total expenditure of $35,000.

C.4 Alternative specifications in PSID

We present several alternative specifications in Table C.3, still using our baseline sample of renters in the PSID.

In column (1), we include liquid wealth as a control (we use the inverse hyperbolic sine transformation to include households with zero wealth). Liquidity constraints feature in some models of nonhomothetic housing demand such as Bilal and Rossi-Hansberg (2021). The estimates are unchanged, suggesting that liquidity constraints are not first order. In column (2), we instrument for expenditure using job tenure. The exclusion restriction is that job tenure affects the housing share only by shifting total expenditure, conditional on controls including family size and age.

Table C.2: Preference estimates, Stone-Geary utility

	(1) OLS	(2) 2SLS
Subsistence requirement, $\underline{q}$	3,810.73	3,071.13
	(193.45)	(330.81)
Asymptotic housing share, κ	0.24	0.26
	(0.01)	(0.01)
R^2	0.17	
First-stage F-stat.		213.5
N	12,351	10,678
No. of clusters	484	217

Note: Housing prices are normalized to mean unity. Instruments in column (2) are household income and housing supply constraints, described in main text. Standard errors clustered at MSA level.

The estimates are similar. Columns (3) and (4) split the sample into movers and non-movers, respectively, in order to explore a key margin of adjustment to housing expenditure. At annual frequency, households can adjust their housing expenditure either by moving or by re-negotiating their rent. The fact that the estimated ϵ in columns (3) and (4) are similar suggest that both margins appear to be operative. Non-movers' housing expenditure is only slightly more inelastic than movers'. Column (5) uses a county-level rental price index from Zillow, a real estate analytics company (Zillow 2017). Reassuringly, the estimates are similar even with different data and a different level of geography. Column (6) does not instrument for price. The results are similar to the baseline, suggesting that endogeneity of prices is not first order. Column (7) shows that the coefficient estimates are robust to excluding demographic controls, which is evidence against households' sorting on variables other than income and price. Column (8) estimates (24) by 2SLS with household fixed effects, and yields very similar results to our fixed-effect GMM estimates. Column (9) repeats 2SLS with household and MSA fixed effects.

C.5 Consumer Expenditure Survey (CEX)

In this section we present results from the Consumer Expenditure Survey (CEX). Reassuringly, all findings are close to our main results.

In the first column of Table C.4, we re-estimate our baseline specification in the CEX. The estimated expenditure elasticity is slightly higher, but the difference is not statistically significant.

C.5.1 Homeowners

Thus far we have focused on renting households because we do not observe expenditure on owner-occupied housing. In this section we explore whether our results extend to homeowners too. An appropriate measure of housing expenditure by homeowners is *rent equivalent*, which is

Table C.3: Preferences, alternative specifications (PSID)
Dependent variable: Log housing share

	(1) GMM	(2) GMM	(3) GMM	(4) GMM	(5) GMM	(6) GMM	(7) GMM	(8) 2SLS	(9) 2SLS
ϵ	-0.303	-0.400	-0.275	-0.363	-0.256	-0.350	-0.308	-0.434	
	(0.037)	(0.128)	(0.043)	(0.045)	(0.043)	(0.030)	(0.030)	(0.140)	
σ	0.523	0.506	0.616	0.369	0.686	0.461	0.530	0.522	
	(0.077)	(0.089)	(0.094)	(0.105)	(0.066)	(0.037)	(0.078)	(0.228)	
Log expenditure								-0.415	-0.430
								(0.148)	(0.137)
Log price								0.457	
								-0.173	
Demographic controls	✓	✓	✓	✓	✓	✓		✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controlling for wealth	✓								
Household FE								✓	✓
MSA FE									✓
Expenditure IV	Income	Tenure	Income	Income	Income	Income	Income	Income	Income
Price	Census	Census	Census	Census	Zillow	Census, no IV	Census	Census	
Sample	Full	Full	Movers	Non-movers	Full	Full	Full	Full	Full
R^2								0.17	0.18
First-stage F-stat.								26.8	34.8
N	10,678	10,353	6,729	3,964	6,572	12,351	10,678	8,670	9,985
No. of clusters	217	217	212	176	192	484	217	197	350

Source: PSID, Census, Zillow, and Saiz (2010).

