

NBER WORKING PAPER SERIES

DYNAMIC URBAN ECONOMICS

Brian Greaney
Andrii Parkhomenko
Stijn Van Nieuwerburgh

Working Paper 33512
<http://www.nber.org/papers/w33512>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025

Parkhomenko acknowledges financial support from the USC Lusk Center for Real Estate. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Brian Greaney, Andrii Parkhomenko, and Stijn Van Nieuwerburgh. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Dynamic Urban Economics

Brian Greaney, Andrii Parkhomenko, and Stijn Van Nieuwerburgh

NBER Working Paper No. 33512

February 2025

JEL No. C63, G11, J61, R10, R21, R23, R31, R52

ABSTRACT

We develop a dynamic urban model combining features of quantitative spatial and macro-housing models. It includes multiple locations, forward-looking households, commuting, costly migration, uninsurable income risk, housing tenure choice, and housing frictions. The model operates in continuous time, with shocks and choices occurring at discrete intervals. This "mixed time" approach enables efficient computation of steady-state equilibria and transition dynamics, even with thousands of location pairs. Using a model of the San Francisco Bay Area, we show how forward-looking behavior, spatial frictions, and transition dynamics reshape estimated effects of spatially heterogeneous shocks and policies, traditionally studied with static models.

Brian Greaney
University of Washington
410 Spokane Lane
Seattle, WA 98105
and University of Washington
bg385@uw.edu

Andrii Parkhomenko
701 Exposition Blvd
Marshall School of Business
University of Southern California
Los Angeles, CA 90089
parkhome@marshall.usc.edu

Stijn Van Nieuwerburgh
Columbia University
Graduate School of Business
Kravis Hall, office 796
665 West 130th Street
New York, NY 10027
and NBER
svnieuwe@gsb.columbia.edu

1 Introduction

The effects of shocks and policies are often unevenly felt across time and space. Since moving is costly and investment in local assets such as housing is widespread, individual welfare effects depend on where people live and work. Therefore, understanding the aggregate and distributional consequences of such shocks and policies requires accounting for geography, moving frictions, and real estate ownership.

The state of the art for analyzing spatially heterogeneous shocks is the quantitative spatial model (QSM) developed by [Ahlfeldt, Redding, Sturm and Wolf \(2015\)](#) and extended in numerous subsequent studies. It has rich geography and some heterogeneity across agents, but is static. While it has been successful in capturing relevant aspects of economic geography, the abstraction from dynamics obscures several important issues.

First, static models are poorly suited to calculate individual welfare effects. The distributional effects of a shock are often at least as important as its aggregate impacts, but correctly estimating these requires accounting for transition dynamics. *Second*, the effects of a shock may be different in the short versus long run. Dynamics are also needed to account for the time horizon of policymakers. The first-best policy implied by a static analysis may be infeasible if voters and politicians only support projects that yield rapid benefits. *Third*, comparative statics ignore history dependence. For historical reasons, there may be a large concentration of economic activity in a location that appears undesirable through the lens of a static model. However, since structures are durable and migration is costly, it may be suboptimal to encourage reallocation. *Fourth*, welfare effects depend on the risks and intertemporal choices households face. People desire to live where they can achieve their preferred tenure and saving, and insure against risk. Welfare effects also depend on real estate ownership. Since real estate investment is a forward-looking decision, a dynamic model is required to properly account for it.

Dynamic models with risk are standard in other fields of economics. For example, macroeconomic models with housing have many of the ingredients needed to study the effects of local shocks and policies. However, they do not account for realistic geography.

In this paper, we build a bridge between static QSMs from urban economics and dynamic macro-housing models. Our model contains a large number of locations, each of which can serve as a residence for some households and a work place for others. Forward-looking households choose their locations of residence and work, taking into account commuting times, amenities, wages, and housing costs. Moving frictions imply that location choices are dynamic. As is standard in macro-housing models, households are finitely-lived, face idiosyncratic income risk, can borrow and save in a risk-free asset,

and choose whether to rent or own housing. Owner-occupied housing is illiquid and requires a downpayment, and homeowners can borrow against their housing wealth subject to a collateral constraint. These ingredients allow us to match key features of homeownership and housing wealth data, as well as capture the welfare effects of house price changes. On the production side, developers in each location build residential and commercial floorspace. Residential floorspace is consumed by households, while commercial floorspace together with labor is used by local firms to produce a traded good. Floorspace rents and prices, as well as wages in every location, are determined in general equilibrium.

While this model is well-suited to analyze spatially heterogeneous shocks and policies, solving it presents several major computational challenges. First, allowing for heterogeneity by age, wealth, illiquid housing, productivity, residence, and workplace implies a large state space. In an urban environment with commuting, the state space increases quadratically with the number of locations. For example, a city with 50 neighborhoods has 2,500 location pairs, each with its own age-wealth-housing-productivity distribution. The state space is even larger when we account for transition dynamics, where time becomes a state variable. Second, as is standard in quantitative urban models, we have to estimate residential and workplace amenities as well as productivities for every location in order to match observed populations, employments, wages, and floorspace prices. As a result, even for a city with a modest number of locations, there are hundreds of parameters to estimate. Finally, housing transaction costs imply that agents face a stopping time problem of when to adjust tenure or house size. This induces kinks in the value function, which complicates computation because the first order conditions are not sufficient.

The first contribution of this paper is to show how these obstacles can be overcome using recent tools from macroeconomics and quantitative urban economics. The key to tractability is to cast the model in continuous time, but only allow shocks and discrete choices to occur at discrete, deterministic time intervals. We refer to this assumption, which builds on [Greaney \(2023\)](#), as “mixed time.” At discrete instances, which we call “shock ages,” households draw idiosyncratic productivity, residence, and workplace preference shocks. They then make discrete choices: residence, workplace, tenure, and owner-occupied house size. Between shock ages, location and owner-occupied housing are fixed, and the agent’s problem reduces to a simple consumption-saving choice.

Separating the timing of discrete versus continuous choices allows us to take advantage of efficient discrete- and continuous-time solution methods. As in quantitative urban models, we assume that location preferences are drawn from an extreme-value distribution, which yields closed-form solutions for value functions and location choice

probabilities at shock ages. Crucially, this assumption also allows us to replace work-place with an alternative state variable, wage, that can be discretized on a grid whose cardinality is independent of the number of locations. This dramatically reduces the state space size when there are many locations. Between shock ages, the consumption-saving choice can be solved efficiently using the continuous-time numerical method developed by [Achdou, Han, Lasry, Lions and Moll \(2022\)](#). This method handles non-convexities without difficulty, and, since time is continuous, the first-order conditions are sufficient.

We calibrate our model to the San Francisco Bay Area, which contains 55 locations or 3,025 location pairs. Our algorithm solves a steady-state equilibrium in 20 seconds on a conventional laptop processor. Importantly, even though the number of location pairs increases quadratically with the number of locations, our algorithm’s computation time increases almost linearly. This opens up the possibility of using our method to solve models with a much larger number of locations.

The second contribution of this paper is to show how accounting for forward-looking behavior, homeownership, spatial frictions, and transition dynamics changes the estimated effects of spatially heterogeneous shocks and policies. To this end, we use the quantitative model of the Bay Area to study two types of policy counterfactuals that are common applications for static QSMs. The first experiment consists of an increase in housing supply (“upzoning”) in locations with below–median construction productivity. The second experiment considers a policy that improves the transportation network by introducing four stations of the planned California High-Speed Rail (HSR). We assume that the HSR will be used for commuting and recalculate the commuting time matrix for the entire Bay Area.

The upzoning experiment attracts more residents and jobs to most upzoned areas: upzoned locations are better places for residents, due to more abundant housing, and for firms, due to more abundant workers nearby. However, due to migration frictions and the durability of structures, spatial reallocation is gradual and the full transition takes about 75 years. The HSR experiment has similar effects on population and employment in treated locations, but for a different reason: locations with stations improve access of workers to jobs, as well as access of employers to workers. As in the upzoning experiment, spatial reallocation is gradual and the transition takes decades to unfold. Moreover, the HSR has non-monotonic effects on employment in some suburban locations. In the first few years, greater access to jobs in the rest of the Bay Area leads to a decline in local employment. But later, as more people move in, local employment partly recovers.

While upzoning improves welfare for the average household, there is a rich distribution of welfare effects. The gains for renters are much larger than for homeowners, whose

housing wealth falls as house prices decline. Older, less productive, and less wealthy homeowners experience sizable welfare losses, as housing accounts for more of their lifetime wealth. The HSR improves welfare for most individuals, especially those living near a station (we do not model the costs of constructing the HSR). The gains are larger for the young, who benefit from better transit for a longer period of time, for homeowners, who benefit from real estate appreciation near HSR stations, and for the most productive workers, who gain most from access to better jobs.

We show that sluggish adjustment to shocks, as well as reallocation costs, mean that welfare gains that account for transition dynamics can be much lower than the gains calculated from comparing two steady states. For example, not accounting for transitions overestimates welfare gains from the upzoning policy by over one-half. There is also rich heterogeneity of impacts across locations, so our results cannot be summarized by a model with stylized geography.

To demonstrate the importance of accounting for homeownership, we build a version of our model with renters only who are exposed to real estate shocks via a local real estate investment trust (REIT). To show the importance of accounting for saving and borrowing decisions, we also build a model with hand-to-mouth households. We then re-run the two counterfactual experiments. The distribution of welfare effects is much larger in our main model. In the upzoning counterfactual, 34% of households experience welfare losses in our main model compared to close to 0% in the no-homeownership and the hand-to-mouth models. In the HSR counterfactual, the numbers are 22% compared to 15-16%. In many locations, over 50% of homeowners experience welfare losses from each policy.

The main reason our model produces a richer distribution of welfare gains is that owner-occupied housing is a spatially undiversified asset that many households own with substantial leverage. For many households, this also the largest asset in their portfolios. Approximating homeownership with a REIT that transfers gains and losses on housing back to the locals falls short of fully capturing the effects of homeownership.

Our findings help explain why many spatial policies that are welfare-improving for an average household often face stiff opposition from local residents and do not get implemented. Our results suggest that urban policy evaluation with a model that abstracts from forward-looking behavior, transition dynamics, realistic geography, and especially homeownership would miscalculate effects that are of central interest to policymakers.

