Why Has House Price Dispersion Gone Up?*

Stijn Van Nieuwerburgh[†] Pierre-Olivier Weill[‡]

January 23, 2010

Abstract

We set up and solve a spatial, dynamic equilibrium model of the housing market based on two main assumptions: households with heterogenous abilities flow in and out metropolitan areas in response to local wage shocks, and the housing supply cannot adjust instantly because of regulatory constraints. In our equilibrium, house prices compensate for cross-sectional productivity differences. We increase productivity dispersion in the calibrated model in order to match the 30-year increase in cross-sectional wage dispersion that we document based on metropolitan-level data. We show that the model quantitatively matches the observed 30-year increase in dispersion of house prices across U.S. metropolitan areas. It is consistent with several other features of the cross-sectional distribution of house prices and wages.

---

[*] First draft: April 2006. We thank Fernando Alvarez, Yakov Amihud, Andy Atkeson, David Backus, Markus Brunnermeier, V.V. Chari, Morris Davis, Matthias Doepke, Xavier Gabaix, Dirk Krueger, Lars Hansen, Jonathan Heathcote, Christian Hellwig, Hugo Hopenhayn, Narayana Kocherlakota, Samuel Kortum, Ricardo Lagos, Robert Lucas, Hanno Lustig, Erzo G.J. Luttmer, François Ortalo-Magné, Christopher Mayer, Ellen McGrattan, Holger Mueller, Torsten Persson, Florian Pelgrin, Andrea Prat, Enrichetta Ravina, Victor Rios-Rull, Jean-Charles Rochet, Esteban Rossi-Hansberg, Robert Shimer, Thomas Sargent, Kjetil Storesletten, Laura Veldkamp, Gianluca Violante, Mark Wright, Randall Wright, and the seminar participants at NYU, IIES in Stockholm, University of Oslo, LSE, HEC Paris, UCLA, Yale University, University of Minnesota, the NBER asset pricing meetings, the University of Chicago, the AEA meetings in Chicago, MIT, the University of Pennsylvania, HEC Lausanne, the Chicago Fed, the NBER public economics and real estate meetings, Minnesota Macro, and Stanford University for comments. We thank the Editor and two anonymous referees for their comments. We acknowledge financial support from the Richard S. Ziman Center for Real Estate at UCLA. Keywords: housing prices, income inequality, housing supply regulation. JEL code: E24, R12, R13.

[†] Department of Finance, Stern School of Business, New York University, and NBER, email: `svnieuwe@stern.nyu.edu`.

[‡] Department of Economics, University of California, Los Angeles, email: `poweill@econ.ucla.edu`.

1 Introduction

This paper sets up and solves a spatial equilibrium model of the housing market. The model is a dynamic version of the canonical compensating-wage differential model of Rosen (1979) and Roback (1982). In contrast with the urban economics tradition of studying house prices in one region given some exogenous outside option of living in the countryside (the "reservation locale"), we model the entire distribution of regions. This has several advantages. First, the model provides predictions for the entire cross-section of income, house prices, and construction. This facilitates comparison with the data, which are realizations of that joint equilibrium distribution. Second, the outside option of living in the reservation locale is no longer exogenous: instead, its value is determined in equilibrium. This is important when studying the effect of changes in aggregate variables on house prices, because they operate precisely through endogenous changes in the value of the reservation locale. Third, our model has the benefit of numerical tractability. This is useful when we solve for transitional dynamics in order to evaluate the effect of changes in aggregate variables on house prices.

We model metropolitan areas as a collection of geographically separated islands, randomly hit by idiosyncratic and persistent productivity shocks in the non-housing sector. Construction firms can build new houses in any metropolitan area, but new construction is irreversible and is subject to supply regulation, implying that the local housing supply cannot adjust instantly in response to a local productivity shock. We assume households differ in their ability and are fully mobile: they can freely move across metropolitan areas, but are constrained to live in the same area they work. The equilibrium provides the endogenous joint dynamics of cross-sectional house prices, construction, labor income, and employment. In particular, more able households sort into more productive regions. House price differentials between metropolitan areas compensate for the income differential of the marginal, lowest-ability household in the location, making him indifferent between staying and moving to the next best metropolitan area. Households also live in smaller and more expensive quarters if they choose to work in higher-income metropolitan areas. Lastly, higher-income metropolitan areas have, on average, a larger housing stock and a larger workforce.

We use this model to study the determinants of the evolution of the regional house price distribution since 1975. While the steep rise of house prices has received a lot of attention in the media and in the literature,[1] one salient feature of that distribution which has received much less attention is the steep rise in the *dispersion* of house prices across regions. Indeed,

---

[1] Case and Shiller (1987), Himmelberg, Mayer, and Sinai (2005), Campbell, Davis, Gallin, and Martin (2009), Gyourko, Mayer, and Sinai (2006) and many others document the historical evolution of house prices at the national and local level.

consider the cross-sectional coefficient of variation (CV), the ratio of the standard deviation to the mean, which is a scale-neutral measure of dispersion. In population-weighted terms, the CV of house prices increased from 0.15 in 1975 to 0.53 in 2007. As we explain below, an increase in the cross-sectional dispersion of regional productivity shocks generates an increase in the dispersion of house prices. Such an increase in the dispersion of productivity is consistent with the increase in the dispersion of labor income (henceforth "wages") that we document in our regional data.[2] However, the increase in the CV of wages is much smaller and increases much less than that of house prices: from 0.08 in 1975 to 0.17 in 2007. Our main quantitative exercise shows that an increase in productivity dispersion can simultaneously generate the 8.6 point increase in the CV of wages and the 38 point increase in the CV of house prices. In fact, our benchmark model predicts a somewhat larger 51 point increase in the CV of house prices.

Our main transitional experiment is to feed in just enough of an increase in productivity dispersion to generate the observed increase in wage dispersion between 1975 and 2007, while keeping both the dispersion of ability and housing supply regulation constant. The increase in productivity dispersion creates flows of workers towards high-productivity metropolitan areas, driving local house prices up because of limited housing supply. Conversely, households flow out of low-productivity areas, driving local house prices down. This increases house price *dispersion*. In addition to explaining the increase in dispersion, our model also generates one third of the observed increase in the average house price level, despite keeping average productivity constant. The level effect arises because house prices are a *convex* function of productivity in the model, or equivalently, because price *differentials* increase with productivity. This effect arises for two reasons. First, as productivity increases, productivity differentials are compensated by housing expenditure differentials for smaller and smaller housing sizes, because households reduce their housing consumption in response to higher prices. Since the price differential is the housing expenditure differential *per unit* of housing consumption, it increases with productivity. A second effect arises because of assortative matching of ability with regional productivity. As productivity increases, the ability of the marginal household increases and so does the wage differential. Since the price differential compensates for the wage differential, the price differential increases with productivity.

The model endogenously generates several features of the joint wage-house price distribution. First, because of assortative matching, it features larger increases in wages at the

---

[2] We use the term "wages" to denote earnings from labor, measured as the product of hours worked and the wage per hour. The Bureau of Economic Analysis similarly refers to earnings as "Wages and salary disbursements."

top of the regional wage distribution. Second, it produces a large increase in the ratio of housing price to construction cost, consistent with the findings of Glaeser, Gyourko, and Saks (2007) and Davis and Heathcote (2007) that the non-structure component of house prices has become more important over the last thirty years. Third, it generates an increasing concentration of people in highly productive regions; the fraction of people working in the highest-wage quintile increases by 8.2 percentage points in model and data. Fourth, it is consistent with the observation that, in repeated cross-sectional regressions of house prices on wages, the coefficient on wages increased over the 1975-2007 period. In other words, a one dollar wage differentials became compensated by a larger and larger price differential.

To understand this last finding, recall that the price differential between two locations of nearby productivity is, to a first-order approximation, proportional to the *constant ability* wage differential that makes the marginal household indifferent between moving locations. The observed wage differential, in contrast, not only depends on the productivity differential but also – through assortative matching – on the ability differential. Consider now an increase in productivity dispersion: productivity differentials rise while ability differentials stay the same, causing a larger percentage increase in price than in wage differentials. Therefore, in the cross-section a one dollar wage differential appears to be compensated by a larger price differential. Note that, if the increase in wage dispersion had been created by an increase in ability dispersion instead, wage differentials would have grown more than price differentials. Hence, the evidence provides indirect support for the mechanism of increasing productivity dispersion, through which we increase wage inequality.

A second transition exercise illustrates that the increase in income inequality is an essential part of the explanation for increasing house price dispersion. To make that point, we attempt to generate the observed increase in house price level and dispersion in a counter-factual economy that experienced no increase in wage dispersion, but only an increase in housing supply regulation. Holding the dispersion of wages constant at its 1975 level, we tighten limits on construction over a thirty-three year period. In order to maximize the impact of supply regulation on prices, we assume that the tightening is more pronounced in high-productivity metropolitan areas. By 2007, the model does predict an increase in the level and dispersion of cross-sectional house prices, but the effects are quantitatively very small. Indeed, the negative impact of regulation on local housing supply is almost completely offset by the equilibrium response of households moving out of tightly regulated areas towards less regulated areas. Because this shifts the local demand down at the same time as the supply, the price impact of supply regulations ends up being quantitatively small.

Related Literature Our model features wage differences across regions, which may reflect productivity gains from agglomeration effects (e.g., Glaeser, Scheinkman, and Schleifer, 1992, 1995). An alternative view in the urban literature is that house price differences reflect differences in amenities and other local traits (e.g, Roback, 1982). We note that our regional productivity process admits a broader interpretation that encompasses both productivity and amenities, which are then reflected into house prices.

Several authors have argued that housing supply regulation is an important determinant of house prices (Glaeser and Gyourko, 2003, 2005; Glaeser, Gyourko, and Saks, 2005, 2007; Quigley and Rosenthal, 2005; Quigley and Raphael, 2005). Both explanations for the increase in level and dispersion of house prices we investigate crucially rely on housing supply regulation. Our quantitative exercise suggests that an initial level of regulation combined with an increase in wage dispersion go a long way towards accounting for the facts.

Our work is related to Gyourko, Mayer, and Sinai (2006), who also study the relationship between the U.S. income distribution and cross-sectional house prices. They provide a two-location model, in which regions differ by housing supply and households differ by income and preference for a particular location. A household lives in the low-supply location if it either has a strong preference for it or a high income. Our paper differs in terms of the economic mechanism -households move for productive rather than preference reasons-, and in terms of methodology. The upside of working with a dynamic and stochastic equilibrium model that is amenable to calibration is that it holds the promise of distinguishing between different mechanisms by looking at their quantitative implications. Spiegel (2001) also studies the link between wages, house prices, and construction in an equilibrium model with a moral hazard friction.

Our model of spatial allocation shares many features with labor search models (Lucas and Prescott, 1974; Alvarez and Veracierto, 1999, 2006) and the spatial allocation model of Shimer (2005). We complement this literature by focusing on a different friction. In our setup, households do not incur any cost when moving between islands. Instead, the flow of households between islands is limited by the supply of housing in each island. Coen-Pirani (2006) works with an island model for studying migration patterns between US states. Eeckhout (2004) uses similar techniques to explain the size distribution of cities.

Our approach to assortative matching of individual ability with regional productivity builds on the model of Sattinger (1993). Recent application of his assignment model to other markets are Gabaix and Landier (2008), Terviö (2008), and Costinot and Vogel (2008). Relative to these papers, we face the additional technical difficulty that the number of households matching with a given region is endogenous because of divisible housing.

Our work connects to the macroeconomics literature that documents increases in wage dispersion at the individual level (e.g. Hornstein, Krusell, and Violante, 2004) and studies its effects on risk-sharing (Krueger and Perri, 2005; Storesletten, Telmer, and Yaron, 2004; Heathcote, Storesletten, and Violante, 2008b; Lustig and Van Nieuwerburgh, 2010) and on asset pricing (Constantinides and Duffie, 1996; Cogley, 2002; Chien and Lustig, 2009; Storesletten, Telmer, and Yaron, 2006; Lustig and Van Nieuwerburgh, 2007). Nakajima (2005) most closely connects these two strands of the literature. He sets up an incomplete markets OLG economy and studies a steady state with low (1967) and one with high individual income inequality (1996). He solves for portfolio allocations between housing and physical capital as well as for equilibrium prices. He finds that the increase in income inequality leads to increased precautionary savings, lower interest rates, and 9% higher house prices. Our model studies the spatial dimension of income and house price inequality in the presence of housing supply restrictions in a complete markets economy.

Finally, our work is complementary to macro and asset pricing models that focus on the role of housing as a consumption good and/or a collateral asset (Iacoviello, 2005; Krueger and Fernández-Villaverde, 2006; Piazzesi, Schneider, and Tuzel, 2006; Lustig and Van Nieuwerburgh, 2005, 2007). In our model, the discount factor is constant across dates and states. An interesting avenue for future work is to incorporate the insights from the asset pricing literature. Recent work by Ortalo-Magné and Prat (2008) along these lines derives equilibrium house and stock prices in a spatial model and shows that a version of the CAPM holds.

The rest of the paper is organized as follows. Section 2 presents our island model. Section 3 calibrates a steady-state of the model to match features of 1975 data. Section 4 provides the quantitative impact of increasing wage dispersion and tightening regulation on prices and the distribution of population and construction. Section 5 concludes.

2 An Island Economy

2.1 The Economic Environment

We start by describing the stochastic environment as well as the technologies for producing housing and non-housing consumption. The next paragraph describes households.

2.1.1 Information and Technology

Time is taken to be discrete and runs forever. The economy is made up of a measure-one continuum of homogenous regions we call islands. At each time $t \in \{1, 2, \ldots\}$, an island's production function of non-housing consumption good is linear in labor with an idiosyncratic productivity $A_t \in \mathbb{R}_+$. We take the productivity process $\{A_t\}_{t=1}^{\infty}$ to be a first-order, N-states, Markov chain with possibly time-dependent support $A_{1t} < A_{2t} < \ldots < A_{Nt}$ and transition function $Q_t(A, \cdot)$. We assume that the productivity process is persistent in the sense that, if $A' > A$, then $Q_t(A', \cdot)$ first-order stochastically dominates $Q_t(A, \cdot)$.[3]

Each island starts at time zero with an initial state $s^0 \equiv (A_0, H_0)$, where A_0 is the initial productivity and $H_0 \in (0, H_{\max})$ the initial housing stock. Although we allow the initial housing stock of an island to be correlated with the initial productivity, we assume that, conditional on A_0, it does not help predict the future path of the productivity.[4] We denote by $\mu_0(ds^0)$ the initial cross-sectional probability measure over initial states s^0. At subsequent times, an island is indexed by its history $s^t = (A^t, H_0)$, where $A^t \equiv (A_0, A_1, \ldots, A_t)$ is the productivity history and H_0 is the initial housing stock. As it is standard, starting from the initial measure μ_0, and using the transition functions $Q_t(A, \cdot)$, one constructs inductively the entire sequence of unconditional probability measures, $\mu_t(ds^t)$, over histories s^t.

Each period, firms purchase construction material in order to construct housing services in the islands of their choosing. A representative construction firm can transform Δ units of construction material into housing consumption according to the Leontief production function $\min\{\Delta, \Pi_t(A_t)\}$, where $\Pi_t(\cdot)$ is some strictly positive bounded function of the current productivity A_t of the island. This function is designed to represent not only technological and physical constraints on construction (such as the amount of constructible land) but also regulatory constraints. One may think of $\Pi_t(A_t)$ as the number of building permits in an island with current productivity A_t. Allowing permits to depend on time and productivity will allow us to capture the commonly held view that regulation became tighter over time, and even more so on highly productive metropolitan areas.

We assume that construction is irreversible and the stock of housing consumption depreciates at rate $\delta \in (0, 1)$. These assumptions are summarized by the constraints

$$\Delta_t(s^t) \geq 0 \tag{1}$$

$$\Delta_t(s^t) \leq \Pi_t(A_t) \tag{2}$$

$$H_t(s^t) = (1-\delta)H_{t-1}(s^{t-1}) + \Delta_t(s^t), \tag{3}$$

where $\Delta_t(s^t)$ denotes the construction flow and $H_t(s^t)$ denotes the housing stock in island s^t. Inequality (1) is the irreversibility constraint, inequality (2) follows from the Leontief

---

[3] This definition of a persistent stochastic process is used, for instance, by Lucas and Prescott (1974).

[4] Formally, $\Pr(A_t \,|\, A_{t-1} \ldots, A_0, H_0) = \Pr(A_t \,|\, A_{t-1}, \ldots, A_0)$. This will imply that, in a dynamic equilibrium, the housing stock does not Granger (1969)-cause productivity.

construction technology, and equation (3) is the law of motion for the housing stock. Lastly, the resource constraint for construction material is

$$\int \Delta_t(s^t)\mu_t(ds^t) \leq M, \tag{4}$$

where M denotes the per-period endowment of perishable construction material.