Note: Column (1) instruments for expenditure using job tenure. Columns (2) and (3) split the sample into households which moved addresses and those which did not. Column (4) use median price per square foot at the county level from Zillow. Column (5) repeats the main GMM specification without instrumenting for price. Column (6) omits demographic controls. Column (7) reports the linearized model with household fixed effects. Column (8) reports the linearized model with household and MSA fixed effects. Price instruments are geographic and regulatory constraints. Standard errors clustered at MSA level.

Table C.4: Preferences (CEX)

Dependent variable: Log housing share

	(1) GMM	(2) GMM	(3) GMM	(4) GMM	(5) GMM	(6) GMM
ϵ	-0.312	-0.253	-0.443	-0.440	-0.312	-0.228
	(0.043)	(0.031)	(0.041)	(0.040)	(0.026)	(0.038)
σ	0.319	0.330	0.114	0.143	0.272	0.338
	(0.198)	(0.146)	(0.155)	(0.157)	(0.109)	(0.147)
Sample	Renters	Pooled	Owners	Owners	Pooled	Owners
Rent measure for owners		Rent equivalent	Rent equivalent	Rent equivalent	Out of pocket	Out of pocket
Including homeowners with second homes				✓		
Demographic controls	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
N	2,995	8,269	5,274	5,659	8,269	5,274

Source: CEX, Census, and Saiz (2010).

Note: Column (1) replicates our baseline specification of Table 1 column (5), using the CEX. Column (2) adds homeowners owning one home, measuring housing expenditure by self-reported rental equivalent. Column (3) restricts to homeowners only. Column (4) includes households who own second homes. Columns (5) and (6) measures housing expenditure as out-of-pocket expenses, defined as mortgage interest, property taxes, insurance, maintenance, and repairs; mortgage principle is excluded. Instrument is log household income. Standard errors clustered at MSA level.

the market rate for the flow of housing services consumed. The PSID consumption module did not elicit rent equivalent until 2017, but rent equivalent is available in all recent waves of the CEX. Therefore we use the CEX to study homeowners.

Column (2) of Table C.4 pools renting and owning households together. The estimate is consistent with significant nonhomotheticity. Restricting attention only to owners (column (3)) yields even stronger nonhomotheticity than the baseline estimate for renters.

In columns (5) and (6), we use an alternative measure of housing expenditure for homeowners, *out-of-pocket expenses*. We define out-of-pocket expenses as the sum of mortgage interest, property tax, insurance, maintenance, and repairs. We omit payments on mortgage principal since these payments are savings, not consumption. Out-of-pocket expenses reflect the user cost of housing, which is equal to the rental value of the house in equilibrium. The estimates are close to our baseline results.

In our main analysis of homeowners, we restrict our sample to households who own a single home, which includes 94% of homeowners in the CEX. The reason is that expenditure on second homes does not reflect the local cost of living, but rather is a luxury more akin to recreation or vacations. That said, it is possible that second homes are a substitute for primary homes in expensive markets: for example, a household could split time between a pied á terre and a large country home. Including the owners of second homes in column (4) leaves our results virtually unchanged.

An alternative to using rent equivalent would be to use home values. Prior work with the PSID (Attanasio and Pistaferri 2016; Straub 2019) imputes rent as a constant fraction of home values, typically six percent (Poterba and Sinai 2008). However, the Residential Financial Survey, used by the BEA to impute rents in the national accounts, shows that the rent-to-value ratio is strongly decreasing in home value. In the CEX, we find that the average rent-to-value ratio is 14% for homes valued at $50,000 compared to 4% for homes valued at $1,000,000. Assuming a constant 6% rent-to-value ratio would bias housing expenditure downward at the bottom and upward at the top of the expenditure distribution. We therefore do not use home values in the analysis.

C.6 Income elasticities from the literature

Table C.5 summarizes estimates of the income elasticity of housing demand from the literature. Controlling for local prices, using expenditure on the right hand side, and accounting for measurement error with an IV are all key in obtaining a consistent estimate of the elasticity.

