Homeownership and the Distributional Effects of Local Shocks

Brian Greaney

Working Paper 2628 **August 2026**

Research Department

https://doi.org/10.24149/wp2628

Homeownership and the Distributional Effects of Local Shocks[*]

Brian Greaney[†]

February 4, 2026

Abstract

I develop a dynamic quantitative spatial framework that combines key features of quantitative spatial and lifecycle housing models to study the distributional effects of spatially heterogeneous shocks. The model features many locations, moving costs, uninsurable income risk, lifecycle dynamics, housing tenure choice, and housing frictions. The model is set in continuous time, which enables efficient computation. I apply the framework to analyze the welfare effects of heterogeneous productivity shocks across U.S. cities. I find that local productivity shocks have important distributional consequences: on average, a 1% shock to local productivity raises residents' welfare by 0.37%. The pass-through from a local productivity shock to welfare varies substantially by age and housing tenure. I show that homeownership plays a central role in spatial redistribution: in an otherwise identical model without homeownership, the average welfare effect of a 1% local productivity shock is just 0.03%. This is because house price changes counteract the welfare effects of wage changes for renters, but augment them for owners. These results suggest that accounting for homeownership is essential for understanding the distributional effects of spatially heterogeneous shocks.

JEL Classification: R12, R23, R00, J61

Keywords: Spatial dynamics, economic geography, migration, homeownership

[*] I am grateful to my advisors Giuseppe Moscarini, Costas Arkolakis, Samuel Kortum, and Michael Peters for their continuous guidance and support throughout this project. I have also benefited tremendously from conversations with Lorenzo Caliendo, Fabian Eckert, Robert Finlay, Joachim Hubmer, Zhen Huo, Tony Smith, Aleh Tsyvinski, Conor Walsh, Elisa Giannone, and participants at the Yale Macro Lunch and International Trade Lunch. The views expressed in this paper are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Dallas or the Federal Reserve System.

[†] Brian Greaney, Federal Reserve Bank of Dallas, brian.greaney@dal.frb.org.

1 Introduction

Economic outcomes have varied tremendously across regions of the United States over the past several decades. While some areas have experienced positive shocks, others have stagnated or declined. There are concerns that these regional disparities may have important distributional consequences. Such worries are magnified by falling migration rates and rapidly rising cost of living in high-income areas. There is a sense that the benefits of prosperous locations are limited to a subset of incumbent residents, while being inaccessible to others.[1]

Although there is reason to believe that these spatial disparities may have distributional effects, their quantitative importance is unclear. It depends on the strength of frictions that link individual outcomes to local economic conditions. Without migration frictions or home-biased investment, the welfare effects of a shock to one area would not vary across space, as utility would equalize across locations. In reality, households face numerous frictions that cause welfare to depend on the local economy. For example, migration entails moving costs and time to look for a new job and housing. Many people own housing, an asset whose value is closely tied to local economic conditions. Such frictions cause spatially heterogeneous shocks to have distributional consequences by concentrating the effects of local shocks on local residents.

Ultimately, the degree to which spatially heterogeneous shocks cause redistribution across households is a quantitative question. Existing quantitative spatial models contain some of the ingredients needed to answer this question. Specifically, they account for realistic geography and migration costs. However, virtually all abstract from homeownership, which is of first-order importance in this context. Most households own a home, and the typical household invests the majority of its wealth in it. Homeownership fundamentally shapes the distributional effects of local shocks. For example, a productivity shock which raises wages in one location also endogenously raises local house prices. If housing is rented, higher housing costs offset the benefits of higher wages, thus mitigating the distributional effects of the shock. If instead housing is owned, incumbent residents are insulated from higher housing costs and benefit from capital gains. In this case, house price changes *exacerbate* the distributional effects of the shock.

In this paper, I develop a dynamic quantitative spatial model that can be used to estimate distributional effects in a setting with rich geography and realistic homeownership. There are many locations, and households face moving costs which limit their ability to respond to shocks via migration. Locations differ by productivity, amenities, and house price elasticity. These features allow the model to match the endogenous responses of populations, wages, and house prices to local shocks. In contrast to standard spatial models, housing can be either rented or owned. There are transaction costs required to buy or sell housing, as well as a collateral constraint that links borrowing capacity to housing wealth. Households are finitely lived, and face uninsurable income risk.

[1] For evidence of regional divergence and a discussion of its distributional consequences, see Moretti (2012). Molloy et al. (2011) and Kaplan and Schulhofer-Wohl (2017) document a long-term decline in internal migration in the United States. Ganong and Shoag (2017) show that high-income cities have experienced a disproportionate increase in house prices during the past thirty years, and argue that this has especially affected low-income households.

These features endogenously generate observed lifecycle patterns of housing investment. They also allow the model to capture the welfare effects of house price changes, which vary by age and tenure.

While the model is well-suited for quantifying the distributional effects of spatially heterogeneous shocks, solving it presents a computational challenge. The combination of incomplete markets, lifecycle dynamics, housing frictions, and geography leads to a large state space. Specifically, there is a joint wealth-housing-income-age distribution in every location. Realistic housing frictions imply that the household's problem is nonconvex, which precludes the use of efficient discrete-time numerical algorithms. It is imperative that the model be solved efficiently, because estimating a quantitative spatial model requires choosing a large number of location-specific parameters. Moreover, analyzing the effects of local shocks when households make dynamic decisions requires computing transition dynamics. The first contribution of this paper is to show that these challenges can be overcome by setting the model in continuous time and taking advantage of numerical techniques recently developed by Achdou et al. (2022). These are much more efficient than discrete-time methods, and can handle nonconvexities without difficulty. The main reason is that, due to fixed housing adjustment costs, the first-order conditions of the household's problem are sufficient in continuous time but not in discrete time. Using these methods expands the types of heterogeneity and frictions that can be considered in spatial models, making it possible to ask new questions and incorporate model features that would otherwise be infeasible.

I calibrate the model to match important features of the data on economic geography and homeownership. I then quantitatively analyze the distributional effects of spatially heterogeneous productivity shocks across U.S. cities. I focus on two shocks. The first reflects heterogeneous exposures to industries with varying productivity changes. The second is chosen to exactly match observed wage changes over the past two decades. The first main result of the paper is that spatially heterogeneous productivity shocks have important distributional consequences. On average, a 1% increase in a city's productivity raises local residents' lifetime utility by 0.37%. The pass-through from local productivity shocks to welfare varies substantially across households. Local productivity shocks affect homeowners more than renters, and middle-aged households more than the young or old. Whereas a local productivity increase raises welfare for most households, it *decreases* welfare for old renters. These households do not benefit from increased wages because they do not participate in the labor market, and are harmed by higher rents.

I next examine the mechanisms that shape the relationship between local shocks and welfare. My first finding is that homeownership plays a critical role. To quantify the importance of homeownership, I repeat the baseline estimation and quantitative analysis when homeownership is prohibited. Without homeownership, the average welfare effect of a 1% local productivity increase is just 0.03%. The fraction of households who are harmed by higher productivity also increases relative to the baseline because there are more old renters. These results suggest that accounting for homeownership is essential when analyzing the distributional effects of spatially heterogeneous shocks. Models that do not include homeownership are likely to underestimate such effects.

In order to better understand the effects of homeownership, I next perform a counterfactual in which first the wage changes and then the house price changes caused by the productivity shocks are shut down. While higher productivity benefits both renters and owners via the wage channel, the effect via the house price channel is positive for owners but negative for renters. This explains why homeowners' welfare is more sensitive to local productivity.

In the final set of results, I investigate the specific mechanisms that cause local productivity shocks to affect renters and owners differently. There are three reasons for this divergence: first, owners receive capital gains when house prices change; second, the pass-through from house prices to effective housing costs differ due to how housing is taxed; and third, transaction costs make moving more expensive for owners. To quantify the relative importance of these channels, I perform a series of counterfactuals in which they are successively shut down. I find that the capital gains channel is by far the most important. Importantly, I find that the welfare effects of house price changes differ substantially between the full model and a version of the model with only one location. Higher house prices affect households both by increasing housing costs, and also by increasing housing wealth. The net welfare effect for a particular household depends on the expected present value of future costs versus capital gains. In the one-location model, higher housing costs cannot be avoided and owners do not expect to realize capital gains until they downsize at the end of life. With multiple locations, higher costs can be avoided by migrating, and capital gains can be realized either by downsizing or by moving to a less expensive location. As a result, house price increases hurt renters less and benefit owners more when there are multiple locations. This highlights the importance of accounting for geography when studying shocks that affect house prices.

Related Literature This paper contributes to a growing literature that uses dynamic spatial models to analyze questions for which space and forward-looking decisions are salient.[2] Solving a general equilibrium model where agents make both forward-looking location and saving choices presents a major computational challenge. The reason is that agents' choices depend on the decisions of agents in all other locations and time periods. This implies a large state space when there are many locations and a non-degenerate wealth distribution. To avoid this challenge, many dynamic spatial models abstract from saving and assume that agents live hand-to-mouth. Important examples include Artuç et al. (2010), Desmet et al. (2018), Caliendo et al. (2019), Eckert and Kleineberg (2019), Allen and Donaldson (2020), Zerecero (2021), Giannone (2022), and Almagro and Domínguez-Iino (2025). An alternative approach is to allow borrowing and saving, but make assumptions about space that imply it is unnecessary to keep track of the wealth distribution in every location. Examples of this approach include Lyon and Waugh (2019) and Bilal and Rossi-Hansberg (2021).[3] Other papers have two types of agents, one that makes forward-looking location choices but cannot save, and another that makes forward-looking saving choices but cannot move. These models are

[2]See Desmet and Parro (2025) for a recent survey of this literature.

[3]Bilal and Rossi-Hansberg (2023) assume migration is costless. Lyon and Waugh (2019) assume locations are infinitesimally small, and that agents who move are randomly assigned a new location.

tractable because the wealth distribution in each location is degenerate. Examples include Kleinman et al. (2023), Cai et al. (2022), Vanhapelto (2022), and Bilal and Rossi-Hansberg (2023). Some recent papers develop dynamic spatial models in which agents jointly make forward-looking location and saving decisions. Dvorkin (2023) uses a model in which agents face multiplicative risk and no borrowing limit, as in Merton (1969). This model is tractable because, even though there is a non-degenerate wealth distribution, location and saving decisions are independent of wealth. Crews (2023) develops a model where agents make forward-looking location and human capital investment choices.

Relative to this literature, the main contribution of this paper is to develop a model in which agents jointly choose location, liquid wealth, and illiquid housing. Accounting for homeownership in a realistic way is essential for analyzing the distributional effects of local shocks, since owner-occupied housing is one of the most important channels through which individual welfare is linked to local economic conditions. In turn, a realistic model of homeownership requires lifecycle dynamics, uninsurable income risk, liquid wealth, and housing frictions, all of which are included in my model. Most closely related to this paper is Giannone et al. (2025). In contemporaneous work, they use a spatial model with homeownership to study the role of credit constraints in shaping migration patterns.[4] Incorporating homeownership is challenging because the state space is large and the household's problem is nonconvex due to fixed housing adjustment costs. I overcome these challenges by using the continuous-time computational methods recently developed by Achdou et al. (2022). These methods are much more efficient than traditional discrete-time methods, especially in the presence of nonconvexities.

My work is also closely related to a literature that explores the distributional effects of uneven local shocks. Important papers include Moretti (2013), Yagan (2014), Diamond (2016), Eckert (2019), Ganong and Shoag (2017), Fogli and Guerrieri (2019), and Notowidigdo (2020). I contribute to this literature by accounting for lifecycle dynamics and homeownership. Since housing is one of the most important channels through which welfare is tied to local economic conditions, and the effects of a local shock vary depending on age and wealth, this enables me to obtain more accurate and disaggregated estimates. An especially relevant paper is Hornbeck and Moretti (2024), who empirically estimate the effects of local productivity shocks on earnings and housing costs for renters and homeowners. Relative to that paper, an advantage of my structural model is that it can be used to compute welfare effects and counterfactuals.

Another set of related papers use incomplete market models with housing to study the welfare effects of house price changes. Examples include Fernandez-Villaverde and Krueger (2011), Berger et al. (2018), Kaplan et al. (2020), and Jones et al. (2022). Also related is Kiyotaki et al. (2011), who study the distributional effects of an aggregate house price shock. I contribute to this literature by accounting for economic geography. This makes it possible to quantify the role of housing in

[4] Since the first draft of this paper was circulated, some other dynamic spatial models with homeownership have been developed. Examples include Luccioletti (2024) and Sun (2025). Greaney et al. (2025) extend the model developed in this paper to include commuting.

shaping the welfare effects of local shocks. Since many shocks that affect house prices have heterogenous impacts across locations, it also enables a more complete understanding of the welfare effects of house price changes. I show that the welfare effects of an aggregate house price shock are substantially different than one whose effects vary across space.

The remainder of the paper is organized as follows. In Section 2, I develop the theoretical framework. Section 3 describes how the model is estimated. In Section 4, I use the model to quantitatively analyze the distributional effects of spatially heterogeneous productivity shocks. Section 5 concludes.

2 Model

In this section, I describe the theoretical framework. The model contains two primary components. The first is geography. There are a large number of locations where households live and work. Locations differ by wage, house price, and amenities. Households can migrate across locations, but must pay a moving cost to do so. Migration costs are necessary for quantifying the effects of local shocks because migration is a potentially important adjustment mechanism, as well as a channel through which the effects of local shocks propagate through space. This component is based on standard quantitative spatial models, reviewed by Redding and Rossi-Hansberg (2017).

The second component is homeownership. Most households own a home, and owner-occupied housing accounts for the majority of the typical household's wealth. Local shocks cause house price changes that affect homeowners' wealth. In order to capture the welfare effects of house price changes, it is necessary to include features that affect housing investment decisions. These include lifecycle dynamics, idiosyncratic risk, tenure choice, and housing illiquidity. Each of these ingredients influence how house price changes affect different types of households. For example, the welfare effect of a house price change depends on how much a household has invested in housing, whether it expects to be a net buyer or seller of housing in the future, and the insurance value of its housing wealth. These ingredients are also necessary to match important features of housing investment behavior, such as the fact that homeownership rates and the fraction of wealth invested in housing varies substantially over the lifecycle. This component is based on standard incomplete market models with homeownership, reviewed by Piazzesi and Schneider (2016).

2.1 Environment

Time t is continuous. The economy is populated by a unit mass of finitely-lived households. Households can live in one of I locations, which are indexed by i. There are two types of goods: a consumption good c and housing. The consumption good is freely traded and serves as the numeraire. Housing is nontradeable across space, and so its price differs by location. Housing can be either rented or owned. House prices are denoted by p_{it} and rental prices by p^r_{it}. All markets are competitive, and there is no aggregate uncertainty.

2.2 Households

Preferences Households derive utility from location-specific amenities A_i, consumption c, and housing services $\mathbf{h}$ according to the flow utility function[5]

$$u(A, c, \mathbf{h}) = \log(A\mathcal{C}(c, \mathbf{h}))$$
$$\mathcal{C}(c, \mathbf{h}) = c^{1-\eta}\mathbf{h}^{\eta}.$$

I refer to the Cobb-Douglas aggregator $\mathcal{C}$, which combines the consumption good c and housing services $\mathbf{h}$, as the "consumption aggregate." Housing services can be obtained either by renting or owning. Let h^r denote the quantity of housing a household rents and h the quantity it owns. Rented housing translates one-for-one into housing services, while each unit of owner-occupied housing provides χ units of services:

$$\mathbf{h}(h^r, h) = h^r + \chi h.$$

The homeownership preference parameter χ captures non-pecuniary benefits of owning housing. Households discount the future at rate ρ and do not have any bequest motive.

Demographics Age is denoted by $j \in [0, J]$, where J is the maximum possible age. Age-j households die at rate $\psi(j)$. When a household dies, it is replaced by a new one. At ages $\mathcal{J}^s = \{0, 1, ..., J-1\}$, which I refer to as "shock ages," households experience idiosyncratic shocks and have the option to move.[6] To be clear, time is continuous—I simply assume that shocks and adjustment opportunities occur at deterministic time intervals. Although individual households experience shocks at discrete intervals, a constant flow of households reaches shock ages at every instant. Hence the densities of all state variables, including those that are adjusted only at shock ages, evolve continuously over time.

Continuous-time spatial models typically assume that shocks and adjustment opportunities arrive according to a Poisson process; in contrast, this paper assumes that they occur at deterministic intervals.[7] This timing structure mirrors the implicit assumption of discrete-time spatial models, where households are allowed to relocate only at the model's frequency. One advantage of the shock-age assumption is that it avoids involuntary immobility driven purely by timing risk. Under Poisson arrival, some households with strong incentives to relocate are prohibited from doing so.

[5] A convenient feature of logarithmic utility is that it allows a closed-form expression for welfare effects.

[6] I assume J is an integer. It is not necessary that shock ages be evenly spaced over the lifecycle; all that is required is that there are a finite number whose timing is known with certainty. In my calibration, shock ages occur annually.

[7] As is standard in quantitative spatial models, I assume that idiosyncratic location preferences have an extreme value distribution. In continuous time, allowing agents to migrate at any time under this assumption would imply continuous migration. For this reason, continuous-time spatial models must restrict the timing of location choices.

With shock ages, households regularly receive the opportunity to move with certainty. As a result, migration choices reflect economic forces—such as migration costs, housing frictions, and lifecycle considerations—rather than random arrival of moving opportunities.

Unlike Poisson arrival, the shock-age assumption implies that agents' decisions depend on the time remaining until the next shock age. This would introduce an additional state variable if agents were infinitely-lived. In lifecycle models, age—which is already a state variable—is a sufficient statistic for time until next shock age. As a result, the shock-age assumption does not expand the state space in this model.

Earnings Households work from birth until the retirement age J^{ret}. They receive a stochastic stream of labor endowments l, which they supply inelastically at wage w_{it} in their location. Labor endowments include both a deterministic lifecycle component $\bar{l}(j)$ and an idiosyncratic component ℓ:

$$\log l = \bar{l}(j) + \ell.$$

ℓ follows an AR(1) process with innovations drawn at shock ages prior to retirement:

$$\ell_{j+1} = \begin{cases} \theta\ell_j + \zeta_{j+1} & \text{if } j < J^{\text{ret}} \\ \ell_j & \text{if } j \geq J^{\text{ret}} \end{cases}, \quad j \in \mathcal{J}^s$$
$$\zeta_j \sim N(0, \sigma^2).$$

Retired households receive a pension benefit $\flat(\ell)$, which is a function of their labor endowment at retirement. For notational convenience, I denote income from earnings and pensions by

$$y_t(i, \ell, j) = \begin{cases} w_{it} l & \text{if } j < J^{\text{ret}} \\ \flat(\ell) & \text{if } j \geq J^{\text{ret}}. \end{cases}$$

Migration Migration between locations is only allowed at shock ages. Prior to making migration decisions, households draw an idiosyncratic location preference ε_i for each location from a Gumbel distribution with scale parameter ν. After observing both earning and location preference shocks, they choose their location. There is a migration cost $\kappa_{ii'}(j)$ required to move from i to i', which may vary both by origin/destination and age. Both location preferences and migration costs are denominated in terms of utility.

The parameter ν determines the responsiveness of population allocations to local prices and amenities. The higher is ν, the more important are idiosyncratic location preferences for location decisions, and the less responsive are migration flows to changes in prices or amenities. From now on, I refer to ν as the "migration elasticity parameter."

Housing Households can only rent or own housing in the location where they live, and they are prohibited from renting and owning simultaneously. Rented house size is restricted to the set $\mathbb{H}^r$, and owner-occupied house size is restricted to the set $\mathbb{H}^o$. Tenure status and owner-occupied housing quantity can only be changed at shock ages. Renters can adjust housing quantity at any time. There are proportional transaction costs required to adjust owner-occupied housing.[8] Whenever a household changes its location or owner-occupied house quantity, it must pay fraction f_s of its old owner-occupied house value and fraction f_b of its new owner-occupied house value. There are no transaction costs required to adjust rented housing. Housing depreciates at rate δ.

Wealth Households can borrow and save in liquid, risk-free assets at rate r.[9] Denote a household's total liquid wealth by b. There is a collateral constraint which restricts borrowing to no more than a fraction ϕ of housing wealth:

$$b \geq -\phi p_{it} h. \tag{1}$$

Total wealth a is the sum of liquid and housing wealth: $a \equiv b + p_{it}h$. Using this notation, the collateral constraint (1) can be written as

$$a \geq (1-\phi) p_{it} h.$$

There is a tax of rate τ_t levied on all income (earnings, pensions, and capital income). Homeowners are allowed to deduct mortgage interest from their taxable income. Between shock ages, wealth evolves according to the budget constraint

$$\dot{a} = (1-\tau_t)[y_t(i,\ell,j) + rb] - c - p^r_{it}h^r - (\delta p_{it} - \dot{p}_{it})h. \tag{2}$$

The first term is after-tax income. The second and third terms are consumption and rental expenditures, respectively. The final term captures depreciation and capital gains on owner-occupied housing. Substituting the definition of a, (2) can be written as

$$\dot{a} = (1-\tau_t)[y_t(i,\ell,j) + ra] - c - p^r_{it}h^r - p^u_{it}h \tag{3}$$

where p^u_{it} is the user cost of housing for owners:

$$p^u_{it} = [(1-\tau_t)r + \delta]p_{it} - \dot{p}_{it}. \tag{4}$$

At shock ages, wealth may change discontinuously due to owner-occupied house size or location adjustments. After adjusting owner-occupied housing from h to h', a household with pre-

[8] In the United States, realtor fees are typically set as a fraction of house value. The macro-housing literature therefore typically assumes proportional transaction fees—see for example Davis and Van Nieuwerburgh (2015) and Piazzesi and Schneider (2016).

[9] As described in Section 2.4, these assets include both a risk-free bond and real estate investment trusts. All liquid assets have return r, so only total liquid wealth b is needed to characterize the household's problem.

adjustment wealth a has wealth

$$a_t^{\mathrm{h}}(a,h,i;h') = \begin{cases} a & \text{if } h' = h \\ a + (1-f_s)p_{it}h - (1+f_b)p_{it}h' & \text{otherwise.} \end{cases}$$

Since households are required to sell any owner-occupied housing when changing locations, wealth and owner-occupied housing immediately after moving from i to i' are

$$a_t^{\mathrm{m}}(a,h,i;i') = \begin{cases} a & \text{if } i' = i \\ a + (1-f_s)p_{it}h & \text{otherwise} \end{cases}$$
$$h^{\mathrm{m}}(h,i;i') = \begin{cases} h & \text{if } i' = i \\ 0 & \text{otherwise.} \end{cases}$$

For notational economy, I denote the set of permitted rental house sizes for a household with owner-occupied house size h by

$$\mathbb{H}^r(h) = \begin{cases} 0 & \text{if } h > 0 \\ \mathbb{H}^r & \text{if } h = 0. \end{cases}$$

At shock ages, households must choose an owner-occupied house size such that the collateral constraint is satisfied after purchase. The set of feasible owner-occupied house sizes for a household with wealth a, owner-occupied house size h, and location i is

$$\mathbb{H}_t^o(a,h,i) = \{h' \in \mathbb{H}^o : a_t^{\mathrm{h}}(a,h,i;h') \geq (1-\phi)p_{it}h'\}.$$

Value Function The household's state variables are wealth a, owner-occupied house size h, location i, idiosyncratic labor endowment ℓ, and age j. Since there is no bequest motive, the value function at the maximum age is $V_t(a,h,i,\ell,J) = 0$. For notational convenience, denote the vector of state variables by $\Omega \equiv (a,h,i,\ell,j)$. Between shock ages ($j \notin \mathcal{J}^s$), the value function satisfies the Hamilton-Jacobi-Bellman (HJB) equation

$$[\rho + \psi(j)]V_t(\Omega) = \max_{c,h^r \in \mathbb{H}^r(h)} u(A_i, c, \mathbf{h}(h^r,h)) + \partial_a V_t(\Omega)\dot{a}_t(\Omega; c, h^r) + \partial_j V_t(\Omega) + \partial_t V_t(\Omega) \tag{5}$$
$$\text{s.t.} \quad \dot{a}_t(\Omega; c, h^r) \geq (1-\phi)\dot{p}_{it}h \ \text{ if } a = (1-\phi)p_{it}h.$$

The left-hand side of (5) is the discounted value function. The first term on the right-hand side is optimally chosen flow utility. The second term captures changes in the value function as the household accumulates or decumulates wealth. The final two terms account for aging and the passage of calendar time. Lastly, the constraint ensures that the collateral limit is obeyed.

I now characterize the household's problem at shock ages by describing each shock-age event in turn. The final decision that shock-age households make is owner-occupied house size. The value of the optimal owner-occupied house size choice for a household with state Ω is

$$V_t^{\mathrm{h}}(\Omega) = \lim_{\iota \downarrow 0} V_{t+\iota}(a_t^{\mathrm{h}}(\Omega; h_t^*(\Omega)), h_t^*(\Omega), i, \ell, j+\iota) \tag{6}$$
$$h_t^*(\Omega) = \operatorname*{argmax}_{h' \in \mathbb{H}_t^o(\Omega)} \lim_{\iota \downarrow 0} V_{t+\iota}(a_t^{\mathrm{h}}(\Omega; h'), h', i, \ell, j+\iota).$$

$V_t^{\mathrm{h}}(\Omega)$ is the value function immediately after making the optimal house choice $h_t^*(\Omega)$ and paying the associated transaction costs. Prior to choosing h, shock-age households choose their location. In Appendix B.2, I show that the expected value of the optimal location choice is

$$V_t^{\mathrm{m}}(\Omega) = \nu \log \sum_{i'} \exp(V_t^{\mathrm{h}}(a_t^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j))^{1/\nu}. \tag{7}$$

Finally, the value function at shock ages is the expected value after integrating over idiosyncratic labor endowment shocks:

$$V_t(\Omega) = \begin{cases} \int V_t^{\mathrm{m}}(a, h, i, \ell', j) f(\ell'|\ell) \mathrm{d}\ell' & \text{if } j < J^{\mathrm{ret}} \\ V_t^{\mathrm{m}}(\Omega) & \text{otherwise} \end{cases} \tag{8}$$

where $f(\ell'|\ell)$ is the conditional probability density of ℓ'.

Density of State Variables Newborn households begin life with no wealth or housing, draw their initial location from a population distribution $P_t^0(i)$, and their initial idiosyncratic labor endowment from the invariant distribution. Between shock ages, the density of state variables $g_t(\Omega)$ satisfies the Kolmogorov Forward (KF) equation

$$\partial_t g_t(\Omega) = -\partial_a[\dot{a}_t(\Omega) g_t(\Omega)] - \partial_j g_t(\Omega) - \psi(j) g_t(\Omega) \tag{9}$$

where $\dot{a}_t(\Omega)$ is the optimally chosen drift of wealth.