2.1.2 Preferences and Endowments

The economy is populated by a measure one continuum of infinitely-lived and fully mobile households with discount factor $\beta \in (0, 1)$. Households have separable utility for non-durable consumption and housing services. Their flow utility for non-durable consumption is taken to be linear, while their flow utility over housing consumption is represented by some strictly increasing, strictly concave, bounded above and twice continuously differentiable function $v : (0, \infty) \to \mathbb{R}$. We assume in addition that $v(\cdot)$ is unbounded below, meaning that $v(h)$ goes to minus infinity as h goes to zero. Lastly, and without further loss of generality since $v(h)$ is bounded above, we assume that $v(h)$ goes to zero as h goes to infinity.[5,6]

We assume that households differ in terms of their ability. Namely, at each time a household supplies inelastically $e \in [\underline{e}, \overline{e}]$ effective units of labor in the island of its choosing, and we let $f(e)$ be the cross-sectional density of effective units of labor in the population. Given firms' linear production function and competition, a household with ability e working in an island with productivity A earns a wage $e \times A$.

Although the assumption of ability differences is not needed for the main qualitative results of the paper, it turns out to be crucial in our quantitative exercise. Introducing ability differences across regions allows us to address a standard self-selection problem when we use wage differentials between regions to infer productivity differentials. Indeed, ability may be imperfectly observable and high ability households may prefer to locate in high productivity areas. This implies that observed wage differentials partly reflect ability differentials, and may be larger than the underlying productivity differentials. The relative size of ability and productivity differentials directly affects the relative size of house price and wage differentials, a key target of our quantitative exercise.

Letting $n_t(e, s^t)$ be the number of households with ability e who choose to live in island

---

[5] An iso-elastic utility function $v(h) = h^{1-\gamma}/(1-\gamma)$ satisfies these parametric assumptions when $\gamma > 1$. Lemma 6 of Appendix A.1 shows that these properties imply that the utility function $v(h)$ satisfies Inada (1963) conditions.

[6] The key implication of quasi-linearity is that the marginal utility of consumption is equated across islands and, in that sense, that our *ex-ante* identical households are fully insured. Appendix B.3 explains how to extend and keep our dynamic model tractable when households have a general convex utility function.

s^t, we have

$$\int n_t(e,s^t)\mu_t(ds^t) = f(e), \tag{5}$$

since the density of households with ability e must be equal to $f(e)$.

A key assumption of our model is that our fully mobile households are constrained to live in the same island they choose to work.[7] In other words, letting $h_t(e,s^t)$ be the housing consumption per household of ability e in island s^t, we have the local resource constraint

$$\int_{\underline{e}}^{\overline{e}} n_t(e,s^t)h_t(e,s^t)\,de \le H_t(s^t). \tag{6}$$

An *allocation* is a collection of measurable functions specifying, for each time $t \in \{1,2,\ldots\}$, each ability e and each island s^t, the number $n_t(e,s^t)$ of households, the housing consumption $h_t(e,s^t)$ per household, the flow $\Delta_t(s^t)$ of construction, and the housing stock $H_t(s^t)$. An allocation is *feasible* if it satisfies the constraints (1)-(6).

2.2 Definition of a Competitive Equilibrium

Every period, a representative competitive construction firm is endowed with construction permits and purchases material at price C_t in order to produce and sell housing consumption in the islands of its choosing.[8] The price of housing consumption in island s^t is denoted $P_t(s^t)$. Hence, the representative construction firm problem is to choose quantities $\Delta_t(s^t)$ of construction material in order to maximize

$$\int \left(P_t(s^t) - C_t\right)\Delta_t(s^t)\mu_t(ds^t), \tag{7}$$

subject to (1)-(2).

We assume that competitive real estate firms purchase the stock of housing consumption in all islands and rent it to households.[9] The rent in island s^t is denoted by $R_t(s^t)$. Clearly,

---

[7] The US Office of Management and Budget defines a metropolitan statistical area (MSA), the empirical counterpart of islands in the model, as "a geographic area consisting of the county or counties associated with at least one core of 50,000 or greater population, plus adjacent counties having a high degree of social and economic integration with the core(s) as measured by commuting ties."

[8] Because of linearity of the construction technology, the distribution of permits across construction firms does not matter. An alternative approach would be to assume that households are endowed with permits to construct on their land and sell them to construction firms. This would deliver the same results.

[9] This assumption is made for expositional simplicity. As it is standard with frictionless housing markets, the same equilibrium price would arise if households were purchasing their homes instead of renting them.

a real estate firm finds it optimal to supply its entire housing stock as long as the rent is strictly positive.

Competition among real estate firms implies that the current price of housing consumption is equal to the rent plus the present value of the price next period, net of depreciation:

$$P_t(s^t) = R_t(s^t) + \beta(1-\delta)E_t\left[P_{t+1}(s^{t+1})\,|s^t\right]. \tag{8}$$

Under the transversality condition

$$\lim_{T\to\infty} \beta^T E_t\left[P_{t+T}\left(s^{t+T}\right) \mid s^t\right] = 0, \tag{9}$$

we obtain the Rosen and Topel (1988) result that a house price is equal to the expected present value of rents net of depreciation

$$P_t(s^t) = E_t\left[\sum_{j=0}^{\infty} \beta^j(1-\delta)^j R_{t+j}(s^{t+j}) \middle| s^t\right]. \tag{10}$$

In Appendix B.2 we provide a complete treatment of the household's inter-temporal problem. Households choose in which location to live and work at any given time, as well as their housing and non-housing consumption in each island. Households receive all profits from the real estate sector: the profit of selling the endowment of construction material to construction firms, the profit of building houses, and the profit of renting out houses. Because of full mobility, the households' inter-temporal problem can be simplified dramatically: it reduces to a sequence of static problems. Every period, a household chooses in which island to work, and how much housing to rent in that island. Namely, given optimal housing consumption choice, a household with ability e who chooses to live in island s^t enjoys the value:

$$u_t(e, s^t) = e \times A_t + \max_{h\geq 0}\left\{v(h) - R_t(s^t)h\right\}. \tag{11}$$

And, of course, a household's optimal location choice is to work and live in any island that yields the maximum value:

$$U_t(e) = \max_{s^t} u_t(e, s^t). \tag{12}$$

A *competitive equilibrium* is a price system and a feasible allocation such that: i) the price and the rent satisfy (10), ii) given the price $P_t(s^t)$ of housing consumption and the price C_t

of construction material, the construction flow $\Delta_t(s^t)$ solves the construction firm's problem, iii) given the rent $R_t(s^t)$, housing consumption $h_t(s^t)$ solves the household's problem and the allocation of households across islands is individually optimal, that is

$$n_t(e, s^t) \geq 0 \quad \text{if} \quad u_t(e, s^t) = U_t(e) \tag{13}$$

$$n_t(e, s^t) = 0 \quad \text{otherwise.} \tag{14}$$

2.3 Equilibrium Characterization

In this subsection we characterize an equilibrium. Section 2.3.1 provides some elementary properties of rents and prices. Section 2.3.2 shows that an equilibrium features assortative matching of ability with productivity. Section 2.3.3 derives the dynamics of the housing stock. Finally, Section 2.3.4 proves that an equilibrium exists and is unique, and provides a numerical algorithm for calculating it.

2.3.1 Rents, Prices, and Housing Consumption

In this section we derive elementary properties of an equilibrium. First, in an island without population, the rent must be equal to zero. Otherwise, a real-estate firm would find it optimal to supply a positive quantity of housing and the market would not clear.[10]

Second, the rent is a function of an island's current productivity, and does not depend on other idiosyncratic characteristics of the island, such as past productivity shocks or the local housing supply. Formally, if the current productivity of an island is equal to A_{it}, the i^{th} element of the productivity grid, then its rents is equal to some R_{it}. Otherwise, if two islands with the same current productivity had different rents, the location choice of households living in the high-rent island would not be optimal: they would prefer to move to the low-rent island where they would earn the same wage (because of equal productivity) but pay a lower rent.

Third, the rent must be increasing in productivity, and strictly increasing across populated islands. Indeed consider two islands with productivity $A_{it} < A_{jt}$. If island A_{it} is populated, then $R_{it} < R_{jt}$. Otherwise the location choice of households in the low productivity island would not be optimal: they would prefer to move to the high-productivity where they would earn a higher wage and pay a lower rent. If island A_{it} is not populated, then $R_{it} = 0$ and, evidently, $R_{it} \leq R_{jt}$.

---

[10] Note that the local housing supply is strictly positive in every island. Indeed, each island starts with a strictly positive housing stock, and the depreciation rate δ is strictly less than one.

Plugging the rent back into the pricing equation (10) and using the Markov property shows that the price $P_t(s^t)$ is a function of the current productivity but does not depend on other idiosyncratic characteristics of an island. In addition, because the rent is increasing in productivity, and the productivity process is persistent the price is also increasing in productivity.

Now, going back to the quasi-linear household's problem of equation (11), it follows that a household's optimal housing consumption does not depend directly on its ability e. Furthermore, because the rent is an increasing function of productivity, it immediately follows that housing consumption is a decreasing function of productivity.

This discussion is summarized in the following proposition:

Proposition 1. *At each time $t \in \{1, 2, \ldots\}$, housing consumption, rent, and price are functions of the islands' current productivity, and do not depend on any other idiosyncratic characteristic of the island. In addition, housing consumption h_{it} is decreasing, rent R_{it} is increasing, and price P_{it} is increasing, with the island's current productivity, A_{it}.*

2.3.2 Assortative Matching

In this paragraph we show that our model formalizes the commonly held view that households with higher ability tend to work in higher productivity locations. This is because the households' objective function is super-modular: the cross-derivative between ability and productivity is positive. This implies that households' location decisions are weakly increasing in ability, i.e. high-ability households find it optimal to locate in higher productivity locations. Formally, we prove the following proposition:

Proposition 2. *At each time, there is an integer $p \in \{1, \ldots, N\}$ and a sequence of ability cutoffs*

$$e_{pt} = \underline{e} < e_{p+1t} < \ldots < e_{Nt} < e_{N+1t} \equiv \overline{e}$$

such that

1. *An island is populated if and only if its productivity is greater than A_{pt}.*
2. *For all $i \in \{p, \ldots, N\}$, households with ability levels $e \in (e_{it}, e_{i+1t})$ strictly prefer to live in islands with current productivity A_{it}.*
3. *For all $i \in \{p, \ldots, N-1\}$, households with ability $e = e_{i+1}$ are indifferent between living in islands with current productivity A_{it} or A_{i+1t}.*

In the proposition, we suppressed the dependence of the integer p on time to simplify notations. The results are illustrated in Figure 1: households with ability in the interval (e_{it}, e_{i+1t}) go work in islands with current productivity A_{it}.

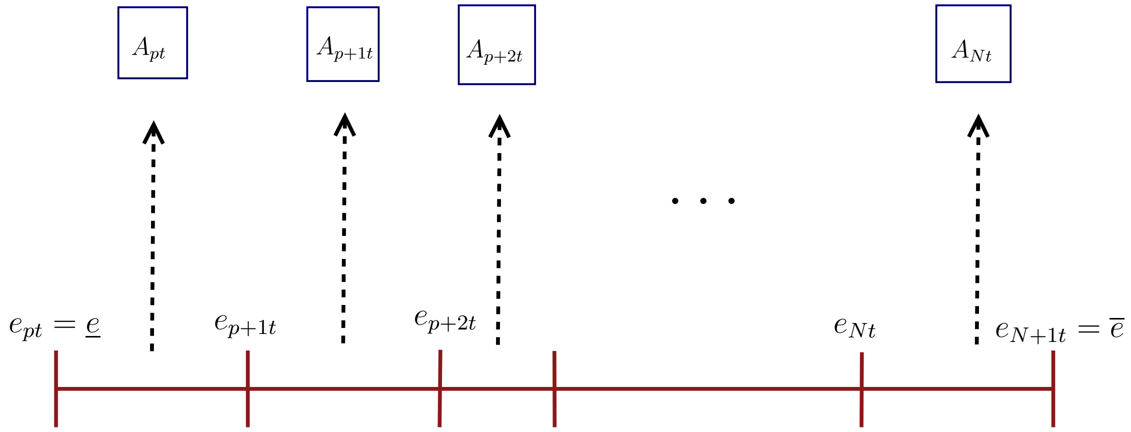

Figure 1: The assignment of abilities to islands

We now provide a method for constructing the cutoffs e_{it} which has two main benefits: it leads to a simple existence proof and it provides an algorithm for calculating the equilibrium assignment of households to islands. We start by letting U_{it} be, for all $i \geq p$, the maximum attainable utility of a household with ability e_{it}. For all $i < p$, we let $e_{it} = e_{pt} = \underline{e}$ and $U_{it} = U_{pt}$. Then, we have

$$U_{it} \geq A_{it} e_{it} + v(h_{it}) - R_{it} h_{it}, \tag{15}$$

with an equality for $i \geq p$, and where h_{it} denote housing consumption in island i at time t. After plugging the first-order condition $R_{it} = v'(h_{it})$ and letting $w(h) \equiv h v'(h) - v(h)$, we obtain that housing consumption is:[11]

$$h_{it} = w^{-1} \left(\max\{A_{it} e_{it} - U_{it}, 0\} \right). \tag{16}$$

Keeping in mind that housing consumption h_{it} only depends on an island's current productivity, we add up the housing market clearing conditions (6) across all islands with current

---

[11] Note that the formula implies that, for $i < p$, the argument of $w^{-1}(\,\cdot\,)$ is zero and the optimal housing consumption is $h_{it} = \infty$. Such housing demand is consistent with optimality because the rent is zero and $v(h)$ is strictly increasing. It is also consistent with market clearing because islands $i < p$ are not populated.

productivity A_{it} to find:

$$n_{it}h_{it} \leq H_{it},$$

with an equality for $i \geq p$, and where H_{it} denotes the *average* housing stock and n_{it} the *average* number of households per island with current productivity A_{it}.[12] Thus, the total number of households living in islands with productivity A_{it} is equal to

$$\mu_{it}n_{it} = \mu_{it}H_{it}/h_{it} = \mu_{it}H_{it}\Phi\left(\max\{A_{it}e_{it} - U_{it}, 0\}\right),$$

where μ_{it} denotes the number of islands with productivity A_{it} and $\Phi(x) \equiv 1/w^{-1}(x)$. From Proposition 2, it follows that the total number of households living in islands A_{it} must be equal to the number of households with ability in between cutoffs e_{it} and e_{i+1t}. Taken together, this gives the difference equation:

$$F(e_{i+1t}) - F(e_{it}) = \mu_{it}H_{it}\Phi\left(\max\{A_{it}e_{it} - U_{it}, 0\}\right), \tag{17}$$

where $F(e)$ is the cumulative distribution function of the ability distribution. This difference equation allows to calculate the sequence of equilibrium cutoffs given a sequence $U_{1t}, \ldots, U_{Nt}$ of maximum attainable utilities.

To calculate the maximum attainable utilities on the right of side of equation (17), we use the indifference conditions of households with abilities $e_{p+1t}, \ldots, e_{Nt}$. Indeed, because of point 3 in Proposition 2, we know that a household with ability e_{i+1t} is indifferent between the following two alternatives. He can work in an island with productivity A_{i+1t} and receive utility U_{i+1t}. Or, he can work in an island with productivity A_{it}, earning a wage $e_{i+1t}A_{it}$, enjoying a quantity h_{it} of housing consumption, and paying a rent $R_{it}h_{it}$. This adds up to a utility $e_{i+1t}A_{it} + v(h_{it}) - R_{it}h_{it} = U_{it} + A_{it}\left(e_{i+1t} - e_{it}\right)$. Equating these two utilities and using equation (15), we obtain the difference equation:

$$U_{i+1t} = U_{it} + \left(e_{i+1t} - e_{it}\right)A_{it}. \tag{18}$$

Note that this difference equation also holds for $i < p$, given our convention that $e_{it} = \underline{e}$ and $U_{it} = U_{pt}$.