Related literature. Our paper is related to the quantitative urban economics literature and the macro-housing literature.¹ The former typically allows for a large number of

¹See [Redding and Rossi-Hansberg \(2017\)](#) for a review of the quantitative spatial literature and [Davis and Van Nieuwerburgh \(2015\)](#) and [Piazzesi and Schneider \(2016\)](#) for reviews of the macro-housing literature.

neighborhoods and commuting decisions, while abstracting from dynamics. Important examples include [Ahlfeldt, Redding, Sturm and Wolf \(2015\)](#) and [Heblich, Redding and Sturm \(2020\)](#). The macro-housing literature has forward-looking agents, housing tenure choice with realistic constraints and frictions, as well as transitional dynamics, but typically abstracts from geography. Main contributions include [Campbell and Cocco \(2007\)](#), [Landvoigt, Piazzesi and Schneider \(2015\)](#), [Favilukis, Ludvigson and Van Nieuwerburgh \(2017\)](#), [Berger, Guerrieri, Lorenzoni and Vavra \(2018\)](#), and [Kaplan, Mitman and Violante \(2020\)](#).

Several recent papers have attempted to bridge the gap between the macro-housing and urban literatures, but have typically done so by imposing restrictive assumptions on individual choices and constraints or geography. [Ortalo-Magné and Prat \(2016\)](#) study a problem where households are exposed to local labor income risk and make a once-and-for-all location choice. Their model solves a rich portfolio choice problem in closed-form, as in [Merton \(1969\)](#), but does not have preferences that admit wealth effects nor allow for recurring consumption and location choices. [Favilukis, Mabile and Van Nieuwerburgh \(2022\)](#) develop a rich macro-housing model to study housing affordability policies, but their model is limited to three locations (city center, city periphery, and an external location). In the quantitative urban tradition, [Takeda and Yamagishi \(2023\)](#) and [Warnes \(2024\)](#) incorporate dynamics in a model of internal city structure with commuting. However, they abstract from savings, housing tenure choice, and floorspace construction.

A major challenge is to allow for forward-looking migration and investment decisions in general equilibrium, since these decisions depend on the choices of all other agents in all future periods. With a large number of locations, this implies an enormous state space.

To avoid this difficulty, most papers with forward-looking migration abstract from consumption-saving decisions. Important examples are [Artuç, Chaudhuri and McLaren \(2010\)](#), [Desmet, Nagy and Rossi-Hansberg \(2018\)](#), [Giannone \(2019\)](#), [Caliendo, Dvorkin and Parro \(2019\)](#), [Eckert and Kleineberg \(2021\)](#), [Zerecero \(2021\)](#), [Martellini \(2022\)](#), [Allen and Donaldson \(2022\)](#), and [Almagro and Domínguez-Iino \(2024\)](#). [Bilal and Rossi-Hansberg \(2021\)](#) allow borrowing and saving, but assume migration is costless, which implies that location is not a state variable.

Other papers allow for both forward-looking migration and saving decisions, but assign these decisions to different types of agents. [Kleinman, Liu and Redding \(2023\)](#), [Cai, Caliendo, Parro and Xiang \(2022\)](#), and [Vanhapelto \(2022\)](#) feature geographically mobile workers who live hand-to-mouth, and capitalists who accumulate wealth but cannot move. These models are tractable because the wealth distribution in each location is degenerate. [Dvorkin \(2023\)](#) extends the approach of [Caliendo, Dvorkin and Parro \(2019\)](#)

to allow for idiosyncratic risk and saving. As in [Merton \(1969\)](#), he assumes that agents face multiplicative risk and no borrowing limit, obtaining closed-form solutions for migration and saving decisions. The key to tractability here is that, although the wealth distribution is non-degenerate, location and portfolio decisions are independent of wealth.

A few recent papers develop models where agents make both forward-looking location choices and intertemporal investment decisions that are subject to constraints. [Crews \(2023\)](#) uses a model in which workers are hand-to-mouth, but invest in human capital subject to a time constraint. Our model allows households to invest in both a liquid asset and illiquid owner-occupied housing, but not human capital. More closely related are [Giannone, Li, Paixao and Pang \(2023\)](#), [Greaney \(2023\)](#), and [Luccioletti \(2023\)](#), who embed lifecycle housing models in a system-of-cities geography. An important limitation is that their models become intractable when the number of locations is large.² This prohibits their application to settings with commuting, in which the state space depends on the number of location *pairs*. Our model not only allows the same agent to make forward-looking location choices and investment decisions subject to constraints, but also can be solved efficiently even when there are thousands of location pairs.

Our paper also contributes to the literature that develops computational methods to solve dynamic spatial models. In particular, our methodological contribution complements that of [Bilal \(2023\)](#). He derives a “Master Equation” that recursively defines equilibrium in many economies where the distribution of state variables is itself a state variable, and analytically derives first- and second-order approximations to it.³ In contrast to our approach, his method can handle aggregate risk.⁴ On the other hand, we highlight two differences between our approaches that make ours more suitable in some contexts. First, our solution method is global, and so can be used to study arbitrarily large shocks. This is important in contexts where individual constraints lead to non-linear effects of shocks. Second, we do not need to assume that policy functions are continuously differentiable with respect to aggregate shocks (see [Bilal \(2023\)](#)’s Assumption 2). This is important because it allows us to incorporate several realistic housing frictions. In our model, construction is irreversible, household borrowing limits depend on house prices, and there are fixed housing adjustment costs.⁵ [Sun \(2024\)](#) develops a deep learning

²[Luccioletti \(2023\)](#) has 12 locations, [Crews \(2023\)](#) has 34, [Giannone, Li, Paixao and Pang \(2023\)](#) have 27, and [Greaney \(2023\)](#) has 50.

³[Bilal and Rossi-Hansberg \(2023\)](#) use this approach to study the effects of climate change in a model with geographically mobile, hand-to-mouth workers and immobile capitalists.

⁴We require that aggregate uncertainty be resolved within a finite time horizon. There *can* be uncertainty about aggregate variables, but the number of future price paths must be finite. See [Steinberg \(2019\)](#) for an example of a shock which induces aggregate uncertainty over a finite horizon.

⁵Irreversible construction is necessary to study situations in which the durability of real estate influences

method to solve a dynamic spatial equilibrium model with a large number of locations. While his model is cast in discrete time, his approach also leverages the advantages of separating choices in time and solving them sequentially.

To our knowledge, we are the first to develop a model that combines forward-looking agents, moving costs, tenure choice, transitional dynamics, and real estate investment into an internal city structure with commuting. We show that these features are indispensable to understand aggregate and distributional effects of shocks and policies that have differential impact across locations and time.

The remainder of the paper is organized as follows. In Section 2, we describe the model. In Section 3, we apply the model to the San Francisco Bay Area. In Section 4, we examine the effects of the transportation and upzoning counterfactuals. Section 5 presents our conclusions.

2 Model

In this section, we present a dynamic quantitative model of an urban area. The key ingredients of the model are (1) many locations, each of which can serve as residence and workplace, (2) commuting between residence and workplace, (3) forward-looking agents, (4) tenure choice with illiquid owner-occupied housing, and (5) endogenous floorspace supply, rents, and prices. The model is set in continuous time.

2.1 Households


The economy is populated by a unit mass of finitely-lived households. There are I locations within a city where households can live and work. A household lives in location i and works in location j .

2.1.1 Timing

Age is indexed by $a \in [0, A]$ and calendar time by t . Households work during ages $a \in [0, A_{\text{ret}})$, are retired during ages $a \in [A_{\text{ret}}, A)$, and die at age A . Both A_{ret} and A are integers. We separate an individual's lifetime into "shock ages," denoted by $\mathcal{A}^s \equiv \{0, 1, \dots, A - 1\}$, and the rest of the lifespan. As described below, households receive a series of shocks

the effects of negative shocks (Glaeser and Gyourko, 2005). It implies that prices are non-differentiable functions of aggregate shocks. The collateral limit and fixed adjustment costs imply that household policies are non-differentiable with respect to aggregate shocks that change house prices.

Figure 1: Household Decision Timeline



at shock ages, as well as make location and owner-occupied house size choices.⁶ The timeline of household decisions is illustrated in Figure 1. The age distribution is uniform. As a result, at any moment in time there is a constant flow of households reaching shock ages. Hence the marginal densities of all state variables, both those that adjust between shock ages and those that only adjust at shock ages, change continuously.

2.1.2 Preferences

Households derive utility from a numeraire consumption good c and housing services $\mathbf{h}$ according to the flow utility function:

$$u(c, \mathbf{h}) = \frac{(c^{1-\eta} \mathbf{h}^\eta)^{1-\gamma}}{1-\gamma}.$$

They also receive utility from liquid wealth b held at death according to a bequest motive of the form used by De Nardi (2004):

$$v(b) = \vartheta_0 \frac{(\vartheta_1 + b)^{1-\gamma}}{1-\gamma}.$$

Households discount the future at rate ρ .

⁶This timing assumption is made primarily for computational reasons, discussed in Section 2.5. In the quantitative model that we describe in Section 3, the interval between shock ages is one year.

2.1.3 Earnings

Working-age households supply one unit of labor inelastically. The log of their individual labor productivity z contains both an idiosyncratic component ζ and a lifecycle component $\bar{z}(a)$:

$$z(\zeta, a) = \zeta + \bar{z}(a).$$

The idiosyncratic component ζ follows the AR(1) process:

$$\zeta_{a+1} = \theta_z \zeta_a + \epsilon_z \text{ with } \epsilon_z \sim N(0, \sigma_z^2),$$

where innovations ϵ_z are drawn by working-age individuals at shock ages.

Earnings are $e^{z(\zeta, a)} w_{jt}$, where w_{jt} is the wage in location j at time t . Retirees receive a pension benefit b . Earnings are subject to a payroll tax of rate τ_z . Payroll tax revenues go toward social security transfers and wasteful government spending. Thus, individual after-tax earnings are characterized by:

$$y_t(j, \zeta, a) = \mathbf{1}(a < A_{\text{ret}})(1 - \tau_z)e^{z(\zeta, a)} w_{jt} + \mathbf{1}(a \geq A_{\text{ret}})b. \quad (2.1)$$

2.1.4 Mobility

Households are only allowed to change their location of residence at shock ages. Immediately after observing their labor productivity ζ , households draw an idiosyncratic preference shock for each residential location ϵ_i^R from a Gumbel distribution with scale parameter ν^R . After observing these shocks, they choose their residential location. When changing residential location from i to i' , they incur moving cost $\mu_{ii'}(a)$ that depends on age. Households in residential location i enjoy a residential amenity E_i^R . Idiosyncratic location preferences, moving costs, and amenities are expressed in utils.

2.1.5 Housing

Housing services can be obtained either by renting or owning. Households can only consume housing in their location of residence. Rental house size h^r is restricted to the set $\mathbb{H}^r$, and owner-occupied house size h is restricted to the set $\mathbb{H}$. The housing services enjoyed by a household with rental housing h^r and owner-occupied housing h are:

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

Housing services per unit vary between rented and owner-occupied housing, and the parameter χ captures the non-pecuniary benefits of ownership. Renting and owning simultaneously is prohibited: owners must have $h^r = 0$ and renters must have $h = 0$.