The first event that occurs at shock ages is the realization of idiosyncratic labor endowment shocks. The density of state variables immediately after labor endowment shocks occur is

$$g_t^{\mathrm{l}}(\Omega) = \begin{cases} \int \lim_{\iota \downarrow 0} g_{t-\iota}(a, h, i, \ell', j-\iota) f(\ell|\ell') \mathrm{d}\ell' & \text{if } j < J^{\mathrm{ret}} \\ \lim_{\iota \downarrow 0} g_{t-\iota}(a, h, i, \ell, j-\iota) & \text{otherwise.} \end{cases} \tag{10}$$

After labor endowment shocks are realized, households draw idiosyncratic location preferences and choose their location. The density of state variables immediately after locations are chosen is

$$g_t^{\mathrm{m}}(\Omega) = \sum_{i'} \int \int \mathbf{1}_t^{\mathrm{m}}(a', h', i'; a, h, i) \pi_t(a', h', i', \ell, j; i) g_t^{\mathrm{l}}(a', h', i', \ell, j) \mathrm{d}a' \mathrm{d}h' \tag{11}$$

where

$$\mathbf{1}_t^{\mathrm{m}}(a',h',i';a,h,i) = \begin{cases} 1 & \text{if } a_t^{\mathrm{m}}(a',h',i';i) = a \text{ and } h^{\mathrm{m}}(h',i';i) = h \\ 0 & \text{otherwise} \end{cases}$$

is an indicator for whether a household with states (a',h',i') has states (a,h) after moving to i and

$$\pi_t(\Omega;i') = \frac{\exp(V_t^{\mathrm{h}}(a_t^{\mathrm{m}}(\Omega;i'),h^{\mathrm{m}}(\Omega;i'),i',\ell,j) - \kappa_{ii'}(j))^{1/\nu}}{\sum_{i''}\exp(V_t^{\mathrm{h}}(a_t^{\mathrm{m}}(\Omega;i''),h^{\mathrm{m}}(\Omega;i''),i'',\ell,j) - \kappa_{ii''}(j))^{1/\nu}} \tag{12}$$

is the fraction of households with state Ω who choose i'.[10] Finally, the density of state variables at shock ages is the density immediately after owner-occupied house size choices have been made:

$$g_t(\Omega) = \int\int \mathbf{1}_t^{\mathrm{h}}(a',h',i,\ell,j;a,h)g_t^{\mathrm{m}}(a',h',i,\ell,j)\mathrm{d}a'\mathrm{d}h' \tag{13}$$

where

$$\mathbf{1}_t^{\mathrm{h}}(a',h',i,\ell,j;a,h) = \begin{cases} 1 & \text{if } a_t^{\mathrm{h}}(a',h',i;h_t^*(a',h',i,\ell,j)) = a \text{ and } h_t^*(a',h',i,\ell,j) = h \\ 0 & \text{otherwise} \end{cases}$$

is an indicator for whether a household with states (a',h',i,ℓ,j) has states (a,h) after choosing owner-occupied housing.

2.3 Production

In each location, a representative firm produces the numeraire good using the technology

$$Y_{it} = Z_{it}L_{it} \tag{14}$$

where

$$L_{it} = \int\int\int\int \exp(\bar{l}(j)+\ell)g_t(\Omega)\mathrm{d}a\mathrm{d}h\mathrm{d}\ell\mathrm{d}j$$

is local labor supply. Since factor prices are set competitively, wages equal local productivity:[11]

$$w_{it} = Z_{it}. \tag{15}$$

Understanding the welfare effects of shocks to local productivities Z_i is the focus of the quantitative analysis in Section 4.

[10] This expression is derived in Appendix B.1. The fact that location preferences have a Gumbel distribution makes it possible to obtain a closed form expression for these probabilities.

[11] In the baseline model, there are no agglomeration or congestion forces in the labor market. In Appendix D.1, I extend the model to consider such forces.

A representative firm produces housing in each location using the technology

$$Y_{it}^h = \overline{\xi} M_{it}^{1-\xi} F_{it}^{\xi} \tag{16}$$

where $\overline{\xi} \equiv (1-\xi)^{1-\xi}\xi^{\xi}$, M is inputs of the numeraire good, and F is permits. I assume that permits are supplied by the government (described below) at price p_{it}^F. The unit cost of construction is therefore $(p_{it}^F)^{\xi}$. I further assume that construction is irreversible.

Total housing demand in location i is

$$H_{it}^d = \int\int\int\int [h_t^r(\Omega) + h] g_t(\Omega) \mathrm{d}a\mathrm{d}h\mathrm{d}\ell\mathrm{d}j. \tag{17}$$

The construction firm chooses construction $Y_{it}^h = 0$ if the construction cost exceeds the house price ($p_{it} < (p_{it}^F)^{\xi}$), and chooses construction $Y_{it}^h = \infty$ if the house price exceeds the construction cost ($p_{it} > (p_{it}^F)^{\xi}$). If $p_{it} = (p_{it}^F)^{\xi}$, it is indifferent how much it produces. In this case, to clear the market it produces the amount demanded by households. Since housing depreciates at rate δ, housing supplies H_{it}^s satisfy

$$\begin{aligned} \dot{H}_{it}^s &= Y_{it}^h - \delta H_{it}^s \\ Y_{it}^h &= \begin{cases} 0 & \text{if } p_{it} < (p_{it}^F)^{\xi} \\ \max\{\dot{H}_{it}^d + \delta H_{it}^d, 0\} & \text{if } p_{it} = (p_{it}^F)^{\xi} \\ \infty & \text{if } p_{it} > (p_{it}^F)^{\xi}. \end{cases} \end{aligned} \tag{18}$$

House market clearing requires that housing demand equals supply: $H_{it}^d = H_{it}^s = H_{it}$. Note that in a stationary equilibrium, $Y_i^h = \delta H_i > 0$ and $p_i = (p_i^F)^{\xi}$.

If housing demand is sufficiently strong that strictly positive construction is demanded ($\dot{H}_{it} > -\delta H_{it}$), house prices equal construction costs: $p_{it} = (p_{it}^F)^{\xi}$. Otherwise, house prices fall below construction costs and are set so that housing demand equals supply with no construction: $p_{it} < (p_{it}^F)^{\xi}$ and $\dot{H}_{it} = -\delta H_{it}$. In equilibrium, house prices never exceed construction costs: $p_{it} \leq (p_{it}^F)^{\xi}$. House market clearing can be summarized by the following complementary slackness condition:

$$\begin{aligned} p_{it} = (p_{it}^F)^{\xi} &\text{ and } Y_{it}^h > 0 \\ \text{or } p_{it} < (p_{it}^F)^{\xi} &\text{ and } Y_{it}^h = 0. \end{aligned} \tag{19}$$

2.4 Rental and Capital Markets

The risk-free return r is set exogenously. In each location, there is a perfectly competitive real estate investment trust (REIT) that owns the local rental housing. The rent charged by REITs is determined by a no-arbitrage condition that ensures the return on rental properties is r:

$$p^r_{it} = (r+\delta)p_{it} - \dot{p}_{it}. \tag{20}$$

REITs are owned by households. Households invest in their local REIT, a "global REIT" that holds all the local REIT shares not owned by local residents, and a risk-free bond. All have return r, so households are indifferent (*ex ante*) what fraction of their liquid wealth is invested in each. I assume that fraction f^{lr} of a household's liquid assets are invested in its local REIT.

Denote local rental housing stocks and liquid assets, respectively, by

$$H^r_{it} = \int\int\int\int h^r_t(a,h,i,\ell,j)g_t(a,h,i,\ell,j)\mathrm{d}a\mathrm{d}h\mathrm{d}\ell\mathrm{d}j$$
$$B^+_{it} = \int\int\int\int \max\{a - p_{it}h, 0\}g_t(a,h,i,\ell,j)\mathrm{d}a\mathrm{d}h\mathrm{d}\ell\mathrm{d}j.$$

The global REIT's investment in i's local REIT is $p_{it}H^r_{it} - f^{\mathrm{lr}}B^+_{it}$. The fraction of liquid assets invested in the global REIT, f^{gr}_t, is chosen so that all rental housing is owned by households:[12]

$$f^{\mathrm{gr}}_t = \frac{\sum_i (p_{it}H^r_{it} - f^{\mathrm{lr}}B^+_{it})}{\sum_i B^+_{it}}.$$

The remainder of liquid wealth is held in the bond.

Although the forward-looking return on rental properties is always r, an unanticipated shock may cause REITs to experience capital gains or losses, which in turn affect households' wealth. Consider a shock that causes house prices to change from p^{ss}_{i0} in the initial steady state to p_{i0}.[13] Denote the proportional change in i's local REIT value by $\hat{v}^{\mathrm{lr}}_i = p_{i0}/p^{\mathrm{ss}}_{i0}$. The proportional change in the global REIT value is

$$\hat{v}^{\mathrm{gr}} = \sum_i f^{\mathrm{gr,ss}}_{i0}\hat{v}^{\mathrm{lr}}_i$$

where

$$f^{\mathrm{gr,ss}}_{i0} = \frac{p^{\mathrm{ss}}_{i0}H^{r,\mathrm{ss}}_{i0} - f^{\mathrm{lr}}B^{+,\mathrm{ss}}_{i0}}{\sum_{i'}(p^{\mathrm{ss}}_{i'0}H^{r,\mathrm{ss}}_{i'0} - f^{\mathrm{lr}}B^{+,\mathrm{ss}}_{i'0})}$$

is the fraction of the global REIT's assets that are invested in local REIT i in the initial steady state. The wealth change experienced by a household with states (a,h,i) when the shock occurs is

$$\mathrm{d}a(a,h,i) = [(\hat{v}^{\mathrm{lr}}_i - 1)f^{\mathrm{lr}} + (\hat{v}^{\mathrm{gr}} - 1)f^{\mathrm{gr,ss}}_0]\max\{a - p^{\mathrm{ss}}_{i0}h, 0\} + (p_{i0} - p^{\mathrm{ss}}_{i0})h.$$

The first term of this expression is the sum of capital gains from local and global REIT holdings. The

[12] I verify *ex-post* that the global REIT invests a strictly positive amount in each local REIT, and that the total real estate investment share satisfies $f^{\mathrm{lr}} + f^{\mathrm{gr}}_t < 1$.

[13] I denote the value of a variable x in the initial steady state by x^{ss}_0, and the value of x when an unexpected shock occurs ($t = 0$) by x_0. For variables that can change discontinuously (e.g. house prices), x^{ss}_0 may not equal x_0.

final term is capital gains on owner-occupied housing.

2.5 Government

I assume that the government sets permit prices according to

$$p_i^F = \overline{p}_i^F \left(\frac{N_i g^N}{N_{i1970}} \right)^{\lambda_i} \tag{21}$$

where N_{i1970} is the local population share in 1970, g^N is aggregate population change since 1970, and N_i is the long-run local population share[14]

$$N_i = \int\int\int\int g(a, h, i, \ell, j) \mathrm{d}a\mathrm{d}h\mathrm{d}\ell\mathrm{d}j.$$

$\overline{p}_i^F$ are exogenous permit price shifters and λ_i are location-specific permit price elasticities. Equations (18) and (19) imply that the long-run elasticity of house price with respect to population is

$$\frac{\partial \log p_i}{\partial \log N_i} = \xi \lambda_i.$$

I set $\lambda_i = \text{ela}_i / \xi$ to match the long-run house price elasticities ela_i estimated by Saiz (2010). In Appendix C, I derive equation (21) directly from Saiz's estimating equation.

My assumptions regarding housing construction and permit prices allow the model to match important features of local housing markets while maintaining tractability. Specifically, the model matches empirical estimates of house price elasticities, which vary substantially across cities. Furthermore, both rental and house prices jump in response to local shocks.[15] Finally, as long as construction is positive, house prices are determined by (21), which is well-behaved. This makes it possible to compute transition dynamics even though there are many price paths to solve for.

The government receives revenue from permit sales, income taxes, and taxing estates. Accidental bequests are taxed at rate 100%, so there are no inheritances. Government revenue is used to fund the pension system and other government spending G, which is spent in a way that does not affect household utility. I assume that G is time-invariant, and that the income tax rate τ_t adjusts as needed so that the government always runs a balanced budget.

[14] In a stationary equilibrium, local populations equal long-run local populations. On the transition path after an unexpected shock, the long-run local population is the one that the economy will eventually converge to.

[15] This would not be the case if permit prices were a function of current population: $p_{it} = \overline{p}_i^F (N_{it} g^N / N_{i1970})^{\lambda_i}$. In response to a positive local productivity shock, N_{it} increases gradually. In this case, p_{it} would increase gradually as well. The rental price expression (20) then implies that rents would jump *down* on impact. Setting permit prices as a function of long-run populations implies that both house and rental prices jump up after a positive shock.

2.6 Definition of Equilibrium

An equilibrium is household value and policy functions $V_t(\Omega)$, $c_t(\Omega)$, $h^r_t(\Omega)$, and $h^*_t(\Omega)$; density of state variables $g_t(\Omega)$; prices w_{it}, p_{it}, p^r_{it}, and p^F_{it}; and tax rate τ_t such that, given an interest rate r, pension benefit $\flat(\ell)$, and other government spending G:

1. Households optimize: the value and policy functions satisfy (5)-(8).
2. The density of state variables is consistent with the policy functions: (9)-(13) are satisfied.
3. Wages are set according to (15).
4. The housing market clears: housing demand (17) equals housing supply (18).
5. Rental prices satisfy (20).
6. Permit prices are set according to (21).
7. The government's budget balances.

A stationary equilibrium is one in which all equilibrium objects are time-invariant.

2.7 Solving the Model

Solving a quantitative spatial model with realistic homeownership presents several computational challenges. First, the state space is large. Second, the discrete choice of whether to rent or own, as well as fixed housing adjustment costs, induce kinks in the value function. This poses serious computational challenges in discrete time.[16] Third, exactly matching geographic data, which is standard in the quantitative spatial economics literature, requires estimating a large number of parameters. In particular, location-specific amenities and productivities must be chosen to match population shares and prices.[17] Finally, since households make dynamic migration and housing decisions, studying the welfare effects of local shocks requires computing transition dynamics. For these reasons, it is imperative that the model be solved efficiently.

I overcome these computational challenges by setting the model in continuous time and using numerical techniques recently developed by Achdou et al. (2022). These techniques leverage four advantages of continuous time that are relevant in this setting.[18] First, the first-order conditions of the household's problem hold with equality and are sufficient at every point in the state space.

[16]For example, since the first order conditions may not be sufficient, the endogenous gridpoint method (Carroll, 2006) cannot be used.

[17]In standard quantitative spatial models, there is no wealth heterogeneity and so amenities can either be read directly from data or dealt with using hat algebra Dekle et al. (2008) or dynamic hat algebra Caliendo et al. (2019) techniques. In my model there is a non-degenerate wealth distribution in each location, so this is not possible.

[18]See Achdou et al. (2022) for a more detailed discussion of the benefits of continuous time. In their Appendix F.1, Achdou et al. (2022) compare computation times using continuous-time versus discrete-time methods.

In a discrete-time version of the model, the first-order conditions would be inequalities due to occasionally binding borrowing constraints. Furthermore, they would not be sufficient due to nonconvexities. The second advantage is that the first-order conditions are "static"—that is, they only involve contemporaneous variables. In a discrete-time version of the model, the first-order conditions would relate marginal utility in one period with the derivative of the value function in the next period. As a result, discrete-time first-order conditions only define optimal policies *implicitly*, and therefore require costly root-finding operations to solve. The third advantage is that, since wealth does not jump discontinuously between shock ages, the HJB equation (5) and KF equation (9) can be represented as sparse matrix equations.[19] These can be solved extremely efficiently, and are also memory-efficient. Finally, the sparse matrix equation that represents the HJB equation is the transpose of the KF equation. As a result, once the HJB equation has been solved, the density of state variables can be computed at virtually no cost.

The first two advantages imply that, given the value function, the household's optimal policy functions can be written in closed form.[20] Given the value function at a shock age, the first three advantages imply that the HJB equation (5) can be efficiently solved backward in time to the previous shock age. At shock ages, shocks and discrete choices occur one at a time. The fact that location preferences have a Gumbel distribution implies that there are closed-form expressions for the value function at shock ages (equations 6-8). As a result, the shock-age value function can be computed efficiently as well. Given the lifecycle structure of the household's problem, there is no need to iterate on the value function. Instead, the equations (5)-(8) can be efficiently solved backward through the household's life to obtain the value and policy functions. Given these, the fourth advantage of continuous time implies that the density of state variables can be obtained efficiently as well.

Using continuous-time computational methods allows me to solve the model more efficiently than would be possible in discrete time, as well as handle the nonconvexities implied by illiquid housing without difficulty. The details of the numerical algorithm, as well as the time required to solve and estimate the model, are described in Appendix E.

3 Estimation

In this section, I describe how I estimate the model. My goal is to choose parameters so that the model matches important features of the spatial data and data on household saving and homeownership. The unit of time used for estimation is one year. The locations are the 50 most populous Metropolitan Statistical Areas (MSAs) in the United States as of the year 2000.

[19] Since I limit all idiosyncratic shocks and discrete choices to shock ages, these matrix equations are tridiagonal.

[20] More specifically, the policy functions can be written as closed-form expressions of the derivative of the value function. These are shown in Appendix E.1. Given the value function, the "upwinding" algorithm described by Achdou et al. (2022) computes the derivative of the value function at virtually no cost.

3.1 Data

The data sources and procedures for constructing the variables used for estimation are described below. Statistics for each city are displayed in Appendix Table 11.

3.1.1 Data Sources

Decennial Census My primary source of geographical data is the 5% public use sample of the 2000 Decennial Census, obtained from IPUMS (Ruggles et al., 2021). In order to be consistent with the model calibration, I keep only individuals between the ages of 20 and 99. I also exclude individuals who are living in group quarters. Results obtained using this dataset are computed using person weights.

American Community Survey I take migration data from the public use sample of the 2005 American Community Survey (ACS). This is the earliest year in which annual migration data at the MSA level is available. I also use data from the public use sample of the 2019 ACS to construct the productivity shock on which my quantitative analysis is based. These data were obtained from IPUMS. As with the Decennial Census data, I exclude individuals outside the ages of 20 and 99 or who live in group quarters. Results obtained using these datasets are computed using person weights.

Bureau of Economic Analysis Regional Accounts I get population in each city from the Bureau of Economic Analysis (BEA) Regional Accounts.[21] I use this data source because it maintains consistent MSA definitions back to 1970. Recall that populations in 1970 and 2000 are needed to compute house prices (see equation 21).

Survey of Consumer Finances The main source of data for household saving and homeownership decisions is the 2001 Survey of Consumer Finances (SCF). The SCF is available on a triennial basis, so these data are based on observations from 1999 to 2001. As with the Census data, I keep only people between the ages of 20 and 99. To be consistent with the model, I also keep only observations with non-negative wealth. Results obtained using this dataset are computed using household weights.

3.1.2 Data Construction

Wages I compute wages in each MSA using the 2000 Decennial Census. The sample of workers consists of individuals between the ages of 20 and 64 who are employed and worked at least 35 hours per week for 40 weeks during the past year. I exclude people who are self-employed, have

[21] The data I use come from Table CAINC1.

farm or business income, whose income is imputed, or have missing data on income, education, or demographics.

An individual's raw wage is their wage and salary income during the past year divided by weeks worked in the last year times usual hours worked per week. In order to control for selection on worker characteristics that my model abstracts from, wages are the residual of a regression that controls for observables that are important determinants of earnings. In particular, I estimate

$$\log w_n = \text{cons} + \mathbf{\Gamma X}_n + \epsilon_n$$

where w_n is the raw wage of person n, cons is a constant, $\mathbf{X}_n$ is a vector of observable characteristics, and ϵ_n is the residual. $\mathbf{X}$ contains age, age squared, a dummy for gender, a series of race dummies, and a series of education dummies. The race dummies are black and white (with all other categories the omitted group) and the education dummies are high school, some college, college, and post-graduate (with less than high school the omitted group). The wage in a location is the average residual wage $\exp(\epsilon_n)$ among workers in that location. Without loss of generality, I scale wages so that the population-weighted average wage across locations is normalized to 1.

House Prices I compute house prices using the 2000 Decennial Census. The house price in each MSA is the median house value reported by homeowners. I exclude values that are top-coded or imputed. Without loss of generality, I normalize the population-weighted average house price across locations to 1.

Migration Rates Migration rates are calculated using the 2005 ACS. The migration rate is the number of individuals who changed MSAs in the last year divided by the number who lived in an MSA both now and one year ago.

Household Portfolios I compute the wealth-income ratio, homeownership rate, housing wealth share, and local REIT investment share using the 2001 SCF. I define wealth as net worth (value of assets minus liabilities), and housing wealth as the value of owner-occupied housing.

3.2 Productivity Shocks

In Section 4, I quantify the distributional effects of spatially heterogeneous productivity shocks by computing the transition dynamics after a shock to local productivities Z_i. I consider two shocks.[22] The first is a shift-share variable that reflects heterogeneous local exposures to industry-specific productivity changes, similar to Bartik (1993). This shock has two main advantages. First, it is a

[22] Both shocks cause productivity Z_i to change (by different amounts) in every location simultaneously. An alternative would be to shock each location separately. I avoid this because it would require solving the model I times, which is prohibitively costly.

plausibly exogenous, empirically important source of local productivity variation. Second, there is an empirical literature that analyzes the effects of this type of shock, which allows me to validate the model by comparing model predictions with empirical estimates. I refer to this as the "Bartik" shock. The second shock is chosen to exactly match observed wage changes over the past two decades. In contrast to the Bartik shock, this shock is not separately identified from other forces that affect wages.[23] The motivation for studying this shock is to understand how observed local wage changes affect different types of households. I refer to this shock as the "Observed Wage" shock.

Denote the proportional change in location i's productivity due to the Bartik shock by $\hat{Z}_i^{\text{btk}}$, and due to the Observed Wage shock by $\hat{Z}_i^{\text{owg}}$. I construct the Bartik shock as follows. First, I compute employment shares f_{is} for every location i and 3-digit industry code s in the year 2000. The Bartik shock in i is the sum of national industry (residual) wage changes between 2000 and 2019, weighted by i's initial employment shares:

$$\hat{Z}_i^{\text{btk}} = \sum_s f_{is} \left(\frac{w_{s2019}}{w_{s2000}} \right).$$

The Observed Wage shock is the observed change in (residual) wages over the same time period:

$$\hat{Z}_i^{\text{owg}} = \frac{w_{i2019}}{w_{i2000}}.$$

In the baseline model (which has no agglomeration or congestion forces in production), the Observed Wage shock induces wage changes that exactly match the data.

I use the time period 2000-2019 to ensure full data availability and to limit the extent to which results are driven by extraordinary disruptions from the mid-2000s housing boom/bust episode and Covid-19 pandemic. In order to focus on distributional effects, I scale both productivity shocks so that the population-weighted average shock is 1.[24] I assume that productivity shocks occur unexpectedly when the economy is in its initial steady state and are permanent. Households have rational expectations, and so correctly anticipate all future prices when a shock occurs.

Table 1 shows correlations between the productivity shocks and local characteristics in the initial steady state. Both shocks are positively correlated with wages and house prices. As a result, they increase the spatial dispersion of wages, consistent with narratives of the "Great Divergence" (Moretti, 2012). Both productivity shocks are also positively correlated with average idiosyncratic labor endowment ℓ, and negatively correlated with homeownership rate and fraction of retirees.

The correlation between the Bartik and Observed Wage shocks is 0.47. The Observed Wage shock is substantially more disperse than the Bartik shock: the standard deviation of $\hat{Z}_i^{\text{owg}}$ is 5.33%, and the standard deviation of $\hat{Z}_i^{\text{btk}}$ is 0.53%. This is because the Bartik shock reflects only one component of

[23]This does not pose a problem for the quantitative exercise performed in Section 4—the model can be used to study *any* local productivity shock, including counterfactual shocks. However, it should be noted that this shock does not correspond to an exogenous productivity shock observed in the data.

[24]The model can be used to study shocks where average productivity changes. Rescaling does not meaningfully affect welfare elasticity estimates (equation 28), since aggregate effects are absorbed by the constant term. The rescaled shocks do have aggregate welfare effects. These are shown in Table 7, which reports both average and median welfare effects.

wage changes. Appendix Table 11 shows the productivity shock in every location.

Table 1: Productivity Shock Correlations

	Log Wage	Log House Price	Average ℓ	Fraction Own	Fraction 65+
$\log \hat{Z}^{\text{btk}}$	0.17	0.27	0.31	-0.11	-0.03
$\log \hat{Z}^{\text{owg}}$	0.21	0.37	0.20	-0.28	-0.06

Notes: This table shows across-location correlations between log productivity shock and various variables in the initial steady state. $\hat{Z}^{\text{btk}}$ is the proportional change in productivity from the "Bartik" productivity shock, and $\hat{Z}^{\text{owg}}$ is the proportional change in productivity from the "Observed Wage" productivity shock.

3.3 Estimation Procedure

In this section I describe my procedure for estimating the model parameters. The parameters related to geographic mobility are selected to match important features of the migration data, while the parameters related to saving and housing are chosen to match household wealth and portfolio data. As is standard in quantitative spatial models, I choose location-specific amenities, productivities, and permit price shifters to exactly match observed population shares, wages, and house prices in each location.

3.3.1 Externally Calibrated Parameters

I take many of the model's parameters either directly from data or existing literature. Table 2 provides a summary of the externally calibrated parameters.

Demographics Households live for at most $J = 80$ years, from age 20 to 100. They retire at age 65 ($J^{\text{ret}} = 45$). I take death rates $\psi(j)$ from Bell and Miller (2005). When a new household is born, it draws its initial location from the current population distribution.

Earning I set the persistence and volatility parameters of the idiosyncratic earning process equal to the values estimated by Floden and Lindé (2001). These are $\theta = 0.9136$ and $\sigma^2 = 0.0426$, respectively. The deterministic lifecycle component of earnings $\bar{l}(j)$ is taken from Hansen (1993). I scale labor endowments so that the median earning of working-age households ($60,899 in the 2001 SCF) in the initial steady state is 1. I assume that the pension benefit is proportional to idiosyncratic labor supply: $\flat(\ell) = \bar{\flat}\ell$. I choose $\bar{\flat}$ so that the ratio of average pension benefit to average earning in the initial steady state is 40%.

Housing I set the housing transaction costs equal to the median costs estimated by Gruber and Martin (2005). The selling cost is $f_s = 7\%$ of the house's value, and the buying cost is $f_b = 2.5\%$. I set the housing depreciation rate to $\delta = 2\%$ and the collateral constraint to $\phi = 0.8$, both of which are within the range of commonly used values discussed by Piazzesi and Schneider (2016).