Taken together, the difference equations (17) and (18) suggest a simple "shooting" algo-

---

[12] Note that, in all islands s^t with current productivity A_{it}, the housing stock $H_t(s^t)$ and the population $n_t(s^t)$ will differ from the averages H_{it} and n_{it}. This is because, unlike housing consumption, the housing stock $H_t(s^t)$ and the population $n_t(s^t)$ depend on the initial housing stock and past productivity realizations.

rithm for calculating the equilibrium assignment of households to islands, given a distribution of housing stocks $H_{1t}, \ldots, H_{Nt}$. Given an initial condition U_{1t} and $e_{1t} = \underline{e}$, we calculate e_{2t} using equation (17), then U_{2t} using equation (18), then e_{3t} using (17), and so on until we obtain the entire sequence $U_{1t}, \ldots, U_{Nt}$ and $e_{1t}, \ldots, e_{Nt}, e_{N+1t}$.[13] Moreover, one can easily show that the terminal cutoff e_{N+1t} is a decreasing function of the initial condition U_{1t}. We use this monotonicity property to prove that there exists a unique U_{1t} such that $e_{N+1t} = \overline{e}$. This discussion is summarized in the following proposition:

Proposition 3 (Equilibrium Assignment)**.** *Given a distribution $H_{1t}, \ldots, H_{Nt}$ of housing stocks, there exists a unique pair of sequences $\{e_{it}\}_{i=1}^{N+1} \in [\underline{e}, \overline{e}]^N$ and $\{U_{it}\}_{i=1}^{N} \in \mathbb{R}^N$ solving the difference equations (17)-(18) with initial condition $e_{1t} = \underline{e}$ and terminal condition $e_{N+1t} = \overline{e}$.*

Another reason why this procedure is computationally convenient is that we always start shooting at the lower bound of the productivity grid, $i = 1$, so that there is no need to guess-and-verify the cutoff p. Instead, p can be calculated in a second step, as the smallest grid point i such that $e_{it}A_{it} > U_{it}$.

2.3.3 Housing Stock

To complete our characterization of an equilibrium, we need to solve for the distribution $H_t(s^t)$ of housing stocks. To that end, we first note that the linearity of the construction firm's objective (7) implies that an optimal construction plan is simply to build $\Pi_t(A_{it}) \equiv \Pi_{it}$ units of housing consumption in every island where $P_t(s^t) > C_t$. Since we proved that $P_t(s^t)$ only depends on the current productivity A_{it} and is increasing, it follows that there is some productivity cutoff A_{ct} such that a construction firm builds a quantity $\Delta_{it} = \Pi_{it}$ of housing if $A_{it} > A_{ct}$, a quantity $\Delta_{ct} \in [0, \Pi_{ct}]$ if $A_{it} = A_{ct}$ and does not construct anything otherwise. Plugging this back into the resource constraint (4) for construction material, we obtain

$$\mu_{ct}\Delta_{ct} + \sum_{i=c+1}^{N} \mu_{it}\Pi_{it} \leq M, \tag{19}$$

with an equality if the following condition is satisfied

$$\sum_{i=1}^{N} \mu_{it}\Pi_{it} > M. \tag{20}$$

---

[13] To make this procedure well defined, we need to artificially extend the domain of the cdf $F(e)$ above the upper bound $\overline{e}$. We deal with this in detail in Appendix A.2.2.

Condition (20), which we assume holds from now onwards, implies that there is a large supply of constructible land. That is, the amount of housing that could be built on all constructible land, on the left-hand side of (20), is greater than the amount of housing that can be built with the available supply M of construction material. Under condition (20), if at the cutoff $\Delta_{ct} < \Pi_{ct}$, then the representative construction firm is indifferent between constructing or not, implying that

$$C_t = P_{ct}, \tag{21}$$

at each time $t \in \{1, 2, \ldots\}$. If $\Delta_{ct} = \Pi_{ct}$, then $C_t = P_{ct}$ is also an equilibrium construction price.

2.3.4 An Algorithm

Taken together, the above paragraphs provide an algorithm for calculating an equilibrium:

1. First, one uses (19) to solve, at each time, for the construction cutoff c and the construction plan $\{\Delta_{it}\}_{i=1}^{N}$.

2. Given the construction cutoffs, one can use the difference equation (3) to calculate the distribution of housing stocks across islands. Note that, unlike rents, prices, and housing consumptions, the local housing stocks will depend on the entire productivity history s^t of an island. However, as Section 2.3.2 made clear, only N moments of the housing stock distribution matters: the average housing stock H_{it} per island with current productivity A_{it}. These N moments jointly solve the difference equation:

   $$H_{it} = (1 - \delta) \sum_{j=1}^{N} \frac{\mu_{jt-1} Q_{t-1}(j, i)}{\mu_{it}} H_{jt-1} + \Delta_{it}, \tag{22}$$

   where the first term on the right-hand side is the (depreciated) average housing stock last period in an island with current productivity A_{it}.

3. Given the distribution of housing stocks H_{it}, one solves for the ability cutoffs e_{it} and the maximum attainable utilities using (17) and (18).

4. Finally, given the ability cutoffs and the maximum attainable utilities, one solves for the housing consumption using (16), for population using $n_{it} = H_{it}/h_{it}$, for rents using $R_{it} = v'(h_{it})$ and for prices using the present value formula (10).

Based on these four steps, we first show:

Proposition 4 (Existence and Uniqueness)**.** *There exists an equilibrium. The equilibrium is unique in the sense that all equilibria share the same rents* R_{it}*, housing consumptions* h_{it}*, population* n_{it}*, housing stocks* H_{it}*, and ability cutoffs* e_{it}*.*

Some equilibrium objects are not uniquely determined: for instance, if $\Delta_{ct} = \Pi_{ct}$, then all construction prices $C_t \in [P_{c-1t}, P_{ct}]$ are consistent with optimality. Also, households at the ability cutoffs e_{it} are indifferent and of measure zero, so their island assignment is indeterminate. These dimensions of indeterminacy, however, do not change the answers of the questions at hand.

Note that the algorithm translates into a fast computational procedure, because the distribution of housing stocks can be characterized before calculating prices. Also, given that the other objects of interest only depend on the current productivity of an island, we do not need to calculate the entire population distribution, $n_t(s^t)$, to calculate population-weighted moments. Instead, it is enough to calculate $n_{it} = \mathbb{E}\left[n_t(s^t) \mid A_{it}\right]$, the average population per island A_{it}. Our discrete state space allows us to calculate expectations and present value quickly, while approximating standard continuous-state processes using the quadrature methods of Tauchen and Hussey (1991). Lastly, the discrete state also speeds up the calculation of households' equilibrium assignment, relative to the ordinary differential equations arising in a continuous-state model. Our algorithm results in quick calculations of transitional dynamics, without using any linearization technique. This turns out to be important for our results, because the price impact of wage dispersion stems from a non-linear convexity effect.

2.4 Convexity

We now show an important property of our model: that increasing productivity dispersion increases house price levels.

Proposition 5 (Convexity)**.** *At each time* $t \in \{1, 2, \ldots\}$*, the rent* R_{it} *is a convex function of an island's current productivity, in that:*

$$\frac{R_{i+1t} - R_{it}}{A_{i+1t} - A_{it}} \tag{23}$$

is increasing in i*.*

The following back-of-the-envelope calculation provides intuition. Consider a household of ability e_{i+1t}, who is indifferent between island A_{i+1t} and island A_{it}:

$$
\begin{aligned}
& e_{i+1t}A_{i+1t} + v(h_{i+1t}) - R_{i+1t}h_{i+1t} = e_{i+1t}A_{it} + v(h_{it}) - R_{it}h_{it} \\
\Rightarrow \quad & R_{i+1t}h_{i+1t} - R_{it}h_{it} = e_{i+1t}\left(A_{i+1t} - A_{it}\right) + v(h_{i+1t}) - v(h_{it}).
\end{aligned}
$$

The equation says that the housing expenditure differential between island i and $i+1$ compensates for the wage differential of the *marginal* household, as well as for the utility differential arising from differential housing consumptions. Now use the last equation to calculate the housing expenditure differential, holding housing consumption constant:

$$
\begin{aligned}
& \left(R_{i+1t} - R_{it}\right)h_{it} = e_{i+1t}\left(A_{i+1t} - A_{it}\right) + v(h_{i+1t}) - v(h_{it}) - R_{it+1}\left(h_{i+1t} - h_{it}\right) \\
\Rightarrow \quad & \left(R_{i+1t} - R_{it}\right)h_{it} \simeq e_{i+1t}\left(A_{i+1t} - A_{it}\right) + \left(v'(h_{i+1t}) - R_{i+1t}\right)\left(h_{i+1t} - h_{it}\right), \\
\Rightarrow \quad & \frac{R_{i+1t} - R_{it}}{A_{i+1t} - A_{it}} \simeq \frac{e_{i+1t}}{h_{it}}, \qquad (24)
\end{aligned}
$$

where the second line follows from a first-order approximation of $v(h_{i+1t}) - v(h_{it})$ and the last line because $R_{i+1t} = v'(h_{i+1t})$.[14] Equation (24) shows that convexity arises for two reasons. The first reason is that, as productivity increase, households respond to higher house prices by reducing their housing consumption: thus, productivity differentials $A_{i+1t} - A_{it}$ are compensated by housing expenditure differentials for smaller and smaller a housing consumption. Since the rent differential, $R_{i+1t} - R_{it}$, is the housing expenditure differential *per unit* of housing consumption, it becomes larger and larger. The second effect arises because ability increases with productivity. Intuitively, the rent differential compensates for the wage differential of the marginal household. But since the ability of the marginal household increases with productivity, the wage differential and the corresponding rent differential becomes larger and larger.

The house price implications of a productivity-induced increase in wage dispersion follow immediately from Proposition 5. Consider a (mean-preserving) increase in productivity dispersion, holding the mapping from productivity to price the same. This mechanically increases wage dispersion. Also, because R_{it} is an increasing function of A_{it}, the rent increases in high-productivity islands and decreases in low-productivity islands. Hence, the cross-sectional dispersion of rents increases. Now, convexity means that the rent increases by more in high-productivity islands than it decreases in low-productivity islands. This creates two level effects. First, the cross-sectional average rent goes up. Second, the house price level increases *in every island.* To understand this second effect, consider the example of

---

[14] Our formal proof does not rely on any approximation. Note also that, if the productivity state were continuous instead of discrete, then equality (24) would hold exactly.

an independent and identically distributed productivity process. That is, every period, the productivity in an island is an independent draw from the cross-sectional distribution. Our pricing equation (8) implies that the price in an island with current productivity A is

$$P_{it} = R_{it} + \frac{E\left[R_{jt+1}\right]}{1-\beta(1-\delta)}, \tag{25}$$

where the expectation is taken with respect to the cross-sectional distribution of productivity. Convexity implies that an increase in productivity dispersion increases the second term in the price equation (25).[15] In words, the house price increases because households anticipate that the rent will increase by more when the island draws a high productivity than it will decrease when it draws a low productivity.

3 Calibration Parameters and Targets

While the previous section establishes the qualitative link between wage and price dispersion, the question remains whether the model can quantitatively generate the observed amount of price dispersion. To that end, we calibrate our model so that the initial steady state of the model matches key moments of the wage and population distribution in 1975. We then engineer an increase in the dispersion of wages of the observed magnitude, and ask whether the model can account for some key features (what we call "targets") of the post-1975 house price distribution. The results of this exercise are in the next section. We start here by describing the moments of the data we are trying to match.

3.1 Targets in the Data

Our goal is to account for the joint distribution of wages and house prices across U.S. metropolitan areas over the last 33 years. First, we briefly describe the wage and price data; details on all data definitions, sources, and construction are relegated to Appendix D. Our sample consists of 330 U.S. metropolitan statistical areas with annual data from 1975 to 2007.

Raw Data Wages are measured using nominal wage per job data available from the Bureau of Economic Analysis Regional Economic Information System (REIS). This is a measure of

---

[15] If the wage process is persistent, then the same effect operates in the long run. Indeed, by ergodicity, the distribution of the wage T periods ahead converges to the cross-sectional distribution as T goes to infinity.

the average annual earnings per employed worker in that region. We also obtain the number of jobs for each metropolitan area from REIS, and use them to calculate population-weighted moments. To calculate the *real* wage per job we deflate the nominal wage per job by a regional cost-of-living index which excludes housing. The index combines data from the Bureau of Labor Statistics to compute year-to-year changes in each MSA, together with data on relative non-housing prices across MSAs from the private data vendor COLI. The base year is 1983-84, when the *average* region has a non-housing price level normalized to 100.

House prices are measured as the nominal median home value. We combine the median single-family home values from the 2000 Census with the Freddie Mac Conventional Mortgage Home Price Index (CMHPI), a repeat-sale house-price index available from 1975 until 2007.[16] Proceeding as with nominal wages, we deflate nominal home values by the non-housing price index to obtain real prices. A balanced panel of prices is only available for a subset of 81 regions. The sample with house price data gradually increases from 81 MSAs in 1975 to 330 MSAs in 1994, and stays constant thereafter. We refer to this growing sample as the unbalanced panel. Figure 2 plots the population-weighted cross-sectional average and coefficient of variation of real wages for the balanced panel of 81 regions (top row) and the unbalanced panel of 330 regions (bottom row). Figure 3 does the same for real home prices. The figures indicates that changes in cross-sectional level and dispersion are similar for both samples.[17] In what follows, we will focus on the unbalanced panel of 330. Both the CV and the level of real wages increase moderately, while the CV and the level of real prices increase strongly.

While the house price series is a constant-quality series, it does not correct for the increase in the quantity of housing services that a typical single-family house provides. We measure this quantity as the average square foot of completed single-family units, for sale inside metropolitan areas. The Census's construction statistics indicate that house size has grown from 1,715 to 2,563 square feet between 1975 and 2007, an average growth rate of 1.256% per year. As explained below, we de-trend house price by size to remove the mechanical increase in house prices that is due to the increase in house size.

Data in the De-Trended Economy. Appendix E explains that, because of the quasi-linear preferences, our model is not consistent with balanced growth in productivity and in

---

[16] The CMHPI is a constant quality house price index. It pertains to single-family properties financed with a mortgage below the conforming loan limit. See Case and Shiller (1987). We use fourth quarter values.

[17] Non-population-weighted moments (not reported here) also display similar increases in level and dispersion.

Figure 2: First and Second Moments of Real Wages in the Data

The top row of the figure plots the population-weighted cross-sectional average, cross-sectional standard deviation, and cross-sectional coefficient of variation of the real wage per job for a balanced panel of 81 metropolitan statistical areas. The bottom panel reports the same moments for an unbalanced panel of regions that grows over time from 81 to 330 metropolitan statistical areas. The real wage per job is calculated as the nominal wage per job divided by the regional non-housing price index.

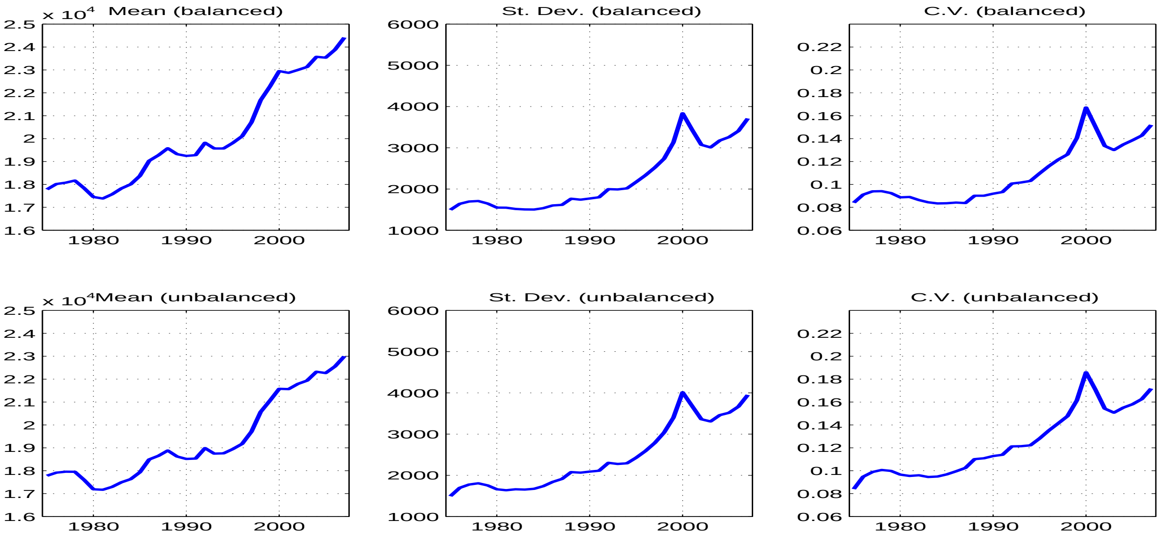

the quantity of housing services.[18] At the same time, the data do not seem consistent with balanced growth either: over our sample period, population-weighted average real wages grew at $g_w = 0.80\%$ per year in the unbalanced panel whereas house size grew at the higher rate of $g_H = 1.256\%$ per year. These observations suggest that our model is better suited to explain de-trended data. This leads us to feed in and confront the model with de-trended data.