Housing is non-tradable across space, so its price varies across locations. Rental prices are denoted by r_{it} and owner-occupied house prices by p_{it} . Housing depreciates at rate δ and owner-occupied housing is subject to a proportional property tax of rate τ_h . Rental housing size can be adjusted freely at any time, while owner-occupied housing can only be adjusted at shock ages. In addition, if a household sells its owner-occupied house, it must pay a fraction ψ of the value of the house in transaction costs. A homeowner who changes residential location must first sell any owner-occupied housing, thereby incurring the transaction cost in addition to the moving cost.

2.1.6 Workplace

Households are allowed to change their work location at shock ages. After choosing residential location and owner-occupied housing, working-age households draw idiosyncratic preference shocks for each workplace location ϵ_j^W from a Gumbel distribution with scale parameter ν^W . After observing these shocks, they choose their work location. Workers who commute from i to j incur a commuting cost d_{ij} . Workers who choose work location j enjoy a workplace amenity E_j^W . Workplace location preferences, commuting costs, and amenities are expressed in utils.

2.1.7 Budget Constraint and Household Portfolio

Households have access to three liquid, risk-free assets: a bond, a residential real estate investment trust (REIT), and a commercial REIT. All liquid assets have (ex-ante) return q , so households are indifferent how their total liquid wealth b is divided among them.⁷ We assume that households invest fractions f^S and f^C of their liquid assets in the residential and commercial REIT, respectively, and hold the remainder of their liquid wealth in the bond.⁸ The interest rate q is set externally.

Whenever households purchase a house, they must satisfy a collateral constraint which prohibits borrowing more than a fraction ϕ of housing wealth:

$$b \geq -\phi p_{it} h. \quad (2.2)$$

⁷REITs are described in Section 2.3. While households always believe that the return on liquid assets will be q with probability 1, an unexpected shock can cause REITs to experience unforeseen capital gains/losses.

⁸For households with non-positive liquid wealth, i.e., $b \leq 0$, $f^S = 0$ and $f^C = 0$.

Between home purchases, if liquid wealth is ever less than or equal to $-\phi p_{it}h$, additional borrowing is prohibited:

$$\dot{b} \geq 0 \text{ if } b \leq -\phi p_{it}h. \quad (2.3)$$

The constraint (2.2) ensures that the loan-to-value (LTV) ratio never exceeds ϕ at origination. If house prices in a location fall, we do not require homeowners who remain in their home to delever to meet the origination LTV limit. However, homeowners can only take on *additional* debt if their LTV ratio is below ϕ . Note that (2.2) and (2.3) imply that renters are not allowed to borrow.⁹

Between shock ages, liquid wealth evolves according to the budget constraint

$$\dot{b} = y_t(j, \zeta, a) + qb - c - r_{it}h^r - (\delta + \tau_h)p_{it}h. \quad (2.4)$$

At shock ages, liquid wealth changes discontinuously for homeowners who change residential location and/or households who adjust owner-occupied house size. Immediately after choosing residential location i' , the liquid wealth and owner-occupied housing of a household with initial states (b, h, i) are:

$$\begin{aligned} \tilde{b}_t^R(b, h, i; i') &= b + \mathbf{1}(i' \neq i)(1 - \psi)p_{it}h, \\ \tilde{h}^R(h, i; i') &= \mathbf{1}(i' = i)h. \end{aligned}$$

After choosing residential location, the household chooses owner-occupied housing. Immediately after adjusting owner-occupied house size to h' , the liquid wealth of a household with state (b, h, i) is:

$$\tilde{b}_t^H(b, h, i; h') = b + \mathbf{1}(h' \neq h)[(1 - \psi)p_{it}h - p_{it}h'].$$

Since the collateral constraint must be satisfied after adjusting house size, the set of permissible owner-occupied house sizes for a household with state (b, h, i) is

$$\mathbb{H}_t(b, h, i) \equiv \{h' \in \mathbb{H} : \tilde{b}_t^H(b, h, i; h') \geq -\phi p_{it}h'\}.$$

Finally, liquid wealth can change discontinuously at the time of an unexpected shock due to REIT capital gains or losses. Let $\hat{v}^m$ denote the proportional change in the sector- m

⁹Our framework can accommodate unsecured borrowing without difficulty. We abstract from it in our baseline model for expositional clarity.

REIT value. Liquid wealth after the shock for a household with liquid wealth b is:

$$\tilde{b}^S(b) = b + \sum_{m \in \{S, C\}} (\hat{\theta}^m - 1) f^m \max\{b, 0\}.$$

2.1.8 Household's Value Function

The household's state variables are liquid wealth b , owner-occupied house size h , residential location i , labor productivity ζ , workplace j (for working-age households), and age a . We denote the vector of state variables by $\Omega \equiv (b, h, i, \zeta, j, a)$. Note that since workplace is freely chosen at shock ages, j is not a state variable at shock ages.

Since the bequest motive depends only on liquid wealth, households optimally sell any owner-occupied housing at death. The value function at the maximum age A is:

$$V_t(b, h, i, \zeta, A) = v(b + (1 - \psi)p_{it}h).$$

Between shock ages, the value function satisfies the Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{aligned} \rho V_t(\Omega) = \max_{c, h^r} u(c, \mathbf{h}(h, h^r)) + \partial_b V_t(\Omega) \dot{b}_t(\Omega, c, h^r) + \partial_a V_t(\Omega) + \partial_t V_t(\Omega) \\ \text{s.t. } \dot{b}_t(\Omega, c, h^r) \geq 0 \text{ if } b \leq -\phi p_{it}h, \text{ and } h^r = 0 \text{ if } h > 0. \end{aligned} \quad (2.5)$$

The first term of the right-hand side of (2.5) is flow utility. The second term captures changes in indirect utility caused by changes in liquid wealth. The final two terms reflect changes in the value function due to aging and the passage of time, respectively.

The value function at shock ages can be characterized recursively as follows. The final decision that shock-age households make is work location. Using properties of the Gumbel distribution, the expected value of the optimal workplace choice for a household with state Ω is:

$$V_t^W(\Omega) = \lim_{\iota \downarrow 0} \begin{cases} v^W \log \left(\sum_j \exp([V_{t+\iota}(b, h, i, \zeta, j, a + \iota) + E_j^W - d_{ij}]/v^W) \right) & \text{if } a < A_{\text{ret}}, \\ V_{t+\iota}(b, h, i, \zeta, a + \iota) & \text{otherwise.} \end{cases} \quad (2.6)$$

Prior to selecting work location, households choose owner-occupied housing h . The value of the optimal owner-occupied housing choice for a household with state Ω is:

$$V_t^H(\Omega) = V_t^W(\tilde{b}_t^H(\Omega; \tilde{h}_t(\Omega)), \tilde{h}_t(\Omega), i, \zeta, a), \quad (2.7)$$

where $\tilde{h}_t(\Omega) = \operatorname{argmax}_{h' \in \mathbb{H}_t(\Omega)} V_t^W(\tilde{b}_t^H(\Omega; h'), h', i, \zeta, a)$. Before choosing owner-occupied housing, households draw residential location preferences and choose their residential location. Homeowners who change residential location sell their house at this time. The expected value of the optimal residential location choice for a household with state Ω is:

$$V_t^R(\Omega) = v^R \log \left(\sum_{i'} \exp \left([V_t^H(\tilde{b}_t^R(\Omega; i'), \tilde{h}^R(\Omega; i'), i', \zeta, a) + E_{i'}^R - \mu_{i'}(a)] / v^R \right) \right). \quad (2.8)$$

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

$$V_t(\Omega) = \begin{cases} \int V_t^R(b, h, i, \zeta', a) f(\zeta' | \zeta) d\zeta' & \text{if } a < A_{\text{ret}}, \\ V_t^R(\Omega) & \text{otherwise} \end{cases} \quad (2.9)$$

where $f(\zeta' | \zeta)$ is the conditional probability density function of individual productivity ζ' .

2.1.9 Initial Conditions

Initial idiosyncratic labor productivity is drawn from the invariant distribution. Households are born with zero wealth—we assume that bequests are not redistributed to other households, but leave the economy. Finally, newborn households freely choose their initial residential and workplace locations after drawing idiosyncratic location preferences.

2.1.10 Density of State Variables

Between shock ages, the density of state variables $g_t(\Omega)$ satisfies the Kolmogorov Forward (KF) equation:

$$\partial_t g_t(\Omega) = -\partial_b[\dot{b}_t(\Omega) g_t(\Omega)] - \partial_a g_t(\Omega), \quad (2.10)$$

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

The density of state variables at shock ages can be characterized as follows. The first shock that households experience is to labor productivity (if of working age). Immediately after productivity shocks have occurred, the density of state variables is:

$$g_t^z(\Omega) = \begin{cases} \sum_j \int \lim_{\iota \downarrow 0} g_{t-\iota}(b, h, i, \zeta', j, a - \iota) f(\zeta | \zeta') d\zeta' & \text{if } a < A_{\text{ret}}, \\ \lim_{\iota \downarrow 0} g_{t-\iota}(b, h, i, \zeta, a - \iota) & \text{if } a \geq A_{\text{ret}}. \end{cases} \quad (2.11)$$

The households then choose their residential location. The density of state variables

after residential locations are chosen is:

$$g_t^R(\Omega) = \sum_{i'} \int \int \mathbf{1}_t^R(b', h', i'; b, h, i) \pi_t^R(b', h', i', \zeta, a; i) g_t^z(b', h', i', \zeta, a) db' dh', \quad (2.12)$$

where

$$\mathbf{1}_t^R(b, h, i; b', h', i') \equiv \mathbf{1}(\tilde{b}_t^R(b, h, i, i') = b' \text{ and } \tilde{h}_t^R(h, i, i') = h')$$

is an indicator function that is 1 if a household with initial states (b, h, i) has liquid wealth b' and owner-occupied housing h' after selecting residential location i' and

$$\pi_t^R(\Omega; i') = \frac{\exp([V_t^H(\tilde{b}_t^R(\Omega, i'), \tilde{h}_t^R(\Omega, i'), i', \zeta, a) + E_{i'}^R - \mu_{ii'}(a)]/\nu^R)}{\sum_{i''} \exp([V_t^H(\tilde{b}_t^R(\Omega, i''), \tilde{h}_t^R(\Omega, i''), i'', \zeta, a) + E_{i''}^R - \mu_{ii''}(a)]/\nu^R)} \quad (2.13)$$

is the fraction of households with state Ω who choose residential location i' .