REIT Investment I calibrate the local REIT investment share f^{lr} as follows. The SCF reports two categories of real estate investment besides primary residence: residential property excluding primary residence, and net equity in non-residential real estate. I refer to the sum of these two categories as "non-primary real estate." I set f^{lr} equal to average non-primary real estate divided by average financial assets. This yields a value of $f^{\mathrm{lr}} = 8.7\%$.[25] The global REIT investment share in the initial steady state is $f_0^{\mathrm{gr,ss}} = 15.6\%$ (recall from Section 2.4 that f_t^{gr} is chosen so that all rental housing is owned by households in the model).

Other I set the risk-free return to $r = 4\%$, which is the standard value mentioned by Davis and Van Nieuwerburgh (2015). This implies the steady-state house price-rent ratio is $1/(r+\delta) = 16.67$ (see equation 20), which is within the historical range of 11-19 reported by Piazzesi and Schneider (2016). Following Herkenhoff et al. (2018), the housing construction parameter ξ is set to 0.38. Total population change in the MSAs in my dataset between 1970 and 2000 is $g^N = 1.3976$. As discussed in Section 2.5, the location-specific permit price elasticities λ_i are chosen to exactly match the long-run house price elasticities estimated by Saiz (2010). Finally, I set other government spending G such that the government's budget balances in the initial steady state with income tax rate $\tau_0^{\mathrm{ss}} = 23.9\%$, which is the average labor income tax in 2000 estimated by McDaniel (2007).

3.3.2 Internally Estimated Parameters

The remaining parameters are internally estimated. These are the discount rate ρ, homeownership preference χ, housing preference weight η, migration costs $\kappa_{ii'}(j)$, migration elasticity parameter ν, permissible house size sets $\mathbb{H}^r$ and $\mathbb{H}^o$, local amenities A_i, local productivities Z_i, and local permit price shifters $\overline{p}_i^F$. All except ν are chosen so that the initial stationary equilibrium of the model matches cross-sectional data from (or as close as possible to) the year 2000. ν is chosen to match an empirical estimate of the long-run labor supply response to MSA-level productivity shocks.

Table 3 provides a summary of the internally estimated parameters. Since the number of estimated parameters equals the number of estimation targets, the model matches all targets to within numerical precision. Below I describe the data moment that is most important for identifying each parameter, but all parameters are estimated jointly. The details of the estimation procedure are described in Appendix E.3.

[25] Implicitly, I am assuming all non-primary real estate is held locally. If households hold non-primary real estate in locations other than their primary residence, this will overstate households' exposure to local real estate. Appendix Figure 10 shows non-primary real estate portfolio shares by net worth decile in the SCF. In Appendix D.4, I assess robustness to alternative assumptions regarding the spatial distribution of non-primary real estate investment.

Table 2: Externally Calibrated Parameters

Parameter	Description	Value	Source
Demographics			
J	Max age	80	Max lifespan 20-100
J^{ret}	Retirement age	45	Retire at 65
$\psi(j)$	Death rates		Bell and Miller (2005)
Earning			
θ	Persistence	0.9136	Floden and Lindé (2001)
σ^2	Volatility	0.0426	Floden and Lindé (2001)
$\bar{l}(j)$	Lifecycle		Hansen (1993)
$\flat(\ell)$	Pension	See text	Replacement ratio 40%
Housing			
f_s	Sell transaction cost	7%	Gruber and Martin (2005)
f_b	Buy transaction cost	2.5%	Gruber and Martin (2005)
δ	Depreciation rate	2%	Piazzesi and Schneider (2016)
ϕ	Collateral constraint	80%	Piazzesi and Schneider (2016)
REITs			
f^{lr}	Local REIT share	8.7%	SCF
f^{gr}	Global REIT share	15.6%	All housing owned by households
Other			
r	Interest rate	4%	Davis and Van Nieuwerburgh (2015)
ξ	Housing construction	0.38	Herkenhoff et al. (2018)
G	Government spending	0.125	See text
λ_i	Permit prices		Saiz (2010)

Notes: This table displays the parameter values which are either taken directly from data, come from existing literature, or are set to standard values. The unit of time used throughout the estimation procedure is one year.

Preferences The discount rate ρ and homeownership preference χ are chosen to match the median wealth-income ratio (2.35) and homeownership rate (71.7%) in the 2001 SCF, respectively. The housing preference weight η is chosen to match the housing expenditure share of 21% reported by Piazzesi and Schneider (2016). Consistent with their methodology, I set housing expenditure for owners equal to imputed rent.

Migration Cost I parameterize the migration cost as a linear function of age and distance:

$$\kappa_{ii'}(j) = \overline{\kappa} + \kappa^{\text{age}} j + \kappa^{\text{dst}} \text{dist}(i, i')$$

where $\text{dist}(i, i')$ is the distance between i and i', provided by Gardner and Hendrickson (2018). The intercept $\overline{\kappa}$ is chosen to match the overall migration rate (3.03%). The age coefficient κ^{age} is chosen to match the difference in migration rates between working-age (3.36%) and retired individuals (1.34%). Finally, the distance coefficient κ^{dst} is chosen to match the relationship between migration probabilities and distances observed in the data. Specifically, I estimate the gravity regression

$$\log \pi_{ii'} = \text{cons} + \mathrm{f}^{\mathrm{o}}_i + \mathrm{f}^{\mathrm{d}}_{i'} + \alpha^{\mathrm{dst}} \log \mathrm{dist}(i, i') + \epsilon_{ii'}$$

where $\pi_{ii'}$ is the fraction of households in i who choose i', cons is a constant, f^{o} are origin fixed effects, f^{d} are destination fixed effects, and ϵ is an error term. I choose κ^{dst} so that the model matches the estimated coefficient $\hat{\alpha}^{\mathrm{dst}} = -0.827$ from the data.

Allowing the migration cost to increase with age is necessary to quantitatively match the observed decline in migration rates over the lifecycle. The model does include some features that cause mobility to decrease with age. For example, younger households have longer to enjoy the benefits of moving to a more desirable location, and older households are more likely to own a home, which raises the cost of moving due to transaction costs. However, there are other important factors which are not included in the model, such as changing family size and local social networks that deepen over time. The parameter κ^{age} is a reduced-form way to capture such forces.

Migration Elasticity I estimate the migration elasticity parameter ν using indirect inference. Specifically, I choose ν to match the long-run effect of local productivity shocks on MSA-level employment estimated by Hornbeck and Moretti (2024, henceforth HM). HM estimate an instrumental variable regression of the form

$$\Delta \log L_i = \text{cons} + \alpha^{\mathrm{mig}} \Delta \log \hat{A}_i + \mathrm{f}^{\mathrm{reg}}_i + \epsilon_i \tag{22}$$

where $\Delta \log L_i$ is the change in log employment of city i between 1980 and 2010, cons is a constant, $\Delta \log \hat{A}_i$ is the predicted change in log productivity of city i between 1980 and 1990,[26] $\mathrm{f}^{\mathrm{reg}}$ are Census region fixed effects, and ϵ_i is the residual. They obtain the empirical estimate $\hat{\alpha}^{\mathrm{mig}} = 4.03$. I use this "identified moment" (Nakamura and Steinsson, 2018) to estimate ν. Specifically, I solve both the initial steady state and the steady state after the Bartik productivity shock in the model, conditional on a guess for ν (and the other estimated parameters). I then estimate the regression

$$\Delta \log L_i = \text{cons} + \alpha^{\mathrm{mig}} \log \hat{Z}^{\mathrm{btk}}_i + \epsilon_i \tag{23}$$

using the model-generated data. I select ν such that the estimated coefficient $\hat{\alpha}^{\mathrm{mig}}$ matches the auxiliary statistic of 4.03. Appendix D.6 examines the sensitivity of results to alternative migration elasticity targets from the literature.

House Sizes I parameterize the set of owner-occupied house sizes $\mathbb{H}^o$ as an evenly spaced grid with 10 nodes. The upper bound $\overline{h}^o$ is chosen so that the population-weighted average maximum house value across cities matches the 90th percentile of the owner-occupied house value distribution. The lower bound $\underline{h}^o$ is chosen so that the model matches the relationship between house price

[26]More specifically, the manufacturing revenue total factor productivity.

and homeownership rate observed in the data. Specifically, I estimate the regression

$$\text{homeownership rate}_i = \text{cons} + \alpha^{\text{h}} \log p_i + \epsilon_i^{\text{h}}. \tag{24}$$

In the data, there is a negative relationship between house price and homeownership rate: $\hat{\alpha}^{\text{h}} = -0.113$. I choose $\underline{h}^o$ to match this coefficient, which I refer to as the "homeownership elasticity." Intuitively, the higher is the minimum house size, the more the homeownership rate decreases in house price, as homeownership becomes less affordable in expensive cities. I estimate a minimum owner-occupied house size of 2.12 and a maximum size of 8.30.[27] I assume that renters can choose any house size less than or equal to the maximum owner-occupied house size: $\mathbb{H}^r = [0, 8.30]$.

Local Parameters The remaining internally estimated parameters are the location-specific amenities A_i, productivities Z_i, and permit price shifters $\overline{p}_i^F$. I choose these so that the initial stationary equilibrium matches population shares, wages, and house prices in each location. Without loss of generality, I normalize the average amenity value to 1. The location-specific parameters for each city are displayed in Appendix Table 12.

As is common in quantitative spatial models, estimating local productivities Z_i and permit price shifters $\overline{p}_i^F$ poses no computational difficulty. Given population allocations and housing demands in the initial steady state (where wages and house prices are those observed in data), the values of Z_i and $\overline{p}_i^F$ that exactly match wages and house prices can be backed out in closed-form using equations (15), (18), and (21).[28] In contrast, it is *not* trivial to estimate local amenities. There is no closed-form relationship between amenities and population allocations, so these cannot be backed out from data. Instead, the amenity values must be estimated numerically.[29] This is one reason why it is imperative that the model be solved efficiently.

3.4 Untargeted Moments

In this section, I evaluate the model's ability to match untargeted moments of interest. Given their importance for my analysis, I focus on statistics related to homeownership, housing investment, migration, spatial sorting, and the dynamic responses of the economy to local productivity shocks.

[27] Recall that both median working-age earning ($60,899 in the 2001 SCF) and population-weighted house price are normalized to 1 in the initial steady state. Hence the average minimum owner-occupied house value is $129,106, and the average maximum is $505,462.

[28] In fact, there is no need to recover the levels of these parameters for estimation or counterfactual analysis. Households do not need to know them to make decisions in the initial stationary equilibrium – all they care about are prices, which exactly match data. And given a guess for population and labor supply allocations on the transition path after a shock, wages and house prices can be computed using the hat algebra techniques of Dekle et al. (2008), which requires only knowledge of how Z_i and $\overline{p}_i^F$ *change* over time.

[29] The dynamic hat algebra techniques developed by Caliendo et al. (2019) cannot be applied in this model because households differ not just by their location, but also by the continuous state variable a.

Table 3: Internally Estimated Parameters

Parameter	Description	Value	Target
ρ	Discount rate	0.021	Median wealth-income: 2.35
χ	Homeown preference	1.13	Homeownership rate: 71.7%
η	House preference	0.21	House expenditure share: 21%
$\overline{\kappa}$	Migration cost	13.62	% migrate: 3.03%
κ^{age}	Migration cost	0.07	% migrate working-retired: 2.02%
κ^{dst}	Migration cost	0.0015	Gravity coefficient: -0.827
ν	Migration elasticity	2.51	Migration elasticity: 4.03
$\underline{h}^o$	Minimum house size	2.12	Ownership elasticity: -0.113
A_i	Amenities	See table 12	Population shares
Z_i	Productivities	See table 12	Wages
$\overline{p}_i^F$	Permit price shifters	See table 12	House prices

Notes: This table displays the parameter values which are internally estimated. All parameters except the migration elasticity parameter ν are chosen so that the initial steady state of the model matches cross-sectional data moments in (or as near as data allows) the year 2000. ν is chosen to match the long-run employment effect of local productivity shocks estimated by Hornbeck and Moretti (2024). The unit of time is one year. Since the number of estimated parameters equals the number of estimation targets, the model matches all targets to within numerical precision. The migration elasticity target is matched to within a tolerance of 5×10^{-3}, wages and house prices are matched exactly, and all other targets are matched to within a tolerance of 5×10^{-4}. The details of the estimation algorithm are described in Appendix E.3.

Homeownership Figure 1a compares homeownership rates over the lifecycle in the model versus data.[30] Recall that the aggregate homeownership rate is an estimation target. Overall, the model does a fairly good job of matching the steady increase in homeownership over the lifecycle. It does predict too low of a homeownership rate at the beginning of the working life (21.6% for 20-29 year-olds in the model versus 36.2% in the data) and correspondingly too high of a homeownership rate at the end of the working life (94.6% for 50-64 year-olds in the model versus 82.8% in the data). Figure 3a compares the homeownership rate across deciles of the wealth distribution in the model versus data. Qualitatively, the model generates the observed monotonically increasing relationship between wealth and homeownership. The model predicts too low of a homeownership rate below the 40th percentile of the wealth distribution, and correspondingly too high of a homeownership rate above the 40th percentile. This is likely due to the fact that there are idiosyncratic reasons why some poor people own and some rich people rent that are not included in the model.

Housing Investment Figure 1b displays the median house wealth share among homeowners for various age groups in the model and data. The model over-predicts housing wealth shares at the beginning and end of life, while under-predicting it for households age 50-64. For 20-29 year-olds, the model predicts a median housing wealth share of 266.3%, whereas in the data the median share is 171.0%. For retired households, the model predicts a housing wealth share of 76.7%, whereas in the data the median share is 65.7%.

[30] The values from Figures 1-3 are shown in Appendix Tables 13-15.

Figure 1: Statistics over the Lifecycle

Notes: This figure compares statistics in the estimated model versus data for various age groups. The data sources are the 2001 Survey of Consumer Finances (homeownership rate, median house wealth share, and median wealth to income ratio) and 2005 American Community Survey (migration rate). See Section 3.1 for more details.

Figure 2a compares selected percentiles of the house wealth share distribution in the model versus data. The model does a fairly good job of matching the house wealth share distribution until the 60th percentile, while predicting too much leverage at the top of the distribution. The discrepancy is largest at the top: the 90th percentile is 324.0% in the model versus 231.1% in the data.

Finally, Figure 3b compares the median housing wealth share across deciles of the wealth distribution in the model versus the data.[31] Qualitatively, the model matches the fact that the housing wealth share is decreasing with wealth. This may seem surprising since the utility function is homothetic. The negative relationship between housing wealth share and wealth is caused by the illiq-

[31] I exclude the first decile of the wealth distribution in Figure 3b for two reasons. First, the homeownership rate for this subgroup is 0 in the model. Second, in the data the homeownership rate is quite low (4.6%) and the median housing wealth share is quite high (3199%). Statistics for the entire wealth distribution are shown in Table 15.

Figure 2: Distributional Statistics

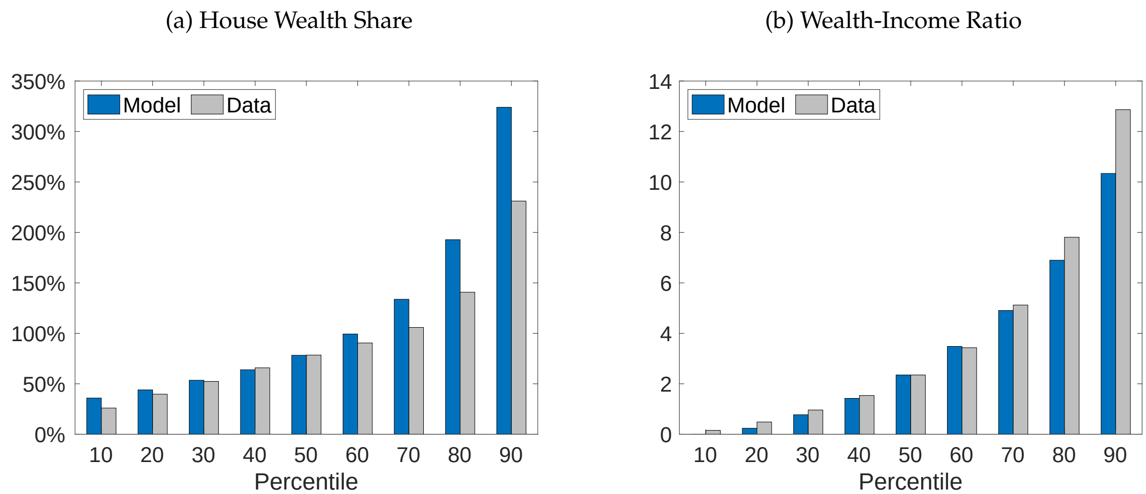

Notes: This figure compares selected percentiles of the house wealth share and wealth-income distribution in the calibrated model versus data. The data source is the 2001 Survey of Consumer Finances. See Section 3.1 for more details.

uidity of owner-occupied housing. Since it is expensive to adjust owner-occupied housing, young (wealth-poor) households invest more in housing than they would in a frictionless environment because they expect to desire a larger house in the future. The opposite is true of old (wealth-rich) households. Quantitatively, the model predicts too high of a housing wealth share at the bottom of the wealth distribution (500% in the 2nd wealth decile versus 313% in the data). The discrepancy between the model and data shrinks with wealth, and the fit is quite close above the 6th decile (where the homeownership rate is highest).

Table 4: Statistics by Tenure

	Model	Data
Migration Rate		
Renters	4.43%	5.99%
Owners	2.48%	1.77%
Median Wealth-Income		
Renters	0.00	0.34
Owners	3.90	3.63

Notes: This table compares statistics by homeownership status in the calibrated model versus data. The data sources are the 2005 American Community Survey (migration rate) and 2001 Survey of Consumer Finances (median wealth-income ratio). See Section 3.1 for more details.

Migration Rate Figure 1c compares the migration rate in the model versus data over the lifecycle. Recall that both the overall migration rate and difference in migration rates between working-age and retired households are estimation targets. As such, it is not surprising that the model does a

Figure 3: Statistics by Wealth Decile

(a) Homeownership Rate

(b) Housing Wealth Share (Owners)

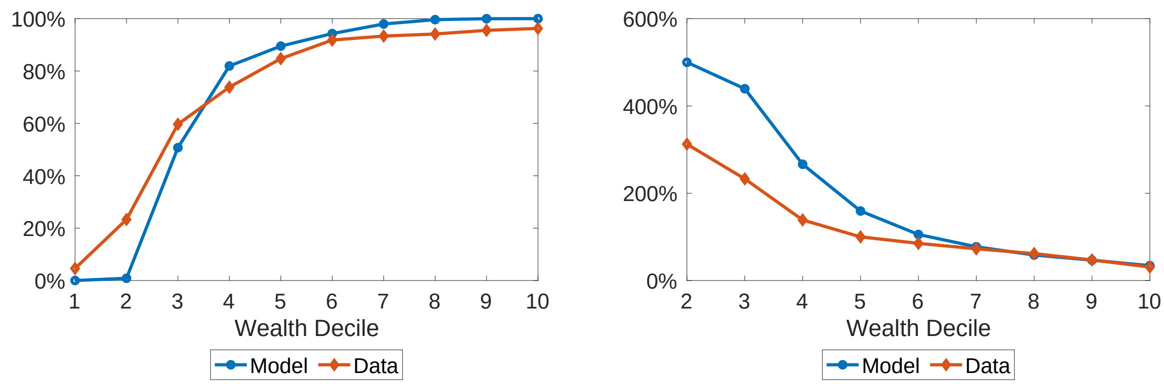

Notes: This figure compares the homeownership rate and median house wealth share (of owners) within deciles of the wealth distribution in the model versus data. The data source is the 2001 Survey of Consumer Finances. See Section 3.1 for more details.

good job of matching the steady decline in migration rates observed over the lifecycle. The top panel of Table 4 compares the migration rates of renters and owners in the model versus the data. The model under-predicts the migration rate of renters (4.4% in the model versus 6.0% in the data) and correspondingly over-predicts the migration rate of owners (2.5% in the model versus 1.8% in the data). However, qualitatively the model does generate the large observed gap in migration rates between renters and owners.

Geographic Patterns The left panel of Figure 4 compares homeownership rates by location in the model versus data. Overall, the model does a fairly good job of matching the spatial distribution of homeownership rates: the correlation between model and data is 63%. The model also does a reasonably good job of matching the across-location dispersion in homeownership rates: the standard deviation of homeownership rates is 8.6% in the model versus 7.1% in the data.

The right panel of Figure 4 compares the fraction of households who are age 65 or older by location in the model versus data. These households are retired in the model. Since the model does not include population growth, it has a higher overall fraction of 65+ than in the data. The model somewhat overpredicts the across-location dispersion in retirement rates: the standard deviation is 3.9% in the model versus 3.0% in the data. The correlation between the model and data is 33%.

Although the model generates positive correlations between observed and model-implied city-level homeownership and retirement rates, the fit is not perfect. The primary focus of this paper is on how different types of households are affected by local productivity shocks. These are determined primarily by household adjustment elasticities (especially migration, tenure, and housing investment) and general-equilibrium price responses, rather than by initial geographic sorting patterns.

In Appendix D.5, I assess robustness of important results to excluding cities with the largest discrepancies between data and model homeownership and retirement rates.

Figure 4: Geographic Patterns: Model vs. Data

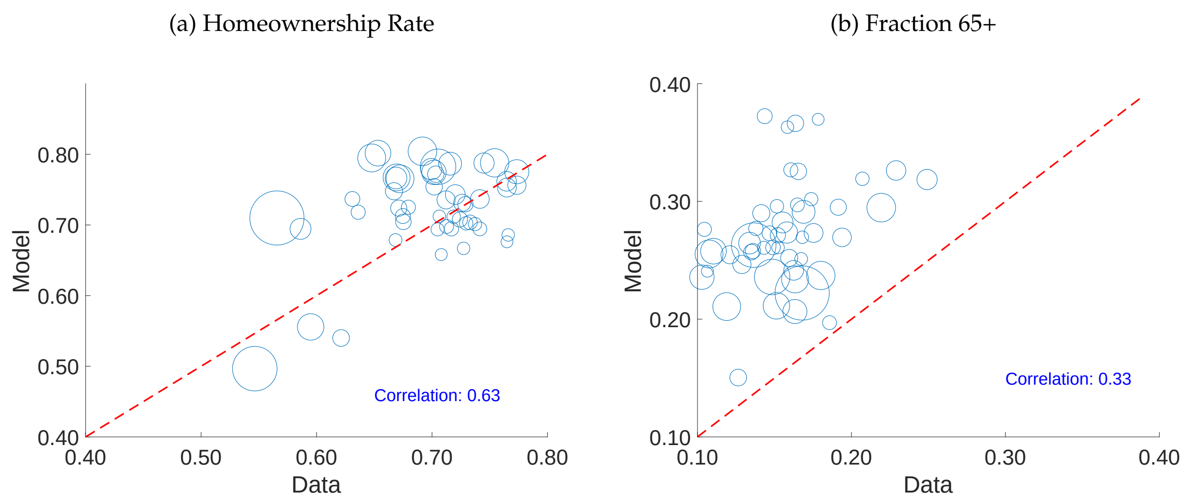

Notes: The left figure compares the homeownership rate of each location in the model versus data. The right panel compares the fraction of households age 65 and over (which corresponds to retirement in the model). The dashed lines are the 45 degree line. The population-weighted correlation between model and data is 63% for homeownership rates and 33% for fraction 65+.

Wealth Distribution Figure 1d displays the median wealth-income ratio over the lifecycle in the model and data. Recall that the median wealth-income ratio is an estimation target. Qualitatively, the model generates the observed increase over the lifecycle. As shown in Figure 2b, the model does a good job of matching the wealth-income distribution until the 90th percentile. The model under-predicts the top decile: the 90th percentile is 10.34 in the model versus 12.87 in the data. This is to be expected given that the model does not contain features that are important for explaining the upper tail of the wealth distribution.[32]

Dynamic Effects of Productivity Shocks The final set of untargeted moments pertain to the dynamic effects of the local productivity shocks $\hat{Z}_i$ described in Section 3.2. I summarize the average effect of a local productivity shock t periods after it occurs by estimating the regression

$$\log \hat{y}_{it} = \Lambda_t + \gamma_t \log \hat{Z}_i + \epsilon_{it} \tag{25}$$

where $\hat{y}_{it} = y_{it}/y_{i0}^{ss}$ is the proportional change in an outcome y relative to its value in the initial steady state, Λ_t is a time fixed effect, and ϵ_{it} is a residual. I estimate (25) using weighted least squares, with weights equal to population shares in the initial steady state.

[32] De Nardi and Fella (2017) examine the mechanisms needed to match important features of the wealth distribution.

Table 5 compares long-run effects in the model versus empirical estimates from Hornbeck and Moretti (2024). Recall that ν is estimated so that the long-run labor supply elasticity with the Bartik shock matches HM's estimate of 4.03. The elasticity is somewhat smaller with the Observed Wage shock (3.33), but still within the 95% confidence interval of $[1.05, 7.01]$. In the baseline model, which abstracts from agglomeration and congestion in production, the elasticity of wages with respect to productivity is 1 at all time horizons ($\hat{\gamma}_t = 1$). This is smaller than HM's estimate of 1.46, but within the 95% confidence interval of $[0.48, 2.44]$. Note that the rental price equation (20) implies that the steady-state price to rent ratio is fixed at $p_i / p_i^r = 1/(r + \delta)$. As a result, the long-run effect on house prices equals the long-run effect on rents. The estimated effect (2.47 with the Bartik shock and 2.81 with the Observed Wage shock) are lower than HM's house price estimate of 3.05, and above their rent estimate of 1.09. While these are inside the 95% confidence interval of $[1.13, 4.97]$ for house prices, they are outside the 95% confidence interval of $[0.15, 2.03]$ for rents.

Since both the Bartik and Observed Wage productivity shocks are scaled so that the population-weighted average productivity change is 0, neither shock generates a large change in aggregate income. As a result, neither shock causes a large change in the income tax rate τ_t. Relative to the baseline tax rate of 23.9%, the tax rate decreases by -0.02% in the long-run after the Bartik shock (to 23.88%) and by -0.49% after the Observed Wage shock (to 23.41%).