D Quantitative Model and Calibration

D.1 Multisector Extension

The model presented in Section 3 has a single sector. In this section we develop an extension which incorporates multiple sectors indexed by $j = 1, ...J$. Doing so has two benefits. First, it

Table C.5: Income elasticities in the literature

Paper	Elasticity	Sample	Local prices?	Expenditure?	IV?
Rosenthal (2014)[a]	-0.88	Renters	✓		
Ioannides, Zabel, et al. (2008)[b]	-0.79	Owners	✓		
Hansen, Formby, Smith, et al. (1996)[c]	-0.73	Renters			
Larsen (2014)[d]	-0.67	Owners			
Zabel (2004)[e]	-0.52	Owners	✓		
Albouy, Ehrlich, and Liu (2016)[f]	-0.28	Renters	✓		
Lewbel and Pendakur (2009)[g]	-0.28	Renters		✓	✓
Attanasio et al. (2012)[h]	-0.22	Both		✓	
Aguiar and Bils (2015)[i]	-0.08	Both		✓	✓
Davis and Ortalo-Magné (2011)[j]	-0.01	Both		✓	✓
Paper benchmark[k]	-0.25	Renters	✓	✓	✓

[a] American Housing Survey, 1985-2011. Table 5, column 1.
[b] American Housing Survey, 1985-1993. Table 5, column 1.
[c] American Housing Survey, 1989. Table 5, column 2, last row.
[d] Norweigian Rental Survey and Consumer Expenditure Survey, 2007. Table 2, row 5.
[e] American Housing Survey, 2001. Table 3, row 3.
[f] US Census, 1970-2014. Table 1, column 3.
[g] Canadian Family Expenditure Surveys, 1969-1996. Median uncompensated elasticity computed using authors' replication file following their Appendix VII.1.
[h] British Household Panel Survey, 1991-2002. Table 4, panel B. Estimates for high- and low-education groups are averaged with weights one-third and two-thirds, respectively.
[i] US CEX, 1980-2010. Table 2, column 1.
[j] US CEX, 1982-2003. Text, page 253.
[k] PSID, 2005-2017.

allows us to consider a particular interpretation of aggregate skill bias A as stemming from differential exposure across skill groups to fast versus slow growing sectors. Second, it motivates the Bartik shift-share instrument used in the model calibration.

We assume that workers are perfectly mobile across sectors within a location, implying that all sectors pay the same wages w_{sn} and w_{un}. Total output in location n, which we denote by Y_n, is a CES aggregate of local produced sectoral outputs denoted y_{jn}

$$Y_n = \left(\sum_j y_{jn}^{\frac{\varphi-1}{\varphi}} \right)^{\frac{\varphi}{\varphi-1}} \tag{67}$$

where φ is the elasticity of substitution between different sectors. Each sector's output is in turn produced according to

$$y_{jn} = \left(\left(q_{jsn} L_{jsn}\right)^{\frac{\rho-1}{\rho}} + \left(q_{jun} L_{jun}\right)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}. \tag{68}$$

where q_{jsn} is the productivity of skilled workers in sector j in location n, L_{jsn} is the supply of such workers, and q_{jun} and L_{jun} are defined analogously.

After some algebra, labor demand can be written

$$w_{in} = L_{in}^{-\frac{1}{\rho}} \left(\sum_j q_{jin}^{\rho-1} \left(w_{sn}^{1-\rho} q_{jsn}^{\rho-1} + w_{un}^{1-\rho} q_{jun}^{\rho-1} \right)^{\frac{\varphi-\rho}{\rho}} \right)^{\frac{1}{\rho}} Y_n^{\frac{1}{\rho}}. \tag{69}$$

for $i = s, u$. To make progress we now make the convenient assumption that the elasticity of substitution between worker types ρ is equal to the elasticity of substitution between sectors φ. With this assumption, labor demand becomes

$$w_{in} = L_{in}^{-\frac{1}{\rho}} Q_{in}^{\frac{\rho-1}{\rho}} Y_n^{\frac{1}{\rho}}, \tag{70}$$

where

$$Q_{in} = \left(\sum_j q_{jin}^{\rho-1} \right)^{\frac{1}{\rho-1}} \tag{71}$$

is a weighted average of group i's sector productivities. Inspecting (2), (3), and (70), we can see that this multi sector model is isomorphic to the one sector model presented in Section 3.