After choosing residential location, households choose owner-occupied housing. The density of state variables after housing adjustments have been made is:

$$g_t^H(\Omega) = \int \int \mathbf{1}_t^H(b', h', i, \zeta, a; b, h) g_t^R(b', h', i, \zeta, a) db' dh', \quad (2.14)$$

where

$$\mathbf{1}_t^H(\Omega; b', h') \equiv \mathbf{1}(\tilde{b}_t^H(b, h, i, \tilde{h}_t(\Omega)) = b' \text{ and } \tilde{h}_t(\Omega) = h')$$

is an indicator function that is 1 if a shock-age household with initial state Ω has liquid wealth b' and owner-occupied housing h' after choosing owner-occupied house size.

The final decision that working-age households make is work location. The density of state variables at shock ages is:

$$g_t(\Omega) = \begin{cases} \pi_t^W(b, h, i, \zeta, a; j) g_t^H(b, h, i, \zeta, a) & \text{if } a < A_{\text{ret}}, \\ g_t^H(\Omega) & \text{otherwise.} \end{cases} \quad (2.15)$$

where

$$\pi_t^W(\Omega; j) = \lim_{i \downarrow 0} \frac{\exp([V_{t+i}(b, h, i, \zeta, j, a) + E_j^W - d_{ij}]/\nu^W)}{\sum_{j'} \exp([V_{t+i}(b, h, i, \zeta, j', a) + E_{j'}^W - d_{ij'}]/\nu^W)} \quad (2.16)$$

is the fraction of households with state Ω who choose work location j .

2.2 Production

There are three sectors of production: the numeraire consumption/investment good, residential floorspace construction, and commercial floorspace construction. Each of these goods is produced in every location by competitive firms.

2.2.1 Traded-Good Firms

The consumption/investment good can be shipped across locations within a city at no cost. Therefore it has a single price, which we normalize to 1. Following [Ahlfeldt et al. \(2015\)](#), we assume that the good is produced using efficiency labor L and commercial floorspace H_C according to the production function

$$Y_{jt} = Z_j L_{jt}^\alpha H_{Cjt}^{1-\alpha},$$

where Z_j is the total factor productivity (TFP) of location j .¹⁰

Firms pay wage w_{jt} per efficiency unit of labor and rent floorspace at rate r_{Cjt} . Factor markets are competitive, so each factor is paid its marginal product:

$$w_{jt} = \alpha Z_j L_{jt}^{\alpha-1} H_{Cjt}^{1-\alpha}, \quad (2.17)$$

$$r_{Cjt} = (1 - \alpha) Z_j L_{jt}^\alpha H_{Cjt}^{-\alpha}. \quad (2.18)$$

Rearranging equation (2.17), labor demand can be written as:

$$L_{jt} = \left(\frac{\alpha Z_j}{w_{jt}} \right)^{\frac{1}{1-\alpha}} H_{Cjt} \quad (2.19)$$

Local labor supply can be computed by integrating over the density of state variables:

$$L_{jt} = \sum_i \int \int \int \int e^{\bar{z}(a)+\zeta} g_t(b, h, i, \zeta, j, a) db dh d\zeta da. \quad (2.20)$$

2.2.2 Developers

There are two types of developers that specialize in building one of the two types of floorspace. Let $m \in \{S, C\}$ index the types of floorspace, with $m = S$ denoting residential

¹⁰Here, Z_j is an exogenous parameter. In Appendix Section B.1, we introduce agglomeration externalities by endogenizing Z_j as a function of local employment density. We show that this has a minor effect on the quantitative results presented in Section 4.

and $m = C$ commercial. Developers construct new floorspace according to the production function

$$Y_{mit}^h = Z_{mit}^h K_{mit},$$

where K is inputs of the numeraire good and Z_{mit}^h is construction productivity, which varies endogenously in a manner that we describe in Section 3.2. When making construction decisions, developers take Z_{mit}^h as given. Construction is irreversible: $Y_{mit}^h \geq 0$.

Residential floorspace demand from renters and owners can be computed by integrating over the density of state variables:

$$H_{Sit} = \sum_j \int \int \int \int [h_t^r(b, h, i, \zeta, j, a) + h] g_t(b, h, i, \zeta, j, a) db dh d\zeta da. \quad (2.21)$$

The population of each location (number of housing units) can be computed as

$$N_{it} = \sum_j \int \int \int \int g_t(b, h, i, \zeta, j, a) db dh d\zeta da. \quad (2.22)$$

Rearranging equation (2.18), commercial floorspace demand can be written as

$$H_{Cjt} = \left[\frac{(1 - \alpha)Z_j}{r_{Cjt}} \right]^{\frac{1}{\alpha}} L_{jt} \quad (2.23)$$

The stock of floorspace in each sector evolves according to the law of motion

$$\dot{H}_{mit} = Y_{mit}^h - \delta H_{mit}, \quad (2.24)$$

so construction demand is $\dot{H}_{mit} + \delta H_{mit}$.

Floorspace construction costs are $1/Z_{mit}^h$. If floorspace price p_{mit} equals construction cost $1/Z_{mit}^h$, developers are indifferent how much to produce. To clear the floorspace market, we assume that they supply the amount of floorspace construction that is demanded. If cost exceeds price, there is no new construction. If price exceeds cost, developers produce an infinite quantity of floorspace, so this cannot be an equilibrium outcome. Formally, construction supply is:

$$Y_{mit}^h = \begin{cases} 0 & \text{if } p_{mit} < 1/Z_{mit}^h, \\ \max\{\dot{H}_{mit} + \delta H_{mit}, 0\} & \text{if } p_{mit} = 1/Z_{mit}^h, \\ \infty & \text{if } p_{mit} > 1/Z_{mit}^h. \end{cases} \quad (2.25)$$

Floorspace market clearing requires that supply equals demand: $Y_{mit}^h = \dot{H}_{mit} + \delta H_{mit}$. The floorspace market clearing can be written as the complementary slackness condition:

$$\begin{aligned} p_{mit} &= 1/Z_{mit}^h \text{ and } Y_{mit}^h > 0, \\ \text{or } p_{mit} &< 1/Z_{mit}^h \text{ and } Y_{mit}^h = 0. \end{aligned} \quad (2.26)$$

That is, if demand is sufficiently strong that there is positive construction, price equals construction cost. If demand is sufficiently weak that no construction is demanded, then prices are no longer determined by construction costs. Instead, market clearing requires that prices be set so that demand falls at the depreciation rate δ , which is the fastest local floorspace quantities can decrease since construction is irreversible. Note that in a steady state, there is always strictly positive construction which exactly offsets depreciation: $Y_{mit}^h = \delta H_{mit} > 0$. Hence, floorspace prices equal construction costs in steady state.

2.3 REITs

Commercial and rented residential units are owned by perfectly competitive real estate investment trusts (REIT). REIT portfolios are perfectly diversified across all locations in the city. REITs purchase floorspace by borrowing from an external credit market at the interest rate q . They rent out space to households and firms at a rent r_{mit} . Since REITs are competitive, rents are set so that the return on real estate investments equal the risk-free rate adjusted for depreciation and property taxes:

$$r_{mit} = (\delta + q + \tau_h)p_{mit} - \dot{p}_{mit}. \quad (2.27)$$

As a result of a shock, REIT portfolios can appreciate or depreciate. This will affect household wealth, as a fraction of household assets are invested in REITs.

2.4 Equilibrium

An *equilibrium* is an allocation, household value function $V_t(\Omega)$, density of state variables $g_t(\Omega)$, wages w_{jt} , floorspace rents r_{mit} , and prices p_{mit} , such that:

1. Households optimize: equations (2.5)–(2.9) are satisfied.
2. The density of state variables is consistent with household optimization: equations (2.10)–(2.16) are satisfied.
3. Firms optimize: equations (2.17)–(2.18) hold.
4. Labor markets clear: the values of equations (2.19) and (2.20) are equal in all locations.

5. Residential and commercial ownership space markets clear: equation (2.26) holds for $m = S, C$ in all locations.
6. Residential and commercial rental space markets clear: equation (2.27) holds for $m = S, C$ in all locations.

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

2.5 Solving the Model

Solving a dynamic quantitative urban model with age, idiosyncratic risk, tenure choice, and illiquid housing presents several computational challenges. First, the state space is large. The main reason is that commuting implies that the size of the state space increases *quadratically* with the number of locations. For example, with the location concept we use in our quantitative analysis, the New York commuting zone has 183 locations, but 33,489 location *pairs*. Each location pair has a separate wealth-housing-productivity-age distribution. The state space is even larger when solving transition dynamics, because time is an additional state variable. Second, as is standard in quantitative urban economics, we choose local amenities and productivities to exactly match population and employment shares, wages, and floorspace prices in every location. As a result, there are many parameters to estimate.¹¹ Finally, fixed housing adjustment costs imply that households face a stopping time problem of when to adjust housing. This induces kinks in the value function where agents adjust housing, which precludes efficient discrete-time algorithms such as the endogenous gridpoint method of [Carroll \(2006\)](#).

A primary contribution of this paper is to develop a method that overcomes these challenges by combining tools from macro and urban economics. The key to tractability is to set the model in continuous time, but only allow idiosyncratic shocks and discrete choices at discrete, deterministic time intervals (shock ages). This timing assumption builds on [Greaney \(2023\)](#).¹² It allows us to make use of both efficient discrete- and continuous-time numerical methods. As is common in quantitative urban models, the fact that location preferences are drawn from an extreme-value distribution allows us to obtain closed-form expressions for the value function and location choice probabilities at shock ages (equations 2.6, 2.8, 2.13, and 2.16). Between shock ages, location and owner-occupied

¹¹Our model delivers closed-form expressions for productivities as functions of populations, labor supplies, and prices. We can therefore read productivities directly off data from a factual equilibrium. In contrast to static QSMs, we do not have closed-form expressions for population and employment shares. As a result, we have to numerically estimate 2I residential and workplace amenities.

¹²[Greaney \(2023\)](#)'s model does not have commuting and has only 50 locations.

housing are fixed, and the agent’s problem reduces to a simple consumption-saving choice (equations 2.5 and 2.10). This can be solved extremely efficiently using continuous-time numerical methods developed by [Achdou, Han, Lasry, Lions and Moll \(2022\)](#).

There are four main advantages of continuous time that are relevant in our setting.¹³ First, the first-order conditions hold with equality and are sufficient at *every* point in the state space. In a discrete-time version of our model, the first order conditions would be inequalities due to occasionally binding borrowing constraints, and they would not be sufficient due to nonconvexities resulting from fixed housing adjustment costs. Second, the first-order conditions include only contemporaneous variables. In discrete time, first-order conditions relate variables in one period to variables in the next period. As a result, optimal policies are defined only implicitly, and must be computed by solving a root-finding problem that typically involves interpolation and computing expectations at every step. This is avoided in continuous time, where the first-order conditions yield explicit expressions for policy functions. Third, since wealth does not jump discontinuously, the HJB and KF equations that define the value function and density of state variables can be represented as sparse matrix equations.¹⁴ This is advantageous because highly efficient routines for solving sparse matrix equations are widely available. Fourth, the matrix that represents the discretized KF equation is the transpose of the matrix that represents the discretized HJB equation. As a result, once the value function has been solved for, the density of state variables can be obtained at virtually no cost. For these reasons, general equilibrium of heterogeneous-agent models can typically be solved much faster in continuous time than in discrete time. The speed gains are greatest in models with nonconvexities, such as ours.