The left panels of Figure 5 display the estimated coefficients $\hat{\gamma}_t$ for house prices. Households respond to shocks by moving out of locations where productivity declines and into locations where productivity increases. This implies a positive correlation between productivity shock and permit price (see equation 21). Since house prices are determined by construction costs in the long run, this further implies a positive long-run relationship between productivity shocks and house price changes. When a shock occurs, demand for housing immediately responds as households react to new information. However, the supply of housing is initially fixed at the initial steady state level. As a result, in the short run house prices may fall below construction costs if demand is sufficiently weak that no new construction is demanded. As in Glaeser and Gyourko (2005), the durability of housing implies that productivity shocks can have an asymmetric effect on house prices and rents in the short run. Specifically, house prices and rents initially fall proportionally more in response to a negative shock than they increase in response to a positive shock. This explains the difference between the short- and long-run effects on house prices. Since the magnitude of the Observed Wage shock is larger than the Bartik shock, the time required for house prices to equal construction costs in all locations is longer for the Observed Wage shock.

The middle panels of Figure 5 display $\hat{\gamma}_t$ for rents, which are related to prices by equation (20). As mentioned above, (20) implies that the long-run rent elasticity equals the long-run house price elasticity. Both prices and rents in a location converge to their long-run level once demand is sufficiently strong that construction is strictly positive.[33] This has occurred everywhere after 1 period with the Bartik shock, and after 8 periods with the Observed Wage shock.

[33] This is due to the assumption that permit prices are a function of *long-run* population levels. This assumption helps generate well-behaved transition dynamics.

Table 5: Long-Run Effects of Productivity Shocks

	Bartik	Observed Wage	Hornbeck and Moretti (2024)
Labor Supply	4.03	3.33	4.03
	(0.10)	(0.11)	(1.52)
Wage	1.00	1.00	1.46
	(0.00)	(0.00)	(0.50)
House Price	2.47	2.81	3.05
	(0.07)	(0.07)	(0.98)
Rent	2.47	2.81	1.09
	(0.07)	(0.07)	(0.48)

Notes: This table compares the average long-run effects of the productivity shock in the model versus the empirical estimates from Hornbeck and Moretti (2024). $\hat{Z}^{\text{btk}}$ is the proportional change in productivity from the "Bartik" productivity shock, and $\hat{Z}^{\text{owg}}$ is the proportional change in productivity from the "Observed Wage" productivity shock. The model estimates are obtained using regression equation (25). The estimates from Hornbeck and Moretti (2024) are displayed in their Table 2. Standard errors are in parenthesis.

The right panels of Figure 5 display $\hat{\gamma}_t$ for populations. Recall that the migration elasticity parameter ν is estimated so that the long-run labor supply elasticity matches Hornbeck and Moretti (2024)'s empirical estimate of 4.03 with the Bartik shock. Since retired households are not directly affected by wages, and face higher moving costs than working-age households, populations respond less to productivity shocks than labor supplies. The long-run population elasticity is 3.31 for the Bartik shock and 2.76 for the Observed Wage shock.

The first two columns of Figure 6 show the percentage change in populations and house prices in every location relative to the initial steady state on the transition paths after the productivity shocks. The last column shows the migration rate. Although the Observed Wage shock induces a much larger population reallocation, it takes roughly the same amount of time (approximately 50 years) for populations to converge to their long-run allocations after both shocks. Intuitively, although the Observed Wage shock induces greater population reallocation (since this shock is more disperse than the Bartik shock), households also have a greater incentive to migrate to positively affected locations in response to this shock. The speed of transitions depends both on migration costs and the variance of idiosyncratic location preference shocks. These are disciplined, in turn, by the observed migration rate and long-run elasticity of population with respect to productivity. Since house prices are forward-looking, they adjust much more quickly to their long-run levels than populations. Finally, after each shock there is an initial increase in the migration rate as households migrate (on average) toward positively affected locations. After this increase, the migration rate gradually converges downward to its long-run value.

Figure 5: Dynamic Effects of Productivity Shock

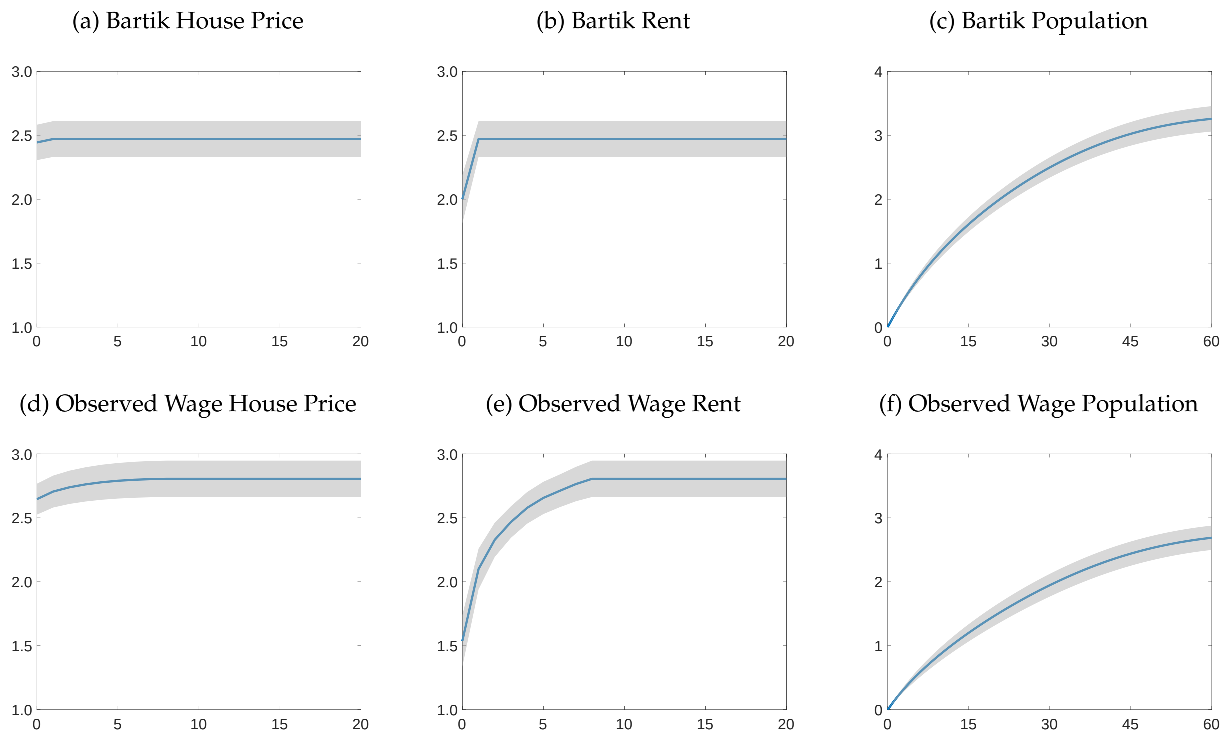

Notes: This figure displays the estimated coefficients $\hat{\gamma}_t$ from the regression (25) for house price, rent, and population. The x axis shows periods since the shock t. The shaded areas represent 95% confidence intervals. All coefficients are estimated using weighted least squares, with weights equal to local population shares in the initial steady state.

4 Quantitative Analysis

In this section, I use the theoretical framework to quantify the distributional effects of spatially heterogeneous productivity shocks. My analysis is focused on how the productivity shocks described in Section 3 affect the welfare of different types of households. The effects of the shocks vary depending on location, age, homeownership status, and wealth. The model accounts for these sources of heterogeneity, which makes it possible to analyze the disaggregate effects of local shocks.

4.1 Measuring Welfare Effects

I measure the welfare effect of a shock as the percentage change in the consumption aggregate $\mathcal{C}$, applied for the remainder of a household's life in the initial steady state, that would make it indifferent whether the shock occurs. Formally, the welfare effect for a household with state Ω is the value $\omega(\Omega)$ that satisfies

$$\mathrm{E}\left\{\int_j^J e^{-\rho(s-j)-\psi(s)}\log(A_s[1+\omega(\Omega)]\mathcal{C}_s)\mathrm{d}s + \sum_{k=\lceil j\rceil}^{J-1} e^{-\rho(k-j)}\Psi(k|j)(\varepsilon_k-\kappa_k)\right\} = V_0(\Omega) \tag{26}$$

Figure 6: Transition Dynamics

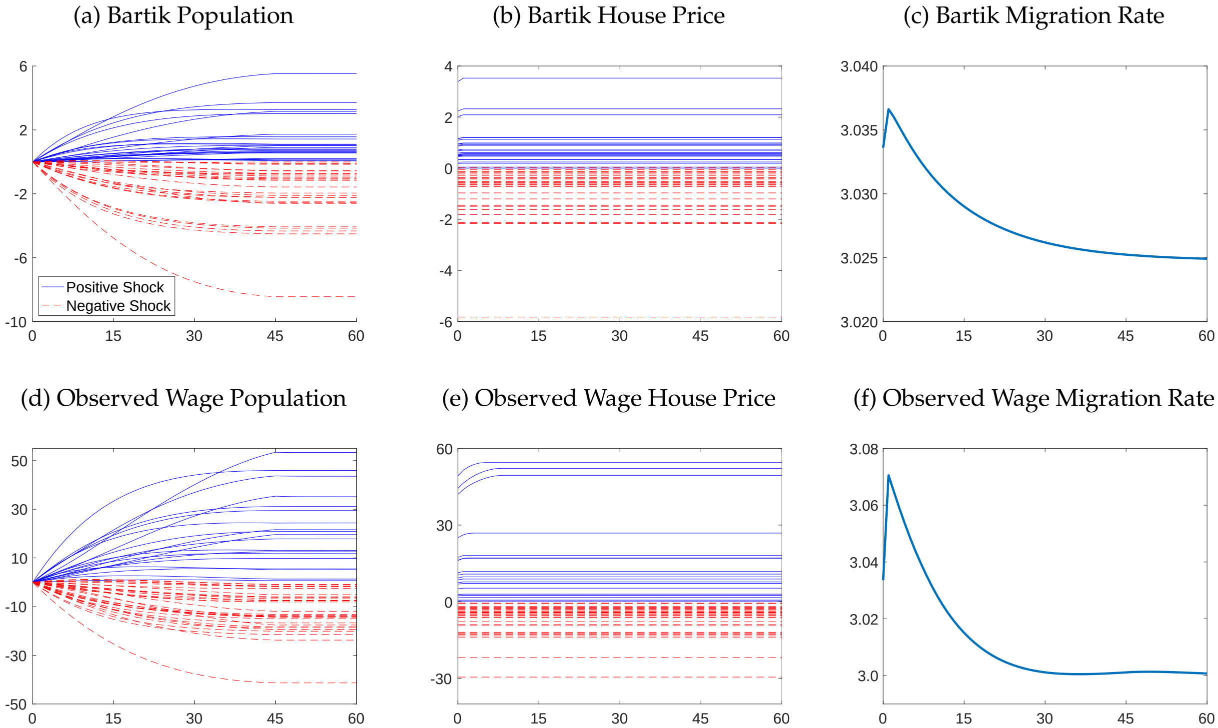

Notes: This figure shows transition dynamics after the Bartik and Observed Wage productivity shocks. Panels (a) and (b) show the percentage change in population and house price relative to the initial steady state, respectively, in each location after the Bartik productivity shock. Solid lines indicate locations that receive a positive productivity shock, while dashed lines indicate locations that receive a negative shock. Panel (c) shows the migration rate after the Bartik productivity shock. Panels (d) - (f) show the analogous figures for the Observed Wage productivity shock.

where A_s and $\mathcal{C}_s$ are amenities and consumption aggregate at age s, $\Psi(k|j)$ is the probability of surviving to age k conditional on living to age j, and ε_k and κ_k are location preference and moving cost incurred at shock age k.[34] The first term in brackets is expected lifetime utility from amenities, consumption, and housing services. The second term is the present discounted value of future location preferences and migration costs. The expectation operator denotes the expected value in the initial steady state. The right-hand side of (26) is the value function immediately after the shock occurs. In Appendix B.3, I show that welfare effects can be written as

$$\omega(\Omega) = \exp\left(\frac{V_0(\Omega) - V_0^{ss}(\Omega)}{\int_j^J e^{-\rho(s-j)-\psi(s)}\mathrm{d}s}\right) - 1. \tag{27}$$

Note that welfare effects vary across households depending on their state variables Ω.

[34] $\lceil x \rceil \equiv \min\{n \in \mathbb{Z} | n \geq x\}$ is the ceiling function, which rounds a real number x up to the nearest integer.

4.2 Quantifying the Distributional Effects of Local Productivity Shocks

I now turn to the main question of the paper: do spatially heterogeneous productivity shocks have unequal welfare effects across households? To answer this question, I use model-generated data to estimate the regression[35]

$$\overline{\omega}_i = \text{cons} + \beta \log \hat{Z}_i + \epsilon(\Omega) \tag{28}$$

where $\overline{\omega}_i$ is the average welfare effect in location i, cons is a constant, and ϵ is a residual. I estimate (28) using weighted least squares, with weights equal to population shares in the initial steady state. The parameter β is a measure of the extent to which local productivity shocks pass through to local residents' welfare. A high value of β suggests that there is a close link between local productivity and individual welfare, while a low value of β suggests that the welfare effects of local productivity shocks are evenly distributed across space. I refer to β as the "welfare elasticity."

The first column of Table 6 shows the estimated welfare elasticity, which is 0.37 for the Bartik shock and 0.39 for the Observed Wage shock. With the Bartik shock, a 1% increase in local productivity increases local residents' welfare by 0.37%. This is the first main result of the paper: there is high pass-through from local productivity shocks to welfare. This suggests that shocks which have heterogeneous effects across space can have important distributional consequences.

The first column of Table 6 also presents the welfare elasticity separately for renters and owners. The elasticity is much higher for owners than renters (0.47 vs. 0.11 for the Bartik shock and 0.51 vs. 0.09 for the Observed Wage shock). Estimates by age group are shown along the rows of Table 6. These entries show that $\hat{\beta}$ is higher for owners than renters within age groups. At the beginning of the lifecycle, the welfare elasticity is increasing with age. This is because owners' welfare is more sensitive to productivity shocks than renters, and the homeownership rate is initially increasing with age (see Figure 1a). At the end of the lifecycle, the welfare elasticity decreases with age. Positive productivity shocks benefit all groups except retired renters. These households are actually *harmed* by higher productivity. They do not benefit from increased wages, and are hurt by higher rents.

The results in Table 6 demonstrate that there is a high average pass-through from local productivity shocks to welfare, and that this elasticity varies considerably among households. In addition to welfare elasticities, the distributional effects of a shock depend on how the shock varies across locations and how households sort across space. Table 7 shows summary statistics of the welfare effect distributions, as well as average welfare effects for different types of households.

The first panel of Table 7 shows how welfare effects vary by idiosyncratic labor endowment ℓ (among working-age households). Both shocks relatively benefit high-productivity workers. This is the result of endogenous sorting patterns. Due to positive complementarity between individual productivity and wage, high-ℓ households sort into high-wage locations. As shown in Table 1,

[35](28) assumes productivity shocks have linear welfare effects. As discussed in Section 3.4, productivity shocks have asymmetric effects on house prices and rents due to housing durability. In Appendix D.2, I explore whether they have asymmetric welfare effects as well. Appendix Figure 11 shows the scatter plot of $\overline{\omega}_i$ versus $\log \hat{Z}_i$ for both shocks.

both productivity shocks are positively correlated with initial wage. As a result, high-productivity households are more likely to live in places that receive favorable shocks.

Table 6: Welfare Regression (28) Estimates

	All Ages	20-29	30-39	40-49	50-64	65+
Bartik						
All Tenures	0.37	0.30	0.47	0.52	0.41	0.23
Renters	0.11	0.22	0.24	0.23	0.05	-0.37
Owners	0.47	0.55	0.55	0.54	0.42	0.38
Observed Wages						
All Tenures	0.39	0.31	0.48	0.53	0.43	0.24
Renters	0.09	0.23	0.23	0.20	0.02	-0.36
Owners	0.51	0.60	0.59	0.57	0.46	0.46

Notes: This table displays the estimated welfare elasticity $\hat{\beta}$ from the regression (28) for various subgroups. "Bartik" refers to the Bartik productivity shock, and "Observed Wage" refers to the Observed Wage productivity shock. Observations are weighted by initial population share.

The second panel of Table 7 shows average welfare effects for working-age households versus retirees. As shown in Table 1, the fraction of retired households is negatively correlated with both shocks. This is due to endogenous sorting of retirees into locations with low prices, which tend to receive low productivity shocks. As a result, working-age households gain more than retirees. The third panel of Table 7 shows average welfare effects for renters versus homeowners. As shown in Table 1, there is a negative correlation between homeownership rate and both productivity shocks. On the other hand, the welfare elasticity is highest for working-age owners (see Table 6), who tend to sort into locations that receive high productivity shocks. The net result is that owners gain slightly more on average than renters with both shocks.

The final panel of Table 7 shows summary statistics of the distribution of welfare effects. The standard deviation of welfare effects from the Bartik productivity shock is 0.24%, and the difference between then 90th and 10th percentiles is 0.35%. For the Observed Wage shock, the standard deviation is 2.57% and the difference between then 90th and 10th percentiles is 4.98%.

4.3 Intuition

In this section, I analyze four special cases of the model to develop intuition for the results shown in the previous section. The first two abstract from borrowing and saving, risk, lifecycle dynamics, and homeownership. When migration costs are infinite (no mobility) or zero (free mobility), it is possible to obtain analytical results in these cases. The last two cases are identical to the baseline model, except that they have one and two locations, respectively.

Table 7: Welfare Effects

	Bartik	Observed Wages
By idiosyncratic labor ℓ		
Percentile 0-31	0.002%	0.211%
Percentile 32-68	0.008%	0.256%
Percentile 69-100	0.013%	0.301%
By age group		
Working-age	0.009%	0.28%
Retired	0.004%	0.19%
By tenure		
Renters	-0.005%	0.24%
Owners	0.013%	0.27%
Distributional Statistics		
Average	0.008%	0.26%
Standard Deviation	0.24%	2.57%
10th Percentile	-0.17%	-2.15%
25th Percentile	-0.06%	-0.81%
50th Percentile	0.04%	-0.01%
75th Percentile	0.11%	0.82%
90th Percentile	0.18%	2.83%

Notes: This table shows the average welfare effects of the productivity shocks for various types of households. The first column shows results for the Bartik productivity shock, and the second column show results for the Observed Wage productivity shock. ℓ is the idiosyncratic component of individual labor endowment. Welfare effects are the percentage change in the consumption aggregate $\mathcal{C}$, applied for the remainder of a household's life in the initial steady state, that would make it indifferent whether the shock occurs.

4.3.1 Special Case 1: No Wealth, Homeownership, Risk, Lifecycle, or Migration

In the first special case, households live hand-to-mouth and cannot own housing or migrate.[36] For simplicity, I also assume that households are infinitely lived, have a constant labor endowment (set to 1), and abstract from taxes. The goal of this example is to illustrate the welfare effects of local productivity shocks when there are no dynamics or migration.

In this case, lifetime utility in the initial steady state is the same for all households in a given location:

$$V_0^{\mathrm{ss}}(i) = \log[A_i w_i (p_i^r)^{-\eta}]/\rho.$$

Since there is no migration, population allocations are not affected by productivity shocks. As a result, the shock has no effect on rents. Lifetime utility after the shock is therefore

$$V_0(i) = \log[A_i \hat{Z}_i w_i (p_i^r)^{-\eta}]/\rho.$$

[36] This would happen endogenously if χ were sufficiently low and ρ and migration costs were sufficiently high.

It follows that the welfare effect of the shock in location i is

$$\overline{\omega}_i = \hat{Z}_i - 1.$$

Taking logs and using the approximation $\log(1 + \overline{\omega}_i) \approx \overline{\omega}_i$ yields

$$\overline{\omega}_i \approx \log \hat{Z}_i.$$

In this case, local productivity shocks pass through approximately one-to-one to welfare: $\beta \approx 1$.

4.3.2 Special Case 2: No Wealth, Homeownership, Risk, Lifecycle; Free Mobility

The second special case is identical to the first, except that there are no migration costs or idiosyncratic location preferences. The goal of this example is to illustrate the welfare effects of local productivity shocks when there are no dynamics or moving frictions. Since households can move freely, lifetime utility is equalized in every location both before and after the productivity shock. Note that lifetime utility may change after the shock, but the change will be the same for all households regardless of their initial location. That is, the estimated constant $\widehat{cons}$ in the welfare regression (28) will in general be nonzero. However, there is no pass-through from local productivity shocks to relative welfare: $\beta = 0$.

4.3.3 Special Case 3: One Location

In the third special case, there is only one location. I examine the effects of two partial equilibrium shocks in this setting: a permanent, 1% increase in wages, and a permanent, 1% increase in house prices and rents. I refer to the first shock as the "wage shock" and the second shock as the "house price shock." The parameters for this example are estimated in the same way as Section 3.3.[37]

The left panel of Figure 7 shows the welfare effect of the wage shock by age group and tenure. The benefits of this shock are decreasing with age. This is because earnings make up a larger fraction of lifetime income for young households. This also explains why renters (who are relatively wealth-poor) benefit more than owners. All age groups prior to retirement are made strictly better off. Since retired households do not participate in the labor market, this shock does not affect them.

The right panel of Figure 7 shows the welfare effect of the house price shock by age group and tenure. Renters of all ages are made worse off due to higher housing costs. Young homeowners are also harmed, but less so than renters. The welfare benefit of the shock is increasing in age for owners, and is positive at the end of the lifecycle. A house price increase affects owners in two ways. First, it harms them by increasing housing costs. This includes transaction and maintenance costs, as well as future home purchases. On the other hand, it benefits them by increasing their housing

[37] The estimated parameter values are shown in Appendix Table 16.

wealth. Housing wealth affects welfare both by providing insurance against adverse shocks (accessible through downsizing or borrowing) and by generating future capital gains. The net welfare effect of a house price change for owners depends on the relative present values of costs versus benefits. Present costs are relatively higher for young homeowners, who are "short" in housing (expect to increase house size in the future) and do not expect to realize capital gains for many years. The opposite is true of old homeowners, which explains why the welfare effect is increasing in age for owners.

Figure 7: One-Location Example

(a) Wage Shock

(b) House Price Shock

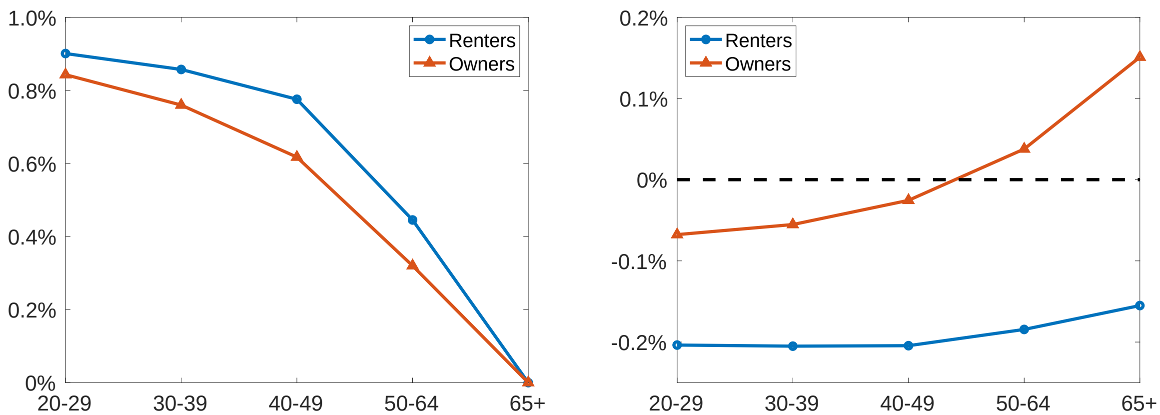

Notes: This figure shows results for the one-location example. The left panel shows the welfare effect of a permanent, 1% increase in wages by age group and tenure. The right panel shows the welfare effect of a permanent, 1% increase in house prices and rents.

Another reason why the welfare effects of the house price shock differ between renters and owners is that there is a gap between the user cost of housing for owners p_t^u and the rental price p_t^r (see equations 4 and 20):

$$p_t^u = p_t^r - \tau_t r p_t.$$

This gap exists because homeowners' imputed rental income is not taxed. It implies that housing is effectively less expensive for owners than renters. In addition, owners' housing costs are less responsive to house price changes than renters'. This is another reason why young homeowners are less harmed by the house price increase than young renters.

The difference in the welfare effects of house price changes for renters versus owners is critical for understanding the main results of the paper. A local productivity shock endogenously raises wages, house prices, and rents. Both renters and owners of working age benefit from higher wages. The welfare gain from higher wages is offset by higher housing costs for renters. Owners benefit from higher house prices due to capital gains, and are more insulated from higher housing costs. As a result, the pass-through from local productivity shocks to welfare is stronger when there is homeownership than it would be if all housing were rented.

4.3.4 Special Case 4: Two Locations

In the final special case, there are two locations. For this example, I examine the effects of a shock that increases location 1's productivity by 1% and leaves location 2's productivity unchanged. I set location-specific parameters equal in both places, so they are identical in the initial equilibrium. The remaining parameters are estimated as in Section 3.3.[38]

Figure 8: Two-Location Example

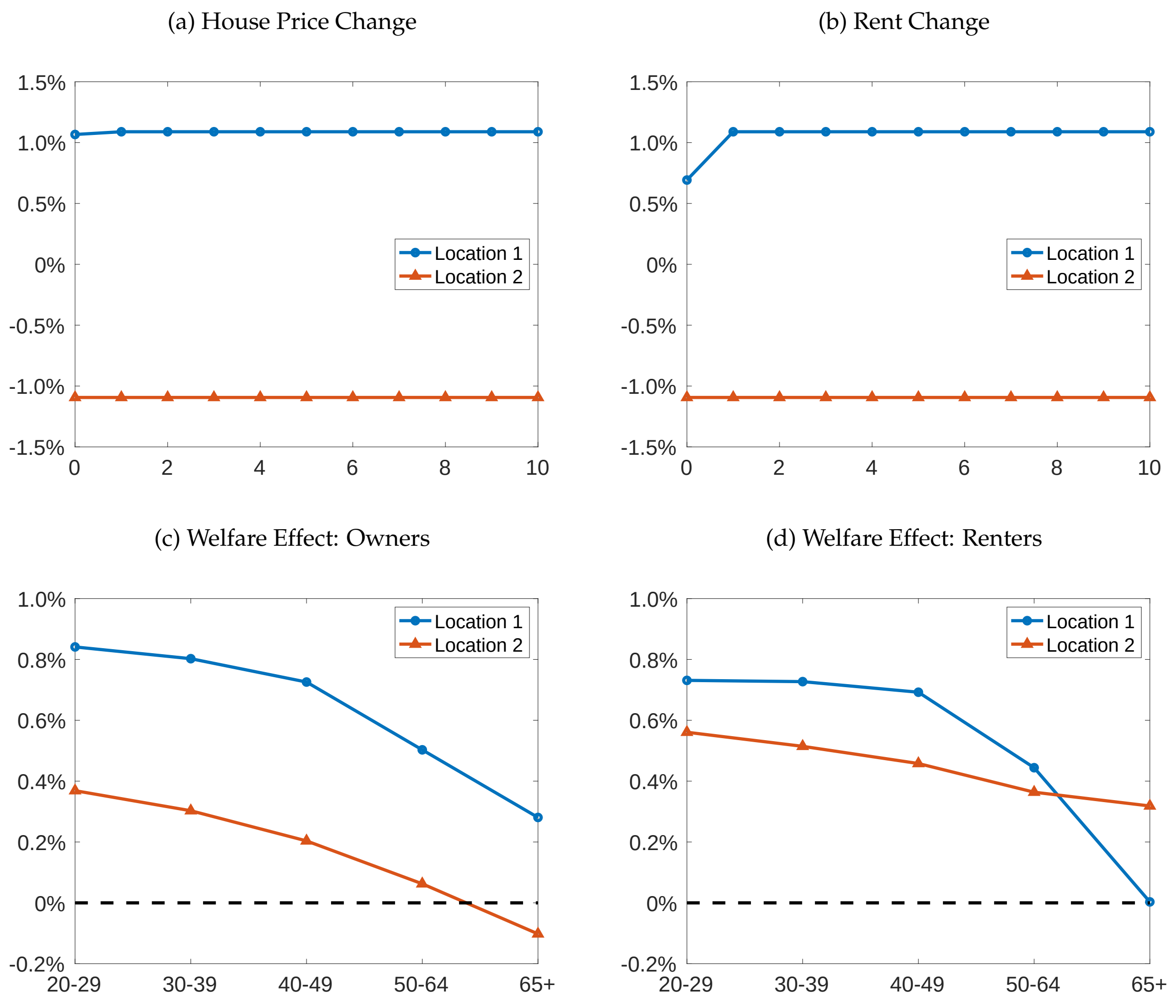

(a) House Price Change

(b) Rent Change

(c) Welfare Effect: Owners

(d) Welfare Effect: Renters

Notes: This figure shows results for the two-location example. This example considers a shock in which location 1's productivity is increased by 1%, while location 2's productivity is unchanged. The top left panel shows the effect of this shock on house prices, and the top right panel shows the effect on rents (expressed as percentage changes relative to the initial steady state). The bottom two panels show welfare effects for owners and renters, respectively.