In the data, we deflate the house price series by the observed size of a house, i.e. by $(1 + g_H)^t$ in year t. This generates a constant-size house price series. We also remove a trend from real wage data. It is important not to remove the entire growth rate g_w because, as will become clear later, the model endogenously generates a trend in wages even when productivity has no trend. If g_m denotes the (endogenous) growth rate of wages in the de-trended model, then we de-trend the real wage at a rate $g_w - g_m$. This guarantees that the de-trended wage grows at the same rate g_m per year in the calibrated model and in the de-trended data.

---

[18] The model is, however, consistent with population growth, i.e. growth in the number of households. Using a standard argument, Appendix E shows how relative prices and per-household quantities are the same in a model with population growth and in an appropriately transformed model without growth.

Figure 3: First and Second Moments of Real Home Prices in the Data

The top row of the figure plots the population-weighted cross-sectional average, cross-sectional standard deviation, and cross-sectional coefficient of variation of real single-family home values in the data for a balanced panel of 81 metropolitan statistical areas. The bottom panel reports the same moments for an unbalanced panel of regions that grows over time from 81 to 330 metropolitan statistical areas. the real home value is computed as the nominal home value divided by the

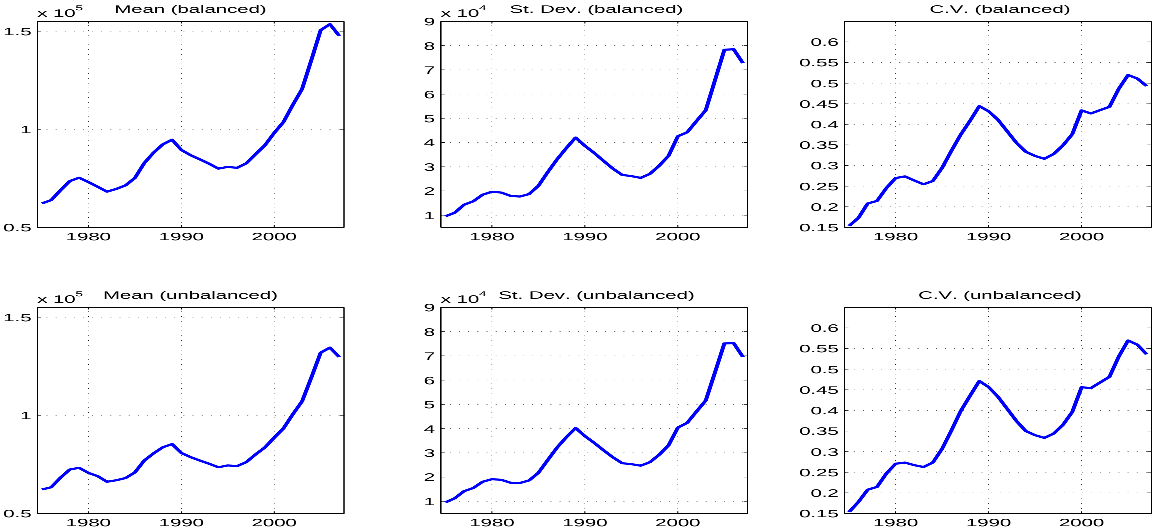

The resulting de-trended, population-weighted average real wage increases from \$17,782 in 1975 to \$19,489 in 2007, an annual change of $g_m = 0.29\%$. The population-weighted CV of real wages increases from 0.084 in 1975 to 0.172 in 2007. The de-trended, population-weighted average real house price increases from \$62,212 in 1975 to \$87,013 in 2007, an annual change of 1.05%. Note in particular that, without the de-trending by house size, house prices go up to about \$150,000 in 2007 (see Figure 3), so that a large fraction of the run-up in levels is indeed accounted for by the increase in house size. The population-weighted CV of real de-trended house prices increases from 0.154 in 1975 to 0.536 in 2007, which is similar to the increase shown in Figure 3. Our main quantitative exercise is to feed in the observed wages into the model and to ask what fraction of the observed increase in level and especially in dispersion of home prices it can explain. In particular, can a small increase in wage dispersion of 8.8 points generate a large increase in house price dispersion of 38.2 points?

Price-Wage Sensitivity In reality, several factors outside of our model presumably contribute to the observed house price dispersion. For example, in addition to productivity differentials, amenity differentials may matter. To quantify the importance of wages differentials in creating price differentials, we compute the following R^2 statistic:

$$R^2 = 1 - \frac{\mathrm{var}(p_{it} - \hat{p}^d_{it})}{\mathrm{var}(p_{it})}, \tag{26}$$

where p_{it} denotes the (real de-trended) house price in region i and period t and $\hat{p}^d_{it}$ denotes a linear projection of real house prices on real wages in the data. This projection, and the associated R^2 in (26), employs all available year-MSA observations, i.e., the entire unbalanced panel, and takes into account the population size of each region. We find that 26.5% of variation in house prices can be explained by variation in wages. This R^2 value is an important target for our model to match.[19] Namely, we will feed in the model observed wage, obtain the model-predicted prices $\hat{p}^m_{it}$, and calculate their R^2 with in equation (26) replacing the linear projection $\hat{p}^d_{it}$ by $\hat{p}^m_{it}$.

In addition, we study the slope b_p from a repeated cross-sectional regression of house prices on wages. These regressions again weigh the importance of each region by its population. In the data, the slope of this regression increases from 0.81 in 1975 to 7.89 in 2007, suggesting that the sensitivity of house prices to wages has increased substantially over time. Explaining the increase in this slope coefficient is another target for our model.

3.2 Calibration

This section discusses the calibration. Table 1 summarizes the parameters and their benchmark values. Our calibration strategy has three components. Five parameters, indicated with an "E" in the first column of the table, are chosen so that the initial steady state of our model replicates key moments of the 1975 data. One parameter, indicated with a "E^*" superscript, is chosen so that we replicate the observed increase in wage dispersion. In order to pick these six parameters, we solve the model repeatedly until the six endogenously generated moments exactly match their counterpart in the data. All other parameters are set "externally" to conventional values. We now describe these choices in more detail.

3.2.1 Preferences

The model is calibrated at annual frequency. We set the households' time discount factor to $\beta = .951$ in order to match the average real interest rates of 5.15% on the conforming

---

[19] In order to focus on the lower-frequency relationship, we run this regression on price and wage data that have been averaged over five-year periods. The R^2 is similar for other horizons. For example, it is 25.6% for annual observations (no averaging) and 27.8% based on ten-year averages.

30-year fixed rate mortgage between 1975 and 2007. This is the most relevant interest rate to use in the present-value formula that pins down the house price. We let households have an iso-elastic utility function $v(h) = \kappa h^{1-\gamma}/(1-\gamma)$ over housing consumption, implying that the price elasticity of housing demand is equal to $-1/\gamma$. Because the micro-level evidence suggests an elasticity of about -0.5, we set $\gamma = 2$.[20] The parameter κ governs the housing expenditure share. We choose κ so that the housing expenditure to income ratio in the model (first averaged across regions, then across time) matches the value of 0.12 in the 2000 Census data.

3.2.2 Productivity and Ability

Productivity Regions differ in their productivity. We choose our finite-state regional productivity process so as to approximate (in the sense of Tauchen and Hussey, 1991) the following geometric random walk with exponential lifetime.

Every period, a measure $\lambda \in (0, 1)$ of new regions is created with an initial log productivity $a_t = \log(A_t)$ drawn from a normal distribution with mean μ_{bt} and standard deviation σ_{bt}. In every subsequent period, a region either disappears with probability λ or survives with probability $1 - \lambda$. In case of survival, it draws a new log productivity

$$a_t = a_{t-1} + \sigma_{at}\varepsilon_t, \tag{27}$$

where ε_t is a standard normally distributed shock. As in Yaari (1965), setting $\lambda > 0$ guarantees the existence of a stationary distribution.[21] Appendix C.1 explains that the cross-sectional distribution of log productivity across islands is not known in closed form (although it can be easily written as a mixture of normal densities) and behaves like a Pareto distribution in its two tails. However, the first and second moment of the cross-sectional productivity distribution can be calculated easily. As explained in the appendix, we then discretize this continuous-state productivity process on $N = 190$ Gaussian quadrature points using the quadrature methods of Tauchen and Hussey (1991). We treat the discretized process as the "true" productivity process, which allows us to apply all the theoretical results of Section 2.

---

[20] Hanushek and Quigley (1980) exploit a natural experiment where a subgroup of 586 low income renters in Phoenix and 799 households in Pittsburgh received rent subsidies ranging from 30-60%, whereas a control group received nothing. They estimate long-run elasticities of -.45 for Phoenix and -.64 for Pittsburgh, based on estimates of how fast the housing demand adjusts towards an equilibrium level in the two years of data.

[21] Although our theoretical section did not consider such exogenous entry and exit of productive locations, it is in fact a straightforward extension. There are only two things that need to be adjusted. First the average housing stock per region of type A_i depreciates faster by a factor $1 - \lambda$. Second the discount factor for the present value formula is also scaled down by $1 - \lambda$.

Thus, in the initial steady state, the productivity process is characterized by four parameters: μ_b, σ_{a0}^2, σ_b^2, and λ. We choose the parameter μ_b to match the 1975 population-weighted average wage per job. We set the variance of productivity innovation σ_{a0}^2 and the variance of productivity at birth σ_b^2 to match the population weighted coefficient of variation of wage per job, with the identifying assumption that initial conditions, represented by σ_b^2, explain half of the variance in productivity. The results turn out to be rather insensitive to this identifying assumption. Lastly, we exogenously fix the death rate λ at 1% per year, which delivers an autocorrelation of wages that is statistically indistinguishable from the data; see Section 4.1 below.

Ability Households differ in their effective units of labor $e \in [\underline{e}, \overline{e}]$.[22] Given that the assumed cross-regional productivity distribution exhibits Pareto behavior in its tails, we chose an ability distribution $f(e)$ with the same properties. Namely, we assume that ability is distributed according to a double-Pareto distribution. Appendix C.1 provides the details. We choose the parameters of this distribution so that, first, ability has a mean normalized to 1. Second, the Pareto coefficient k_e, which relates inversely to the cross-sectional standard deviation of ability, is such that we match the 1975 sensitivity of prices to wages b_{p0} of 0.81. Since we simultaneously match the dispersion of wages in 1975, by matching b_{p0} we also match the observed covariance between prices and wages. In other words, the fraction of price dispersion that is explained by wages in 1975, according to a naive ordinary least squares regression, is the same in the model and in the data. Of course, this measure could be biased either upward of downward because of omitted variables that directly impact house prices and, at the same time, are related to wages. Section 4.6.1 proposes a sensitivity analysis: Instead of matching the fraction of price dispersion explained by wages in 1975, we match the full price dispersion in 1975. This calibration implies a worse fit for the subsequent evolution of house price dispersion, but a better fit for the evolution of house price levels.

To understand more precisely the relationship between ability differentials and the sensitivity b_{p0} of house prices to wages, note that the price differential between two regions is determined by the wage differential of the *marginal household*. I.e., the wage decline that a household would incur by moving from its current region to the next-highest productivity region, holding –of course– its ability constant. In short, the house price differential reflects a constant-ability wage differential. The key observation is that the *observed* wage differential may be larger than the price differential because it not only reflects productivity differentials, but also the ability differentials of households. In particular, the more cross-sectional

---

[22] We implicitly assume that all members of the household share the same ability.

dispersion in ability there is (lower k_e), the smaller price differentials are relative to wage differentials. Cross-sectionally, this results in a lower sensitivity, b_{p0}, of house prices to wages (see Appendix B.1). When there are no ability differences (k_e is very large), the sensitivity of prices to wages is at its highest.

The 1975 data suggest a sensitivity of prices to wages requiring a Pareto coefficient of $k_e = 17.89$. This value implies a cross-sectional standard deviation of ability of 0.079. The variance $0.0062 = 0.079^2$ of ability of we use is not excessive. Appendix B.4 spells out an argument which shows that the difference between the overall cross-sectional variance of individual wages and the variance of wages that are averaged at the regional level is an upper bound on the cross-sectional variance of ability. Based on micro data from Heathcote, Storesletten, and Violante (2008a) and our regional data, we find that our ability variance is thirteen times smaller than this upper bound.

Transition Exercise In our main exercise below, we engineer an increase in the cross-regional wage dispersion of the same magnitude as in the data. The discipline in our transition experiment comes from assuming that the entire increase is generated by an increase in the log productivity dispersion over time. The dispersion of ability (i.e., k_e) and all other parameters stay constant during the transition. To keep things simple, we assume a linear increase in the productivity dispersion between the initial steady state (1975) and period 32 of the transition (2007). From period 33 (2008) onwards, we assume that productivity dispersion stays constant at its final steady state value. We choose the final steady state value so that we exactly hit the coefficient of variation of the real wage per job in 2007.

3.2.3 Construction Technology

We set the housing depreciation rate $\delta = 0.016$. This is the average depreciation rate over the last 35 years, calculated as the ratio of depreciation at current cost and the current cost net stock of residential fixed assets from the Fixed Asset Tables provided by the Bureau of Economic Analysis. See also Davis and Heathcote (2007).

We set the yearly endowment M_t of construction material so that, year-by-year, the aggregate housing supply per household in the model, H_t, matches the de-trended house size per household we observe in the data. For the years 1975-2007, we feed in the observed de-trended size. After 2007, the de-trended size equals the initial steady state value of 1,715 square feet. This procedure amounts to exogenously fixing the total quantity of square feet in the economy and letting the equilibrium endogenously allocate these square feet across regions.[23]

The last object we need to calibrate is the permit function $\Pi_t(A_t)$ which measures the maximum amount of construction per period in a region with productivity A_t. We start by assuming a constant permit function, i.e. $\Pi_t(A_t) = \pi_a$. This captures the notion that housing supply regulation is no tighter in some metropolitan areas than in others.[24] Because the parameter π_a determines the distribution of housing across islands, it indirectly governs the distribution of households across islands. Indeed, a larger π_a allows firms to construct more housing in high-wage areas, which in turns increases the population in these areas. This observation motivates us to choose π_a in order to match the 1975 concentration of jobs in high-wage metropolitan areas, as follows.[25] Each year, we sort the largest sample of regions we can find, into (equal-sized) wage quintiles and compute the fraction of jobs in each quintile. The data indicate an increase in the fraction of jobs that are concentrated in the highest wage quintile (Q5) from 64.9% in 1975 to 73.1% in 2007 (see Appendix D.5). We choose the parameter π_a in order to match the 1975 Q5 number.

3.2.4 Demographics

In order to control for the effect of demographics on house prices, we feed into our model the observed 1975-2007 data for the growth rate in the number of households g_{Nt} as well as for the number of jobs per household. Appendix E.3 provides the details. The growth rate in the number of households enters in the depreciation rate of the de-trended housing stock: $(1-\delta)/(1+g_{Nt})/(1+g_H)$.[26] Finally, the number of jobs is relevant for household earnings, which is the product of the real wage per job and the number of jobs per household.

---

[23] Of course, if we hold all other parameters the same, reducing M_t results in a smaller aggregate housing supply and raises house prices. In our benchmark calibration we find that a 10% decrease in the aggregate housing stock increases house prices by about 10%. See Appendix B.6 for details.

[24] While this assumption is uncontroversial as a description of the early 1970s, some have argued that housing supply restrictions have become tighter and more widespread over time. In Section 4.7, we allow for such a change and find that the quantitative effects of tightening regulation had negligible effects on equilibrium house prices.

[25] The strategy of calibrating $\Pi(A_t)$ directly to regulation data, instead of relying on its indirect impact on the population distribution, is not feasible. While there exist indices of housing supply constraints at the metropolitan level (e.g., Malpezzi (1996) and Saks (2005)), they have no time-series dimension. In addition, there is no natural mapping between such ordinal measures and our quantity constraint $\Pi(A_t)$.

[26] In the data, the number of households has grown faster than the population (1.53% average growth per year versus 1.05%) because the number of persons per household has declined (-0.46% growth). We feed in the faster growth rate in the number of households and thus capture its effect on house prices.

4 Quantitative Results

In this section we investigate the effects of feeding in the model the progressive increase in wage dispersion we observed in the data. We study the economy's transition from 1975 until 2007, and ultimately towards the new steady-state. In the figures we present below, the red dashed line denotes the initial 1975 steady state, the green dash-dotted line denotes the final steady state, the blue solid line denotes the transition path, and the dashed red line with circles denotes the data.