The main advantage of limiting idiosyncratic shocks and discrete choices to shock ages is that it allows us to replace workplace with a lower-dimensional state variable. Since workplace is freely chosen at shock ages, which are also when workplace amenities are enjoyed and commuting costs are paid, the only workplace characteristic that is relevant between shock ages is wage. As a result, we can replace the discrete state variable workplace with the continuous state variable wage in the HJB equation (2.5). The value of workplace j in equation (2.6) is the value of earning j ’s wage plus its workplace amenity minus its commuting cost. Wage can be discretized on a grid whose cardinality is independent of the number of locations, which keeps the state space size manageable even when there are many location pairs. Suppose instead that the opportunity to change workplace arrived according to a Poisson process.¹⁵ In that case, households would be uncertain how

¹³See [Achdou, Han, Lasry, Lions and Moll \(2022\)](#) for a more detailed discussion.

¹⁴Since we limit idiosyncratic shocks to shock ages, in our model these matrix equations are tridiagonal.

¹⁵This assumption is made for location choices by [Crews \(2023\)](#), [Bilal \(2023\)](#), and [Bilal and Rossi-Hansberg](#)

long they would have to remain in their current workplace. Since households care about the entire future path of earnings, current wage would not be a sufficient state variable for workplace.¹⁶ As a result, workplace would be a state variable for working-age agents, and the state space size would increase quadratically in the number of locations instead of linearly.

Another attractive feature of the shock age assumption is that it is realistic. It is also common: discrete-time models implicitly assume that shocks and choices are made with discrete frequency. While in reality shocks and moves can happen at any time, migration is infrequent and labor and mortgage contracts often last for a year or more. In continuous time, it is impractical to allow migration at any time because extreme-value location preference shocks would induce continuous migration. If moving opportunities instead arrived according to a Poisson process, some households with a strong desire to move would not receive the opportunity to do so. With our assumption, households regularly receive the opportunity to move with certainty.

Computational time. We demonstrate the performance of our algorithm by solving stationary equilibrium in commuting zones (CZs) with varying numbers of locations: Portland (18 locations and 324 pairs), Seattle (33 locations and 1,089 pairs), San Francisco (55 locations and 3,025 pairs), Los Angeles (123 locations and 15,129 pairs), and New York (183 locations and 33,489 pairs).¹⁷ Figure 2 shows the relationship between the number of locations and time to solve equilibrium.¹⁸ Solving the model for a small CZ such as Portland takes just 7.6 seconds, while New York takes 2.91 minutes. Solving it for the San Francisco Bay Area, the CZ that we use for our quantitative analysis in Sections 3 and 4, takes 20 seconds. Importantly, even though the size of the state space grows quadratically with the number of locations, computational time grows less than quadratically. This is due to the fact that wage is a sufficient state variable for workplace between shock ages. As a result, it is possible to use our model to study cities with tens of thousands of location pairs. Appendix C spells out the details of the computational algorithm.


(2023) in their continuous-time spatial models.

¹⁶In our model, households care about the path of wages until the next shock age. We approximate the transition path of wages with a step function in which wages are constant between shock ages, so current wage is a sufficient state variable for workplace. The accuracy of this approximation is increasing in the frequency of shock ages.

¹⁷That is, we compute the value function and density of state variables given prices and parameters. Locations and model parameters are described in Section 3.

¹⁸For this demonstration, we used a MacBook Pro laptop computer with a 1.4 GHz Intel core i5 processor.

Figure 2: Number of locations and computational time



Note: The solid line shows the relationship between computational time required to solve the stationary equilibrium and the number of locations in an urban area. The lower dotted line shows extrapolated (using the two smallest CZs) computational time if it were a linear function of the number of locations. The upper dotted line shows extrapolated time if it were a quadratic function of the number of locations.

3 Application

In this section, we show how our model can be applied to study policy counterfactuals in an urban area with many locations. We focus on the San Francisco Bay Area, which has 55 locations.¹⁹

3.1 Data

Model locations. Following [Delventhal and Parkhomenko \(2024\)](#), we define a model location as the intersection of a Census Public Use Microdata Area (PUMA) and a county.²⁰ The San Francisco Bay Area, as defined by the Association of Bay Area Governments, contains 55 model locations that cover nine counties in the Bay Area: Alameda, Contra Costa, Marin, Napa, San Francisco, San Mateo, Santa Clara, Solano, and Sonoma.²¹ This

¹⁹We used the same data sources to demonstrate the computational efficiency of our method for other cities with a smaller or larger number of locations in Section 2.5.

²⁰PUMA is the smallest geography for which individual-level data is publicly available. The Census Bureau designs PUMAs to have between 100,000 and 200,000 residents. In densely populated areas, where there are many PUMAs to a county, each PUMA is a model location. This allows us to take advantage of geographically-detailed data and study patterns within metro areas. In rural areas, where there may be several counties in a single PUMA, each county is a model location.

²¹For more details, see the website of the Association of Bay Area Governments: <https://abag.ca.gov/>. We chose not to use the San Francisco Commuting Zone (CZ), as defined in [Tolbert and Sizer \(1996\)](#), because it is separate from the San Jose CZ. At the same time, merging San Francisco and San Jose CZs would require us to include several counties to the south of the area which appear to be far from the rest of the Bay Area. We also chose not to use the San Francisco Combined Statistical Area, as it includes several counties to the east of the area which are also quite far from the Bay Area. Instead, our definition of the Bay Area is a contiguously developed area around San Francisco, San Pablo, and Suisun Bays. It has a combined population of 7.8 million people according to the 2020 Census.

number of locations implies 3,025 residence-workplace location pairs. Appendix Table [A.1](#) lists model locations and some of their characteristics.

Characteristics of residents. To obtain citywide characteristics of residents we use the 2012–2016 five-year sample of the American Community Survey (ACS) microdata. We restrict our attention to heads of household between 20 and 84 years old who do not live in group quarters. To compute data moments that describe labor market outcomes, we focus on 20–64 year olds, who report income and hours, work at least 35 hours a week and at least 27 weeks a year, are not self-employed, and earn at least half of the federal minimum wage. To compute data moments that describe housing market outcomes, we focus on 20–84 year olds and drop observations that report living in mobile homes, trailers, boats, tents, and farmhouses.

Residents, jobs, and commuting. To obtain information on resident population, job counts, and commuting flows, we turn to the LEHD Origin-Destination Employment Statistics (LODES) database, taking averages across 2012–2016. LODES provides workplace and residence job counts separately by education level or by industry at the Census block level, which we aggregate to the level of model locations.

Wages. We use the Census Transportation Planning Products (CTPP) database and the ACS data for 2012–2016 to obtain estimates of average wage for each location. We use the data reported for the period from 2012 to 2016, and estimate wage indices for each location after controlling for the effects of age, gender, race, industry, occupation, and education. Appendix [A](#) provides more details.

Floorspace rents and prices. To obtain local prices of residential floorspace, we aggregate the zip-code level Zillow Home Value Index for years 2012 to 2016 to the level of model locations. Since in the steady state price-rent ratios do not differ across locations, we obtain local rents by dividing prices by the model-implied price-rent ratio for the Bay Area: 26.4. See Section [3.2](#) below for more information.

To obtain commercial real estate rents, we build hedonic indices for each model location using the transaction-level data of leases on office, retail, and industrial properties for the period from 2012 to 2016 from the data provider Compstak. The correlation between commercial and residential rents across model locations is 90%. Appendix [A](#) contains more details.

Commute times. We use the Census Transportation and Planning Products (CTPP) data to estimate the commute times between each pair of model locations. The CTPP

database reports commuting time data for origin-destination pairs of Census tracts, and is tabulated using ACS data. Appendix Section A contains more details on how we process this data.

3.2 Model Parameters

Next we describe how we obtain the values of parameters. When discussing internally calibrated parameters, we mention which data moment is central in determining that parameter, but we calibrate all parameters jointly. Parameter values are shown in Table 1.

Life cycle. Workers are born at age 20, retire at age 65, and die at age 85. The interval between shock ages $a \in \mathcal{A}^s$ is one year. In Appendix Section B.2, we consider a version of the model where the interval between shock ages is half a year. This does not change the results of the counterfactual experiments that we describe later in Section 4.

Preferences. We calibrate the weight of housing in utility η to match the median rent-to-earnings ratio of 0.24, calculated from the 2012–2016 ACS data for the San Francisco Bay Area. The calibrated η is 0.3281. We calibrate the preference for homeownership χ to match the homeownership rate. In the Bay Area, the homeownership rate is 54.2% and the calibrated value of χ is 1.0506.

The Gumbel scale parameters ν^R and ν^W of the distributions of residence and workplace preference shocks, respectively, determine the relative importance of location fundamentals versus idiosyncratic preferences when households make location choices. The parameter ν^R governs how much sorting by income there is across locations. If idiosyncratic preferences are strong, then there should be a lot of income mixing within locations and average incomes across locations will not differ much. On the other hand, if the preferences are weak, then there should be strong segregation by income and average incomes across locations will differ substantially. Therefore, we first compute log average hourly earnings of residents in each location.²² Then we calibrate ν^R to match the variance across locations of log average earnings (0.01043) and obtain $\nu^R = 2.0833$.

The parameter ν^W determines how far workers are willing to commute. If idiosyncratic workplace preferences are strong, then workers should be relatively insensitive to commute times when choosing their workplace. Otherwise, workplace choices should be sensitive to commute times. Thus, we calibrate ν^W to match the 90th percentile of commuting times for commutes 90 minutes or less and obtain $\nu^W = 0.2241$.²³

²²We remove the effects of age, gender, industry, and years of schooling from log earnings prior to calculating the average in a location.