The top left panel of figure 8 shows the effect of the shock on house prices, expressed as the percentage change relative to the initial steady state. In the long run, the shock raises location 1's population

[38]For this example, I set house price elasticities in each location equal to the population-weighted average from the baseline (0.74). The estimated parameter values are shown in Appendix Table 16.

by 1.48% and lowers location 2's population by the same amount. Given the calibrated house price elasticity, this implies that construction cost increases by 1.09% in location 1 and decreases by the same amount in location 2. These correspond to the long-run house price changes. However, in the short run, house price does not equal construction cost in location 1. At the new construction cost, location-1 housing demand would initially be lower than the initial steady state housing quantity. Since construction is irreversible, to clear the market, the location-1 house price is initially below construction cost. The gap between construction cost and house price lasts for one period. After that, positive construction is again demanded due to increased population and housing depreciation. In contrast, the location-2 house price immediately falls to the new construction cost. As a result, the shock initially lowers location 2's house price more than it increases location 1's. The top right panel of Figure 8 shows the effect on rents, which are related to prices by equation (20). Finally, since the shock raises earnings in location 1, the tax rate required to balance the government's budget falls from 23.90% in the initial steady state to 23.77% in the long run.

The bottom panels of Figure 8 show welfare effects by age group and tenure.[39] For both locations and tenures, welfare effects are decreasing with age. Holding rents and house prices constant, the shock benefits working-age households by increasing location 1's wage. Younger workers benefit more from this channel because they have more remaining working years. In location 2, young workers also benefit more than old workers because they face lower migration costs to access the higher wage in location 1. Both working-age renters and owners benefit more in location 1 than location 2, as they receive higher wages without having to incur moving costs. Even though retirees are not affected by wage changes, old owners are made better off in location 1 due to higher house prices and worse off in location 2 due to lower house prices. The opposite is true for old renters: retired renters in location 1 are harmed due to higher rents, whereas retired renters in location 2 benefit from lower rents.

Table 8 compares welfare elasticities in the baseline model (with the Bartik productivity shock) versus the two-location example. The main qualitative insights from the baseline model also apply in the two-location example: the average welfare elasticity is high, and the elasticity is much higher for owners than renters. The main difference is that the welfare elasticity for renters is nearly 2/3 lower in the two-location example (0.04 with two locations versus 0.11 in the baseline). Geography is an important reason for this discrepancy. In the full model, locations differ on several dimensions – for example, wages, house prices, and amenities. Migration after a shock is undesirable if there are no nearby locations with attributes a household finds attractive. In the two-location example, locations are initially identical. Migration is therefore a more effective adjustment mechanism, especially for renters (who tend to be young and need not pay housing transaction costs to move). As a result, the welfare elasticity is lower in the two-location example. This difference highlights the role of geography in shaping the welfare effects of local shocks.

[39]The values from Figure 8 are shown in Appendix Table 17.

Table 8: Two-Location Welfare Elasticities

	All Ages	20-29	30-39	40-49	50-64	65+
Baseline						
All	0.37	0.30	0.47	0.52	0.41	0.23
Renters	0.11	0.22	0.24	0.23	0.05	-0.37
Owners	0.47	0.55	0.55	0.54	0.42	0.38
Two-Location Example						
All	0.34	0.24	0.42	0.50	0.43	0.17
Renters	0.04	0.17	0.21	0.24	0.08	-0.32
Owners	0.46	0.47	0.50	0.52	0.44	0.38

Notes: This table compares the estimated welfare elasticity from equation (28) for various subgroups in the baseline model (with the Bartik productivity shock) versus those from the two-location example.

4.4 Mechanisms

In this section, I perform a series of counterfactuals that isolate the effects of important mechanisms. The goal of these experiments is to develop a more complete understanding of the channels through which local productivity shocks affect welfare for different types of households.

4.4.1 Homeownership

The results shown in Table 6 suggest that homeownership plays a key role in determining the distributional effects of local productivity shocks. Quantitative spatial models typically do not account for homeownership, or if they do, they incorporate it in a stylized manner. In this section, I explore how omitting homeownership affects conclusions about spatial redistribution. To do this, I analyze a version of the model in which homeownership is prohibited. I estimate the parameters of this model using the same procedure employed for the baseline, and compute transition dynamics after the same productivity shock.[40] I then estimate the welfare regression (28) as above.

The first two panels of Table 9 compare estimates from the baseline model with those from the no-ownership model. Without homeownership, the welfare elasticity is nearly 0 (0.03, versus 0.37 in the baseline). The welfare elasticity is decreasing in age, and is negative for the 50-64 and 65+ age groups. While the average elasticity is smaller than for renters in the baseline model (which is 0.11), the elasticities are similar within age group. The overall difference is in part accounted for by a compositional effect: whereas the fraction of renters decreases with age in the baseline model, it is constant (100%) in the no-ownership model. Since the welfare elasticity is lowest for old renters in both models, the average elasticity is lower in the no-ownership model.

[40]The estimated parameter values are shown in Appendix Table 16.

Table 9: Counterfactual Welfare Elasticities

	All Ages	20-29	30-39	40-49	50-64	65+
Baseline						
All	0.37	0.30	0.47	0.52	0.41	0.23
Renters	0.11	0.22	0.24	0.23	0.05	-0.37
Owners	0.47	0.55	0.55	0.54	0.42	0.38
No Ownership						
All	0.03	0.22	0.23	0.18	-0.01	-0.33
Wage Changes Only						
All	0.32	0.44	0.50	0.47	0.28	0.00
Renters	0.38	0.44	0.52	0.56	0.43	0.00
Owners	0.29	0.47	0.49	0.46	0.28	0.00
House Price Changes Only						
All	0.05	-0.16	-0.04	0.04	0.13	0.23
Renters	-0.27	-0.23	-0.29	-0.34	-0.37	-0.37
Owners	0.17	0.07	0.06	0.07	0.14	0.38
No Capital Gains						
All	-0.00	0.22	0.23	0.17	-0.05	-0.43
Renters	0.11	0.22	0.24	0.23	0.05	-0.37
Owners	-0.04	0.23	0.22	0.16	-0.05	-0.44
Equalize House Price Sensitivity						
All	0.33	0.27	0.42	0.46	0.35	0.18
Renters	0.09	0.20	0.21	0.19	0.03	-0.37
Owners	0.41	0.50	0.50	0.49	0.36	0.31
No Transaction Costs						
All	0.36	0.29	0.44	0.49	0.40	0.23
Renters	0.11	0.22	0.24	0.23	0.05	-0.37
Owners	0.44	0.52	0.51	0.50	0.40	0.40

Notes: This table compares the estimated welfare elasticity from equation (28) for various subgroups in the baseline model with those of the counterfactuals described in Section 4.4. This table shows results for the Bartik productivity shock. Results for the Observed Wage productivity shock are shown in Appendix Table 18.

The disparity between the baseline and no-ownership models further highlights the central role of homeownership in shaping the welfare effects of local shocks. These results suggest that models which abstract from homeownership are likely to underestimate the distributional effects of spatially heterogeneous shocks. This is further demonstrated in Table 10, which compares distributional statistics in the baseline model versus the model with no ownership. For both shocks, the standard deviation of welfare effects is roughly twice as large in the baseline as the no-ownership model (0.24% vs. 0.13% for the Bartik shock and 2.57% vs. 1.26% for the Observed Wage shock).

Table 10: Distribution of Welfare Effects

	Standard Deviation	p10	p25	p50	p75	p90
Bartik Shock						
Baseline Model	0.24%	-0.17%	-0.06%	0.04%	0.11%	0.18%
No Ownership	0.13%	-0.11%	-0.04%	0.01%	0.05%	0.11%
Observed Wage Shock						
Baseline Model	2.57%	-2.15%	-0.81%	-0.01%	0.82%	2.83%
No Ownership	1.26%	-0.82%	-0.18%	0.32%	0.84%	1.83%

Notes: This table compares summary statistics of the welfare effect distribution in the baseline model versus the model without homeownership. "Bartik" refers to the Bartik productivity shock, and "Observed Wage" refers to the Observed Wage productivity shock.

4.4.2 Labor versus Housing Market

The next two experiments are partial-equilibrium exercises that are designed to elucidate the separate roles of labor and housing markets. In the first counterfactual, "wage changes only," wages evolve as in the baseline transition, while rents, house prices, and taxes are fixed at their initial steady state values. The third panel of Table 9 displays the estimated welfare elasticities from this counterfactual. Shutting down rental and house price changes increases the welfare elasticity for renters (0.38, versus 0.11 in the baseline) and lowers it for owners (0.29, versus 0.47 in the baseline). This confirms the intuition mentioned above: while endogenous house price changes mitigate the welfare effects of wage changes for renters, they augment them for owners.

In the second experiment, "house price changes only," rents and house prices evolve as in the baseline transition, while wages and taxes are fixed at their initial steady state values. Without wage changes, the average welfare elasticity falls substantially, from 0.37 to 0.05. In this counterfactual, the only effect of a local productivity increase for renters is to raise housing costs, so their welfare elasticity is negative. The elasticity is positive for owners, but much lower than the baseline (0.17 versus 0.47).

Figure 9 shows the estimated welfare elasticities by age and tenure for these counterfactuals. The differences relative to Figure 7 highlight the importance of geography in determining the welfare effects of local shocks. Whereas the welfare benefit of a wage increase is strictly decreasing in age in the one-location example, the benefit is hump-shaped in the full model. Since migration costs are increasing in age, in the multi-location model young households are less affected by local wage changes because they can easily migrate to strong labor markets. As a result, in the full model the welfare effect of wage changes is initially increasing in age. Similar logic explains why the harm from higher rents is increasing in age in the full model, but not the one-location model. Another qualitative difference is that whereas house price increases hurt young households in the one-location model, they benefit owners of all ages when there are multiple locations. There are two reasons for this discrepancy. First, whereas higher housing costs cannot be avoided when there is only one location, young owners in the multi-location environment can migrate to places with low housing costs. Second, in the multi-location model, capital gains are realized not only when changing house size, but also when changing location. The first difference reduces the expected increase in lifetime costs caused by higher local house prices, while the second increases the expected present value of capital gains. This explains why young homeowners benefit (on average) from higher local house prices in the full model.

Figure 9: Partial Equilibrium Welfare Elasticities

(a) Wage Changes Only

(b) House Price Changes Only

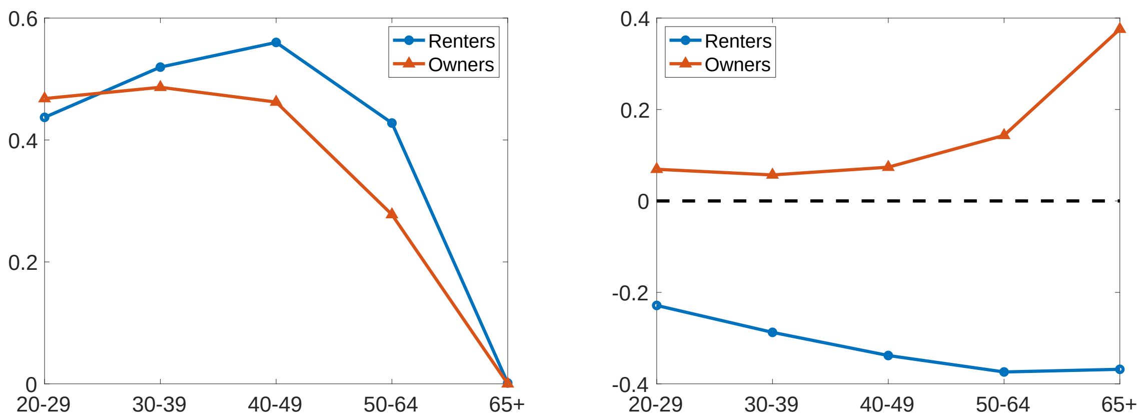

Notes: The left panel shows the estimated welfare elasticity by age group and tenure in the "wage changes only" counterfactual. In this counterfactual, wages evolve as in the baseline transition, while rents, house prices, and taxes are fixed at their initial steady state levels. The right panel shows the estimated welfare elasticity by age group and tenure in the "house price changes only" counterfactual. In this counterfactual, rents and house prices evolve as in the baseline transition, while wages and taxes are fixed at their initial steady state levels.

4.4.3 Renters versus Owners

In the final set of counterfactuals, I explore why the welfare effects of local shocks differ so much between renters and owners. Two of the reasons for this discrepancy are discussed in Section 4.3.3: owners receive capital gains when house prices change, and their housing costs are less sensitive to house price changes. Another difference is that while both renters and owners face utility migration costs, owners who move also incur housing transaction costs. These reduce the benefits of moving, making owners more exposed to local shocks. In this section, I perform a series of partial-equilibrium counterfactuals that seek to quantify the relative importance of these channels.[41]

No Capital Gains The first counterfactual shuts down capital gains. In this counterfactual, liquid wealth is adjusted to offset capital gains that are received when the shock occurs. Specifically, post-

[41] By "partial equilibrium," I mean that prices and taxes evolve as in the baseline transition.

shock liquid wealth is set to $b_0 = b_0^{\text{ss}} - (p_{i0} - p_{i0}^{\text{ss}})h$. This ensures that total wealth is the same immediately before and after the shock occurs.

The results of this counterfactual, referred to as "no capital gains" are shown in the fifth panel of Table 9. Since renters do not receive capital gains, they are unaffected. In contrast, the welfare elasticity for owners falls dramatically, from 0.47 in the baseline to -0.04. The magnitude of this decrease highlights the paramount role of capital gains. Without capital gains, higher housing costs are enough to make old homeowners suffer when local productivity increases. The welfare elasticity remains positive for younger homeowners, but is substantially reduced.

Equalize House Price Sensitivity The next counterfactual equalizes the sensitivity of housing costs to house prices for renters and owners. As discussed in Section 4.3.3, the pass-through from house prices to housing costs is higher for renters than owners. Specifically, the difference in the derivative of per-unit housing cost with respect to house price between renters and owners is (see equations 4 and 20)

$$\frac{dp_{it}^r}{dp_{it}} - \frac{dp_{it}^u}{dp_{it}} = \tau_t r.$$

This difference is due to the fact that imputed rental income is not taxed.[42] In this counterfactual, I impose an additional tax of $\tau_t r(p_{it} - p_{i0}^{\text{ss}})h$ on homeowners. This leaves the gap between rental price and user cost in the initial steady state unchanged, but ensures that the response of housing costs to house prices is the same for renters and owners after the shock.

The results of this counterfactual, referred to as "equalize house price sensitivity" are shown in the sixth panel of Table 9. Making housing costs more responsive to house prices reduces the welfare elasticity for owners (from 0.47 in the baseline to 0.41). However, the effect is modest compared to that of shutting down capital gains. The welfare elasticity is modestly reduced for renters as well (from 0.11 to 0.09). This is because some renters will eventually own housing in their current location. When they do transition to homeownership, the higher house price sensitivity will attenuate welfare effects relative to the baseline. This mechanism is less relevant for old renters (who are less likely to purchase a home in the future), which explains why they are less affected.

No Transaction Costs The final counterfactual shuts down housing transaction costs immediately after the shock. As in the baseline transition, homeowners receive capital gains from initial house price changes. However, immediately afterward, all housing wealth is converted to liquid wealth (without transaction fees). In effect, owners are "transformed" into renters with the same amount of wealth they had prior to the transformation. As a result, in the immediate aftermath of the shock, all households of a given age face the same moving costs regardless of their initial tenure.

[42] If imputed rental income (net of depreciation and capital gains) were taxed, the user cost of housing would equal the rental price. In this case, owners would pay an additional tax of $\tau_t r p_{it} h$. Since imputed rental income is not taxed, the user cost of housing is less than the rental price: $p_{it}^u = p_{it}^r - \tau_t r p_{it}$.

The final row of Table 9 shows the results of this counterfactual, referred to as “no transaction costs.” Renters are unaffected by this counterfactual, because all their wealth is already liquid. For owners, the welfare elasticity decreases only slightly relative to the baseline (from 0.47 to 0.44). While housing transaction costs do play a role in making owners’ welfare more sensitive to local productivity shocks than renters, the effect is quantitatively small.

5 Conclusion

Many shocks have uneven effects across space. There are numerous reasons why individual welfare is tied to local economic conditions. Two of the most important are migration frictions, which make it difficult to move across labor markets, and owner-occupied housing, which links financial wealth to local house prices. Accounting for these channels is necessary to understand the distributional effects of spatially heterogeneous shocks.

In this paper, I develop a dynamic quantitative spatial model and use it to quantify the distributional effects of spatially heterogeneous productivity shocks. The model incorporates realistic homeownership into an otherwise standard quantitative spatial model. This poses a computational challenge, which I overcome using continuous-time computational techniques. The computational challenge is worth confronting because homeownership fundamentally changes the relationship between local shocks and individual welfare. In the baseline model, the average pass-through from a local productivity shock to welfare is 0.37. Without homeownership, the pass-through is only 0.03. Accounting for realistic household heterogeneity makes it possible to quantify how the welfare effects of local shocks vary across different types of households. I find that both the magnitude and, in some cases, the sign of such effects vary across households of different age groups and tenure. These results suggest that spatially heterogeneous shocks can have important distributional consequences, and that much of this redistribution is due to homeownership.

References

ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2022): "Income and wealth distribution in macroeconomics: A continuous-time approach," *The Review of Economic Studies*, 89, 45–86. [3, 5, 16, 17, 70, 71]

ALLEN, T. AND D. DONALDSON (2020): "Persistence and path dependence in the spatial economy," Tech. rep., National Bureau of Economic Research. [4]

ALMAGRO, M. AND T. DOMÍNGUEZ-IINO (2025): "Location sorting and endogenous amenities: Evidence from amsterdam," *Econometrica*, 93, 1031–1071. [4]

ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade shocks and labor adjustment: A structural empirical approach," *American economic review*, 100, 1008–45. [4]

BARTIK, T. J. (1993): "Who benefits from local job growth: migrants or the original residents?" *Regional studies*, 27, 297–311. [19]

BELL, F. C. AND M. L. MILLER (2005): "Life tables for the United States social security area 1900-2100," *Social Security Administration Publications*. [21, 23]

BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2018): "House prices and consumer spending," *The Review of Economic Studies*, 85, 1502–1542. [5]

BILAL, A. AND E. ROSSI-HANSBERG (2021): "Location as an asset," *Econometrica*, 89, 2459–2495. [4]

——— (2023): "Anticipating Climate Change Across the United States," Tech. rep., National Bureau of Economic Research. [4, 5]

CAI, S., L. CALIENDO, F. PARRO, AND W. XIANG (2022): "Mechanics of spatial growth," Tech. rep., National Bureau of Economic Research. [5]

CALIENDO, L., M. DVORKIN, AND F. PARRO (2019): "Trade and labor market dynamics: General equilibrium analysis of the china trade shock," *Econometrica*, 87, 741–835. [4, 16, 25]

CARROLL, C. D. (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics letters*, 91, 312–320. [16]

CREWS, L. G. (2023): "A Dynamic Spatial Knowledge Economy," Working Paper. [5]

DAVIS, M. A. AND S. VAN NIEUWERBURGH (2015): "Housing, Finance, and the Macroeconomy," in *Handbook of Regional and Urban Economics*, vol. 5, 753–811. [9, 22, 23]

DE NARDI, M. AND G. FELLA (2017): "Saving and wealth inequality," *Review of Economic Dynamics*, 26, 280–300. [30]

DEKLE, R., J. EATON, AND S. KORTUM (2008): "Global rebalancing with gravity: Measuring the burden of adjustment," *IMF Economic Review*, 55, 511. [16, 25]

DESMET, K., D. K. NAGY, AND E. ROSSI-HANSBERG (2018): "The geography of development," *Journal of Political Economy*, 126, 903–983. [4]

DESMET, K. AND F. PARRO (2025): "Spatial Dynamics," Tech. rep., National Bureau of Economic Research. [4]

DIAMOND, R. (2016): "The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000," *American Economic Review*, 106, 479–524. [5]

DUSTMANN, C., U. SCHÖNBERG, AND J. STUHLER (2016): "The impact of immigration: Why do studies reach such different results?" *Journal of Economic Perspectives*, 30, 31–56. [69]

DVORKIN, M. A. (2023): "Heterogeneous agents dynamic spatial general equilibrium," Working Paper. [5]

ECKERT, F. (2019): "Growing apart: Tradable services and the fragmentation of the US economy," *Unpublished Manuscript, Yale University*. [5]

ECKERT, F. AND T. KLEINEBERG (2019): "Can We Save the American Dream? A Dynamic General Equilibrium Analysis of the Effects of School Financing on Local Opportunities," *Unpublished Manuscript, Yale University*. [4]

FERNANDEZ-VILLAVERDE, J. AND D. KRUEGER (2011): "Consumption and saving over the life cycle: How important are consumer durables?" *Macroeconomic dynamics*, 15, 725–770. [5]

FLODEN, M. AND J. LINDÉ (2001): "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics*, 4, 406–437. [21, 23]

FOGLI, A. AND V. GUERRIERI (2019): "The end of the American dream? Inequality and segregation in US cities," *NBER Working Paper*. [5]

GANONG, P. AND D. SHOAG (2017): "Why has regional income convergence in the US declined?" *Journal of Urban Economics*, 102, 76–90. [2, 5]

GARDNER, J. AND J. R. HENDRICKSON (2018): "If I leave here tomorrow: an option view of migration when labor market quality declines," *Southern Economic Journal*, 84, 786–814. [23]

GIANNONE, E. (2022): "Skilled-biased technical change and regional convergence," Working Paper. [4]

GIANNONE, E., Q. LI, N. PAIXAO, AND X. PANG (2025): "Unpacking Moving," *Working Paper*. [5]

GLAESER, E. L. AND J. GYOURKO (2005): "Urban decline and durable housing," *Journal of political economy*, 113, 345–375. [31]

GREANEY, B., A. PARKHOMENKO, AND S. VAN NIEUWERBURGH (2025): "Dynamic Urban Economics," Tech. rep., National Bureau of Economic Research. [5]

GRUBER, J. W. AND R. F. MARTIN (2005): "Precautionary Savings and the Wealth Distribution with Illiquid Durables," *SSRN Electronic Journal*. [22, 23]

HANSEN, G. D. (1993): "The cyclical and secular behaviour of the labour input: Comparing efficiency units and hours worked," *Journal of Applied Econometrics*, 8, 71–80. [21, 23]

HERKENHOFF, K. F., L. E. OHANIAN, AND E. C. PRESCOTT (2018): "Tarnishing the golden and empire states: Land-use restrictions and the US economic slowdown," *Journal of Monetary Economics*, 93, 89–109. [22, 23, 63]

HORNBECK, R. AND E. MORETTI (2024): "Estimating who benefits from productivity growth: local and distant effects of city productivity growth on wages, rents, and inequality," *Review of Economics and Statistics*, 106, 587–607. [5, 24, 26, 31, 32, 63, 64, 69, 70]

JONES, C., V. MIDRIGAN, AND T. PHILIPPON (2022): "Household Leverage and the Recession," *Econometrica*, 90, 2471–2505. [5]

KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2020): "The housing boom and bust: Model meets evidence," *Journal of Political Economy*, 128, 3285–3345. [5]

KAPLAN, G. AND S. SCHULHOFER-WOHL (2017): "Understanding the long-run decline in interstate migration," *International Economic Review*, 58, 57–94. [2]

KIYOTAKI, N., A. MICHAELIDES, AND K. NIKOLOV (2011): "Winners and Losers in Housing Markets," *Journal of Money, Credit and Banking*, 43, 255–296. [5]

KLEINMAN, B., E. LIU, AND S. J. REDDING (2023): "Dynamic spatial general equilibrium," *Econometrica*, 91, 385–424. [5]

KOPECKY, K. A. AND R. M. SUEN (2010): "Finite state Markov-chain approximations to highly persistent processes," *Review of Economic Dynamics*, 13, 701–714. [71]

LUCCIOLETTI, C. (2024): "Should Governments Subsidize Homeownership?" Working Paper. [5]

LYON, S. AND M. E. WAUGH (2019): "Quantifying the losses from international trade," *Unpublished Manuscript*. [4]

MCDANIEL, C. (2007): "Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003," *Manuscript, Arizona State University*, 19602004. [22]

MELO, P. C., D. J. GRAHAM, AND R. B. NOLAND (2009): "A meta-analysis of estimates of urban agglomeration economies," *Regional science and urban Economics*, 39, 332–342. [63]

MERTON, R. C. (1969): "Lifetime portfolio selection under uncertainty: The continuous-time case," *The review of Economics and Statistics*, 247–257. [5]