We now use this multisector model to show how differential growth across sectors relates to changes in aggregate skill bias A. We begin with the convenient assumption that the q_{jin} can be decomposed as follows

$$q_{jin} = q_i q_n q_j q_{ji}. \tag{72}$$

That is, some locations n, skill types i, and sectors j are uniformly more productive, and different skill types may have different productivities in different sectors, captured by q_{ji}. Now let us suppose sectors experience different levels of productivity growth, denoted by $\Delta \log q_j$. With these assumptions, one can show that

$$\Delta \log A \simeq \sum_j \left(\ell_{js} - \ell_{ju}\right) \Delta \log q_j, \tag{73}$$

where $\Delta \log A$ is the change in log aggregate skill bias and $\ell_{ji} \equiv \frac{\sum_n L_{jin}}{\sum_j \sum_n L_{jin}}$ is the share of national employment of type i in sector j. Skilled workers experience faster productivity growth because their employment is concentrated in sectors which experience relatively fast growth. Note that this does not imply that productivity growth is the same in all locations; some locations will grow faster than others, but only because they are relatively more skill intensive. A natural example of this dynamic would be business services. These sectors are skill intensive, and have experienced fast productivity growth since 1980 because they are particularly able to take advantage of advances in information and communications technology (Eckert, Ganapati, and Walsh 2022).

Of course, local productivities q_{jin}, need not have this special structure. More generally, we will have

$$\Delta \log A \simeq \sum_n \lambda_n \sum_j \left[\ell_{jsn} \left(\Delta \log q_{jsn} - \Delta \log q_{jun}\right)\right] + \left[\Delta \log q_{jun} \left(\ell_{jsn} - \ell_{jun}\right)\right], \tag{74}$$

where the μ_n are 1980 employment weights. This expression shows that movement in aggregate skill bias comes from faster growth in productivity for skilled workers in a given sector and location (the first term in square brackets) and greater exposure to fast growing sectors or regions (the second term in square brackets), averaged across space using 1980 employment weights.

The multi sector model can also motivates the use of Bartik instruments. These instruments, denoted by Z_{in}, are defined by

$$Z_{in} = \sum_j \ell_{jin} g_j \tag{75}$$

where $\ell_{jin} \equiv \frac{L_{jin}}{\sum_j L_{jin}}$ is the share of type i working in sector j in location n and g_j is aggregate growth in wages in sector j. In order to be valid instruments, we need the Bartik shifters to be relevant, i.e, correlated with changes in local labor demand for type i, Q_{in}. Differencing Q_{in} over time, we obtain

$$\Delta \log Q_{in} = \sum_j \ell_{jin} \Delta \log q_{jin}. \tag{76}$$

Assuming that local changes in productivity in sector j are correlated with national changes in productivity, and thus with aggregate wage growth in sector j, the relevance of the Bartik instrument follows immediately.

D.2 Calibration

Migration elasticity

We calibrate θ by requiring our model to match the results of Hornbeck and Moretti (2022), following Greaney (2023). That paper estimates the causal effect of TFP shocks on employment and wages. We mimic their setting by shutting down all shocks other than shocks to productivity and then repeating their regressions using the output of our model. Our target is the ratio of the effect on employment to the effect on wages — the long run elasticity of employment to wages. We target of $3.35/1.54 = 2.18$ (see Table 3, Column (8) of Hornbeck and Moretti (2022)). Formally we proceed as follows:

(i) Guess θ

(ii) Invert the model in 1980 and 1990 to obtain fundamentals $\left(A_{in}^t, B_{in}^t\right)_{i,n}, \left(\Pi_n^t\right)_n, \left(L_i^t\right)_i$ for $t = 1980, 1990$.

(iii) Solve the model with fundamentals $\left(A_{in}^{90}, B_{in}^{80}\right)_{i,n}, \left(\Pi_n^{80}\right)_n, \left(L_i^{80}\right)_i$ to obtain $\left(\hat{l}_{in}^{90}, \hat{w}_{in}^{90}\right)_{i.n}$.

(iv) Define $L_n^{80} = \sum_i l_{in}^{80}$, $W_n^{80} = \sum_i l_{in}^{80} w_{in}^{80} / \sum_i l_{in}^{80}$ and $\log Z_n^{80} = \sum_i l_{in}^{80} \log A_{in}^{80} / \sum_i l_{in}^{80}$ and likewise for $\hat{L}_n^{90}$, $\hat{W}_n^{90}$ and $\log \hat{Z}_n^{90}$

(v) Estimate the models below by OLS, weighting by 1980 employment:

$$\log \hat{L}_n^{90} - \log L_n^{80} = \pi^L \left(\hat{Z}_n^{90} - Z_n^{80}\right) + v_n^L$$
$$\log \hat{W}_n^{90} - \log W_n^{80} = \pi^W \left(\hat{Z}_n^{90} - Z_n^{80}\right) + v_n^W.$$

The fact that we only study changes between 1980 and 1990 is innocuous, because our model has no transitional dynamics.