4.1 Wages

Our calibration parameters are picked so that, in the initial steady state of the model, the population-weighted cross-sectional mean and CV of the real wage per job is the same as in the 1975 data. Then, along the transition path, we linearly increase the dispersion of productivity by increasing the innovation variance of productivity shocks over 32 periods; see Appendix C.1.1 for the details. We pick the path of productivity dispersion so that the population-weighted CV of real wages in period 32 of the model's transition is the same as in the 2007 data. The procedure is summarized in Figure 4. The left panel plots the exogenous cross-sectional standard deviation of productivity we feed in the model: it increases from 0.0063 in the initial steady state to 0.1029 in period 32, and then stays constant until the final steady state. The right panel shows that this indeed allows us to match the 8.8 point increase in the CV in the data (red solid line). The middle panel shows that we also match the increase in the cross-sectional average wage between 1975 and 2007 as part of our detrending procedure.[27] Note that, even though the standard deviation of log productivity is held fixed after 2007 (period 32 of the transition), the mean and CV of wages continue to rise as the economy converges towards the final steady state. As more construction continues to take place in the newly productive regions, the population continues to relocate there. Table 2 shows the key moments in the data (Panel A) and in our benchmark model (Panel B). Rows 1 and 2 contain the moments for wages. Appendix B.5 shows that the model fits additional moments of the wage data beyond the mean and CV.

---

[27] Since we do not attempt to explain the cross-sectional (co-)movement of wages and house prices at business cycle frequencies, we do not try to match the entire time series of the average or CV of wages. In reality, other factors such as unemployment or interest rates, whose dynamics our model abstracts from, undoubtedly affect house prices.

Figure 4: Increasing the Wage Dispersion

The left panel plots the cross-sectional standard deviation of log productivity that arises from our calibration (exogenous). The middle panel plots the equilibrium population-weighted average of the (endogenous) real wage while the right panel plots the equilibrium population-weighted coefficient of variation. The red dashed line is the initial steady-state, the green dash-dotted line is the final steady-state, and the solid line (without markers) denotes the first 200 years along the transition path. This figure is for our benchmark calibration. In the middle and right panels, the dashed red line with circles plots the data from 1975 until 2007.

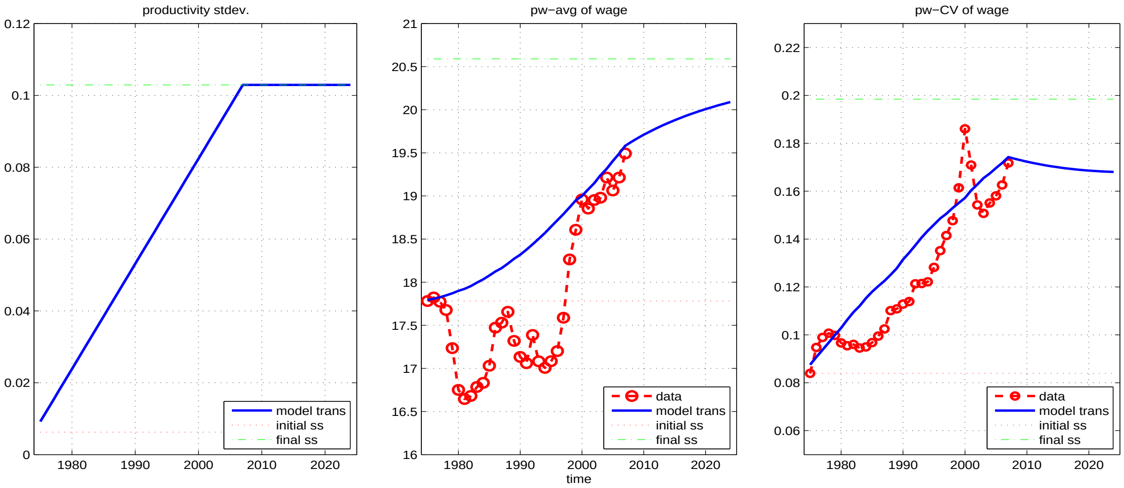

4.2 House Prices

Our main object of interest is the post-1975 evolution of house prices. In particular, the central question of our paper is whether the modest increase in wage dispersion (population-weighed CV increases by 8.8 points) can generate a large increase in house price dispersion (population-weighed CV increases by 38.2 points)? Figure 5 shows the model's predictions for the population-weighed average (left panel) and population-weighed CV (right panel) of house prices in the initial and final steady states (dashed lines), as well as along the transition path (solid line). Both the level and the CV of house prices are predicted to continue rising towards the new steady state after 2007.

The results show that our model features enough amplification to turn the modest increase in productivity dispersion into a large increase in house price dispersion. Our benchmark calibration generates an increase in the population-weighed CV of house prices of 51 points from 0.022 in the initial steady-state to 0.532 thirty-two periods into the transition, see Table 2, Panel B, Row 3. In the data, the CV increases from 0.154 in 1975 to 0.536 in 2007 (Panel A, Row 3). Thus the model is able to account for the observed cross-sectional dispersion in house prices in 2007. Because its initial steady state CV of house prices is lower than the observed 1975 value, the increase exceeds that in the data. This low initial CV is a direct consequence of the low observed 1975 price-wage sensitivity we match in our initial steady-state as part of the calibration. Section 4.6.1 contains a version of the model and Panel C contain calibration results when we match the 1975 CV of house prices instead.

Because of the convexity effect discussed in Section 2.4, the increase in dispersion also generates a moderate increase in the population-weighted average house price, from \$55,719 in the initial steady state to \$62,571 after 32 periods of transition (Row 4). In the data, average house prices increase from \$62,212 to \$87,013. Thus, while the model accounts for all of the increase in the dispersion of house prices, it can only account for one-third of the increase in house prices (11% vs 33% increase).

The above results are *population* statistics of house prices and wages because they are derived from and calculated for a model with a continuum of regions. While our continuum of regions model matches the key features of the observed wage distribution, we consider the additional exercise of feeding in the observed wage data from our unbalanced panel of 330 regions and computing *sample* statistics. More precisely, we evaluate the equilibrium price-wage function at the observed wage data for each region and each period. Having fed in the observed wages, we can then recompute the model's implications for house prices at the observed region-year observations. The results are almost identical to the population moments from the previous paragraph. First, the CV of house prices increases from 0.022 in

Figure 5: House Prices

The left panel plots the population-weighted cross-sectional mean of the real median home value, which re refer to as house price. It is calculated as the square foot housing price multiplied by the detrended housing size. The right panel plots the corresponding population-weighted coefficient of variation, the ratio of standard deviation to the mean. In both panels, the red dashed line denotes the initial 1975 steady state. The green dash-dotted line denotes the final steady state which is reached well beyond 2007. The blue solid line denotes the transition from the initial steady state to the final steady state. The dashed red line with circles plots the data from 1975 until 2007.

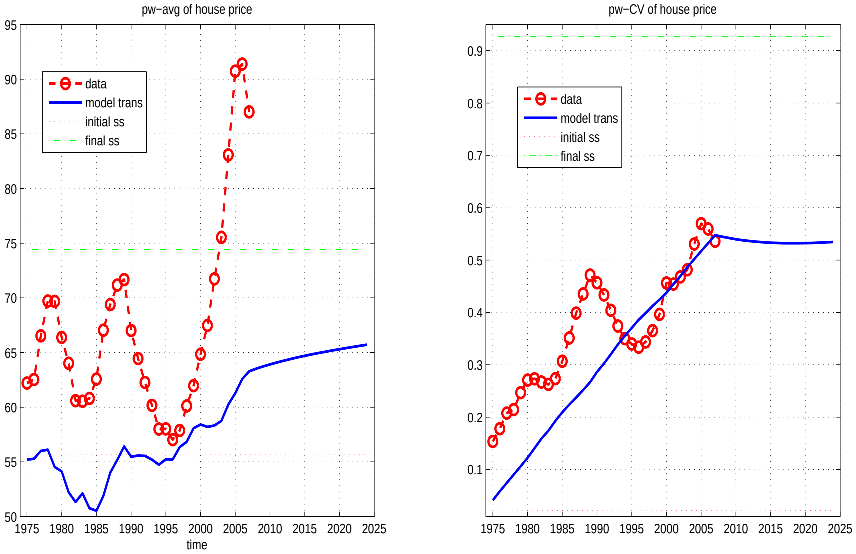

the initial steady state to 0.545 thirty-two periods into the transition. Second, the average house price level increases from \$55,755 to \$62,478. The reason for this close correspondence between population and sample moments is that the population distribution for wages does a good job describing its sample counterpart; see Section 4.1.

Our benchmark model generates a strong increase in the ratio of the average house price relative to its construction cost.[28] This ratio increases from 1.75 in the initial steady state to 2.62 thirty-two periods into the transition, and continues to climb to its final steady state value of 3.65 afterwards. A similarly strong increase in the non-structure component of house prices is present in national and regional data (Davis and Heathcote (2007) and Glaeser et al. (2007)). Davis and Heathcote (2007) emphasize the value of land; similarly Glaeser, Gyourko, and Saks (2007) emphasize the value of the right to build on that land. The empirical evidence suggests that the non-structure component accounts for about half of the value of the U.S.

---

[28] The assumption of a centralized market for construction material implies that construction costs are the same in every island in the model. Although this implication of the model is violated in the data, Davis and Palumbo (2008) show that it is a reasonable approximation: very little of cross-sectional variation in housing prices is due to variation in construction costs.

Figure 6: Price-Wage Sensitivity

Each period we run a cross-sectional regression of real house prices on real wages. The slope coefficient b_p is computed as the population-weighted covariance of prices and wages, divided by the population-weighted variance of wages. The dashed line with circles denotes the time series for b_p in the data, while the solid line (no markers) denotes the same slope in the model. The model's slope is computed by feeding in the observed wage data, and evaluating them at the equilibrium price function.

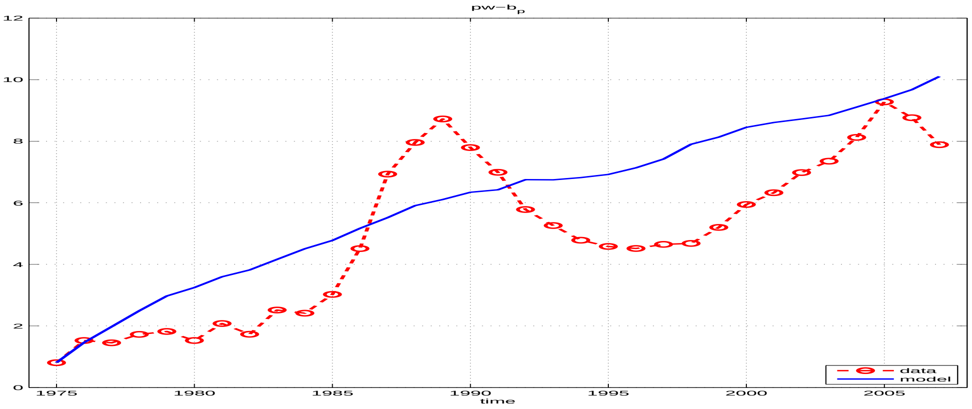

housing stock, and it has risen much faster than the structure component since the 1970s. In our model, this non-structure component is the shadow price of an additional construction permit. As a region becomes more productive and thus more attractive to households, this shadow price increases. Hence, our model is able to account for the increase in the house price to construction cost ratio despite the fact that the quantity of construction permits stays constant.

4.3 Sensitivity of Prices to Wages

The model captures several features of the observed relationship between wages and house prices. Our calibration is designed to match the population-weighted sensitivity coefficient $b_{p0} = 0.81$ that arises from running a cross-sectional regression of house prices on wages in 1975. Second, the increase in productivity dispersion generates an increase in that sensitivity coefficient over time from 0.81 in the initial steady state to 9.82 thirty-two periods into the transition. The wage-feeding exercise described above generates an increase from 0.81 to 10.10. In the data, the increase in sensitivity is similar: from 0.81 in 1975 to 7.89 in 2007. Figure 6 plots the population-weighted sensitivity coefficient b_p in the data (solid red line with dots) and in the model (solid blue line).

The wage-feeding exercise suggests an important specification test for our model: does it predict the right amount of sensitivity of house prices to wages as measured by the R^2

statistic in equation (26)? In our benchmark model, we obtain an R^2 statistic of 31.0%, while the value in the data is 26.5%. So, wages account for abound 30% of the variation in house prices across regions in both model and data. This R^2 statistics of 31% arises because of two features of the calibration. First, we choose ability dispersion in order to match the initial sensitivity of house prices to wages, b_{p0}. Indeed, a model without ability dispersion would lead to model-implied house prices that are extremely sensitive to wages. The slope b_p from a repeated cross-sectional regression of house prices on wages, averaged over time, would be 12.22, compared to 5.39 in our benchmark model. It would result in prices that vary too much in the cross-section: the R^2 statistics would be -31.5%. The second feature of the calibration that matters for the R^2 statistics of 30% is the endogenous increase in the price-wage sensitivity, b_p, along the transition path. If we were to keep the sensitivity constant along the transition, the R^2 would drop from 30% to about 0.06%. In short, while (constant) ability dispersion is instrumental in obtaining the right *level* of price-wage sensitivity, the increase in productivity dispersion is instrumental in obtaining the right *increase* in price-wage sensitivity.

The increasing sensitivity provides indirect evidence for an increasing productivity dispersion as the root cause for the increase in the wage dispersion. The alternative explanation is that the increase in wage dispersion is due to highly-skilled workers becoming relatively more productive. In our model, this would amount to an increase in ability dispersion. However, as explained in the second paragraph of Section 3.2.2, an increase in wage dispersion engineered through an increase in ability dispersion would result in a *reduction* in the cross-sectional sensitivity of house prices to wages. The data show an increasing sensitivity.

Finally, our model predicts a convex relationship between houses prices and productivity. In our calibrated model, this translates into a convex relationship between house prices and wages. We find direct evidence for a similar convex relationship in the data: a cross-sectional regression of house prices on demeaned wages and squared demeaned wages generates not only significant linear but also quadratic coefficients.

4.4 Population Dynamics

Our calibration guarantees we match the 64.88% share of population that lives and works in the 20% highest-wage regions (Q5). In the benchmark model, Q5 drops from 64.88% to 46.95% in the first period of the transition. It then rises gradually to 55.18% over the next 32 periods. Although the initial drop is counterfactual (we explain why it occurs below and propose a resolution in Section 4.6.2) the 8.2% increase between 1976 and 2007 is similar (identical) to the 9.4% (8.2%) increase in the data between 1976 (1975) and 2007. The increase in population concentration after 1976 is made possible by increased construction in high-productivity regions. In contrast, low-productivity regions see no construction, a declining housing stock because of depreciation, and they lose population. This construction pattern facilitates population concentration in the highest-wage quintile. Figure 7 shows that the population further concentrates towards the highest productivity regions as the economy moves towards the final steady-state, at which point Q5 is 80.50%. Even though there are no more exogenous changes to the productivity process after 2007 and the construction threshold has reached its steady state value, the housing distribution continues to adjust towards its steady state. This continued population concentration towards high-wage regions explains why the population-weighted average and CV of wages and house prices in Figures 4 and 5, respectively, continue to increase in the model after 2007.

Figure 7: Population Distribution

This figure plots the population distribution by wage quintile in the benchmark model. We use the model with an increasing productivity dispersion to generate a population time-series for each MSA. We sort the MSAs into five equally sized wage bins and calculate the ratio of the number of people in each quintile to the number of people in the economy (normalized to 1). The graph shows the distribution in the initial steady state (1975, left bars), after 32 years (2007, middle bars) and in the final steady state (right bars).

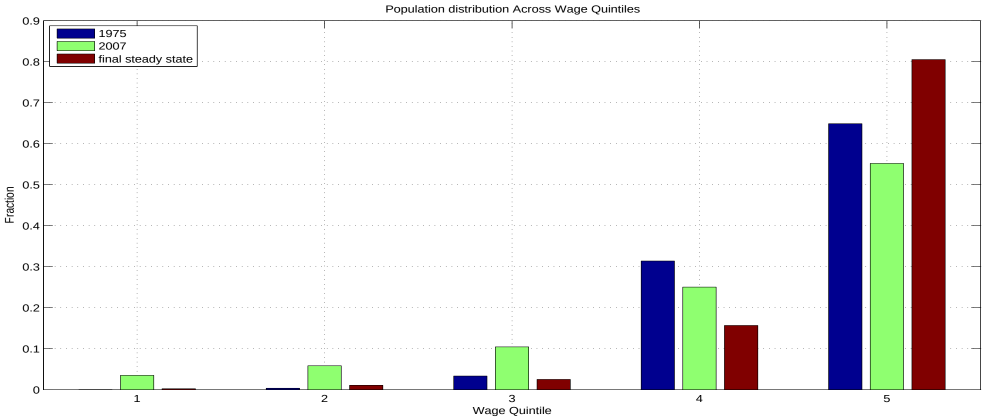

As mentioned above, one problem with the benchmark calibration is the large initial drop in Q5 between the initial steady state and the first period of the transition. This drop is an artefact of the specific mechanics driving the increase in cross-sectional wage dispersion. Indeed, in order to obtain a gradual, linear increase in the dispersion of log productivity, plotted in the left panel of Figure 4, we need a large jump in the *innovation* variance of log productivity in period 1, σ_{a1}. This high innovation variance acts as a big shock to the wage distribution between the initial steady state and period 1. Previously high-productivity, high-population regions may draw a very negative productivity innovation which puts them in a lower wage quintile and vice versa. This breaks the strong association between population and wages, resulting in the initial drop in Q5.[29] One consequence of these productivity dynamics is that the *rank correlation of wages* between adjacent years is counter-factually low. For instance, the rank correlation between the initial steady state and the first period of the transition is 68.8% compared to a rank correlation of 96.8% in the 1975-76 wage data. This rank correlation gradually increases along the transition path and exceeds 95% only twenty years into the transition. In the data it is above 95% throughout. In Section 4.6.2 below, we consider an alternative calibration that increases productivity dispersion in a rank-preserving way. Its prediction for the fraction of jobs in the highest wage quintile matches the data.