MOLLOY, R., C. L. SMITH, AND A. WOZNIAK (2011): "Internal migration in the United States," *Journal of Economic perspectives*, 25, 173–96. [2]

MORETTI, E. (2012): *The new geography of jobs*, Houghton Mifflin Harcourt. [2, 20]

——— (2013): "Real wage inequality," *American Economic Journal: Applied Economics*, 5, 65–103. [5]

NAKAMURA, E. AND J. STEINSSON (2018): "Identification in macroeconomics," *Journal of Economic Perspectives*, 32, 59–86. [24]

NOTOWIDIGDO, M. J. (2020): "The incidence of local labor demand shocks," *Journal of Labor Economics*, 38, 687–725. [5]

PIAZZESI, M. AND M. SCHNEIDER (2016): "Housing and Macroeconomics," in *Handbook of Macroeconomics*, vol. 2, 1547–1640. [6, 9, 22, 23]

REDDING, S. J. AND E. ROSSI-HANSBERG (2017): "Quantitative spatial economics," *Annual Review of Economics*, 9, 21–58. [6]

ROUWENHORST, K. G. (1995): "Asset Pricing Implications of Equilibrium Business Cycle Models," in *Frontiers of business cycle research*, Princeton University Press, 294–330. [71]

RUGGLES, S., S. FLOOD, S. FOSTER, R. GOEKEN, J. PACAS, M. SCHOUWEILER, AND M. SOBEK (2021): "Integrated public use microdata series, USA: Version 11.0 [database]," *Minneapolis: University of Minnesota*. [18]

SAIZ, A. (2010): "The geographic determinants of housing supply," *The Quarterly Journal of Economics*, 125, 1253–1296. [15, 22, 23, 61, 62]

SUN, J. (2025): "The Distributional Consequences of Climate Change: The Role of Housing Wealth, Expectations, and Uncertainty," Tech. rep., Working Paper. [5]

VANHAPELTO, T. (2022): "House Prices and Rents in a Dynamic Spatial Equilibrium," Working Paper. [5]

YAGAN, D. (2014): "Moving to opportunity? Migratory insurance over the Great Recession," *Job Market Paper*. [5]

ZERECERO, M. (2021): "The birthplace premium," Working Paper. [4]

A Additional Tables and Figures

Table 11: MSA Statistics

City	Population Share	Wage	House Price	$\hat{Z}^{\text{btk}} - 1$	$\hat{Z}^{\text{owg}} - 1$
Atlanta-Sandy Springs-Roswell, GA	0.0262	0.8217	0.6850	-0.0003	-0.0494
Austin-Round Rock, TX	0.0088	0.7599	0.6850	0.0066	0.0516
Baltimore-Columbia-Towson, MD	0.0171	0.8090	0.6850	-0.0005	0.0704
Birmingham-Hoover, AL	0.0069	0.7439	0.4110	0.0026	-0.0124
Boston-Cambridge-Newton, MA-NH	0.0320	0.8651	1.0047	0.0016	0.0391
Buffalo-Cheektowaga-Niagara Falls, NY	0.0080	0.7100	0.4110	-0.0097	-0.0001
Charlotte-Concord-Gastonia, NC-SC	0.0102	0.7698	0.5937	-0.0048	-0.0019
Chicago-Naperville-Elgin, IL-IN-WI	0.0554	0.8542	0.8220	-0.0020	-0.0413
Cincinnati, OH-KY-IN	0.0130	0.7766	0.5024	-0.0041	-0.0002
Cleveland-Elyria, OH	0.0153	0.7667	0.5937	-0.0045	-0.0561
Columbus, OH	0.0107	0.7684	0.5937	-0.0021	-0.0353
Dallas-Fort Worth-Arlington, TX	0.0343	0.8097	0.5937	0.0010	-0.0291
Denver-Aurora-Lakewood, CO	0.0145	0.8105	0.8220	0.0019	0.0387
Detroit-Warren-Dearborn, MI	0.0256	0.8747	0.5937	-0.0222	-0.1274
Hartford-West Hartford-East Hartford, CT	0.0086	0.8647	0.6850	0.0035	-0.0574
Houston-The Woodlands-Sugar Land, TX	0.0290	0.8046	0.5024	0.0099	-0.0130
Indianapolis-Carmel-Anderson, IN	0.0094	0.7599	0.5024	-0.0016	-0.0242
Jacksonville, FL	0.0080	0.7311	0.5024	0.0043	0.0046
Kansas City, MO-KS	0.0103	0.7612	0.5024	-0.0041	-0.0273
Knoxville, TN	0.0062	0.6507	0.4110	-0.0028	0.0174
Las Vegas-Henderson-Paradise, NV	0.0098	0.7897	0.6850	-0.0020	-0.0485
Los Angeles-Long Beach-Anaheim, CA	0.0873	0.8476	2.2834	0.0008	-0.0100
Louisville/Jefferson County, KY-IN	0.0085	0.7195	0.5024	-0.0101	-0.0305
Memphis, TN-MS-AR	0.0076	0.7573	0.5024	-0.0053	-0.0336
Miami-Fort Lauderdale-West Palm Beach, FL	0.0361	0.7512	0.5937	0.0026	-0.0739
Milwaukee-Waukesha-West Allis, WI	0.0072	0.7773	0.6850	-0.0074	-0.0629
Minneapolis-St. Paul-Bloomington, MN-WI	0.0147	0.8153	0.5937	0.0001	0.0133
Nashville-Davidson–Murfreesboro–Franklin, TN	0.0090	0.7408	0.5937	-0.0096	-0.0115
New Orleans-Metairie, LA	0.0088	0.7020	0.5024	0.0094	-0.0043
New York-Newark-Jersey City, NY-NJ-PA	0.1291	0.9392	1.6441	0.0021	0.0023
Oklahoma City, OK	0.0069	0.6688	0.4110	-0.0016	0.0781
Orlando-Kissimmee-Sanford, FL	0.0120	0.7072	0.5024	-0.0025	-0.0568
Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	0.0354	0.8434	0.5937	0.0022	-0.0170
Phoenix-Mesa-Scottsdale, AZ	0.0230	0.7666	0.5937	0.0036	-0.0188
Pittsburgh, PA	0.0168	0.7049	0.4110	0.0030	0.0341
Portland-Vancouver-Hillsboro, OR-WA	0.0130	0.7785	0.8220	0.0004	0.0521
Providence-Warwick, RI-MA	0.0115	0.7555	0.6850	-0.0013	0.0270
Raleigh, NC	0.0063	0.7799	0.6850	0.0076	-0.0164
Richmond, VA	0.0069	0.7613	0.5937	-0.0042	-0.0027
Riverside-San Bernardino-Ontario, CA	0.0200	0.8009	0.5937	-0.0079	-0.0221
Rochester, NY	0.0078	0.7295	0.4110	-0.0109	-0.0401
St. Louis, MO-IL	0.0157	0.7593	0.5024	-0.0023	-0.0296
San Antonio-New Braunfels, TX	0.0120	0.6663	0.4110	-0.0002	0.0177
San Diego-Carlsbad, CA	0.0193	0.7839	1.2787	0.0036	0.0775
San Francisco-Oakland-Hayward, CA	0.0308	0.9721	2.2834	0.0035	0.1313
San Jose-Sunnyvale-Santa Clara, CA	0.0123	1.0475	2.2834	0.0108	0.1388
Seattle-Tacoma-Bellevue, WA	0.0207	0.8325	1.0047	0.0017	0.1482
Tampa-St. Petersburg-Clearwater, FL	0.0181	0.7139	0.4110	0.0020	-0.0145
Virginia Beach-Norfolk-Newport News, VA-NC	0.0099	0.6673	0.5937	-0.0023	0.0734
Washington-Arlington-Alexandria, DC-VA-MD-WV	0.0341	0.8786	1.0047	0.0002	0.0447

Notes: This table displays statistics for each city in the dataset. $\hat{Z}^{\text{btk}}$ is the proportional change in productivity from the "Bartik" productivity shock, and $\hat{Z}^{\text{owg}}$ is the proportional change in productivity from the "Observed Wage" productivity shock.

Table 12: Estimated Location-Specific Parameters

City	Amenity	Productivity	Permit Price Shifter
Atlanta-Sandy Springs-Roswell, GA	1.0119	0.8217	0.1562
Austin-Round Rock, TX	0.9867	0.7599	0.1353
Baltimore-Columbia-Towson, MD	0.9744	0.8090	0.2414
Birmingham-Hoover, AL	0.8407	0.7439	0.0723
Boston-Cambridge-Newton, MA-NH	1.0862	0.8651	0.7146
Buffalo-Cheektowaga-Niagara Falls, NY	0.8893	0.7100	0.1187
Charlotte-Concord-Gastonia, NC-SC	0.9289	0.7698	0.1588
Chicago-Naperville-Elgin, IL-IN-WI	1.0940	0.8542	0.3750
Cincinnati, OH-KY-IN	0.9094	0.7766	0.1358
Cleveland-Elyria, OH	0.9746	0.7667	0.3088
Columbus, OH	0.9256	0.7684	0.1862
Dallas-Fort Worth-Arlington, TX	1.0378	0.8097	0.1011
Denver-Aurora-Lakewood, CO	1.0376	0.8105	0.1927
Detroit-Warren-Dearborn, MI	0.9261	0.8747	0.2517
Hartford-West Hartford-East Hartford, CT	0.8594	0.8647	0.3081
Houston-The Woodlands-Sugar Land, TX	0.9971	0.8046	0.0686
Indianapolis-Carmel-Anderson, IN	0.8883	0.7599	0.1380
Jacksonville, FL	0.9236	0.7311	0.0376
Kansas City, MO-KS	0.9206	0.7612	0.1332
Knoxville, TN	0.9005	0.6507	0.0505
Las Vegas-Henderson-Paradise, NV	0.9918	0.7897	0.0171
Los Angeles-Long Beach-Anaheim, CA	1.4968	0.8476	1.7810
Louisville/Jefferson County, KY-IN	0.9132	0.7195	0.1386
Memphis, TN-MS-AR	0.8771	0.7573	0.1075
Miami-Fort Lauderdale-West Palm Beach, FL	1.1284	0.7512	0.0073
Milwaukee-Waukesha-West Allis, WI	0.9156	0.7773	0.3110
Minneapolis-St. Paul-Bloomington, MN-WI	0.9521	0.8153	0.1268
Nashville-Davidson–Murfreesboro–Franklin, TN	0.9373	0.7408	0.1319
New Orleans-Metairie, LA	0.9697	0.7020	0.0986
New York-Newark-Jersey City, NY-NJ-PA	1.2628	0.9392	2.8874
Oklahoma City, OK	0.9337	0.6688	0.0694
Orlando-Kissimmee-Sanford, FL	1.0071	0.7072	0.0111
Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	0.9890	0.8434	0.2284
Phoenix-Mesa-Scottsdale, AZ	1.0736	0.7666	0.0397
Pittsburgh, PA	0.9749	0.7049	0.1273
Portland-Vancouver-Hillsboro, OR-WA	1.1023	0.7785	0.1434
Providence-Warwick, RI-MA	1.0053	0.7555	0.2998
Raleigh, NC	0.8869	0.7799	0.1167
Richmond, VA	0.8849	0.7613	0.1714
Riverside-San Bernardino-Ontario, CA	1.0236	0.8009	0.0134
Rochester, NY	0.8720	0.7295	0.0827
St. Louis, MO-IL	0.9573	0.7593	0.1526
San Antonio-New Braunfels, TX	1.0229	0.6663	0.0575
San Diego-Carlsbad, CA	1.2268	0.7839	0.1110
San Francisco-Oakland-Hayward, CA	1.2114	0.9721	2.8206
San Jose-Sunnyvale-Santa Clara, CA	1.0346	1.0475	1.7429
Seattle-Tacoma-Bellevue, WA	1.1332	0.8325	0.2218
Tampa-St. Petersburg-Clearwater, FL	1.0074	0.7139	0.0128
Virginia Beach-Norfolk-Newport News, VA-NC	1.0375	0.6673	0.0799
Washington-Arlington-Alexandria, DC-VA-MD-WV	1.0581	0.8786	0.5034

Notes: This table displays the estimated amenity A_i, productivity Z_i, and permit price shifter $\overline{p}_i^F$ for each city. These are chosen to match population shares, wages, and house prices in the initial stationary equilibrium.

Table 13: Values from Figure 1

Age Group	Homeownership Rate Model	Homeownership Rate Data	House Wealth Share Model	House Wealth Share Data	Migration Rate Model	Migration Rate Data	Wealth-Income Model	Wealth-Income Data
20-29	21.6%	36.2%	266.3%	171.0%	5.9%	6.6%	0.01	0.47
30-39	69.2%	63.1%	146.8%	117.7%	3.9%	3.7%	1.05	1.12
40-49	88.5%	76.3%	84.5%	81.5%	2.8%	2.1%	2.58	2.06
50-64	94.6%	82.8%	53.8%	68.0%	2.1%	1.7%	4.68	3.48
65+	75.2%	80.3%	76.7%	65.7%	1.5%	1.3%	6.22	7.30

Notes: This table displays the values shown in Figure 1. It compares statistics in the estimated model versus data over the lifecycle. The data sources are the 2001 Survey of Consumer Finances (for the homeownership rate, median house wealth share, and median wealth to income ratio) and 2005 American Community Survey (for the migration rate). See Section 3.1 for more details.

Table 14: Values from Figure 2

Percentile	House Wealth Share Model	House Wealth Share Data	Wealth-Income Model	Wealth-Income Data
10	35.9%	26.0%	0.00	0.16
20	44.0%	39.7%	0.24	0.49
30	53.5%	52.4%	0.78	0.96
40	63.9%	65.9%	1.42	1.54
50	78.3%	78.5%	2.35	2.35
60	99.3%	90.5%	3.48	3.43
70	133.8%	105.9%	4.91	5.12
80	192.8%	140.8%	6.90	7.81
90	324.0%	231.1%	10.34	12.87

Notes: This table displays the values shown in Figure 2. It compares selected percentiles of the house wealth share and wealth to income ratio distributions in the estimated model versus the data. The data source is the 2001 Survey of Consumer Finances. See Section 3.1 for more details.

Table 15: Values from Figure 3

Wealth Decile	Homeownership Rate Model	Data	House Wealth Share Model	Data
1	0.0%	4.6%		3198.7%
2	0.9%	23.3%	500.0%	312.5%
3	50.7%	59.7%	439.4%	233.0%
4	81.9%	73.8%	266.4%	138.7%
5	89.5%	84.7%	159.3%	100.0%
6	94.3%	91.8%	105.5%	85.1%
7	98.0%	93.3%	77.3%	72.7%
8	99.6%	94.1%	58.4%	61.7%
9	100.0%	95.5%	46.6%	47.1%
10	100.0%	96.3%	34.1%	30.9%

Notes: This table displays the values shown in Figure 3. It compares the homeownership rate and median house wealth share (among owners) within deciles of the wealth distribution in the calibrated model versus the data. The data source is the 2001 Survey of Consumer Finances. See Section 3.1 for more details.

Table 16: Estimated Parameter Values from Alternate Models

Parameter	Baseline	One Location	Two Locations	No Ownership	Agglomeration and Congestion
ρ	0.021	0.021	0.021	0.017	0.021
χ	1.13	1.03	1.08		1.14
η	0.21	0.21	0.21	0.22	0.21
$\overline{\kappa}$	13.62		6.06	14.22	14.84
κ^{age}	0.07		0.07	0.06	0.07
κ^{dst}	0.0015		0	0.0016	0.0017
ν	2.51		2.24	2.59	2.73
$\mathbb{H}^o$	{2.12,...,8.30}	{2.12,...,8.30}	{2.12,...,8.30}		{2.12,...,8.30}
$\mathbb{H}^r$	[0,8.30]	[0,8.30]	[0,8.30]	[0,8.30]	[0,8.30]

Notes: This table displays the internally estimated parameter values for the baseline model (column 2), one-location example (column 3), two-location example (column 4), no-ownership model (column 5), and model with agglomeration and congestion forces in production (column 6). All models use the same estimation procedure, which is described in Section 3.3. The one- and two-location models are described in Section 4.3, the no-ownership model is described in Section 4.4, and the agglomeration and congestion model is described in Appendix D.1. In the two-location model, the fixed component of migration costs $\overline{\kappa}$ is not separately identified from the distance component κ^{dst}, so without loss of generality I set $\kappa^{\text{dst}} = 0$. It is not possible to estimate the minimum owner-occupied house size to match the observed relationship between house prices and homeownership rates in the one- or two-location models, so I set the minimum house size equal to the value from the baseline estimation in these models. In the no-ownership model, the set of rental house sizes is set equal to that from the baseline estimation.

Table 17: Values from Figure 8

	Location 1		Location 2	
Age Group	Renters	Owners	Renters	Owners
20-29	0.73%	0.84%	0.56%	0.37%
30-39	0.73%	0.80%	0.51%	0.30%
40-49	0.69%	0.73%	0.46%	0.20%
50-64	0.44%	0.50%	0.36%	0.06%
65+	0.00%	0.28%	0.32%	-0.10%

Notes: This table displays the values for the bottom panels of Figure 8. These show the welfare effects of the two-location example by tenure and age group. In this example, the productivity of location 1 is permanently increased by 1%, while the productivity of location 2 is unchanged.

Table 18: Results of Counterfactuals for the Observed Wage Shock

	All Ages	20-29	30-39	40-49	50-64	65+
Baseline						
All	0.39	0.31	0.48	0.53	0.43	0.24
Renters	0.09	0.23	0.23	0.20	0.02	-0.36
Owners	0.51	0.60	0.59	0.57	0.46	0.46
No Ownership						
All	0.04	0.25	0.24	0.18	-0.02	-0.30
Wage Changes Only						
All	0.37	0.58	0.59	0.53	0.30	-0.00
Renters	0.47	0.58	0.66	0.67	0.43	0.00
Owners	0.33	0.57	0.56	0.51	0.29	0.00
House Price Changes Only						
All	0.04	-0.23	-0.08	0.02	0.14	0.24
Renters	-0.34	-0.31	-0.37	-0.40	-0.38	-0.36
Owners	0.19	0.05	0.04	0.07	0.17	0.46
No Capital Gains						
All	-0.04	0.23	0.20	0.12	-0.11	-0.46
Renters	0.09	0.23	0.23	0.20	0.02	-0.36
Owners	-0.09	0.22	0.19	0.11	-0.12	-0.49
Equalize House Price Sensitivity						
All	0.33	0.28	0.43	0.46	0.36	0.19
Renters	0.07	0.20	0.20	0.18	0.01	-0.36
Owners	0.44	0.54	0.53	0.50	0.38	0.39
No Transaction Costs						
All	0.39	0.31	0.48	0.52	0.43	0.26
Renters	0.09	0.23	0.23	0.20	0.02	-0.36
Owners	0.52	0.59	0.58	0.57	0.46	0.48

Notes: This table compares the estimated welfare elasticity from equation (28) for various subgroups in the baseline model with those of the counterfactuals described in Section 4.4. This table shows results for the Observed Wage productivity shock. Results for the Bartik productivity shock are shown in Table 9.

Figure 10: Local Real Estate Investment by Wealth Decile

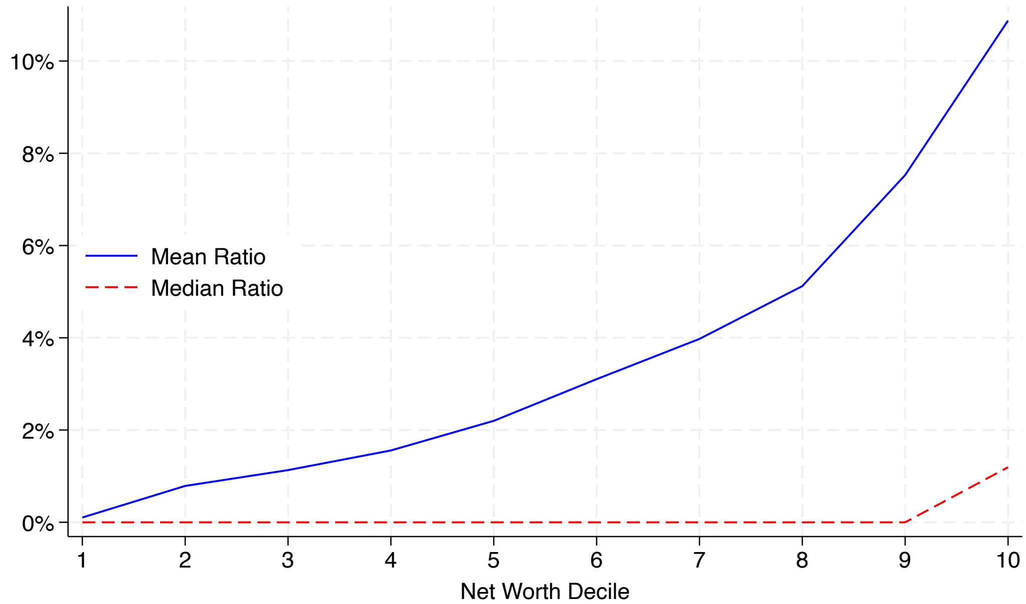

Notes: This figure shows the fraction of assets invested in non-primary real estate by net worth decile in the 2001 Survey of Consumer Finances. "Non-primary real estate" is the sum of "residential property excluding primary residence" and "net equity in non-residential real estate." The "Mean Ratio" series shows average non-primary real estate divided by average assets. The "Median Ratio" series shows median non-primary real estate divided by median assets. The economy-wide ratio of average non-primary real estate to average assets is 8.7%. This is the value I use for the local REIT share f^{lr}.

Figure 11: Scatter Plot of the Welfare Regression (28)

(a) Bartik

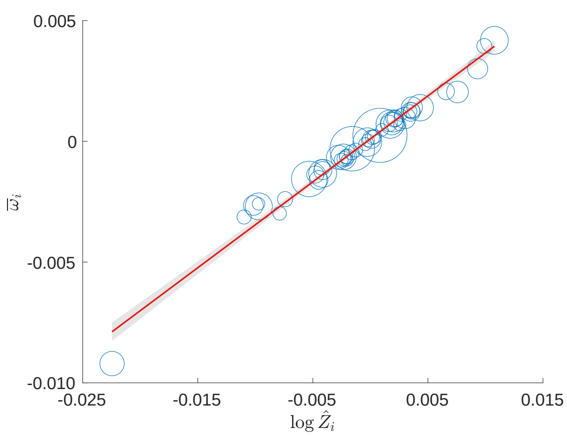

(b) Observed Wages

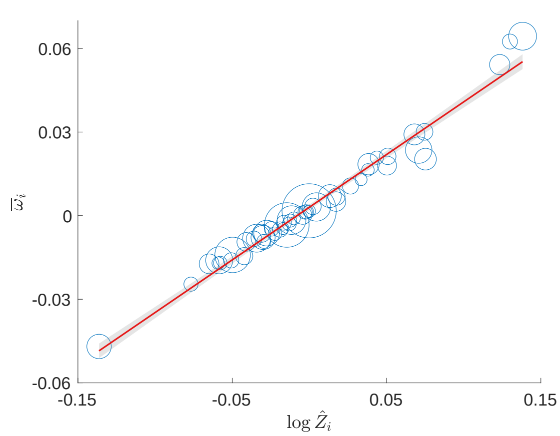

Notes: These figures show scatter plots of the welfare elasticity regression (28). $\overline{\omega}_i$ is the average welfare effect in location i, weighted by the density of state variables. $\hat{Z}_i$ is the proportional change in productivity from a shock. The left panel shows welfare effects from the Bartik productivity shock, and the right panel shows welfare effects from the Observed Wage productivity shock. The solid red lines are best fit lines, and the shaded areas are 95% confidence intervals.