²³Out of 139 mln commuters we observe in the nationwide 2012–2016 LODS data, 9.8 mln travel between

Table 1: Model parameters

Parameter	Description	Value	Target or source	Value
Internally calibrated				
ρ	discount factor	0.0215	median wealth-earnings ratio	1.647
η	weight of housing in utility	0.3281	median rent-to-earnings	0.240
χ	preference for homeown.	1.0506	homeownership rate	0.542
ϑ_0	bequest motive	0.5113	homeown., ages 80–84 vs 20–24	0.642
ν^R	Gumbel scale, resid. shocks	2.0833	100× variance of log earnings	1.043
ν^W	Gumbel scale, work. shocks	0.2241	p90 commute time, minutes	67.0
κ	cost of commuting	0.0092	commuting gravity coeff.	-0.0408
μ^0	moving cost, intercept	12.5598	cross-county migration rate	0.0224
μ^a	moving cost, age coeff.	0.1524	migration, ages 20–24 vs 80–84	0.0853
Externally calibrated				
q	interest rate	0.02	data	
γ	risk aversion	2	standard value	
ϑ_1	bequest motive curvature	0.01	see the text	
θ_z	labor prod. persistence term	0.9136	Floden and Lindé (2001)	
σ_z^2	labor prod. variance term	0.0426	Floden and Lindé (2001)	
α	labor share in production	0.82	Valentinyi and Herrendorf (2008)	
δ	housing depreciation rate	0.0110	BEA and Davis et al. (2021)	
ψ	housing transaction fee	0.06	standard value	
τ_h	property tax	0.0071	Brookings data	
τ_z	payroll tax	0.179	OECD	
ϕ	collateral constraint	0.8	standard value	
$\min \mathbb{H}$	smallest own. size	4.0781	median wage / p5 own. value	
$\max \mathbb{H}$	largest own. size	10.7834	p90 own. / p10 own. size	
$\max \mathbb{H}^r$	largest rental size	5.3917	p90 rental size / p10 own. size	
f^S	share of resid. REIT	0.0845	SCF data	
f^C	share of comm. REIT	0.0574	SCF data	

Note: The table describes model parameters. See the text for more details.

Migration. We parameterize the moving cost as a linear function of age:

$$\mu_{i'i''}(a) = \mathbf{1}(i' \neq i)(\mu^0 + \mu^a a). \quad (3.1)$$

locations that are over 3 hours apart. Due to reasons outlined in [Graham et al. \(2014\)](#), many of these long commutes arise due to errors in assigning work or residence locations. We therefore truncate observations with commute times greater than 90 minutes from our average commute time calculations.

In the ACS data, we can identify migration across the nine counties within the Bay Area, but not across PUMAs. Thus, we calibrate the moving cost intercept μ^0 to match the annual migration rate between locations that belong to different counties within the Bay Area. In 2012–2016, this migration rate was 2.24 percent.²⁴ We calibrate the age coefficient μ^a to match the difference in migration rates between 20–24 year olds and 80–84 year olds, which is equal to 8.53 percentage points. We obtain $\mu^0 = 12.5598$ and $\mu^a = 0.1524$.

Commuting. We parameterize the transportation cost as a linear function of travel time,

$$d_{ij} = \kappa \times \text{time}_{ij}, \quad (3.2)$$

where time_{ij} is the time in minutes required to travel from location i to location j . Parameter κ measures the sensitivity of individual utility to the time spent commuting. As is standard in the literature, we first estimate the “gravity regression,”

$$\ln N_{ij} = \varsigma \text{time}_{ij} + \varphi_i^R + \varphi_j^W + \varepsilon_{ij}, \quad (3.3)$$

where N_{ij} is the fraction of workers who live in location i and work in j , and φ_i^R and φ_j^W are residence and workplace fixed effects. Then, we calibrate κ such that the estimated coefficient ς is the same in the model as in the data, and obtain $\kappa = 0.0092$.

Income. The lifecycle component of labor productivity $\bar{z}(a)$ is taken from Hansen (1993). Without loss of generality, we scale $\bar{z}(a)$ so that median earning is 1. The parameters of the stochastic process for idiosyncratic labor productivity ζ , the persistence term θ_z and the variance term σ_z , are taken from Floden and Lindé (2001). We discretize ζ using Rouwenhorst’s method (Rouwenhorst 1995, Kopecky and Suen 2010) with 5 gridpoints.

Bequests. We calibrate the strength of the bequest motive ϑ_0 to match the percentage-point difference in homeownership rates between 80–84 year-olds and 20–24 year-olds. In San Francisco, this difference is 64.2 percentage points and our calibrated ϑ_0 is equal to 0.5113. We set the curvature parameter ϑ_1 to 0.01, a small positive number that ensures that the bequest function is well-defined for zero bequests and is approximately log-linear in the size of the bequest.

Housing. We calculate the housing depreciation rate δ as follows. From the BEA Depreciation Estimates for 2016, we divide the value of current-cost depreciation of residential fixed assets owned by households by the value of the current-cost net stock of

²⁴Migration across county borders is relatively uncommon. The migration rate for any residential move within the Bay Area is higher at 9.7 percent.

residential fixed assets, and obtain the annual depreciation rate of 0.0239. This estimate excludes the value of land embedded in house values. Since land does not depreciate, we multiply the structures depreciation rate by $1 - 0.5393$, where 0.5393 is the average share of land in house values in the Bay Area, from [Davis, Larson, Oliner and Shui \(2021\)](#)'s tract-level estimates. This yields $\delta = 0.011$.

Next, we describe how we build the sets of possible sizes of owner-occupied and rental units, $\mathbb{H}$ and $\mathbb{H}'$. We let $\mathbb{H}$ be discrete and $\mathbb{H}'$ continuous.²⁵ For owner-occupied houses, we build an equally-spaced grid with six support points. We calibrate the smallest and the largest value in the grid as follows. The smallest value is set to the ratio of the value of the house in the 5th percentile of the value distribution in the Bay Area to median labor earnings (equal to 1 in our model). We use ACS to build the distribution of house values and control for year of construction and type of structure. To obtain the largest value, we multiply the smallest value by the ratio of the 90th to the 10th percentiles of owner-occupied house sizes in the San Francisco metropolitan area from the 2015 American Housing Survey (AHS) data.. Our calibrated smallest and largest owner-occupied house sizes are 4.0781 and 10.7834.²⁶

The lower bound of the size of a rental unit is zero. We calibrate the maximum size to match the ratio of the 90th percentile of rental units to the 10th percentile of owner-occupied units in the 2015 American Housing Survey (AHS) data, equal to 1.32. Our calibrated largest rental unit size is 5.3917.

REIT ownership. The shares of REITs in household liquid assets, f^S and f^C , are calculated as follows. From the 2016 Survey of Consumer Finances (SCF) data, we estimate that the average household owns \$50,280 of residential real estate other than primary residence (\$49,432 directly and \$848 via REITs), as well as \$34,186 of commercial real estate (\$29,462 directly and \$4,724 via REITs). An average household owns \$595,348 of total assets excluding primary residence.²⁷ This implies that an average household has $f^S = 8.45\% = \$50,280/\$595,348$ of her assets in residential real estate other than primary residence and $f^C = 5.74\% = \$34,186/\$595,348$ in commercial real estate.

Taxes. The property tax rate τ_h is calculated as follows. First, we use the county-level data on taxes paid as a share of home value from Brookings for the nine counties in the

²⁵For computational reasons, it is convenient when the distribution of owner-occupied dwelling sizes is discrete because the size is a state variable, and when the distribution of rental dwelling sizes is continuous because a renter's housing consumption can be solved analytically from first-order conditions.

²⁶Without loss of generality, we normalize the population-weighted average house price to 1 (which corresponds to median earning).

²⁷The numbers for the *median* household are much smaller.

Bay Area.²⁸ Then, we use county population levels from the 2010 Census and find the weighted-average property tax rate of 0.71%. We use the payroll tax rate reported by the OECD for the U.S. The average tax rate for the period 2012–2016 was 17.9%.²⁹

Interest rate and discount factor. We model the city as a small open economy and set the interest rate at $q = 0.02$ which corresponds to the average 10-year real interest rate from 1962 to 2024. Using the calibrated values of δ and τ_h and equation (2.27), this interest rate results in the price-rent ratio of 26.4.³⁰ We calibrate the discount factor ρ to match the median wealth-income ratio of 1.647 from the 2016 SCF, and obtain the value of 0.0215.

Production. Valentinyi and Herrendorf (2008) estimate the share of land and structures in production to be 0.18. Since the numeraire production technology is constant-returns-to-scale in floorspace and labor, we set $\alpha = 0.82$.

Local amenities and traded-good productivity. Local residential and employment amenities, E_i^R and E_j^W , are calibrated to match local residential population and employment. Local labor productivity in the traded-good sector Z_j is calibrated to match average wages for workers employed in location j . The values of location-specific parameters are listed in Appendix Table A.1.

Construction. The construction productivity functions Z_{mit}^h are chosen to match empirical estimates of floorspace supply elasticities. Let $r_{mit}(p_{mit})$ denote the floorspace rent when floorspace price is p_{mit} , determined by equation (2.27). In the residential segment, we assume that Z_{Sit}^h is such that rent when construction is strictly positive (so that price equals construction cost) is a log-linear function of number of housing units:

$$r_{Sit}(1/Z_{Sit}^h) = \bar{r}_{Si} \left(\frac{N_{it}}{N_{i0}} \right)^{1/\xi_i^S}. \quad (3.4)$$

In the previous expression, ξ_i^S is the housing unit supply elasticity in the residential market, $\bar{r}_{Si}$ is the exogenous component of residential construction productivity, and N_{i0} is population in the initial steady state. Recall that construction irreversibility implies that construction may not be strictly positive for part of the transition path after an unexpected shock, if demand is sufficiently weak. In this case, $p_{Sit} < 1/Z_{Sit}^h$, and the equilibrium rent r_{Sit}

²⁸Brookings property tax map: <https://www.brookings.edu/articles/map-property-taxes-in-your-county/> (accessed on September 11, 2023).

²⁹See <https://stats.oecd.org/index.aspx> (variable “Average income tax rate” located in “Public Sector, Taxation and Market Regulation”/“Taxation”/“Taxing Wages”; accessed on September 14, 2023).

³⁰Using individual observations from the 2012–2016 ACS data, we estimated hedonic price and rent indices, and calculated the price-rent ratio of 26.7, close to the value in the model.

will not equal $r_{Sit}(1/Z_{Sit}^h)$. However, since construction is always strictly positive in steady state, equation (3.4) ensures that residential rents are log-linear functions of populations in the long run. This is also true of prices, since the long-run price/rent ratio is $1/(\delta + q + \tau_h)$ (see equation 2.27).

We calculate ξ_i^S by combining 10-year tract-level unit supply elasticities from [Baum-Snow and Han \(2023\)](#) and 30-year MSA-level elasticities from [Saiz \(2010\)](#). The productivity shifter $\bar{r}_{Si}$ is chosen to match observed rents r_{Si}^{data} in the initial steady state: $\bar{r}_{Si} = r_{Si}^{\text{data}}$.