(vi) Calculate π^L / π^W.

(vii) Update θ until π^L / π^W converges to the target value.

This procedure yields $\theta = 0.285$.

Elasticity of substitution in production

The production side of the model is standard and we externally calibrate $\rho = 3.85$ to match Card (2009).[34] That paper estimates the elasticity of substitution between workers of different skill groups at the MSA level using immigration as an instrument for labor-supply changes. The elasticity is larger than canonical estimates from Katz and Murphy (1992) and Acemoglu and Autor (2011), who report values close to 1.6. However, Katz and Murphy (1992) estimate an aggregate production function on time-series data, whereas Card (2009) estimates a city-level production

[34] See Table 5, column (7), in Card (2009) for the negative inverse elasticity of $-1/\rho = -0.26$.

function on cross-sectional data. Studies estimating the elasticity of substitution at the city level tend to find values between 3 and 5 (Bound et al. 2004; Beaudry, Doms, and Lewis 2010; Baum-Snow, Freedman, and Pavan 2018; Eckert, Ganapati, and Walsh 2022).[35]

Tax system

We use data from the 1981/91/2001/11 waves of the PSID (each containing summary information on the *prior* year's income). Using the same sample restrictions as in section 2, we run the PSID data through the NBER's TAXSIM program. For each household, pre-tax income is computed as adjusted gross income minus Social Security transfers. Post-tax income is computed as pre-tax income minus federal and state taxes (including payroll taxes) plus Social Security transfers. We estimate (32) in logs by pooled OLS over the four periods. Our estimated $\hat{\tau}$ is 0.174 (robust s.e. 0.003). The R^2 of the regression is 0.98, suggesting that, despite its parsimony, a log-linear tax equation is a good approximation to the actual tax system in the United States. Our estimate is close to Heathcote, Storesletten, and Violante (2017), who estimate $\hat{\tau} = 0.181$.

Housing Supply Elasticities

Equation 40 tells us the change in prices between any two equilibria as a function of changes in the shifter $\hat{\Pi}_n$, changes in housing expenditure in location n, which we denote by $\hat{HD}_n$, and the housing supply elasticity γ_n. We take the observed equilibrium in 1980 as the baseline and the observed equilibrium in 2010 as the 'new' equilibrium, so that (40) gives us an expression for changes in local price index over time. Taking logs,

$$\log \hat{p}_n = \gamma_n \log \hat{HD}_n + \log \hat{\Pi}_n. \tag{77}$$

Following Saiz (2010), we parameterize γ_n as a function of geographical and regulatory constraints,

$$\gamma_n = \gamma + \gamma_L UNAVAL_n + \gamma_R WRLURI_n. \tag{78}$$

$UNAVAL_n$ is a measure of geographic constraints from Saiz (2010) and $WRLURI_n$ is the Wharton Residential Land Use Regulation Index developed by Gyourko, Saiz, and Summers (2008). Substituting the expression for γ_n into (77) yields an estimating equation for γ, γ_L and γ_R,

$$\log \hat{p}_n = (\gamma + \gamma_L UNAVAL_n + \gamma_R WRLURI_n) \log \hat{HD}_n + \log \hat{\Pi}_n. \tag{79}$$

We interpret $\log \hat{\Pi}_n$ as an unobserved shock to housing supply in location n.

Saiz (2010) reports values of land unavailability $UNAVAL_n$ and regulatory constrains $WRLURI_n$ for a subset of MSAs. After dropping those for which these measures are missing, we are left with 193 MSAs. Prices p_n are obtained from hedonic regressions in the Census data as

[35] An exception is Diamond (2016), who estimates an elasticity close to 1.6 in line with the time-series results.

described in the text. We use Census data on employment, wages, and (19) to construct housing expenditure $\sum_i \eta_{in} e_{in} l_{in}$ for each MSA. Finally, following Diamond (2016), we instrument for housing expenditure using a Bartik shifter Z_{int}, where the shares are a region's industrial composition in 1980, and the shift is change in average wages nationwide (excluding the region itself). We also use the interactions of Z_{int} with $UNAVAL_n$ and $WRLURI_n$) as instruments. Table D.1 reports the result of estimating (79) by 2SLS. For the 193 locations with complete data, we then define