B Mathematical Derivations

B.1 Location Probability Equation (12)

For notational clarity, throughout this Appendix I omit time subscripts. Denote the difference, net of location preferences, in the value of choosing location i' and i'' for a household with state Ω by

$$\bar{\varepsilon}_{i'i''}(\Omega) \equiv V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega;i'),h^{\mathrm{m}}(\Omega;i'),i',\ell,j) - V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega;i''),h^{\mathrm{m}}(\Omega;i''),i'',\ell,j) + \kappa_{ii''}(j) - \kappa_{ii'}(j).$$

The probability that the household chooses i' is

$$\pi(\Omega;i') = \int_{-\infty}^{\infty} f(\varepsilon_{i'}) \prod_{i'' \neq i'} F(\bar{\varepsilon}_{i'i''}(\Omega) + \varepsilon_{i'}) \mathrm{d}\varepsilon_{i'}.$$

The cumulative distribution and probability density functions of location preferences, respectively, are[43]

$$\begin{aligned} F(\varepsilon) &= \exp(-\exp(-\varepsilon/\nu - \overline{\gamma})) \\ f(\varepsilon) &= (1/\nu)\exp(-\varepsilon/\nu - \overline{\gamma})\exp(-\exp(-\varepsilon/\nu - \overline{\gamma})). \end{aligned}$$

Substitute into the previous expression, using $\bar{\varepsilon}_{i'i'}(\Omega) = 0$:

$$\begin{aligned} \pi(\Omega;i') &= \int_{-\infty}^{\infty} (1/\nu)\exp\left(-\varepsilon_{i'}/\nu - \overline{\gamma} - \exp(-\varepsilon_{i'}/\nu - \overline{\gamma}) - \sum_{i'' \neq i'} \exp(-[\bar{\varepsilon}_{i'i''}(\Omega) + \varepsilon_{i'}]/\nu - \overline{\gamma})\right) \mathrm{d}\varepsilon_{i'} \\ &= \int_{-\infty}^{\infty} (1/\nu)\exp\left(-\varepsilon_{i'}/\nu - \overline{\gamma} - \sum_{i''} \exp(-[\bar{\varepsilon}_{i'i''}(\Omega) + \varepsilon_{i'}]/\nu - \overline{\gamma})\right) \mathrm{d}\varepsilon_{i'}. \end{aligned}$$

Define $\zeta_{i'} \equiv \varepsilon_{i'}/\nu + \overline{\gamma}$. Then

$$\begin{aligned} \pi(\Omega;i') &= \int_{-\infty}^{\infty} (1/\nu)\exp\left(-\zeta_{i'} - \sum_{i''} \exp(-\bar{\varepsilon}_{i'i''}(\Omega)/\nu - \zeta_{i'})\right) \mathrm{d}\varepsilon_{i'} \\ &= \int_{-\infty}^{\infty} (1/\nu)\exp\left(-\zeta_{i'} - \exp(-\zeta_{i'}) \sum_{i''} \exp(-\bar{\varepsilon}_{i'i''}(\Omega)/\nu)\right) \mathrm{d}\varepsilon_{i'}. \end{aligned}$$

Define $\lambda_{i'}(\Omega) \equiv \log \sum_{i''} \exp(-\bar{\varepsilon}_{i'i''}(\Omega)/\nu)$. Then

$$\begin{aligned} \pi(\Omega;i') &= \int_{-\infty}^{\infty} \exp(-\zeta_{i'} - \exp(-(\zeta_{i'} - \lambda_{i'}(\Omega))))\mathrm{d}\zeta_{i'} \\ &= \exp(-\lambda_{i'}(\Omega)) \int_{-\infty}^{\infty} \exp(-(\zeta_{i'} - \lambda_{i'}(\Omega)) - \exp(-(\zeta_{i'} - \lambda_{i'}(\Omega))))\mathrm{d}\zeta_{i'} \end{aligned}$$

[43] $\overline{\gamma} \equiv \int_{-\infty}^{\infty} x \exp(-x - \exp(-x))\mathrm{d}x$ is the Euler-Mascheroni constant, which is added so that $\mathrm{E}(\varepsilon_i) = 0$.

where I have used $d\zeta_{i'} = (1/\nu)d\varepsilon_{i'}$. Define $y_{i'}(\Omega) \equiv \zeta_{i'} - \lambda_{i'}(\Omega)$. Then

$$\pi(\Omega; i') = \exp(-\lambda_{i'}(\Omega)) \int_{-\infty}^{\infty} \exp(-y_{i'}(\Omega) - \exp(-y_{i'}(\Omega))) \mathrm{d}y_{i'}(\Omega).$$

Since $\exp(-y_{i'}(\Omega) - \exp(-y_{i'}(\Omega)))$ is a probability density function (specifically, the Gumbel distribution with location parameter 0 and scale parameter 1), it follows that

$$\begin{aligned}
\pi(\Omega; i') &= \exp(-\lambda_{i'}(\Omega)) \\
&= \sum_{i''} \exp\left(\frac{V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) - [V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i''), h^{\mathrm{m}}(\Omega; i''), i'', \ell, j) - \kappa_{ii''}(j)]}{\nu} \right) \\
&= \frac{\exp(V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j))^{1/\nu}}{\sum_{i''} \exp(V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i''), h^{\mathrm{m}}(\Omega; i''), i'', \ell, j) - \kappa_{ii''}(j))^{1/\nu}}.
\end{aligned}$$

B.2 Expected Value of Optimal Location Choice (7)

The expected value of the optimal migration choice is

$$\begin{aligned}
V^{\mathrm{m}}(\Omega) &= \mathrm{E}_{\varepsilon} \max_{i'} [V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) + \varepsilon_{i'}] \\
&= \sum_{i'} \int_{-\infty}^{\infty} [V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) + \varepsilon_{i'}] f(\varepsilon_{i'}) \prod_{i'' \neq i'} F(\bar{\varepsilon}_{i'i''}(\Omega) + \varepsilon_{i'}) \mathrm{d}\varepsilon_{i'}.
\end{aligned}$$

Follow the steps of Appendix B.1 to write this as

$$\begin{aligned}
&\sum_{i'} \int_{-\infty}^{\infty} [V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) + \varepsilon_{i'}](1/\nu) \exp\left(-\varepsilon_{i'}/\nu - \overline{\gamma} - \sum_{i''} \exp(-[\bar{\varepsilon}_{i'i''}(\Omega) + \varepsilon_{i'}]/\nu - \overline{\gamma}) \right) \mathrm{d}\varepsilon_{i'} \\
&= \sum_{i'} \exp(-\lambda_{i'}(\Omega)) \int_{-\infty}^{\infty} [V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) + \varepsilon_{i'}] \exp(-y_{i'}(\Omega) - \exp(-y_{i'}(\Omega))) \mathrm{d}y_{i'}(\Omega) \\
&= \sum_{i'} \exp(-\lambda_{i'}(\Omega)) \int_{-\infty}^{\infty} [V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) + (y_{i'}(\Omega) - \overline{\gamma} + \lambda_{i'}(\Omega))\nu] \exp(-y_{i'}(\Omega) - \exp(-y_{i'}(\Omega))) \mathrm{d}y_{i'}(\Omega) \\
&= \sum_{i'} \exp(-\lambda_{i'}(\Omega)) \left[V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) + (\lambda_{i'}(\Omega) - \overline{\gamma})\nu + \nu \int_{-\infty}^{\infty} y_{i'}(\Omega) \exp(-y_{i'}(\Omega) - \exp(-y_{i'}(\Omega))) \mathrm{d}y_{i'}(\Omega) \right] \\
&= \sum_{i'} \exp(-\lambda_{i'}(\Omega)) \left[V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) + \lambda_{i'}(\Omega)\nu \right] \\
&= \sum_{i'} \pi(\Omega; i') [V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j) - \nu \log \pi(\Omega; i')]
\end{aligned}$$

where I have used the definition of $\overline{\gamma}$ in the second to last step. Since

$$-\nu \log \pi(\Omega; i') = -[V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i'), h^{\mathrm{m}}(\Omega; i'), i', \ell, j) - \kappa_{ii'}(j)] + \nu \log \sum_{i''} \exp(V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega; i''), h^{\mathrm{m}}(\Omega; i''), i'', \ell, j) - \kappa_{ii''}(j))^{1/\nu}$$

This simplifies to

$$\begin{aligned}&\sum_{i'}\pi(\Omega;i')\nu\log\sum_{i''}\exp(V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega;i''),h^{\mathrm{m}}(\Omega;i''),i'',\ell,j)-\kappa_{ii''}(j))^{1/\nu}\\=&\nu\log\sum_{i'}\exp(V^{\mathrm{h}}(a^{\mathrm{m}}(\Omega;i'),h^{\mathrm{m}}(\Omega;i'),i',\ell,j)-\kappa_{ii'}(j))^{1/\nu}.\end{aligned}$$

B.3 Derivation of Welfare Measure (27)

Recall the definition of the welfare measure $\omega(\Omega)$ (equation 26):

$$\mathrm{E}\left\{\int_j^J e^{-\rho(s-j)-\psi(s)}\log(A_s[1+\omega(\Omega)]\mathcal{C}_s)\mathrm{d}s+\sum_{k=\lceil j\rceil}^{J-1}e^{-\rho(k-j)}\Psi(k|j)(\varepsilon_k-\kappa_k)\right\}=V_0(\Omega).$$

This can be written as

$$\int_j^J e^{-\rho(s-j)-\psi(s)}\log(1+\omega(\Omega))\mathrm{d}s+\mathrm{E}\left\{\int_j^J e^{-\rho(s-j)-\psi(s)}\log(A_s\mathcal{C}_s)\mathrm{d}s+\sum_{k=\lceil j\rceil}^{J-1}e^{-\rho(k-j)}\Psi(k|j)(\varepsilon_k-\kappa_k)\right\}=V_0(\Omega).$$

Note that the term in brackets is simply the value function in the initial steady state. Therefore the previous expression can be written as

$$\int_j^J e^{-\rho(s-j)-\psi(s)}\log(1+\omega(\Omega))\mathrm{d}s+V_0^{\mathrm{ss}}(\Omega)=V_0(\Omega).$$

Rearrange to obtain the welfare measure:

$$\omega(\Omega)=\exp\left(\frac{V_0(\Omega)-V_0^{\mathrm{ss}}(\Omega)}{\int_j^J e^{-\rho(s-j)-\psi(s)}\mathrm{d}s}\right)-1.$$

C Derivation of the Permit Price Equation (21)

Saiz (2010) estimates equations of the form[44]

$$\Delta\log p_i=\mathrm{ela}_i\Delta\log H_i+\mathbf{\Gamma}\boldsymbol{X}_i+\epsilon_i \tag{29}$$

where Δ denotes the change between 1970 and 2000, H_i is number of housing units in city i, $\boldsymbol{X}_i$ is a vector of controls, and ϵ_i is a residual. Assuming that H_i is proportional to population level $\mathbb{N}_i$, (29) implies

$$p_{i2000}=\overline{p}_i\left(\frac{N_{i2000}\mathbb{N}_{2000}}{N_{i1970}\mathbb{N}_{1970}}\right)^{\mathrm{ela}_i}$$

[44]See equation (3) of Saiz (2010).

where $\overline{p}_i \equiv p_{i1970} \exp(\boldsymbol{\Gamma X}_i + \epsilon_i)$, N_{it} is i's population share in year t, and $\mathbb{N}_t = \sum_i \mathbb{N}_{it}$ is total population in year t. I calibrate the initial steady state of my model to match cross-sectional data from the year 2000, and abstract from population growth. Since I do not consider any counterfactuals that would change $\overline{p}_i$, within the context of the model $\overline{p}_i$ is constant and the total population level is fixed at $\mathbb{N}_{2000}$. The previous expression thus implies that in the long-run, local house prices and populations are related by

$$p_i = \overline{p}_i \left(\frac{N_i g^N}{N_{i1970}} \right)^{\text{ela}_i}$$

where $g^N \equiv \mathbb{N}_{2000}/\mathbb{N}_{1970}$ is population change since 1970. Equations (18) and (19) imply that the house price in steady state is[45]

$$p_i = (p_i^F)^{\xi}.$$

In order to match the long-run elasticities of house prices with respect to population ela_i estimated by Saiz (2010), I thus set

$$p_i^F = \overline{p}_i^F \left(\frac{N_i g^N}{N_{i1970}} \right)^{\lambda_i}$$

where $\overline{p}_i^F \equiv \overline{p}_i^{1/\xi}$ and $\lambda_i \equiv \text{ela}_i/\xi$.

D Robustness and Extensions

D.1 Agglomeration and Congestion Forces in Production

In the baseline model, there are no agglomeration or congestion forces in production. In this appendix, I examine how incorporating such forces affects the main results. To do this, I consider the alternative production function

$$\begin{aligned} Y_{it} &= \mathcal{Z}_{it}(L_{it}) L_{it}^{\alpha} K_{it}^{\vartheta} \mathcal{F}_{it}^{1-\alpha-\vartheta} \\ \mathcal{Z}_{it}(L_{it}) &= Z_{it} L_{it}^{\varsigma} \end{aligned} \tag{30}$$

where K is capital, $\mathcal{F}$ is a fixed factor of production, $\mathcal{Z}$ is endogenous productivity, and ς is the elasticity of local productivity with respect to employment. I assume that firms take productivity as given when choosing employment. I also assume that factor prices are set competitively (as in the baseline model), the return on capital R_t is set exogenously, and $\mathcal{F}$ is owned by absentee owners. If $\alpha + \vartheta < 1$, there is a congestion force in production due to the fixed factor. If $\varsigma > 0$, there is an agglomeration force in production because productivity is increasing in labor supply. Wages are

$$w_{it} = \overline{Z}_{it} L_{it}^{\frac{\alpha+\vartheta+\varsigma-1}{1-\vartheta}} \tag{31}$$

[45] Since housing stocks are constant in steady state ($\dot{H}_i = 0$), construction is positive and so house prices equal construction costs.

where

$$\overline{Z}_{it} \equiv \left[\alpha Z_{it} \left(\frac{\vartheta}{\alpha R_t} \right)^{\vartheta} \mathcal{F}_{it}^{1-\alpha-\vartheta} \right]^{\frac{1}{1-\vartheta}}$$

is a function of parameters. Following Herkenhoff et al. (2018), I set the labor share to $\alpha = 0.66$ and the capital share to $\vartheta = 0.29$, so that the fixed factor share is 0.05. For the agglomeration parameter ς, I use the mean estimate of 0.058 reported by Melo et al. (2009) in their meta-analysis of estimates of urban agglomeration economies. Since $(\alpha + \vartheta + \varsigma - 1)/(1 - \vartheta) = 0.0113 > 0$, agglomeration forces outweigh congestion in this calibration. The remaining parameters are estimated in the same way as the baseline model. The final column of Table 16 compares the estimated parameters from the model with agglomeration and congestion versus the baseline.

Table 19: Results with Agglomeration and Congestion Forces in Production

	Baseline	Agglomeration and Congestion	Hornbeck and Moretti (2024)
Long-Run Elasticities			
Labor*	4.03	4.03	4.03
	(0.104)	(0.11)	(1.52)
Wage	1.00	1.04	1.46
	(0.000)	(0.00)	(0.50)
House Price	2.47	2.49	3.05
	(0.071)	(0.07)	(0.98)
Rent	2.47	2.49	1.09
	(0.071)	(0.07)	(0.48)
Welfare			
Welfare Elasticity	0.375	0.379	
Standard Deviation of Welfare Effects	0.239%	0.242%	

* Targeted in estimation.

Notes: This table compares results in the baseline model (column 1) versus the model with agglomeration and congestion (column 2). The results in this table are from the Bartik productivity shock. The top panel shows long-run elasticities with respect to local productivity shocks, estimated using equation (25). The third column shows corresponding empirical estimates from Hornbeck and Moretti (2024). Standard errors are in parenthesis. The bottom panel shows the welfare elasticity, estimated using equation (28), and the standard deviation of welfare effects.

Table 19 compares key results from the baseline model versus the model with agglomeration and congestion. The results of this table are from the Bartik productivity shock. The first panel shows the average long-run effect of the shock on labor supplies, wages, house prices, and rents, as estimated in regression (25). Recall that both the baseline model and the model with agglomeration and congestion are estimated to match Hornbeck and Moretti (2024)'s empirical estimate of the long-run labor supply elasticity (4.03). The effect on wages is greater with agglomeration and congestion than the baseline (estimated coefficient of 1.04, versus 1 in the baseline), and closer to Hornbeck and Moretti's estimate of 1.46. The long-run house price and rent elasticity are little changed (2.49 with congestion and agglomeration versus 2.47 in the baseline).

The bottom panel of Table 19 shows the estimated welfare elasticity from equation (28) and the standard deviation of welfare effects. Quantitatively, agglomeration and congestion forces have little effect on the welfare elasticity (0.375 in the baseline versus 0.379 with agglomeration and congestion). The same is true of the dispersion of welfare effects (standard deviation of 0.239% in the baseline versus 0.242% with agglomeration and congestion).

These results suggest that agglomeration forces strengthen the pass-through from local productivity shocks to welfare, and increase the distributional effects of spatially heterogeneous productivity shocks. Agglomeration forces may be too weak in the baseline model, as evidenced by the fact that the long-run wage elasticity is lower in the baseline model than Hornbeck and Moretti (2024)'s empirical estimate. However, the main results from the baseline model are not substantially changed by incorporating agglomeration and congestion forces calibrated to standard values.

D.2 Asymmetric Effects of Productivity Shocks

The welfare elasticity regression (28) assumes a linear relationship between productivity shocks and welfare effects. As discussed in Section 3.4, productivity shocks have asymmetric effects on house prices and rents in the short run due to housing durability. In this Appendix, I explore whether they have asymmetric welfare effects as well.

I test for asymmetric welfare effects by estimating the regression

$$\overline{\omega}_i = \text{cons} + \beta_1 \log \hat{Z}_i + \beta_2 \max\{\log \hat{Z}_i, 0\} + \epsilon(\Omega). \tag{32}$$

(32) is equivalent to the baseline welfare elasticity regression (28) except that it allows negative and positive productivity shocks to have different effects. A significant estimate for β_2 suggests that productivity shocks have asymmetric welfare effects. $\beta_2 > 0$ indicates that positive productivity shocks benefit households proportionally more than negative shocks hurt them, and vice versa.

Table 20: Asymmetric Effects of Productivity Shocks

	Welfare	Rent	House Price
Bartik	-0.04	-1.69***	-0.35
	(0.02)	(0.17)	(0.21)
Observed Wage	0.08***	-1.90***	0.15
	(0.02)	(0.13)	(0.18)

Notes: The first column shows the estimated coefficient $\hat{\beta}_2$ from the asymmetric welfare elasticity regression (32). The second and third columns show the estimated coefficient $\hat{\gamma}_{2,0}$ from the regression (33) for rents and house prices, respectively. Standard errors are in parenthesis. * denotes significance at the 10% level, ** significance at the 5% level, and *** significance at the 1% level.

The first column of Table 20 shows the estimated coefficients $\hat{\beta}_2$ from (32). Standard errors are in parentheses. $\hat{\beta}_2$ is statistically insignificant with the Bartik shock, and significantly positive (at the 1% confidence level) with the Observed Wage shock. While we cannot reject the null hypothesis of no asymmetric welfare effects with the Bartik shock, there is evidence that positive productivity shocks benefit households proportionally more than negative productivity shocks hurt them with the Observed Wage shock.

In order to understand the sources of asymmetric welfare effects, I next examine the extent of asymmetric effects on rents and house prices.[46] To do this, I estimate a modified version of (25) that allows for asymmetry:

$$\log \hat{y}_{it} = \Lambda_t + \gamma_{1,t} \log \hat{Z}_i + \gamma_{2,t} \max\{\log \hat{Z}_i, 0\} + \epsilon_{it}. \tag{33}$$

As suggested by the intuition in Section 3.4, I focus on short-term effects $\hat{y}_{i0}$ (i.e. the proportional change immediately after the shock occurs). The second column of Table 20 shows the estimated coefficients $\hat{\gamma}_{2,0}$ for rents. For both the Bartik and Observed Wage shocks, $\hat{\gamma}_{2,0}$ is significantly negative for rents at the 1% confidence level. That is, negative productivity shocks reduce rents on impact proportionally more than positive shocks increase them. All else equal, lower rents benefit both renters and owners (due to lower user cost of housing). Hence the asymmetric response of rents is a source of asymmetric welfare effects. By itself, this channel makes households benefit proportionally more from a positive productivity shock than they are harmed by a negative productivity shock.

The third column of Table 20 shows the estimated coefficients for house prices. There is little evidence of asymmetry for house prices: $\hat{\gamma}_{2,0}$ is insignificantly negative with the Bartik shock, and insignificantly positive with the Observed Wage shock.[47] Recall that durable housing only implies asymmetric effects on rents and house prices in the short run. Since house prices are the discounted value of future rents, there is less asymmetry in short-run price effects than short-run rent effects. Since homeowners benefit from capital gains, by itself the house price channel could cause negative productivity shocks to hurt some owners proportionally more than positive productivity shocks help them (especially old owners, who care relatively more about capital gains and relatively less about housing costs). This channel is offset by the rent channel. In the case of the Bartik shock, the net effect is that there is little asymmetry in welfare effects. In the case of the Observed Wage shock, the net effect is that households gain proportionally more in locations with positive productivity shocks than they lose in locations with negative productivity shocks. Appendix Figure 11 shows the full scatter plots of average welfare effects versus productivity shocks.

[46] Recall that wages respond one-to-one to productivity shocks in the baseline model. Therefore productivity shocks have symmetric effects on wages.

[47] It may seem surprising that $\hat{\gamma}_{2,0}$ is (insignificantly) positive with the Observed Wage shock (that is, house prices increase proportionally more in locations with positive shocks than they decrease in locations with negative shocks). This across-location relationship is possible because of spatial heterogeneity in house price elasticities. In a *given* location, a negative shock reduces house price proportionally more than a positive shock increases it.

D.3 Number of Locations

In this appendix I examine the sensitivity of key results to the number of locations I. In each case, I use the I largest cities in the dataset, and estimate parameters using the same procedure as the baseline (described in Section 3.3). For this exercise, I focus on results from the Bartik productivity shock. In addition to the baseline case of $I = 50$, I also consider $I = 10, 25$, and 75.

The top row of Table 21 shows the welfare elasticity from equation (28) for various values of I. For the set of values considered, the elasticity varies between 0.34 ($I = 10$) and 0.38 ($I = 25$). The second and third rows of Table 21 show the welfare elasticities separately for renters and owners. For both groups, the welfare elasticity is fairly stable. For renters, it varies between 0.11 ($I = 25$) and 0.13 ($I = 10$). For owners, it varies between 0.41 ($I = 10$) and 0.47 ($I = 25$ and 50). For each value of I, the welfare elasticity is roughly 3-4 times larger for owners than renters.

These results suggest that the elasticity of welfare with respect to local productivity shocks is not very sensitive to the number of locations, at least for $I \geq 10$. Of course, the overall distributional effects of a shock depend not just on the welfare elasticity, but also on the shock each location receives, and how household characteristics covary with the shock. This is demonstrated in the fourth row of Table 21, which shows the standard deviation of welfare effects. The standard deviation with 10 locations is less than half of that from the baseline model. On the other hand, the ratio of standard deviation of welfare effects with versus without homeownership is fairly stable.

The appropriate economic geography for analyzing a particular shock depends on how the effects of the shock vary across space. Some shocks require more locations to properly model than others. A benefit of the efficient computational techniques developed in this paper is that they make it possible to analyze the distributional effects of spatially heterogeneous shocks in settings with rich geography.

Table 21: Results by Number of Locations I

	$I = 10$	$I = 25$	$I = 50$	$I = 75$
Welfare Elasticity: All	0.34	0.38	0.37	0.35
Welfare Elasticity: Renters	0.13	0.11	0.11	0.12
Welfare Elasticity: Owners	0.41	0.47	0.47	0.44
Welfare Effect Std. Dev.	0.10%	0.23%	0.24%	0.25%
Welfare Effect Std. Dev, No Ownership	0.06%	0.12%	0.13%	0.14%

Notes: This table displays results for varying numbers of locations I. In the baseline model, $I = 50$. The results in this table pertain to the Bartik productivity shock. The welfare elasticity is the estimated coefficient $\hat{\beta}$ from equation (28). "No Ownership" refers to the No Ownership model, described in Section 4.4.

D.4 Local Ownership of Rental Housing

In the baseline model, households invest fraction f^{lr} of their liquid assets in local rental housing and fraction f_t^{gr} in a spatially diversified global REIT. The local REIT portfolio share f^{lr} is calibrated to match SCF data (as described in Section 3.3.1) and f_t^{gr} is chosen so all rental housing is owned by households in the model. A challenge for this calibration strategy is that the SCF does not report ownership of local residential rental properties. It does include two categories of real estate investment besides primary residence (residential property excluding primary residence, and net equity in non-residential real estate), but there is no way of knowing whether these are held locally and whether these are residential rental properties. Furthermore, the SCF does not provide information about indirect investments in local residential real estate. As a result, the data does not provide a clear picture of how much households invest in local versus spatially diversified rental properties.

In this appendix I perform a robustness check where all rental housing is locally owned. Specifically, I set $f_t^{\text{gr}} = 0$ and f_{it}^{lr} such that all rental housing in each location is owned by local households. Note that for this exercise, the local REIT share varies both across time and locations. The motivation for this exercise is to understand how assumptions about local and global REIT portfolio shares affect the main results of the paper.

Table 22 compares results from the Bartik productivity shock in the model with full local ownership versus the baseline. Relative to the baseline, the model with full local ownership has a higher elasticity of welfare with respect to local productivity shocks (0.40 in the local ownership model versus 0.37 in the baseline). The elasticity increases by similar amounts relative to the baseline for both renters and owners. The standard deviation of welfare effects across all households is also slightly higher in the model with full local ownership.

The reason that the welfare elasticity is higher in the full local ownership model is that households are more exposed to local real estate values. In the baseline calibration, households invest $f^{\text{lr}} = 8.7\%$ of their liquid assets in local REITs. In the full local ownership model the average local REIT share is 24.3%.[48] While this does strengthen the link between individual welfare and local economic conditions relative to the baseline model, the difference is not large. The reason is simply that households make much larger investments in owner-occupied housing than rental properties, even in the full local ownership model. In the model (both the baseline and with full local ownership) the total value of owner-occupied housing is approximately 5.3 times the total value of rented housing. In both the data and model, owner-occupied housing is by far the largest investment (often levered) for the typical household. Although the data provides limited information regarding how much households invest in local non-owner-occupied real estate, it is much less than they invest in owner-occupied real estate even under the assumption that all real estate is locally owned.[49] While other

[48] Across locations, local REIT shares vary from 15.9% to 47.1% in the full local ownership model. The across-location standard deviation is 6.8%.

[49] Owner-occupied housing is likely even more important relative to other real estate investments for the *typical* household. In the 2001 SCF, 71.7% of households own a home, while 18.0% of households own *any* other real estate. While the

real estate holdings play a role, the main investment channel that links individual welfare to local economic conditions is owner-occupied housing.

Table 22: Results for Model with Full Local Ownership

	Baseline	Full Local Ownership
Welfare Elasticity: All	0.37	0.40
Welfare Elasticity: Renters	0.11	0.12
Welfare Elasticity: Owners	0.47	0.50
Welfare Effect Std. Dev.	0.24%	0.25%

Notes: This table displays results for the model with full local ownership versus results for the baseline model. In the full local ownership model, the global REIT portfolio share f_t^{gr} is set to 0 and the local REIT portfolio share f_{it}^{lr} is set so that all rental housing is owned by local households. In the baseline model, the local REIT share (which does not vary across time or locations) is $f^{\text{lr}} = 8.7\%$ and the global REIT share is $f^{\text{gr}} = 15.6\%$. In both models, all rental housing is owned by households in the model.

D.5 Robustness to Excluding Worst-Fit Cities

In the last part of Section 3.4, I assess the model's ability to match untargeted geographic sorting patterns. Figure 4 shows observed and model-predicted homeownership and retirement rates across cities. Although these are positively correlated, the fit is not perfect. In this appendix, I assess robustness to excluding cities with the largest deviations between observed and model-implied homeownership and retirement rates.

To implement this robustness check, I re-estimate the welfare elasticity regression (28) after excluding the cities with the largest discrepancies between the model and data. Specifically, I first exclude the 10% worst-fitting cities in terms of homeownership, and then then 10% worst-fitting in terms of fraction of residents age 65 and older. As in Section 4.2, I estimate the welfare regression first for all households, and then separately for renters and owners.

Table 23 reports the results of this robustness check. Excluding the worst-fit cities has very little impact on the estimated welfare elasticities. In both cases, the welfare elasticity for the entire population and for homeowners changes by less than 0.004. The largest change is the welfare elasticity for renters when the worst-fit cities by homeownership rate are excluded. In this case, the welfare elasticity declines by 0.007. In all cases, the key result that the average welfare elasticity is high, and is much higher for owners than renters, remains unchanged.

owner-occupied housing wealth share is high for the bulk of the population, non-owner-occupied real estate investments are much more skewed toward the top of the wealth distribution: see Figures 3 and 10.

Table 23: Welfare Elasticities Excluding Worst-Fit Cities

	Baseline	Drop 10% Worst-Fit: Homeownership Rate	Drop 10% Worst-Fit: Fraction 65+
Welfare Elasticity: All	0.375	0.372	0.376
Welfare Elasticity: Renters	0.115	0.107	0.114
Welfare Elasticity: Owners	0.466	0.469	0.467

Notes: This table reports the estimated welfare elasticities from regression (28) in the baseline specification and when the 10% worst-fitting cities in terms of homeownership (column 3) and fraction of households age 65 and older (column 4) are excluded. Rows report estimates for all households, and separately for renters and homeowners. The results in this table are from the Bartik productivity shock.