In the commercial segment, we set Z_{Cit}^h so that, when construction is strictly positive, there is a log-linear relationship between commercial rents and commercial floorspace quantity H_{Cit} :

$$r_{Cit}(1/Z_{Cit}^h) = \bar{r}_{Ci} \left(\frac{H_{Cit}}{H_{Ci0}} \right)^{1/\xi_i^C}. \quad (3.5)$$

We infer $\bar{r}_{Ci}$ in the same way as for the residential sector but using the Compstak commercial rent data. Due to the absence of local estimates of commercial floorspace supply elasticities, we calculate ξ_i^C by combining 10-year tract-level residential floorspace supply elasticities from [Baum-Snow and Han \(2023\)](#) and 30-year MSA-level elasticities from [Saiz \(2010\)](#).³¹ Appendix A provides more details and Appendix Table A.1 list the values of elasticities and construction productivities.

3.3 Non-targeted Moments


This section describes the baseline model's implications for commuting, income, wealth, homeownership, and migration.

Commuting. While we observe the full commuting flow matrix for San Francisco, our calibration only targets the elasticity of flows with respect to commute times and the number of residents and workers in each location. Panel A of Figure 3 shows that our model matches commuting flows quite well. The correlation between log shares in the model and in the data is 0.89. Panel B shows that our model also produces a distribution of commuting times that is similar to the one in the data.

Income and wealth. In our calibration, we match average wages by workplace and the variance of wages across residential locations. Panel C of Figure 3 shows that our model also matches the distribution of average incomes across residential locations well, with a correlation of 0.79, albeit with a smaller variance than in the data. Our model also

³¹Many of the same regulatory and administrative processes at the local level affect both residential and commercial development.

Figure 3: Non-targeted moments



Note: The maps figure shows several non-targeted moments, in the model and in the data. See the text for more details.


produces realistic patterns for wealth inequality. As panel D shows, the distribution of wealth by age follows a hump-shaped profile and wealth levels for homeowners are much higher than for renters.³² Furthermore, Figure 4 shows that in our model, as in the data, both financial and housing wealth are the highest in San Francisco, Silicon Valley, and San Rafael.

Homeownership. Our calibration targets the overall homeownership rate and the percentage-point difference in homeownership rates between 80–84 year-olds and 20–24 year olds. Panel E of Figure 3 shows that our model is successful in generating a realistic life-cycle evolution of homeownership rates.

Migration. We target the overall migration rate across counties within the Bay Area and the difference in migration rates between ages 20–24 and 80–84. As shown in panel F of Figure 3, our model matches the entire age profile of migration quite well. Importantly, our model generates the flattening of the migration profile with age.

³²Figure 9 in Davis and Van Nieuwerburgh (2015) shows that these relationships are similar in the data.

Figure 4: Spatial distribution of wealth



Note: The maps show the average financial and housing wealth relative to median annual wage in each location. Since renters hold zero housing wealth, housing wealth is calculated for homeowners only.

4 Policy Experiments

In this section, we use our model to investigate the long-run and transitional effects of two illustrative policy experiments that are commonly studied using static QSMs. We then discuss why dynamics and rich spatial heterogeneity are necessary to obtain many of the important results of these experiments.

4.1 Effects of Housing and Transportation Policies


We use the model to run two counterfactual experiments: (1) an increase in housing supply in the most constrained locations and (2) an improvement the transportation network. In both experiments, we assume that the population of the Bay Area is fixed, i.e., it is a closed city. These experiments are not meant to serve as evaluations of specific policies, but rather as illustrations of what our dynamic model with homeownership can offer compared to traditional static QSMs.

4.1.1 Counterfactual Experiments

Upzoning. In our first counterfactual scenario, housing supply increases. Since housing supply is endogenous, the increase in housing supply is engineered through an increase in the productivity of the residential construction sector which can represent, for example, a relaxation of zoning constraints, often referred to as upzoning.³³ In particular,

³³Zoning constraints or other administrative barriers to housing supply are commonplace in U.S. cities (Gyourko and Molloy, 2015).

Figure 5: Counterfactual changes in housing productivity and travel times



Note: Panel A shows the increase in housing construction productivity Z_{Si0}^h in the upzoning counterfactual. Panel B shows the locations of HSR stations in the HSR counterfactual, as well as the reduction in outgoing travel times in minutes for each location, weighted by pre-HSR commuting flows. Panel C shows the reduction in incoming travel times, weighted by pre-HSR commuting flows. Bordered locations are the locations where upzoning takes place (panel A) or where HSR stations are located (panels B and C).

we first calculate the median housing productivity Z_{Si0}^h in the Bay Area and then, for each location below the median, we increase its Z_{Si0}^h 10% toward the Bay Area median.³⁴ Locations with low calibrated Z_{Si0}^h are those where prices are relatively high, which may indicate insufficient supply of housing. Panel A of Figure 5 shows that upzoning is largely concentrated around Silicon Valley and San Francisco—areas that are notorious for the difficulty to build and for high prices.

High-Speed Rail. In our second counterfactual scenario, the transportation infrastructure is improved. In particular, we simulate the construction of the Bay Area section of the California High-Speed Rail (HSR).³⁵ We do not model construction costs. It envisages four stations in the Bay Area: San Francisco, Millbrae/SFO, San Jose, and Gilroy. The projected travel time between San Francisco and Millbrae/SFO is 10 minutes, between Millbrae/SFO and San Jose is 20 minutes, and between San Jose and Gilroy is 18 minutes.

To simulate the HSR in our model, we calculate the counterfactual commuting time matrix as follows. First, we adjust travel times between model locations that receive an HSR station by adding 2 minutes for each stop, 5 minutes waiting time for each trip, as

³⁴A common alternative upzoning counterfactual is to increase supply elasticities ξ_i^S in low-elasticity locations (see for example [Hsieh and Moretti 2019](#)). Solely changing elasticities, without another spatially heterogeneous shock, has limited effects. For simplicity, we instead change construction productivities. See [Greaney \(2024\)](#) for further discussion of housing supply elasticity counterfactuals.

³⁵[Fajgelbaum et al. \(2023\)](#) also study the impact of the California HSR in a static quantitative spatial model of California.

well as extra time needed to travel between residence within the model location and the station and between the station and workplace destination. This extra time corresponds to the average travel time for trips within the model location from the CTPP. Then, we construct alternative routes between each pair of locations that use the HSR network. Finally, the counterfactual travel time is the minimum between the travel time on a route that uses HSR and the pre-HSR travel time.


The introduction of the HSR lowers travel time for 659 out of 3,025 location pairs, which represent 4.7% commuting flows in the benchmark economy. Panels B and C of Figure 5 show the reduction of outbound and inbound travel times, i.e., the travel times experienced by residents and workers in a given location, weighted by pre-HSR commuting flows. Time savings are the largest in locations with stations but remain substantial for many neighboring areas. Due to the asymmetric nature of the commuting time matrix and differences in the number of residents and jobs in each location, the magnitudes of changes in outgoing and incoming travel times may differ in a given location.

In these counterfactuals, the upzoning and the introduction of the HSR are unanticipated shocks. As soon as they occur, agents correctly anticipate the entire future path of prices (wages, prices, and rents in both the residential and commercial property sectors) in every location. Appendix Section C details the computation of the transition path. Also, in each experiment local productivity in the tradeable sector is fixed. In Appendix B.1, we demonstrate that endogenizing productivity to allow for agglomeration effects does not lead to meaningfully different quantitative results.

4.1.2 Transitional and Long-Run Spatial Effects

Upzoning. The increase in construction productivity allows developers to build more housing in upzoned locations. As can be seen in panel A of Figure 6, this attracts more residents from locations where upzoning did not take place in the long run. However, not all upzoned locations gain residents. For example, because upzoning is more aggressive in Silicon Valley (see panel A of Figure 5), it draws some of its new residents from East Bay where upzoning is more moderate. Panel B of Figure 6 shows the long-run relocation of jobs to the areas around Silicon Valley and San Francisco, which happens because firms in those areas have access to new workers who moved to these areas in response to greater housing supply. Long-run movements of residents and jobs can also be analyzed

Figure 6: Long-run spatial effects of upzoning



Note: The maps show the long-run percentage changes in residents and jobs in the upzoning counterfactual. Bordered areas represent the locations where upzoning takes place.

using a static QSM.³⁶ Our dynamic approach, however, also allows looking at transitional dynamics.


Panels A and B in Figure 7 show how the number of residents and jobs evolve along the transition path in each location. Due to moving costs, housing transaction costs, and durability of structures, it takes on average 8.9 and 8.3 years for the number of residents and jobs, respectively, to move halfway to the new steady state, and about 75 years for the full transition.

Residential prices and rents fall in upzoned areas on impact and remain lower in the long run, as demonstrated in panels D and E of Figure 7. As residents move away from non-treated locations, prices and rents there fall as well, and every single location in the Bay Area experiences a long-run decline in house prices. As the number of jobs around upzoned locations goes up, the demand for commercial real estate increases and so do commercial prices and rents, as shown in panels G and H of Figure 7.

Price-rent ratios exhibit a non-monotonic pattern. In the residential segment, the price-rent ratios in most upzoned locations jump on impact but then gradually decrease (panel F in Figure 7). This is because, as more residents move in each year following the zoning reform, residential rents increase after a large initial drop, while prices are forward-looking and immediately reflect the entire expected future rent path. Moreover,

³⁶This does not mean that comparing two equilibria of a static model will yield the same results as comparing two steady states of the dynamic model. Our model has risk-averse agents who make forward-looking choices, which affects the steady state. The standard static QSM is not nested by our model. Comparing our results to those from a static model is further complicated by the fact that a static model would have different parameter values, some of which would have to be calibrated using different moments than the ones we use.

Figure 7: Transitional effects of upzoning




Note: The figure shows the evolution of residents, jobs, homeownership rates, real estate prices, rents, and price-rent ratios over the transition in the upzoning counterfactual. Blue lines represent the model locations where housing supply productivity was increased. Gray lines represent other locations.

by making housing supply more abundant and lowering house prices, the policy promotes homeownership across the Bay Area. Panel C shows that homeownership rates go up in most locations, upzoned or not.

In the commercial real estate segment, the price-rent ratio first jumps in many locations in anticipation of a gradual increase in jobs. Then it slowly falls and returns to the initial level (panel I in Figure 7).

High-Speed Rail. The construction of the HSR reduces time required to travel to and from locations with stations, as well as many nearby areas. Thus, residents who live

Figure 8: Long-run spatial effects of HSR



Note: The maps show the long-run percentage changes in residents and jobs in the HSR counterfactual. Bordered areas represent the locations where HSR stations are built.


in those locations have better access to jobs in the Bay Area, while employers in those locations have better access to workers.