Table D.1: Housing Supply Elasticity Estimates
Dependent variable: Log price change, 1980-2010

γ	0.209
	(0.069)
γ_L	0.090
	(0.055)
γ_R	0.230
	(0.057)

Source: Census. Robust standard errors in parentheses.

$$\gamma_n = \gamma + \gamma_L UNAVAL_n + \gamma_R WRLURI_n.$$

Of the remaining locations, 50 are the nonmetro portions of states and 26 are MSAs for which $UNAVAL_n$ and $WRLURI_n$ are not available. For the 26 MSAs, we define γ_n to be the median among the 193 MSAs with complete information. For the 50 state residuals, we set γ_n to the lowest value among the 193 MSAs with complete information, on the assumption that supply is likely to be more elastic in nonmetro areas.

E Counterfactual

E.1 Skill Premia

In the body of the paper, we focus on the spatial distribution of skilled versus unskilled workers, and in particular variation in the log ratio of skilled to unskilled employment across space. But our model also makes predictions for wages. We focus on the local skill premium ω_{nt}, defined as the log ratio of skilled to unskilled wages in location n in year t. In 1980 local skill premia were moderately closely correlated with skill ratios. Column (1) of Table E.1 shows that this relationship became much tighter by 2010, with the correlation rising from 0.208 to 0.618. Column (2) shows that this was largely driven by fast growth in the skill premium in initially skill intensive locations by fixing μ_{nt} at its 1980 level in each location. The correlation still rises substantially, reaching 0.559 in 2010.

The columns labeled 'Model' show how these correlations evolve in our main counterfactual exercise, in which we hold A_t constant at its 1980 level and let all other fundamentals evolve as they did in the data. (1) shows that in the absence of increases in A_t, the correlation between ω_{nt} and μ_{nt} would have risen slightly faster, and (2) shows this is true even fixing local skill ratios at their 1980 levels. The final row shows that increases in A_t explain a slightly *negative* fraction of the observed increase in the correlation between ω_{nt} and μ_{nt}. Intuitively, the mechanism we focus on pushes skilled workers towards expensive cities. The CES production function (29) then implies that, absent shocks to local skill bias a_{nt}, skill premia in these locations will fall. Since these locations are typically skill intensive – as we can see from the columns labelled 'Data' in Table E.1 – this tends to push down the correlation between ω_{nt} and μ_{nt}.

Table E.1: Skill Ratios and Skill Premia

	(1) $\rho\left(\omega_{nt}, \mu_{nt}\right)$		(2) $\rho\left(\omega_{nt}, \mu_{n0}\right)$	
	Data	Model	Data	Model
1980	0.208	0.208	0.208	0.208
2010	0.618	0.621	0.559	0.578
% of observed change accounted for	-	-0.61%	-	-5.35%

Note: The columns labelled (1) show the correlation of local log skill premia ω_{nt} with log skill ratios μ_{nt}, weighting locations by their 1980 employment shares. The first of these uses the data and the second shows these correlations in the counterfactual in which A_t is held constant at its 1980 level. The columns labelled (2) show the same correlations, but use skill ratios measured in 1980 rather than contemporaneously. The final row shows the difference between data and counterfactual, expressed as a percentage of the growth in the data between 1980 and 2010. Note that a negative value here indicates that increases in A_t have pushed these correlations down, while they have risen in the data.

E.2 Covariance of skill ratios and expenditure shares

In Panel (e) of Figure 3, we vary the covariance between expenditure shares η_{sn} and skill ratios μ_n by creating many simulated datasets. We now describe how we simulate these datasets in detail. For each simulated dataset, we copy wages, prices and total MSA populations from the data exactly. The only objects we change are the log skill ratios μ_n, which we construct using

$$\mu_n = a + bm_n + c\eta_{sn} \tag{80}$$

where a, b and c are parameters and $m_n \sim \mathcal{N}(0, 1)$. We choose c to vary the covariance between μ_n and η_{sn}, and then choose a and b to match two salient features of the data: a is chosen to match the overall skill share ϕ; and b is chosen to match the variance of log skill ratios across space $\mathcal{M}$, which we have used as our main measure of sorting throughout. To create the plot in Panel (e), we choose 10 values of c evenly spaced between 0 and 10.6, with the upper bound chosen to generate a covariance between μ_n and η_{sn} roughly 20% larger than the value observed in the data. For each value of c, we take 10 draws of m_n. For each draw of m_n, we construct a new μ_n, calculate its covariance with η_{sn} and re-run our main counterfactual exercise. The result is the 100 simulations plotted in Panel (e).