D.6 Migration Elasticity Target

As discussed in Section 3.3.2, the baseline estimation targets the long-run migration elasticity of 4.03 estimated by Hornbeck and Moretti (2024). This target is well-suited to my estimation strategy because it uses city-level data, focuses on long-run responses (30 years), and estimates the effects of plausibly exogenous local productivity shocks. This aligns closely with my quantitative analysis.

Nonetheless, the literature reports a variety of alternative migration elasticity estimates. In their survey of the literature, Dustmann et al. (2016) report estimates ranging from 1.5 to 5. Most of these use much shorter time horizons than Hornbeck and Moretti (2024), and so are not directly comparable (and are likely lower than the long-run elasticity). In this appendix, I analyze how the migration elasticity target affects results by re-estimating the model to match a lower migration elasticity target drawn from this range. Specifically, I re-estimate the model using the procedure described in Section 3.3.2, with the migration elasticity target replaced with 3.25.

The results of this exercise are shown in Table 24. A lower migration elasticity implies less directed migration in response to local shocks. This affects welfare elasticities in two ways. First, migration is a less effective adjustment mechanism. All else equal, this increases welfare elasticities. Second, it implies that house prices and rents are less responsive to local shocks. This is shown in the fifth row of Table 24: lowering the migration elasticity target reduces the average long-run house price and rent elasticity from 2.47 to 2.08. This dampens capital gains for owners and weakens the rent-offset channel for renters. All else equal, this lowers the welfare elasticity for owners and raises it for renters.

The net effect is to raise the estimated welfare elasticity for renters from 0.11 to 0.15 and to lower it for owners from 0.47 to 0.44. Both the welfare elasticity among all households and the standard deviation of welfare effects are nearly unchanged. Although the estimated welfare elasticities are quantitatively affected by the migration elasticity target, the main qualitative results of the paper are unchanged. Lowering the migration elasticity target reduces the gap between the welfare elasticities of owners and renters, but the difference is still large.

Table 24: Results with Alternative Migration Elasticity Target

	Baseline Target (4.03)	Alternative Target (3.25)
Migration Elasticity Parameter ν	2.51	3.68
Welfare Elasticity: All	0.37	0.36
Welfare Elasticity: Renters	0.11	0.15
Welfare Elasticity: Owners	0.47	0.44
House Price Elasticity	2.47	2.08
Welfare Effect Std. Dev.	0.24%	0.23%

Notes: This table reports baseline results (column 1) and results when the alternative migration elasticity is targeted (column 2). The baseline migration elasticity target, which is taken from Hornbeck and Moretti (2024), is 4.03. The alternative migration elasticity is 3.25. In both cases, all parameters are estimated as described in Section 3.3.2. The results in this table are from the Bartik productivity shock.

E Computational Appendix

In this appendix, I describe the procedures for numerically solving and estimating the model. Section E.1 describes how to solve the household's problem and the density of state variables. In Section E.2, I describe how I solve for stationary equilibria. The estimation algorithm is presented in Section E.3. Section E.4 describes the algorithm for computing transition dynamics after an unexpected shock. Finally, Section E.5 shows the amount of time required to estimate the model and compute counterfactuals.

E.1 Household's Problem

Overview In this Section, I describe the algorithm for solving the household's value function V and the density of state variables g. The algorithm is based on the upwinding method described by Achdou et al. (2022). It can be used to solve V and g in either a stationary equilibrium or at a point in time along a transition path. At a high level, the algorithm proceeds as follows. The value function at the maximum age J is 0. Starting from this boundary condition, we solve the Hamilton-Jacobi-Bellman equations (5)-(8) backward from age J, time $t+J$ to age 0, time t (in a stationary equilibrium, the value function does not vary with time). At each age/time, we use the upwinding algorithm to compute the derivative of the value function with respect to wealth. Given this, we can recover the household policies in closed-form using the first-order conditions of the household's problem. The policy functions at each age/time, which are needed to solve g, are saved. Once we have iterated backward over the entire lifecycle in this fashion, we have the value and policy functions.

The density of state variables at birth (described in Section 2), is a known function of parameters and the age-0 value function. Starting from this boundary condition, we solve the Kolmogorov

Forward equations (9)-(13) forward from age 0, time t to age J, time $t+J$. This step again uses the upwinding algorithm and the policy functions computed when solving V. Once we have iterated forward over the entire lifecycle in this fashion, we have the density of state variables.

Details It is convenient to replace the state variable a (total wealth) with the alternative state variable $x \equiv b + \phi p_{it} h$. x is liquid wealth held in excess of the collateral constraint. I refer to x as "voluntary wealth." Using this variable is helpful because the collateral constraint can be written as $x \geq 0$ (see equation 1). This allows the voluntary wealth grid to be the same for all agents, regardless of owner-occupied house size.[50]

I discretize voluntary wealth x using an unevenly spaced grid that concentrates nodes near the borrowing limit $x = 0$ (where the curvature of the value and policy functions is highest). I discretize age j using an evenly spaced grid with interval length 1 and the AR(1) process for the idiosyncratic component of labor endowments ℓ using Rouwenhorst's method (Rouwenhorst, 1995 and Kopecky and Suen, 2010). This method produces a transition matrix that I use to compute the expectation over labor endowment innovations at shock ages.

Given the value function at the maximum age $V_t(x, h, i, \ell, J) = 0$, I solve the Hamilton-Jacobi-Bellman (HJB) equation (5) using the finite differences method described by Achdou et al. (2022) to obtain the value function immediately after the last shock age, $\lim_{\iota \downarrow 0} V_{t-1+\iota}(x, h, i, \ell, J-1+\iota)$. The upwinding method yields the derivative of the value function with respect to voluntary wealth $\partial_x V(\Omega)$. The first-order conditions of the HJB equation are (the second applies only to renters):

$$c = \begin{cases} \frac{1-\eta}{\partial_x V_t(\Omega)} & \text{if } x > 0 \\ \min\left\{\frac{1-\eta}{\partial_x V_t(\Omega)},\ (1-\tau_t)[y_t(\Omega) + rb] - p^r_{it} h^r - \delta p_{it} h + \phi \dot{p}_{it} h\right\} & \text{otherwise} \end{cases}$$

$$h^r = \min\left\{\frac{\eta}{p^r_{it} \partial_x V_t(\Omega)},\ \overline{h}^o\right\}.$$

Next, I evaluate equations (6)-(8) to obtain the value function at the last shock age, $V_{t-1}(x, h, i, \ell, J-1)$. This requires evaluating the value function off of the discrete x grid, which I do using linear interpolation. I iterate backward in this manner from age J to 0 to obtain the value and optimal policy functions throughout the lifecycle.

The density of state variables is discretized on the same grids as the policy and value functions. The density of state variables at age 0, $g_t(x, h, i, \ell, 0)$, is as described in Section 2. Given the density of state variables at age 0, I again use the upwinding algorithm to solve the Kolmogorov Forward (KF) equation (9) for the density of state variables immediately before the first shock age, $\lim_{\iota \downarrow 0} g_{t+1-\iota}(x, h, i, \ell, 1-\iota)$. I then evaluate equations (10)-(13) to obtain the density of state variables at the first shock age, $g_{t+1}(x, h, i, \ell, 1)$. I use linear interpolation to assign mass that falls between voluntary wealth gridpoints to the adjoining nodes. I iterate forward in this manner from age 0 to J

[50] Recall that the collateral constraint requires $a \geq (1-\phi) p_{it} h$.

to obtain the density of state variables throughout the lifecycle.

E.2 Computing Stationary Equilibrium

Overview In this appendix, I describe the algorithm for computing stationary equilibria. A stationary equilibrium is one in which all equilibrium objects (listed in Section 2.6) are time-invariant. Computing a stationary equilibrium requires finding an tax rate τ, vector of wages w_i, and vector of house prices p_i such that the government's budget balances and labor and housing markets clear in every location. At a high level, the algorithm for computing stationary equilibrium proceeds as follows. First, guess the income tax rate, wages, and house prices (given house prices, rents are implied by equation 20). Then use the algorithm described in Appendix E.1 to solve the household's value function and density of state variables. Integrate over the density of state variables to compute tax revenue, labor supplies, and housing demands. Given these, compute equilibrium condition violations. We use a fixed-point iteration algorithm to search for a vector of tax rates and prices that balance the government's budget and clear all markets. In practice, this algorithm converges more quickly than derivative-based root finding algorithms, especially when the number of locations is large.

Details I compute stationary equilibria using the following fixed-point iteration algorithm:

1. Guess the income tax rate τ, wages w_i, and house prices p_i. In steady state, rents are $p_i^r = (r+\delta)p_i$ (see equation 20)

2. Given the guess for tax rate and prices, compute the density of state variables using the algorithm described in Section E.1. Integrate over this density to compute labor supplies L_i, populations N_i, housing stocks H_i, aggregate income INC, and aggregate accidental bequests.

3. Denote government revenue from sources other than income taxes (i.e. permit sales and accidental bequests) by REV^{oth} and expenditure (pensions and other government spending G) by EXP. In equilibrium, the government's budget must balance:

$$REV^{oth} + \tau INC = EXP$$
$$\Rightarrow \tau = (EXP - REV^{oth})/INC$$

   In equilibrium, wages must satisfy equation (31) (which collapses to the baseline wage equation (15) when $\alpha + \vartheta = 1$ and $\varsigma = 0$):

$$w_i = \overline{Z}_i L_i^{\frac{\alpha+\vartheta+\varsigma-1}{1-\vartheta}}.$$

   Finally, in a stationary equilibrium, construction is strictly positive and so house prices equal construction costs. By equations (19) and (21), equilibrium house prices satisfy

$$p_i = [\overline{p}_i^F (N_i g^N / N_{i1970})^{\lambda_i}]^{\xi}.$$

Check for convergence. If

$$\begin{aligned}
&|REV^{oth} + \tau INC - EXP| < \varepsilon \\
\text{and } &\max_i |w_i - \overline{Z}_i L_i^{\frac{\alpha+\vartheta+\varsigma-1}{1-\vartheta}}| < \varepsilon \\
\text{and } &\max_i |p_i - [\overline{p}_i^F (N_i g^N / N_{i1970})^{\lambda_i}]^{\xi}| < \varepsilon
\end{aligned}$$

for the numerical tolerance parameter $\varepsilon < 0$, stop. Otherwise, update the guess for tax rate, wages, and house prices using

$$\begin{aligned}
\tau &= \tau + \nabla((EXP - REV^{oth})/INC - \tau) \\
w_i &= w_i + \nabla(\overline{Z}_i L_i^{\frac{\alpha+\vartheta+\varsigma-1}{1-\vartheta}} - w_i) \\
p_i &= p_i + \nabla([\overline{p}_i^F (N_i g^N / N_{i1970})^{\lambda_i}]^{\xi} - p_i)
\end{aligned}$$

where $\nabla \in (0, 1]$ is a damping parameter and return to step (2).[51]

E.3 Estimation

This appendix describes the procedure for estimating the internally estimated parameters discussed in Section 3.3.2. The estimation procedure has a nested loop structure. In the inner loop, all the internally estimated parameters except the migration elasticity parameter ν are estimated, conditional on a value for ν. Since these parameters are chosen to match targets in the initial steady state, the inner loop involves solving the initial stationary equilibrium. The outer loop estimates ν, conditional on all other parameters. Since ν is chosen to match the long-run migration elasticity, the outer loop involves solving the long-run stationary equilibrium after the Bartik productivity shock. Since the number of estimated parameters equals the number of estimation targets, the estimation procedure matches every target to within numerical precision.

E.3.1 Inner Loop: Estimate Parameters Conditional on ν

Conditional on a value of ν, the remaining internally estimated parameters are ρ, χ, η, $\overline{\kappa}$, κ^{age}, κ^{dst}, $\underline{h}^o$, and $\{A_i, Z_i, \overline{p}_i^F\}_{i=1}^I$. The estimation target for the amenity vector $\{A_i\}_{i=1}^I$ is the vector of observed population shares. Since population shares always sum to 1, a normalization of the amenity vector is required. Without loss of generality, I set the amenity in location 1 to $A_1 = I - \sum_{i=2}^I A_i$, so that

[51] I use $\varepsilon = 10^{-4}$ and $\nabla = 0.1$.

the average amenity equals 1. The estimation targets for the vectors of productivities and permit price shifters $\{Z_i, \overline{p}_i^F\}_{i=1}^I$ are the vectors of observed wages and house prices. Households do not care about these parameters *per se*: all they care about are wages and house prices. In the estimation procedure, I solve the household's problem conditional on observed wages and house prices. As shown below, the $\{Z_i, \overline{p}_i^F\}_{i=1}^I$ values that exactly match observed wages and house prices can then be backed out in closed-form.

The inner-loop algorithm is as follows:

1. Guess a vector $\Theta \equiv (\rho, \chi, \eta, \overline{\kappa}, \kappa^{\text{age}}, \kappa^{\text{dst}}, \underline{h}^o, \{A_i\}_{i=2}^I)$. Set $A_1 = I - \sum_{i=2}^I A_i$, so that the average amenity is normalized to 1.

2. Given the parameters (Θ, A_1, ν) and observed wages, and house prices $\{w_i, p_i\}_{i=1}^I$, a stationary equilibrium can be computed using the algorithm described in Appendix E.2. Define the objective function $F(\Theta)$, which returns the vector of differences between estimation targets and corresponding model-implied moments for every element of Θ. That is, $F(\Theta)$ is a function of $I+7$ parameters that returns $I+7$ deviations from estimation targets.

3. Find a vector Θ^* that satisfies
$$\max |F(\Theta^*)| < \epsilon$$
for a numerical tolerance parameter $\epsilon > 0$.[52]

4. Step (3) yields populations and labor supplies $\{N_i, L_i\}_{i=1}^I$ in the initial steady state. Use equations (15), (19), and (21) to back out productivities and permit price shifters that exactly match observed wages w_i^{data} and house prices p_i^{data} in the initial steady state:
$$Z_i = w_i^{\text{data}}$$
$$\overline{p}_i^F = \left(\frac{N_i g^N}{N_{i1970}}\right)^{-\lambda_i} (p_i^{\text{data}})^{1/\xi}$$
Note that in the initial steady state, $p_i = (p_i^F)^\xi$ since construction is strictly positive in all locations.

Conditional on a value of ν, the inner loop finds a parameter vector that matches all estimation targets except the long-run migration elasticity to within numerical precision ϵ.

E.3.2 Outer Loop: Estimate ν

The outer loop of the estimation procedure uses bisection to find a value of ν that matches the long-run migration elasticity. In each iteration of the outer loop, the remaining parameters of the model are estimated. As a result, once the outer loop converges, the estimation procedure is complete.

[52]In my numerical implementation, I use Matlab's fsolve function to find Θ^* and set $\epsilon = 5 \times 10^{-4}$.

The outer loop uses an objective function $f^{\text{nu}}(\nu)$ that measures the distance between observed and model-implied migration elasticity. $f^{\text{nu}}(\nu)$ is defined as follows. For a given value of ν, estimate the remaining parameters of the model using the inner loop algorithm described in Appendix E.3.1. Then use the algorithm described in Appendix E.2 to compute the long-run steady state after the Bartik productivity shock. Using the model-generated data from the initial and long-run steady states, estimate $\hat{\alpha}^{\text{mig,model}}$ from the auxiliary regression (23). The objective function is

$$f^{\text{nu}}(\nu) = \hat{\alpha}^{\text{mig,data}} - \hat{\alpha}^{\text{mig,model}}$$

where $\hat{\alpha}^{\text{mig,data}} = 4.03$.

The outer-loop algorithm is as follows:

1. Find a bracket $(\underline{\nu}, \overline{\nu})$ such that $0 < \underline{\nu} < \overline{\nu}$, $f^{\text{nu}}(\underline{\nu}) < 0$, and $f^{\text{nu}}(\overline{\nu}) > 0$.
2. Set $\nu = (\underline{\nu} + \overline{\nu})/2$. If $|f^{\text{nu}}(\nu)| < \epsilon^{\text{nu}}$ for a numerical tolerance parameter $\epsilon^{\text{nu}} > 0$, stop.[53] Otherwise, proceed to step (3).
3. Set

$$\begin{aligned} \underline{\nu} &= \nu \quad \text{if } f^{\text{nu}}(\nu) < 0 \\ \overline{\nu} &= \nu \quad \text{if } f^{\text{nu}}(\nu) > 0 \end{aligned}$$

   and return to step (2).

As long as f^{nu} is continuous (which it appears to be in practice), the outer loop will converge to a value of ν that matches the migration elasticity to within ε^{nu}. Since all other parameters are estimated each time f^{nu} is evaluated, at that point all estimation targets are matched to within the specified numerical tolerance and the estimation procedure is complete.

E.4 Transition Dynamics

Overview In this section, I describe the algorithm for computing transition dynamics after an unexpected shock. I assume that the economy is in its initial steady state at $t = 0$, when the shock occurs. The economy eventually converges to a new steady state. I assume (and verify ex-post) that the economy has approximately converged to its long-run stationary equilibrium by time T, where T is large but finite. The outline of the algorithm for computing transition dynamics is as follows. First, compute the initial and long-run equilibria using the algorithm described in Appendix E.2. This yields the value function at time T, V_T, and the density of state variables at time 0, g_0. Next, guess paths of tax rates, wages, and house prices. Starting with V_T, iterate backward from $t = T$ to $t = 0$ on the HJB equations (5)-(8) as described in Appendix E.1 to obtain the value and policy

[53] In my numerical implementation, I set $\epsilon^{\text{nu}} = 5 \times 10^{-3}$.

functions along the transition path. Next, beginning with the initial density of state variables g_0, iterate forward on the Kolmogorov forward equations (9)-(13) from $t = 0$ to $t = T$ as described in Appendix E.1 to obtain the density of state variables along the transition path. Integrate over the density of state variables to obtain market clearing violations along the entire path. We iterate on paths of tax rates and prices using a fixed-point algorithm very similar to the one used to compute stationary equilibria described in Appendix E.2 until paths of tax rates and prices that balance the government's budget and clear all markets along the transition path are found.

Details The algorithm for computing transition dynamics is as follows:

1. Compute the initial ($t = 0$) and final ($t = \infty$) steady states using the algorithm described in Section E.2.

2. Guess paths of income tax rates τ_t, wages w_{it}, house prices p_{it}, and an indicator that is 1 if construction is strictly positive and 0 otherwise, $\mathbf{1}_{it}^{\text{cns}}$. This guess should converge to the final steady state values well before $t = T$.[54]

3. Given the guess for house prices p_{it}, compute rents p_{it}^r using equation (20).

4. Given the value function at time t and prices at $t-1$, compute the value and policy functions at time $t-1$ as described in Appendix E.1. Iterating backwards from $t = T$ (when the value function $V_T = V_\infty$, computed in step 1, is known) to $t = 0$ yields the value and policy functions along the entire transition path.

5. The density of state variables immediately before the shock, g_0, is computed in step 1. As shown in Section 3, house prices jump immediately when a shock occurs. This implies that voluntary wealth x jumps as well (both for owners and renters with positive assets, due to REIT holdings). The new values of x will, in general, not lie on the discretized grid. Use linear interpolation to assign mass that falls between voluntary wealth gridpoints to the adjoining nodes. This yields the discretized density of state variables at $t = 0$ (immediately after the shock). Given the density of state variables at t, compute the density at $t+1$ as described in Appendix E.1. Iterating forward from $t = 0$ to $t = T$ yields the density of state variables along the entire transition path.

6. Integrate over the density of state variables to compute the paths of labor supplies L_{it}, populations N_{it}, housing stocks H_{it}, aggregate income INC_t, and accidental bequests. In equilibrium, the government's budget balances:

$$REV_t^{oth} + \tau_t INC_t = EXP.$$

[54]Since construction is strictly positive in steady state, $\mathbf{1}_{i\infty}^{\text{cns}} = 1$.

In addition, wages satisfy equation (31) (which collapses to the baseline wage equation (15) when $\alpha + \vartheta = 1$ and $\varsigma = 0$):

$$w_{it} = \overline{Z}_{it} L_{it}^{\frac{\alpha+\vartheta+\varsigma-1}{1-\vartheta}}.$$

Conditional on the construction indicator guess $\mathbf{1}_{it}^{\text{cns}}$, the house market clearing conditions (18) and (19) imply that equilibrium house prices and demands satisfy

$$\begin{aligned} p_{it} &= (p_{it}^F)^{\xi} \text{ if } \mathbf{1}_{it}^{\text{cns}} = 1 \\ H_{it} &= H_{it}^{\text{nc}} \text{ if } \mathbf{1}_{it}^{\text{cns}} = 0 \end{aligned}$$

where

$$H_{it}^{\text{nc}} = \begin{cases} H_{i0}^{\text{ss}} & \text{if } t = 0 \\ (1-\delta) H_{it-1} & \text{otherwise} \end{cases}$$

is the supply of housing when there is no construction and permit prices are (see equation 21):

$$p_{it}^F = \overline{p}_i^F \left(\frac{N_{i\infty} g^N}{N_{i1970}} \right)^{\lambda_i}.$$

Check for convergence conditional on $\mathbf{1}_{it}^{\text{cns}}$. If

$$\begin{aligned} &\max_t |REV_t^{oth} + \tau_t INC_t - EXP| < \varepsilon \\ &\max_{i,t} |w_{it} - \overline{Z}_{it} L_{it}^{\frac{\alpha+\vartheta+\varsigma-1}{1-\vartheta}}| < \varepsilon \\ \text{and } &\max_{i,t|\mathbf{1}_{it}^{\text{cns}}=1} |p_{it} - (p_{it}^F)^{\xi}| < \varepsilon \\ \text{and } &\max_{i,t|\mathbf{1}_{it}^{\text{cns}}=0} |H_{it} - H_{it}^{\text{nc}}| < \varepsilon \end{aligned}$$

for the numerical tolerance parameter $\epsilon > 0$, proceed to step (7). Otherwise, update the guess for tax rates, wages, and house prices using

$$\begin{aligned} \tau_t &= \tau_t + \nabla((EXP - REV_t^{oth})/INC_t - \tau_t) \\ w_{it} &= w_{it} + \nabla(\overline{Z}_{it} L_{it}^{\frac{\alpha+\vartheta+\varsigma-1}{1-\vartheta}} - w_{it}) \\ p_{it} &= \begin{cases} p_{it} + \nabla[(p_{it}^F)^{\xi} - p_{it}] & \text{if } \mathbf{1}_{it}^{\text{cns}} = 1 \\ p_{it} + \nabla\left(\frac{H_{it}}{H_{it}^{\text{nc}}} p_{it} - p_{it}\right) & \text{if } \mathbf{1}_{it}^{\text{cns}} = 0 \end{cases} \end{aligned}$$

where $\nabla \in (0, 1]$ is a damping parameter and return to step (3).

7. Update the construction indicator. First compute construction demand

$$Y^h_{it} = \begin{cases} H_{it} - H^{\mathrm{ss}}_{i0} & \text{if } t = 0 \\ H_{it} - H_{it-1} + \delta H_{it-1} & \text{otherwise} \end{cases}.$$

Then update the construction indicator using

$$\mathbf{1}^{\mathrm{cns}'}_{it} = \begin{cases} 1 & \text{if } \mathbf{1}^{\mathrm{cns}}_{it} = 0 \text{ and } p_{it} > (p^F_{it})^{\xi} \\ 0 & \text{if } \mathbf{1}^{\mathrm{cns}}_{it} = 1 \text{ and } Y^h_{it} < 0 \\ \mathbf{1}^{\mathrm{cns}}_{it} & \text{otherwise} \end{cases}.$$

In the first case, the price at which housing demand equals supply with no construction exceeds the construction cost. This is not an equilibrium outcome, as construction firms would strictly prefer to construct housing. In the second case, when house price equals construction cost, negative construction is demanded. This is also not an equilibrium outcome because it violates the irreversible construction constraint. In these cases, and only these cases, the construction indicator needs to be changed.

8. If the construction indicator has converged, that is, $\mathbf{1}^{\mathrm{cns}'}_{it} = \mathbf{1}^{\mathrm{cns}}_{it}$ for all i and t, stop. Otherwise, set $\mathbf{1}^{\mathrm{cns}}_{it} = \mathbf{1}^{\mathrm{cns}'}_{it}$ and return to step (3).

E.5 Computational Time

Table 25 displays the amount of time required to perform each of the steps of the quantitative analysis. These include estimating the model and computing the long-run steady state and transition dynamics after the productivity shocks. The first two rows of Table 25 also show the amount of time required to solve a partial equilibrium steady state and transition (that is, solve the household's problem and compute the density of state variables given parameters and prices). The most time-intensive operation is computing transition dynamics. It takes a little less than 1 hour to compute the transition after the Bartik productivity shock, and about 3.5 hours to compute the transition after the Observed Wage productivity shock. Computing the transition path after the Observed Wage shock takes longer because it is a larger shock, so the algorithm described in Section E.4 requires more iterations to converge.

Table 25: Computational Time

	Time
Partial Equilibrium Steady State	5.58 seconds
Partial Equilibrium Transition	5.68 minutes
Estimation	1.08 hours
Bartik Transition	0.95 hours
Observed Wage Transition	3.39 hours

Notes: This table shows the amount of time required to complete each of the steps in the quantitative analysis. "Partial Equilibrium" refers to the amount of time required to solve the household's value function and density of state variables given parameters and prices. All computations were performed in MATLAB R2025a on an Amazon Web Services EC2 m5.16xlarge instance (64 vCPUs, 256 GB RAM).

Table 26 shows how computation time varies with the number of locations I. The first column shows the amount of time required to solve a partial equilibrium steady state relative to the amount of time required in the baseline case of $I = 50$ (5.58 seconds: see Table 25). The second column shows the relative time required to solve a partial equilibrium transition (5.68 minutes with $I = 50$). For the set of I values considered, computational time for both operations scales roughly linearly with the number of locations.

Table 26: Computational Time Relative to 50-Location Model

I	Steady State	Transition
10	0.19	0.21
25	0.46	0.50
50	1.00	1.00
75	1.51	1.57

Notes: This table shows computation time for various numbers of locations I, relative to the baseline case of $I = 50$. The first column shows the amount of time required to compute a partial equilibrium steady state (that is, the amount of time required to solve the household's problem and compute the density of state variables given parameters and prices). The second column shows the relative time required to solve a partial equilibrium transition path. In the baseline case of $I = 50$, it takes 5.58 seconds to solve a partial equilibrium steady state and 5.68 minutes to solve a partial equilibrium transition path (see Table 25). All computations were performed in MATLAB R2025a on an Amazon Web Services EC2 m5.16xlarge instance (64 vCPUs, 256 GB RAM).