4.5 Mobility

The increasing dispersion in productivity causes migration from previously productive to newly productive regions. The migration pattern and the magnitude of the migration rate predicted are similar in model and data.

To measure migration in the data, we use U.S. Census data for the in-migration and out-migration between 1995 and 2000, available for 271 MSAs. Net migration is defined as the difference between in- and out-migration. We focus on the sub-population of young (25-39), single, college-educated because this group is more likely to move for productive reasons, and therefore, to more closely approximate the agents in our model. We sort regions into 25 wage-per-job bins and compute net migration rates for each bin. We adjust for population growth by scaling the population in 2000 so that it is the same as in 1995. We measure net migration in the same way in the model; Appendix C.2 contains the details.

When we compare model to data, we find a similar pattern: out-migration from low-wage areas and in-migration into high wage areas. The top-10% lowest-wage regions see an out-migration of about 7% in both model and data. The top-10% highest-wage regions see an in-migration of about 1.5% in both model and data. Figure 8 shows the annual net migration rate in the benchmark model and in the data. These results suggest that the assumption of frictionless reallocation does not lead to excessive migration. This is due to

---

[29] Since we assume that the innovation variance reaches its steady state value from period 33 onwards, the linear increase in the standard deviation of log productivity also requires that the innovation variance increases between period 1 and period 32, and that it overshoot its final steady state value. Precisely, the time path for the innovation standard deviation of log productivity in our benchmark calibration is as follows: 0.0004 (initial steady state), 0.0068 (period 1), 0.008 (period 2), gradually increasing to 0.0263 (period 32), then constant at 0.0103 (from period 33 until final steady state).

the high persistence of wages. Indeed, consider the equilibrium of a version of our model with constant wages: in that extreme case, nobody would find it optimal to move despite perfect mobility.

Figure 8: Net Migration Rates

The model plots net migration rates in the benchmark model (left bars) and in the data (right bars). Each pair of bars represents the net migration (in-migration minus out-migration) between 1995 and 2000 of a group of regions with similar real wages. In particular, we form 25 groups of regions, sorted by their real wage from lowest to highest. The data is from the U.S. Census for young, single, college-educated persons. The model computes migration rates in the same way as in the data.

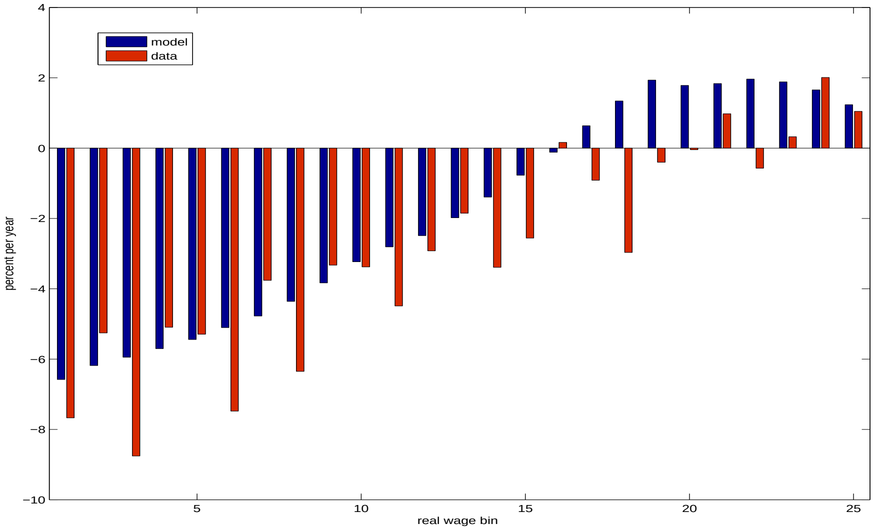

4.6 Robustness

In this section, we discuss two alternative calibrations and their implications for wages and house prices.

4.6.1 Alternative 1: Calibrating to Initial CV of House Prices

First, we explore a calibration in which the initial steady state CV of house prices matches the 1975 value of 0.154 in the data. See Table 2, Panel C. This is an alternative to matching the 1975 sensitivity of prices to wages b_p in Panel B. The calibration, which continues to match the level and CV of wages in 1975 and 2007, features less ability dispersion and more regional productivity dispersion than the benchmark. In particular, the cross-sectional standard deviation of ability is 0.056 compared to 0.079 in the benchmark calibration (k_e

is 25.28 compared to 17.89), the initial productivity dispersion is higher (0.039 compared to 0.006), and it rises to a higher value in the final steady state (0.135 versus 0.103). This calibration predicts a rise in the population-weighted CV of house prices of 53 points, similar to the 51 point increase in the benchmark. It generates a substantially larger increase in house price levels: a 23% increase compared to an 11% increase in the benchmark and a 33% increase in the data. The downside to matching the initial CV of house prices is that we overstate the initial sensitivity coefficient: $n_{p0} = 5.37$ versus 0.81 in the 1975 data. Both models feature a similar increase in this sensitivity coefficient. Hence, this alternative calibration introduces excess sensitivity of house prices to wages: a year-by-year cross-sectional regression of house prices on wages delivers a slope coefficient of 8.2 (averaged over time) inside the model and 3.2 in the data. The excess sensitivity problem is still less pronounced than in the model without ability dispersion though. The R^2 for house prices in equation (26) is 19% in this model compared to 26.5% in the data and 31% in the benchmark model.

4.6.2 Alternative 2: Rank-Preserving Increase in Wage Dispersion

Second, we explore a calibration in which the increase in the standard deviation of productivity dispersion is engineered in a different way. As explained above, in order to generate a linear increase in dispersion, we need that the innovation variance jumps from the initial steady state to the first period of the transition, then gradually rises until period 33, and then jumps back down to the final steady state value. The initial jump in innovation variance introduces too much "mixing" in the cross-sectional productivity distribution and leads to a counter-factually low rank correlation between city-specific wages in adjacent periods. As an alternative, we consider a lower time path for the innovation standard deviation of log productivity: a linear rise from its initial steady state to its final steady state value over 33 periods. Because this lower time path of productivity shocks, by itself, generates a smaller increase in productivity dispersion than in our benchmark, we add a second engine of productivity dispersion. We deterministically increase the productivity of regions above the average and, vice versa, decrease the productivity of regions below the average. Importantly, unlike random productivity shocks, these deterministic changes preserve the rank of regions in the productivity distribution. Formally, the law of motion for log productivity becomes:

$$a_t = \text{avg}(a_{t-1}) + \rho_t \left(a_{t-1} - \text{avg}(a_{t-1})\right) + \sigma_{at}\varepsilon_t,$$

where $\text{avg}(a_{t-1})$ denotes the cross-sectional average productivity. Setting $\rho_t > 1$ allows us to increase dispersion in a rank-preserving way; Appendix C.1 explains the mechanics in detail.

As in the benchmark model, we match the 1975 sensitivity coefficient b_{p0}. The calibration is essentially identical to the benchmark model.

The main moments of interest are discussed in Panel D of Table 2. First, the model continues to generate a large increase in the CV of house prices for a given increase in the CV of wages. The increase is 60 points compared to 51 points in the benchmark. Second, it generates a larger increase in house prices: 19% increase compared to 11% in the benchmark. Third, the sensitivity coefficient increases strongly from 0.81 in 1975 to 11.95 in 2007. Fourth, and most significantly, this calibration generates population dynamics close to those observed in the data. The fraction of people working and living in the top-20% regions in terms of wage is 64.88% in the initial steady state and rises to 74.77% after 32 periods. This is close to the 73.09% in the 2007 data. This calibration avoids the steep drop in $Q5$ in the first period of the transition, which we noted for the benchmark model. Instead, the 1976 value for $Q5$ in the model is 64.62%, close to the initial steady state value. The population then gradually relocates towards the newly productive regions. In the final steady state, $Q5$ reaches a value of 80.37%, similar to the benchmark model. The key difference with the benchmark model, therefore, is the transition path of $Q5$. Because it avoids the initial drop in population, this model generates a higher increase in population-weighted house prices and a higher population-weighted sensitivity of prices to wages. By the same token, this version of our model matches the rank correlation of wages between adjacent years. It is 99.55% on average in the model and 99.22% on average in the data. Finally, the model generates an R^2 statistic of 25.1%, close to the 26.5% number in the data.

4.7 Increase in Regulation

In this section, we use our model to pursue an alternative explanation for the increase in the level and dispersion of house prices: housing supply regulation has become tighter over time and has gradually spread geographically (from the coastal areas inland). To keep matters as simple as possible, we hold all parameters of the benchmark calibration fixed with one exception. Instead of increasing the dispersion of productivity over time, we hold it fixed at its initial steady state level. The resulting initial steady state wage distribution is the same as in our benchmark model and matches the observed 1975 wage distribution as before. Instead, we tighten regulation. We let the number of permits at time t in a region with productivity A be

$$\Pi_t(A) = \pi_a \left(\frac{A}{A_{\min}} \right)^{\phi_t}.$$

The elasticity parameter $\phi_t = 0$ in 1975 as in our benchmark model. We let it decrease from 0 to $\phi_t = -0.5$ between periods one and thirty-two of the transition. We keep it constant at -0.5 after 2007 (periods thirty-three and later). Figure 9 illustrates how lowering the supply elasticity parameter ϕ reduces the number of permits, and more so in highly productive regions. This rotation captures the stylized fact that supply regulation gradually tightened over the last three decades, especially among more productive regions such as the coastal metropolitan areas.

Figure 9: Tightening Housing Supply Regulation

This figure plots the permit function $\Pi(A) = \pi_a \left(A/A_{\min}\right)^{\phi}$. The top line denotes the situation in 1975 when $\pi_a = .30026$ and $\phi = 0$. The bottom line denotes the situation in 2007 and beyond when $\pi_a = .30026$ and $\phi = -0.5$. In the years between 1975 and 2007, ϕ decreases linearly from 0 to -0.5, so that the permit function gradually rotates from the top line to the bottom line.

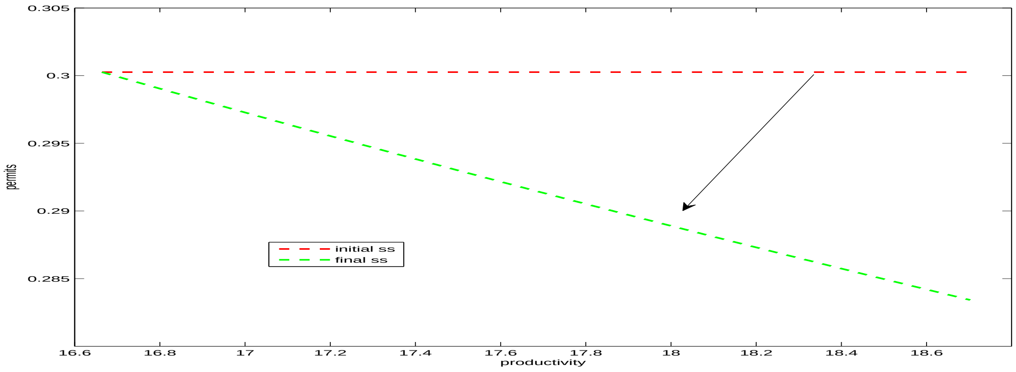

Panel E of Table 2 shows the results of the regulatory tightening exercise. We find that decreasing building permits has quantitatively minor effects on average house prices and on the dispersion of house prices. While both increase, the increases are quantitatively small compared to those we found in our benchmark exercise. The same is true for the price-wage sensitivity which increases only slightly over time. The population in Q5 is also almost constant. Finally, our R^2 goodness-of-fit metric for house prices is only 5%, one-fifth of its value in the data and one-sixth of its value in the benchmark model. The intuition for the small impact of regulation is simple. While tighter regulation reduces the supply of houses in high-wage metropolitan areas, the equilibrium response of labor is to move out, thereby effectively reducing the housing demand in those same areas. The net effect is a tiny increase in price. A similar intuition is at work in the closed city model of Arnott and MacKinnon (1977) and the open city model of Aura and Davidoff (2008). We have explored alternative values for the supply elasticity parameter ϕ (-3, -1, -0.1, and even +1). The results were quantitatively similar across cases because of the endogenous response of mobility to the various regulatory changes. In the same vein, tightening regulation alongside an increase in wage dispersion delivers the same quantitative results as in our benchmark calibration with constant regulation. Impediments to labor mobility, absent from the model, may slow down the reduction in housing demand, but are unlikely to reverse it. These results suggest that an increase in wage dispersion is an important ingredient to generate a quantitatively meaningful increase in house price level and dispersion.

4.8 Rental Prices

In the model, the price of a house equals the present discounted value of the rents. An alternative to testing the model's implications for house prices would be to test its implications for rents. After all, the spatial equilibrium model also predicts a relationship between rents and wages. We collected nominal rent data from the Fair Market Rents database, as detailed in Appendix D.3. As we did with nominal house prices, we deflate them by the regional non-housing CPI as well as by the trend in house size. Census data suggest that the size of multi-family homes, which is likely to be rental housing, grew at the same rate as single-family housing, which is likely to be owner-occupied. The rental data are only available in 1982 and from 1984 until 2007. As the last two rows of Table 2, panel A show, de-trended real rents seem to have fallen from \$4,220 per year in 1982 to \$3,420 in 2007. The CV increased only moderately, from 0.153 in 1982 to 0.190 in 2007. Our model generates a moderate increase in average rents and a large increase in the CV of rents, just as with house prices. Therefore, while the model can account for the observed increase in house price dispersion, it produces an increase in rent dispersion that is too large relative to the data. This is akin to the excess volatility puzzle, according to which equity prices are too volatile relative to their underlying dividends.

A potential explanation for this divergence is that the cash flows entering the present value formula for house prices are unlikely to be the rents we measure in the data. Glaeser (2007a) make several compelling empirical observations suggesting that house price and rent series can be best understood as the costs of two different types of housing, reflecting different demands on two related, but not directly comparable, markets.[30] This market segmentation causes a (severe) selection problem when comparing the present-value of observed rents to ob-

---

[30] Rents in our model are then to be interpreted as the per-period user cost of owner-occupied housing. Since there are no regional data on the user cost, and since single-family ownership price data are of high quality, it seems natural to test the model using house price data instead. Like us, the bulk of the spatial location literature derives implications for (implicit) rents, but almost always tests them on owner-occupied house price data (recent examples are Gyourko, Mayer, and Sinai, 2006; Glaeser, 2007b).

served house prices. This selection problem could help explain our empirical observation that owner-occupied house price dispersion increased much more than rent dispersion: Indeed, the increase in income inequality, the key driving force of our model, was most pronounced in the top half of the income distribution, a group that is more likely to be composed of homeowners. One way to address selection would be to study housing units that are both for rent and for sale: unfortunately there exists no such regional panel data set for the United States, but other countries or certain regions with the U.S. may have such data. Another way to go would be to develop a model where agents choose to self-select into the rental or the ownership market. This extension is left for future research.

5 Conclusion

Our paper provides a new general equilibrium framework for analyzing the joint dynamics of regional income, house prices, and housing quantities. It extends the Rosen-Roback spatial equilibrium model along several dimensions in order to establish closer contact with the data. We used our framework to study the quantitative effect of wage dispersion and housing supply regulation for the regional house price level and its dispersion. The model accounts for several features of the joint price-wage distribution. Faced with an increase in the productivity dispersion across metropolitan areas, households choose to reallocate from lower towards higher-productivity metropolitan areas. This pushes up house prices in high-wage areas. The observed increase in wage dispersion is sufficient to generate the observed increase in the house price dispersion across metropolitan areas.