E.3 Alternative Measures of Sorting

In the body of the paper we focus on a particular measure of spatial sorting by skill, the variance of the log skill ratio

$$\mathcal{M} = \frac{1}{2}\text{Var}(\mu_n).$$

Table E.2: Alternative Measures of Sorting

	(1) $\mathcal{M}$	(2) $\mathcal{T}$	(3) $\mathcal{D}$	(4) $\mathcal{N}$
Data: Change 1980-2010, %	32.8	34.5	18.9	8.3
Counterfactual: Change 1980-2010, %	23.7	25.6	15.4	5.9
% of observed change accounted for	27.4%	25.9%	19.2%	28.8%

Note: Each column reports the results of the main counterfactual in which all fundamentals evolve as they did in the data, apart from aggregate skill bias A_t, which is held constant at its 1980 level. The first column uses our baseline measure of sorting and reproduces the results from Column (1) of Table 5. Column (2) uses the Theil index from (81). Column (3) uses the dissimilarity measure from (82). Column (4) uses the ninety-ten difference from (83). In each case, the first row reports the percentage change 1980-2010 in the data, and the second the percentage change in the counterfactual. The final row reports the difference between data and counterfactual, expressed as a percentage of the increase in the data, and corresponds to the share of the increase in each measure accounted for by the increase in A_t.

We now consider three alternative measures of sorting. First, the Theil index for a non-negative variable x with weights λ_n is defined as

$$\mathcal{T} = \sum_i \lambda_n \left(\frac{x_n}{\bar{x}}\right) \log\left(\frac{x_n}{\bar{x}}\right) \tag{81}$$

where $\bar{x}$ is the weighted average of x_n. We use this as a measure of sorting by setting $x_n = \exp(\mu_n)$, where μ_n is the log-skill ratio, and weight by 1980 employment. Second, the dissimilarity index for two populations u and s, spread over geographical units indexed by n is given by

$$\mathcal{D} = \frac{1}{2} \sum_n |\ell_{sn} - \ell_{un}| \tag{82}$$

where ℓ_{sn} is location n's share in overall skilled employment and ℓ_{un} is defined analogously. This measure has been used recently by Fogli and Guerrieri (2019) to study within-city segregation by income. Third, we define the 90-10 difference of the log skill ratio as

$$\mathcal{N} = \mu^{90} - \mu^{10} \tag{83}$$

where μ^{90} is the ninetieth percentile of the log skill ratio distribution and μ^{10} is defined analogously. The 90-10 difference captures the extent of dissimilarity between a relatively skilled versus relatively unskilled location, and thus can be thought of as a measure of how intensely the two skill groups cluster.

The first row of Table E.2 reports the percentage change in each measure of sorting between 1980 and 2010. All measures of sorting increased, although the magnitude varies across measures. The second row shows how each measure of sorting evolves in our main counterfactual exercise, corresponding to Column (1) of Table 5, and the final row expresses this as a share of the observed increase that is explained by rising aggregate skill bias. The share explained is quite similar across different measures, although a little lower for the dissimilarity index. We conclude that our results are robust to using alternative measures of spatial sorting by skill.

E.4 Alternative Parametrization of Preferences

We recalibrate our model to Price Independent Generalized Linear (PIGL) utility, a leading case of nonhomothetic preferences (Boppart 2014; Eckert and Peters 2023). PIGL admits a closed form for the indirect utility function (26),

$$v_{in} = \frac{1}{\varepsilon}(e_{in}^{\varepsilon} - 1) - \frac{\Omega}{\varsigma}(p_n^{\varsigma} - 1)$$

for parameters $0 < \varepsilon < \varsigma < 1$ and $\Omega > 0$. By Roy's identity, the housing share is

$$\eta_{in} = \Omega e_{in}^{-\varepsilon} p_n^{\varsigma} \tag{84}$$

Taking logs, adding a time subscript, and interpreting the scalar Ω as an idiosyncratic household demand shifter Ω_{int}, (84) is equivalent to the linearized estimating equation (21) for NHCES utility. The income elasticity is ε and the price elasticity is ς, which correspond to β and ψ, respectively, in (22). We can therefore read the parameters directly off column (4), Table 1, setting $\varepsilon = 0.248$ and $\varsigma = 0.390$. After recalibrating the full model we find that the skill premium explains 25.9% of the increase in sorting since 1980, comparable to our baseline results. We conclude that our findings are not sensitive to the parametrization of utility.