The same thirty years since 1975 also saw a tightening of housing supply regulation, especially in the coastal areas. One might think that the supply effect induced by this regulatory tightening could, in and of itself, account for the increase in house price level and dispersion. However, because the equilibrium response of households to move out of the more tightly regulated areas, the house price effects of tightening supply restrictions are small. So, while supply constraints are important, the increase in wage dispersion is an essential part of the explanation.

The model's prediction of an increasingly strong cross-sectional sensitivity of house prices to wages is consistent with the data. It suggests that increasing dispersion of regional productivity, as opposed to an increasing dispersion in the ability of households, underlies the changes in spatial location, wage, and house price patterns we have observed over the last three decades.

Table 1: Benchmark Calibration

The six parameters with a notation "E" next to them are determined by repeatedly solving the model until six moments in the data, listed in the last column of the corresponding rows, are matched exactly. One of the six parameter has a "E^*" next to indicate that it governs a 2007 moment of the data, The other ones, without a star, govern features of the 1975 data. While none of the parameters solely pin down the moment mentioned on the same row, that parameter is in practice the key parameter for matching that moment. the abbreviation "pw-CV" stands for population-weighted coefficient of variation and "ss" for steady state.

	Parameter	*Description*	*Benchmark*	*Source/Matches*
		Preferences		
	β	*time discount factor*	0.9510	*historical avg. 30-yr fixed-rate mortgage rate*
	γ	*inverse elasticity of housing demand*	2	*Hanushek & Quigley (1980)*
E	κ	*weight on housing*	7.0248	*average housing expenditure share*
		Productivity and Ability Processes		
E	$\bar{A}$	*mean productivity level*	17.654	*1975 pw-avg. of real wage per job*
E	σ_{a0}	*productivity dispersion in initial ss*	0.0063	*1975 pw-CV of real wage per job*
E^*	σ_{aT*}	*productivity dispersion in final ss*	0.1029	*2007 pw-CV of real wage per job*
	λ	*death rate of regions*	0.01	*consistent with AC of wages*
	σ_{bt}	*productivity dispersion at birth*	*see text*	
E	k_e	*governs ability dispersion*	17.889	*1975 pw- sensitivity of house prices to wages*
		Technology		
	δ	*depreciation rate of housing stock*	0.0160	*Bureau of Economic Analysis*
	M_t	*construction material*	*see text*	*house size from Census of Construction*
E	π_a	*building permits*	0.3003	*1975 fraction in the top-20% wage regions*

Table 2: Summary Statistics Data and Model

The table reports the cross-sectional average and coefficient of variation, defined as the ratio of the standard deviation to the mean, of the real wage per job (Rows 1 and 2) and the real median home value (Rows 3 and 4) across US metropolitan areas. Row 5 reports the sensitivity of real house prices to real wages, b_p, measured as the time-series average of the slope coefficient of a cross-sectional linear regression of real house prices on real wages. Row 6 reports the fraction of jobs in the highest wage quintile, $Q5$. Rows 7 and 8 report the cross-sectional mean and CV of the real rent. Row 9 reports the R^2 statistics of equation 26. In the data panel (A), this is simply the R^2 of a regression of observed prices on wages, measuring the fraction of observed house prices variation accounted for by wages in the data. In the model panel (B to E), the R^2 measures the empirical variation that can be accounted for by model-generated prices, i.e. the prices obtained by feeding in the model the wages we observed in the data. The cross-sectional mean and CV moments are population-weighted, as well as the sensitivity coefficient b_p. Panel A is for the data. Real wages, prices, and rents are obtained by dividing the nominal amounts by the regional CPI ex-shelter. The means represent thousands of 1983 dollars; the coefficient of variation is unit-free. The sample is the unbalanced panel of 330 metropolitan statistical areas for which we have house price data. Population weights are computed as the number of jobs in that region relative to the full sample. House prices and wages are detrended as explained in the main text. Panel B is for our benchmark model. The first column (1975) refers to the initial steady state, the second column (2007) to period 32 of the transition path towards the new steady-state. Panel C is for an alternative calibration that targets the 1975 CV of house prices instead of the 1975 sensitivity coefficient b_p. Panel D is for an alternative calibration that engineers the increase in wage dispersion in a rank-preserving way. Panel E investigates a tightening of housing supply regulation. The numbers denoted by a star are values in 1982 instead of 1975; 1982 is the first observation on rental data. The corresponding numbers in the models denote period 8 of the transition instead of the initial steady state.

	A: Data		B: Benchmark		C: Altern. 1		D: Altern. 2		E: Regulation	
	1975	2007	1975	2007	1975	2007	1975	2007	1975	2007
Mean wage	17.78	19.49	17.78	19.49	17.78	19.69	17.78	20.32	17.78	17.78
CV wage	0.084	0.172	0.084	0.172	0.084	0.172	0.084	0.172	0.084	0.084
Mean hp	62.21	87.01	55.72	62.57	55.67	70.15	57.65	69.77	55.72	55.80
CV hp	0.154	0.536	0.022	0.532	0.154	0.671	0.021	0.620	0.022	0.025
b_p	0.81	7.89	0.81	9.82	5.61	13.60	0.81	11.95	0.81	0.91
Q_5	64.88	73.09	64.88	55.18	64.88	61.73	64.88	74.77	64.88	64.89
Mean rent	4.22*	3.42	4.45*	5.00	4.52*	5.51	4.57*	5.03	4.39*	4.09
CV rent	0.153*	0.190	0.126*	0.515	0.252*	0.651	0.094*	0.602	0.021*	0.025
R^2	26.46		30.96		19.06		25.08		5.61	

References

Fernando Alvarez and Marcelo Veracierto. Labor market policies in an equilibrium search model. *NBER Macroeconomic Annual*, pages 265–303, 1999.

Fernando Alvarez and Marcelo Veracierto. Fixed-term employment contracts in an equilibrium search model. Working Paper, University of Chicago, 2006.

Richard J. Arnott and James G. MacKinnon. Measuring the cost of height restrictions with a general equilibrium model. *Regional Science and Urban Economics*, pages 359–375, 1977.

Saku Aura and Thomas Davidoff. Supply constraints and housing prices. *Economic Letters*, 99: 275–277, 2008.

Sean D. Campbell, Morris A. Davis, Joshua Gallin, and Robert F. Martin. What moves housing markets: A variance decomposition of the rent-price ratio. *Journal of Urban Economics*, 66:90 – 102, 2009.

Karl E. Case and Robert J. Shiller. Prices of single-family homes since 1970: New indexes for four cities. *New England Economic Review*, pages 46–56, September/October 1987.

YiLi Chien and Hanno Lustig. The wealth distribution and aggregate risk. *Review of Financial Studies*, 2009. Forthcoming.

Daniele Coen-Pirani. Understanding gross worker flows across U.S. states. Working Paper, Carnegie Mellon University, August 2006.

Tim Cogley. Idiosyncratic risk and the equity premium: Evidence from the consumer expenditure survey. *Journal of Monetary Economics*, 49:309–334, 2002.

George M. Constantinides and Darrell Duffie. Asset pricing with heterogeneous consumers. *Journal of Political Economy*, 104:219–240., 1996.

Arnaud Costinot and Jonathan Vogel. Matching and inequality in the world economy. Working Paper, MIT and Columbia University, 2008.

Morris A. Davis and Jonathan Heathcote. The price and quantity of residential land in the united states. *Journal of Monetary Economics*, 54:2595 – 2620, 2007.

Morris A. Davis and Michael G. Palumbo. The price of residential land in large us cities. *Journal of Urban Economics*, 63:352 – 384, 2008.

Jan Eeckhout. Gibrat's law for (all) cities. *American Economic Review*, 94(5):1429–1451, 2004.

Xavier Gabaix and Augustin Landier. Why has CEO pay increased so much? *Quarterly Journal of Economics*, 123:49–100, 2008.

Edward L. Glaeser and Joseph Gyourko. The impact of zoning on housing affordability. *Economic Policy Review*, 9(2):21–39, 2003.

Edward L. Glaeser and Joseph Gyourko. Urban decline and durable housing. *Journal of Political Economy*, 113(2):345–375, 2005.

Edward L. Glaeser, José Scheinkman, and Andrei Schleifer. Growth in cities. *Journal of Political Economy*, 100(6):331–370, 1992.

Edward L. Glaeser, José Scheinkman, and Andrei Schleifer. Economic growth in a cross-section of cities. *Journal of Monetary Economics*, 36:117–143, 1995.

Edward L. Glaeser, Joseph Gyourko, and Raven E. Saks. Why is Manhattan so expensive? *Journal of Law and Economics*, 48(2):331–370, 2005.

Edward L. Glaeser, Joseph Gyourko, and Raven E. Saks. Why have house prices gone up? *American Economic Review Papers and Proceedings*, 95(2):329–333, 2007.

Joseph Glaeser, Edward L. Gyourko. Arbitrage in the housing market. Working Paper, Harvard University and Wharton School of Business, December 2007a.

Joseph Glaeser, Edward L. Gyourko. Housing dynamics. Working Paper, Harvard University and Wharton School of Business, May 2007b.

Clive W. J. Granger. Investing causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(4):424–438, 1969.

Joseph Gyourko, Christopher Mayer, and Todd Sinai. Superstar cities. Working Paper, NBER, July 2006.

Eric A. Hanushek and John M. Quigley. What is the price elasticity of housing demand? *The Review of Economics and Statistics*, 62(3):449–454, 1980.

Jonathan Heathcote, Kjetil Storesletten, and Gianluca Violante. The macroeconomic implications of rising wage inequality in the U.S. Working Paper, Federal Reserve Bank of Minneapolis, University of Oslo, and NYU, 2008a.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Insurance and opportunities: A welfare analysis of labor market risk. *Journal of Monetary Economics*, 55:501–525, 2008b.

Charles Himmelberg, Christopher Mayer, and Todd Sinai. Assessing high house prices: Bubbles, fundamentals, and misperceptions. *Journal of Economic Perspectives*, 19:67–92, 2005.

Andreas Hornstein, Per Krusell, and Giovanni L. Violante. The effects of technical change on labor market inequalities. In Philippe Aghion and Steven Durlauf, editors, *Handbook of Economic Growth*. Elsevier Science, Amsterdam, 2004.

Matteo Iacoviello. House prices, borrowing constraints, and monetary policy in the business cycle. *American Economic Review*, 95(3), June 2005.

Ken-Ichi Inada. On a two-sector model of economic growth: Comments and generalization. *Review of Economic Studies*, 30(2):119–127, 1963.

Dirk Krueger and Jesús Fernández-Villaverde. Consumption over the life cycle: Facts from the consumer expenditure survey. *Review of Economics and Statistics*, 89:552–565, 2006.

Dirk Krueger and Fabrizio Perri. Does income inequality lead to consumption inequality? Evidence and theory. *The Review of Economic Studies*, 74:163–193, 2005.

Robert E. Lucas and Edward C. Prescott. Equilibrium search and unemployment. *Journal of Economic Theory*, 7(2):188–209, 1974.

Hanno Lustig and Stijn Van Nieuwerburgh. Housing collateral, consumption insurance and risk premia: An empirical perspective. *Journal of Finance*, 60(3):1167–1219, 2005.

Hanno Lustig and Stijn Van Nieuwerburgh. Can housing collateral explain long-run swings in asset returns? Working Paper, UCLA and NYU Stern, June 2007.

Hanno Lustig and Stijn Van Nieuwerburgh. How much does household collateral constrain regional risk sharing. *Review of Economic Dynamics*, Forthcoming, 2010.

Stephen Malpezzi. Housing prices, externalities, and regulation in U.S. metropolitan areas. *Journal of Housing Research*, 7(2):209–241, 1996.

Makoto Nakajima. Rising earnings instability, portfolio choice, and housing prices. Working Paper, University of Illinois, Urbana-Champaign, 2005.

Francois Ortalo-Magné and Andrea Prat. Spatial asset pricing: A first step. Working Paper, University of Wisconsin and LSE, 2008.

Monika Piazzesi, Martin Schneider, and Selale Tuzel. Housing, consumption and asset pricing. *Journal of Financial Economics*, Forthcoming, 2006.

John Quigley and Steven Raphael. Regulation and the high cost of housing in California. *American Economic Review Papers and Proceedings*, 95(2):323–328, 2005.

John Quigley and Larry Rosenthal. The effects of land use regulation on the price of housing: What do we know? what can we learn? *Cityscape*, 8(1):69–138, 2005.

Jennifer Roback. Wages, rent, and the quality of life. *Journal of Political Economy*, 90(2):191–229, 1982.

Sherwin Rosen. Wage-based indexes of urban quality of life. In Mieszkowski and Straszheim, editors, *Current Issues in Urban Economics*. Johns Hopkins University Press, Baltimore, 1979.

Sherwin Rosen and Robert Topel. Housing investment in the United States. *Journal of Political Economy*, 96(4):718–740, 1988.

Raven E. Saks. Job creation and housing construction: Constraints on metropolitan area employment growth. Federal Reserve Board Finance and Economics Discussion Series No. 2005-49, September 2005.

Michael Sattinger. Assignment models of the distribution of earnings. *Journal of Economic Literature*, 31:831–880, 1993.

Robert Shimer. Mismatch. *American Economic Review*, 97:1074–1101, 2005.

Matthew Spiegel. Housing returns and construction cycles. *Real Estate Economics*, 29(4):521–551, 2001.

Kjetil Storesletten, Chris Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *The Journal of Monetary Economics*, 51:609–633, 2004.

Kjetil Storesletten, Chris Telmer, and Amir Yaron. Asset pricing with idiosyncratic risk and overlapping generations. Working Paper, University of Oslo, February 2006.

George Tauchen and Robert Hussey. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica*, 59(2):371–396, 1991.

Marko Terviö. The difference that CEOs make: An assignment model approach. *American Economic Review*, 98:642–668, 2008.

Menahem E. Yaari. Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32:137–150, 1965.

A Proofs

A.1 Preliminary Results

The following Lemma compiles technical results which are used in the following subsection.

Lemma 6. *Consider some strictly increasing strictly concave, and twice continuously differentiable function* $v : (0, \infty) \to \mathbb{R}$*. Suppose that* $v(h)$ *goes to minus infinity as* h *goes to zero, and that* $v(h)$ *goes to zero as* h *goes to infinity. Then*

1. *The derivative v'(h) goes to infinity as h goes to zero, and goes to zero as h goes to infinity.*
2. *The function* $hv'(h)$ *goes to zero as* h *goes to infinity.*
3. *The function* $w(h) \equiv hv'(h) - v(h)$ *is continuous and strictly decreasing, goes to zero as* h *goes to infinity, and goes to infinity as* h *goes to zero.*
4. *The function* $\Phi(x) = 1/w^{-1}(x)$ *is continuous and strictly increasing. It can be extended by continuity at zero with* $\Phi(0) = 0$*. It goes to infinity as* x *goes to infinity.*
5. *The function* $R(x) \equiv v' \circ w^{-1}(h)$ *is increasing, convex, continuous, goes to zero as* x *goes to zero and goes to infinity as* x *goes to infinity.*
6. *Consider any density* $g(A)$ *such that, for all* $x \in \mathbb{R}$*,*

$$G(x) = \int_{A_{\min}}^{A_{\max}} \Phi\left(\max\{A - x, 0\}\right) g(A)\, dA < \infty.$$

Then, the function $G(x)$ *is continuous.*

Proof.

1. For any $h_1 > h_2$, concavity implies that $v'(h_2)(h_1 - h_2) \geq v(h_1) - v(h_2)$. Therefore, $v'(h_2)h_1 \geq v'(h_2)h_2 + v(h_1) - v(h_2) \geq v(h_1) - v(h_2)$. Letting h_2 go to zero in the inequality implies that $v'(h_2)$ goes to infinity as h_2 goes to zero. Second, since $v'(h)$ is positive and decreasing, it has some positive limit $\underline{v}'$ as h goes to infinity. Since $v(h)$ is concave, then for all $h_1 > h_2$, $0 \geq v(h_1) \geq v(h_2) + v'(h_1)(h_1 - h_2) \geq v(h_2) + \underline{v}'(h_1 - h_2)$. Letting h_1 go to infinity shows that $\underline{v}' = 0$. Therefore, $v'(h)$ goes to zero as h goes to infinity.
2. Rearranging the previous inequality implies that

$$v(h_1) + h_2 v'(h_1) - v(h_2) \geq h_1 v'(h_1) \geq 0.$$

Letting h_1 go to infinity shows that $-v(h_2) \geq \limsup_{h \to \infty} hv'(h) \geq 0$ for all h_2. Letting h_2 go to infinity shows that $hv'(h)$ also goes to zero as h goes to infinity.

3. Consider the function $w(h) \equiv hv'(h) - v(h)$. The above results show that $w(h)$ goes to zero as h goes to infinity. Because $w'(h) = hv''(h) < 0$, it follows that $w(h) \geq 0$. Lastly, since $w(h) \geq -v(h)$, letting h go to zero shows that $w(h)$ goes to infinity as h goes to zero.

4. The previous paragraph implies that the function $\Phi(x) = 1/w^{-1}(x)$ is well defined. It is continuous, increasing, goes to zero as x goes to zero, and to infinity as x goes to infinity. Lastly, consider the function $R(x)$ is increasing because both $v'(x)$ and $w^{-1}(x)$ are decreasing. Point 1 and 3 of the Lemma imply that it is goes to zero as x goes to zero, and to infinity as x goes to infinity.

5. In order to prove that $R(x)$ is convex, note that

$$\begin{aligned} R'(x) &= \frac{v'' \circ w^{-1}(x)}{w' \circ w^{-1}(x)} \\ &= \frac{v'' \circ w^{-1}(x)}{w^{-1}(x) \times v'' \circ w^{-1}(x)} \\ &= \frac{1}{w^{-1}(x)}, \end{aligned}$$

   where the second line follows from the fact that $w'(h) = hv''(h)$. Since $w^{-1}(x)$ is decreasing, it follows that $R'(x)$ is increasing, which establishes convexity.

6. Pick any $x \in \mathbb{R}$ and some $\eta > 0$. Then that, for all $y \in [x - \eta, x + \eta]$

$$\begin{aligned} |G(x) - G(y)| &\leq \int_{A_{\min}}^{\alpha} \big|\Phi\left(\max\{A - x, 0\}\right) - \Phi\left(\max\{A - y, 0\}\right)\big| g(A)\, dA \\ &\quad \int_{\alpha}^{A_{\max}} \big|\Phi\left(\max\{A - x, 0\}\right) - \Phi\left(\max\{A - y, 0\}\right)\big| g(A)\, dA \\ &\leq \int_{A_{\min}}^{\alpha} \big|\Phi\left(\max\{A - x, 0\}\right) - \Phi\left(\max\{A - y, 0\}\right)\big| g(A)\, dA \\ &\quad 2\int_{\alpha}^{A_{\max}} \Phi\left(\max\{A - x + \eta, 0\}\right) g(A)\, dA \end{aligned}$$

   where the second inequality follows because $\Phi(x)$ is decreasing. Now, because $G(x-\eta) < \infty$, it follows that for all $\varepsilon > 0$ there exists some $\alpha > 0$ such that the second integral on the right-hand side is less than $\varepsilon/2$. Since the function $\Phi\left(\max\{z, 0\}\right)$ is uniformly continuous over the compact $[0, \alpha - x + \eta]$, there exists some $\eta' < \eta$, such that $|x - y| < \eta'$ implies that $|\Phi\left(\max\{A - x, 0\}\right) - \Phi\left(\max\{A - y, 0\}\right)| < \varepsilon/2$. Plugging this back into the first integral on the right-hand side shows that $|x - y| < \eta'$ implies that $|G(x) - G(y)| < \varepsilon$.

□

A.2 Proofs of the results in the text

A.2.1 Proof of Proposition 2

We proceed in five steps. First, we let $\varepsilon_{it} \in [\underline{e}, \overline{e}]$ be the set of households who find it optimal to locate in an island with current productivity A_{it}. Then, we have

Result 1. *If, for some j, $F(\varepsilon_{jt}) = 0$, then $\varepsilon_{it} = \emptyset$ for all $i < j$.*

To prove the first statement, consider an island A_{it} with a measure $F(\varepsilon_{it}) = 0$ of households. Then the rent must be $R_{it} = 0$. For all $j < i$, we then have that $R_{jt} \geq 0 = R_{it}$ and $A_{jt} < A_{it}$. Thus, island j is strictly less attractive to household than island i and, consequently, is not populated, i.e. $\varepsilon_{jt} = \emptyset$. We then show that:

Result 2. *Consider $i \neq j$. Then $\varepsilon_{it} \cap \varepsilon_{jt}$ is either empty or is a singleton.*

Indeed, household $e \in \varepsilon_{it} \cap \varepsilon_{jt}$ is indifferent between island i and j if and only if $eA_{it} + v(h_{it}) - R_{it}h_{it} = eA_{jt} + v(h_{jt}) - R_{jt}h_{jt}$. Since there is a unique e solving this equation, the result follows. We then have:

Result 3. *If, for some i, $F(\varepsilon_{it}) > 0$, then $F(\varepsilon_{jt}) > 0$ for all $j > i$.*

Note that, because of result 2, all households e in the set ε_{it} live in islands of type i. Thus, islands of type i must be inhabited by a positive measure of households, so housing demand is non zero and $R_{it} > 0$. Now, if an island $j > i$ were populated by a measure zero of households, then $R_{jt} = 0$ and households $e \in \varepsilon_{it}$ would strictly prefer it over island i because of its lower rent and strictly higher wage, which would contradict optimality. Let p be the smallest integer such that $F(\varepsilon_{it}) > 0$ and let e_{it} be the infimum of ε_{it} for all $i \geq p$. This infimum is well defined since, by the previous result then ε_{it} is not empty for all $i \geq p$. For $i = N+1$, we define $e_{N+1t} \equiv \overline{e}$. Next, we show:

Result 4. *Suppose $A_{it} < A_{jt}$. If household $e \in [\underline{e}, \overline{e}]$ weakly prefers island j to i, then all households $e' > e$ strictly prefer j to i.*

This results follows because the household's objective function is super-modular. Since:

$$eA_{jt} + v(h_{jt}) - R_{jt}h_{jt} > eA_{it} + v(h_{it}) - R_{it}h_{it} \geq 0 \Leftrightarrow e(A_{jt} - A_{it}) \geq v(h_{it}) - v(h_{jt}) - R_{it}h_{it} + R_{jt}h_{jt}, \quad (28)$$

and $A_{jt} > A_{it}$, then the inequality is strict for all $e' > e$. Equipped with this result, we obtain:

Result 5. *The lowest ability cutoff is $e_{pt} = \underline{e}$ and, for all $i \geq p$, $\varepsilon_{it} = [e_{it}, e_{i+1t}]$.*

First note that, the sets ε_{it} are increasing. Otherwise, suppose we had $j > i$, $e \in \varepsilon_{it}$, $e' \in \varepsilon_{jt}$ and $e' < e$. Then e' weakly prefers j to i and so, by result 4, e would strictly prefer j to i, a contradiction. It then follows that the sequence e_{it} is weakly increasing. Now all $e \in (e_{it}, e_{i+1t})$ must belong to ε_{it}. Otherwise, if some $e \in (e_{it}, e_{i+1t})$ belonged to ε_{jt} for some $j < i$, then we could find some $e' \in \varepsilon_{it} < e$. By result 4, e would strictly prefer i to j, which is a contradiction. Also, if e belonged to ε_{jt} for $j > i$, then $e_{jt} \geq e_{i+1t} > e$ which contradicts the fact that e_{jt} is the infimum of ε_{jt}. By a similar line of argument, we can show that any $e \notin (e_{it}, e_{i+1t})$ cannot belong to ε_{it}. This shows that $(e_{it}, e_{i+1t}) \subseteq \varepsilon_{it} \subseteq [e_{it}, e_{i+1t}]$. It then follows that $e_{it} < e_{i+1t}$ because otherwise ε_{it} would have measure zero. Since, in an equilibrium, all household must live in some location, we must also have that $\varepsilon_{pt} = \underline{e}$. Lastly, letting $e \to e_{i+1t}$ from the left and from the right in equation (28) for i and $i+1$, we obtain that household e_{i+1t} is indifferent between island i and island $i+1$, so both e_{it} and e_{i+1t} belong to the set ε_{it}.

A.2.2 Proof of Proposition 3

In this proof we suppress time subscripts to simplify notation. We proceed in two steps. First, we show that the system of difference equations (17)and (18) has a unique solution with $e_{1t} = \underline{e}$ and $e_{N+1t} = \overline{e}$. Then, we show that the unique solution of the difference equation is indeed the basis of an equilibrium assignment of households to islands.

Step 1. To solve the system of difference equations, it is useful to consider the following change of variables:

$$e = F^{-1}\left(\min\{1-q, 1\}\right) \equiv \psi(q), \tag{29}$$

where $q_i \in (-\infty, 1]$. The variable q is a "generalized percentile," such that when $q \leq 0$, then $e = \overline{e}$ and when $q = 1$, $e = \underline{e}$. In terms of q_i, the system of difference equation becomes:

$$\begin{aligned} q_i - q_{i+1} &= \mu_i H_i \Phi\left(\max\{\psi(q_i)A_i - U_i, 0\}\right) && (30) \\ U_{i+1} &= \left(\psi(q_{i+1}) - \psi(q_i)\right) A_i + U_i && (31) \end{aligned}$$

One first sees that a pair of sequences $\{U_i\}_{i=1}^N$ and $\{e_i\}_{i=1}^{N+1}$ solves (17)-(18) with $e_1 = \underline{e}$ and $e_{N+1} = \overline{e}$ if and only if the pair of sequences $\{U_i\}_{i=1}^N$, $q_i = \{1 - F(e_i)\}_{i=1}^{N+1}$ solves (30) and (31) with $q_1 = 1$ and $q_{N+1} = 0$.

Now, given an initial condition U_1 and $q_1 = 1$, we show that $U_i - \psi(q_i)A_i$ is strictly increasing in the initial condition U_1, q_i is increasing in the initial condition U_1, and strictly increasing if $q_i < 1$. We proceed by induction. Clearly, the property is true for $i = 1$. Now suppose it is true for all $j \leq i$. We have that

$$q_{i+1} = q_i - \mu_i H_i \Phi\left(\max\{\psi(q_i)A_i - U_i, 0\}\right).$$

Then there are three cases to consider. If $q_i < 1$, then q_{i+1} is strictly increasing in U_1 given that q_i and $U_i - \psi(q_i)A_i$ are strictly increasing in U_1, and $\Phi(\,\cdot\,)$ is an increasing function. The second case is if $q_i = 1$ and $q_{i+1} < 1$, then $\psi(q_i)A_i - U_i > 0$ and so q_{i+1} is strictly increasing in U_1. Lastly, if $q_i = 1$ and $q_{i+1} = 1$, then $\psi(q_i)A_i - U_i \leq 0$, and we obtain that q_{i+1} is increasing in U_1. Now turn to:

$$U_{i+1} - \psi(q_{i+1})A_{i+1} = U_i - \psi(q_i)A_i + (A_i - A_{i+1})\psi(q_{i+1}). \tag{32}$$

The result follows by the induction hypothesis and because $A_i - A_{i+1} < 0$ and $\psi(q)$ is decreasing.

We thus have that q_{N+1} is increasing, and strictly increasing if $q_{N+1} < 1$. Moreover, for all $U_1 > \overline{e}A_N$, $U_i \geq U_1 > \overline{e}A_N \geq e_i A_i$ for all i, so $q_i = 1$ for all i. On the other hand

$$q_{i+1} - q_i < -\min_{j \in \{1,\ldots,N\}}\{\mu_j H_j\} \times \Phi\left(\max\{\psi(q_1)A_1 - U_1, 0\}\right),$$

because, from equation (32), the sequence $\psi(q_i)A_i - U_i$ is increasing in i. Since $\psi(q_1) = \underline{e}$, it follows that, as U_1 goes to minus infinity, $q_{i+1} - q_i$ goes to minus infinity, and so does q_{N+1} Taken together, these properties show that there exists a unique U_1 such that $q_{N+1} = 0$.

Step 2. We now verify that the solution we constructed is indeed the basis of an equilibrium. That is, given $\{U_i\}_{i=1}^N$ and $\{e_i\}_{i=1}^{N+1}$ we let:

$$\begin{aligned} p &= \min\{i \in \{1,\ldots,N\} \,:\, e_i A_i - U_i > 0\} \\ n_i &= H_i \Phi\left(\max\{e_i A_i - U_i, 0\}\right) \\ h_i &= H_i / n_i \\ R_i &= v'(h_i). \end{aligned}$$

The labor market clears by construction of n_i and the housing market clears by construction of h_i. The rent R_i makes it optimal for a household in island i to consume h_i. All we need to verify is that, for all $i \geq p$, households $e \in [e_i, e_{i+1}]$ find it indeed optimal to live in island i. We start by noting that:

$$\begin{aligned} U_i(e) &\equiv eA_i + v(h_i) - R_i h_i = eA_i + v(h_i) - v'(h_i)h_i = eA_i - w(h_i) \\ &= eA_i - \max\{e_iA_i - U_i, 0\} = (e - e_i)A_i + \min\{e_iA_i, U_i\}. \end{aligned}$$

For $\underline{e} = e_p$, we have that, for $i < p$, $U_i(\underline{e}) = 0 + \underline{e}A_i$, since $U_i = U_p \geq \underline{e}A_p > \underline{e}A_i$. For $i = p$, then $U_p(\underline{e}) \geq \underline{e}A_p$ by definition of p. Thus, $\underline{e}$ prefers island p to any island $i < p$, and so does any household $e \geq \underline{e}$ because location choices are monotonic in ability. Thus, islands $i < p$ are not populated. Now, for any $i \geq p$, we have that $U_i \leq e_iA_i$ because of (17) and the fact that $e_{i+1} > e_i$. Thus:

$$\begin{aligned} U_{i+1}(e) &= (e - e_{i+1})A_{i+1} + U_{i+1} = (e - e_{i+1})A_{i+1} + U_i + (e_{i+1} - e_i)A_i \\ &= U_i(e) + (e - e_{i+1})(A_{i+1} - A_i) = U_p(e) + \sum_{j=p}^{i}(e - e_{j+1})(A_{j+1} - A_j), \end{aligned}$$

where the second equality follows by equation (18), the third equality by definition of $U_i(e)$, and the last equality by iterating backward until $j = p$. The terms of the sum are positive if and only if $e \geq e_{j+1}$. It thus follows that a household finds it optimal to locate in the largest j such that $e \geq e_j$. In other words, households in $[e_i, e_{i+1}]$ find it optimal to locate in islands of type i.

A.2.3 Proof of Proposition 4

The result follows from the four steps outlined in Section 2.3.4. Equation (19) uniquely determine, for each time, a construction cutoff c and a construction plan $\{\Delta_{it}\}_{i=1}^N$. Next, given the cutoffs and the construction plans, the difference equation (22) uniquely determine, for each time, the average housing stock per island with current productivity A_{it}, $\{H_{it}\}_{i=1}^N$. Then, the proof of Proposition 2 delivers, at each time, the unique sequence of ability cutoffs $\{e_{it}\}_{i=1}^{N+1}$ and of maximum attainable utilities $\{U_{it}\}_{i=1}^N$. The housing consumption per household, $\{h_{it}\}_{i=1}^N$ is given by equation (16), and the population weights by the market clearing condition $n_{it} = H_{it}/h_{it}$. The rent is given by $R_{it} = v'(h_{it})$ and the price is given by calculating present values.

A.2.4 Proof of Proposition 5

We start from the definition

$$\begin{aligned} U_{it} &= e_{it}A_{it} + \max_{h \geq 0}\{v(h) - R_ih\} \\ &\equiv e_{it}A_{it} - \theta(R_i) \end{aligned} \tag{33}$$

where $\theta(R) \equiv \min_{h \geq 0}\{Rh - v(h)\}$. Note that, for each h, the function $h \mapsto Rh - v(h)$ is positive, increasing, and affine. Being the upper envelope of such a family of functions, $\theta(R)$ is increasing, and concave. We then write:

$$\frac{R_{i+1t} - R_{it}}{A_{i+1t} - A_{it}} = \frac{R_{i+1t} - R_{it}}{\theta(R_{i+1t}) - \theta(R_{it})} \times \frac{\theta(R_{i+1t}) - \theta(R_{it})}{A_{i+1t} - A_{it}}. \tag{34}$$

The first term is increasing in i because, as argued above, the function $\theta(R)$ is increasing and concave. The second terms is also positive and increasing. Indeed, using (33), we have that $\theta(R_{it}) = U_{it} - e_{it}A_{it}$. Therefore:

$$\begin{aligned}
\theta(R_{i+1t}) - \theta(R_{it}) &= e_{i+1t}A_{i+1t} - U_{i+1t} + U_{it} - e_{it}A_{it} \\
&= e_{i+1t}A_{i+1t} - U_{it} - (e_{i+1t} - e_{it})A_{it} + U_{it} - e_{it}A_{it} \\
&= e_{i+1t}(A_{i+1t} - A_{it}),
\end{aligned}$$

where the second line follows from the indifference equation (18). Thus, the second term of (34) is simply equal to e_{i+1t} which is increasing because of assortative matching of ability with productivity.