E.5 Endogenous Amenities

Diamond (2016) shows the importance of endogenous amenities for understanding the location choices of skilled and unskilled workers. In that paper, skilled workers impose a positive externality on other skilled workers, which motivates sorting. In this subsection we consider how our results might change in the presence of endogenous amenities.

We incorporate endogenous amenities into the quantitative model of Section 3. Following Diamond (2016) we model amenities as

$$B_{in} = \overline{B}_{in} \mathcal{B}_n^{\nu_i} \tag{85}$$

$$\mathcal{B}_n = \left(\frac{\ell_{sn}}{\ell_{un}} \right)^{\zeta} \tag{86}$$

Amenities are the product of an exogenous component $\overline{B}_{in}$ and an endogenous component $\mathcal{B}_n$. Diamond (2016) finds that the endogenous component is increasing in the skill ratio and is valued more by skilled types, so $\zeta > 0$ and $\nu_s > \nu_u > 0$. We assume $(\nu_s - \nu_u)\,\zeta < 1 + \frac{1-\tau}{\theta\rho}$ to ensure that endogenous amenities are not attractive enough to generate perfect sorting.

The first result extends Proposition 4 to endogenous amenities. Even with endogenous amenities, sorting does not depend on aggregate skill bias when preferences are homothetic.

Proposition 6 *Consider the quantitative model with endogenous amenities. Suppose $\epsilon = 0$ so that preferences are homothetic. Then, $\mathcal{M}$, the level of sorting, does not depend on aggregate skill bias A.*

The proof is similar to the proof of Proposition 4. From location choice (28) and endogenous amenity preferences (85), the log skill ratio can be written

$$\mu_n = \varkappa + \frac{1}{\theta} \log \frac{e_{sn}}{e_{un}} + \log \frac{\overline{B}_{sn}}{\overline{B}_{un}} + (\nu_s - \nu_u)\,\mathcal{B}_n \tag{87}$$

for some $\varkappa$ that does not vary by location. As in the benchmark model, homotheticity causes the price indices to drop out of (87). Substituting expenditures (32), wages (31) and (30), and the endogenous amenity (86) gives, after some algebra,

$$\mu_n = \tilde{\vartheta}_0 + \tilde{\vartheta}_1 \log a_n + \tilde{\vartheta}_2 \log \frac{\overline{B}_{sn}}{\overline{B}_{un}} \tag{88}$$

Table E.3: Model calibration, endogenous amenities

Parameter	Value
θ	0.40
	(0.07)
ν_s	0.77
	(0.24)
ν_u	0.27
	(0.18)
ζ	2.00
	(0.96)
N	166

Note: Robust standard errors in parentheses. Parameters estimated via GMM using replication package from Diamond (2016).

where $\tilde{\vartheta}_1$ and $\tilde{\vartheta}_2$ depend only on $\theta, \rho, \tau, \nu_s, \nu_u$, and ζ.

Second, we show quantitatively that, in the presence of nonohomothetic preferences, endogenous amenities amplify the response of sorting to skill bias. Using the replication package from Diamond (2016), we estimate the labor supply parameters $(\theta, \nu_s, \nu_u, \zeta)$ in a three equation system: employment growth as a function of real wage growth and changes in amenities for the two types, and changes in amenities as a function of changes in the skill ratio. We use the same instruments as Diamond (2016) and look at the single difference 1980-2000. Table E.3 reports the estimated parameters, which are comparable to those reported by Diamond (2016) and which satisfy the parametric restriction imposed earlier.

Repeating the main counterfactual, we find that the shock to aggregate skill bias explains 53.5% of the increase in sorting. Intuitively, endogenous amenities feed back to strengthen agglomeration for skilled types. A shock to skill bias increases skilled types' willingness to live in expensive cities. As skill ratios rises there, endogenous amenities make those cities even more attractive to skilled workers, increasing the skill ratio further.