As a result, the four locations with HSR stations experience an increase in residents, as shown in panel A of Figure 8. However, as we can see in panel A of Figure 9, these changes only happen gradually due to moving costs, housing transaction costs, and durability of structures. The full transition unfolds over 75 years, and it takes on average 9.3 and 14.5 years for the number of residents and jobs, respectively, to move halfway to the new steady state. The magnitude of long-run population growth in locations with stations ranges from 2.6% in San Jose to 9.8% in San Francisco. One important reason for this heterogeneity is relatively high estimated housing supply elasticity in downtown San Francisco. This allows it to have a lower price growth (see panel D) and a higher population growth than in other locations.³⁷

The HSR has an interesting effect on the distribution of jobs in the Bay Area. With the exception of Millbrae/SFO, all treated locations lose jobs on impact, as can be seen in panel B of Figure 9. However, as time passes and new residents move into locations with stations, jobs partly recover. For example, Gilroy loses 3% of jobs on impact but the long-run job decline is only 2.1%. Gilroy is relatively unproductive and some residents who held local jobs before the HSR switch to more attractive jobs elsewhere and use the HSR to get to work. Many locations that do not have HSR stations but are relatively close also gain jobs. For example, the location to the south of Millbrae/SFO that contains San Mateo (and a large area of mostly undeveloped land) experiences a nearly 16% long-run increase

³⁷In our calibration, the elasticity is 1.47 in the downtown San Francisco location where the HSR station is built and 0.31 in the location where the San Jose station is built.

Figure 9: Transitional effects of HSR



Note: The figure shows the evolution of residents, jobs, homeownership rates, real estate prices, rents, and price-rent ratios over the transition in the HSR counterfactual. Thick blue lines represent the four model locations that contain the HSR stations. Thin gray lines represent other locations.

in jobs (see panel B of Figure 8) because its workers experience the largest reduction in travel times of all locations (see panel B of Figure 5). Other factors that help San Mateo gain more jobs than any other location are its high productivity and high elasticity of commercial floorspace supply.

The HSR project has a sizable effect on real estate markets, especially in locations with HSR stations. House prices jump immediately by between 2 and 6.9 percent in anticipation of future growth of rents (panel D of Figure 9). In the long run, price appreciation reaches 2.8–8.8 percent. The variation in magnitudes reflects different population dynamics and housing supply elasticities across treated locations. Rents adjust more gradually but also

end up higher in the long run (panel E). As a result, the price-rent ratio in treated locations jumps in the first few years (panel F). This rise in price-rent ratios lowers homeownership. While workers, especially young ones, are eager to move closer to the HSR, higher price-rent ratios often mean that they choose to rent. As panel C shows, homeownership rates fall by nearly 6 percentage points in San Francisco in the first 20 years before partly recovering in the long run. At the same time, the outflow of residents from several other locations leads to lower price-rent ratios and higher homeownership there.

Panels G to I of Figure 9 show the evolution of commercial real estate prices, rents, and price-rent ratios. Since moving jobs is costless, both prices and rents reach their new long-run levels in just a few years.

4.1.3 Welfare Analysis

Next, we compute welfare gains from each of the policies. We measure welfare effects as the percentage change in the consumption aggregate $c^{1-\eta}h^\eta$, applied in the initial steady state for the remainder of a household's lifetime, that would make it indifferent between remaining in the initial steady state and experiencing the counterfactual transition.

Aggregate welfare gains. Table 2 shows that upzoning leads to an aggregate welfare gain of 0.5%. The HSR results in an aggregate welfare gain of 0.4%. We also compute welfare gains by ignoring the transition, simply comparing the pre-policy and the new long-run steady states. In the case of the HSR, the difference between the two measures is negligible. However, in the upzoning counterfactual, not accounting for the transition overestimates welfare gains by over one-half. This is because the expansion of housing supply leads to sizable wealth losses among homeowners in the upzoned locations which lowers their utility at the beginning of the transition due to legacy real estate exposure. Another reason why accounting for transition dynamics lowers welfare gains is that it is costly to relocate to locations that offer higher utility due to moving and housing transaction costs. This is also true in the HSR counterfactual, where welfare gains would be larger in the absence of moving and transaction costs.

Our dynamic setting with rich agent heterogeneity allows us to dissect welfare gains along several dimensions. Figure 10 shows welfare effects by age, worker productivity, liquid wealth, tenure status, and residential location before the shock took place.

Distributional welfare gains: Upzoning. In the upzoning counterfactual, welfare effects are positive for renters and young owners, and negative for older owners, as shown in panel A of Figure 10. Many homeowners, especially the older ones, lose housing wealth as a result of housing supply expansion. Homeowners' welfare gains depend little

Table 2: Aggregate welfare effects

Welfare gain, %	Upzoning	HSR
Accounting for transition	0.50	0.40
Comparing steady states	0.78	0.40

Note: The table describes aggregate welfare gains in the two counterfactual experiments. Welfare gain is measured as the percentage change in the consumption aggregate $c^{1-\eta}\mathbf{h}^\eta$, applied for the remainder of a household's life, that all households would need to receive in the initial equilibrium to equalize average lifetime utility pre- and post-counterfactual.

on whether their location was upzoned or not. Due to spatial equilibrium effects, non-upzoned locations lose residents and also see a decline in home values. Welfare gains for owners monotonically decrease with age, because older owners have fewer working-age years to use the labor market as insurance against the negative housing wealth shock. Young homeowners gain because they tend to own smaller houses and upzoning allows them to move to a bigger house at a lower cost.


Renters win not only because rents fall but also because lower prices make it easier for them to become homeowners. The gains among renters increase with age for those who live in upzoned locations, and fall for those who live elsewhere. Older renters are less mobile. They are more likely to remain in their current locations and reap the benefits of the increased housing supply in upzoned neighborhoods.

Panel C demonstrates the role of wealth. The least wealthy owners lose the most. They are highly leveraged and have little liquid wealth, thus even a modest negative shock to house values will have large effects on their welfare. The wealthiest owners, especially those in upzoned locations, also lose more from upzoning because they tend to have a larger amount of housing wealth before the shock. Wealthier renters benefit more because, as prices fall, a dollar of wealth goes farther and allows them to purchase a larger house.

Distributional welfare gains: HSR. In the HSR experiment, welfare gains are high for young individuals and even higher for the 30–40 year olds. Then they fall with age until retirement, as shown in panel B of Figure 10. This is because the youngest individuals can enjoy the benefits of the HSR for a longer period of time. However, young workers are also more geographically mobile and are more likely to move away from neighborhoods with stations due to idiosyncratic preference shocks.

Due to the appreciation of house prices due to greater connectivity in many locations, homeowners in locations with HSR stations are the biggest winners. This is especially

Figure 10: Distributional welfare effects



Note: The figure shows counterfactual welfare gains by age (panels A and B) and total wealth (panels C and D) in each counterfactual. Welfare gains are shown separately for owners (left side of each panel) and renters (right side), as well as separately for residents of treated and non-treated locations. Welfare gain is measured as the percentage change in the consumption aggregate $c^{1-\eta}h^\eta$, applied for the remainder of a household's life in the initial steady state, that would make it indifferent between remaining in the initial steady state or experiencing the counterfactual transition.

true for retirees who are unlikely to move to a different location and can reap the benefits of higher house values. At the same time, old renters who live in locations with HSR stations lose because they do not benefit from greater job connectivity but their rents still go up.

Proximity to the HSR matters too, and welfare gains are concentrated among residents who live in locations where HSR stations are built. However, the gains are also substantial for those who live in other locations that are fewer than 20 minutes away and are smaller for those who live farther out.

Panel D shows that, while owners of all wealth levels gain, the least wealthy owners in locations with HSR stations benefit the most. They are leveraged and even small positive

shocks to house values have large effects on their welfare. The wealthiest owners also experience sizable gains. They tend to own large houses in the pre-HSR economy, and when a station is built in their neighborhood, they experience a wealth windfall. Gains for renters do not depend much on wealth.

4.2 Why Dynamics, Many Locations, and Homeownership Matter

4.2.1 Why Dynamics and a Large Number of Locations Are Required

There is a sizable literature that has studied versions of our counterfactuals in different models. Policy counterfactuals with improving infrastructure have been predominantly studied using static QSMs with a large number of locations (Severen, 2021; Allen and Arkolakis, 2022; Tsivanidis, 2023; Chen et al., 2024; Fajgelbaum et al., 2023). Policy counterfactuals with increasing housing supply within a city have been studied either using static frameworks (Allen et al., 2016; Acosta, 2022) or using dynamic frameworks with a small number of within-city locations (Favilukis et al., 2022).

Our policy experiments demonstrate how a dynamic model produces important results that cannot be obtained from a static model. First, a static model cannot produce transitional dynamics that arise from individual choices. As we showed in Figures 7 and 9, transitions can be long and sometimes non-monotonic. They allow for a policy impact analysis at any time horizon, rendering our model more suitable for policymaking. The standard approach in the literature of comparing two static equilibria would miss the fact that the influx of residents into treated and some non-treated locations happens very gradually. Prices and rents adjust at different paces so that the policy has different effects on renters and owners. Moreover, disregarding transitions may result in over- or underestimation of welfare gains, as was the case in our upzoning experiment.

Second, a static model cannot accommodate risk and intertemporal choices such as saving and housing tenure choice. But these choices are intertwined with location choices. Workers choose where to live and work not only based on wages, housing costs, commuting costs, and amenities, as in the static model, but also take into account which locations will allow them to make optimal tenure choices, saving decisions, and provide better insurance against labor income risk.

Third, the dynamic nature of our model results in rich agent heterogeneity by age, productivity, housing tenure, housing and non-housing wealth. Modeling these characteristics in a static model would require far-fetched assumptions. However, as we demonstrated in Figure 10, this heterogeneity is crucial for understanding the welfare implications of spatial policies.

Our policy counterfactuals also show why a large number of locations is necessary; one cannot simply collapse the geography into a small number of locations such as the treated and the non-treated. For example, the HSR produces a non-trivial adjustment of travel times throughout the entire transportation network, even though stations are built in just four locations. As a result, there is large variation in how different locations respond to the policy. A simpler model with two locations, treated and non-treated, will miss such spatial spillover effects. It would also be unable to explain why some non-treated neighborhoods benefit and others lose from the policy. Similarly, in the upzoning counterfactual, jobs change not only based on the treatment status but also based on the proximity of a given area to an upzoned location.

4.2.2 Why Homeownership Is Required

Adding homeownership to dynamics and a large number of locations adds complexity. It also necessitates modeling saving and borrowing decisions.³⁸ So it is reasonable to ask how important this extra addition is.

To demonstrate the importance of homeownership, we build a version of our model which shuts down homeownership. By setting the ownership preference parameter $\chi = 0$, we ensure that no one chooses to own a house of any size. We continue to allow individuals to save in the risk-free asset. We fix the parameters $(\rho, \vartheta_0, \vartheta_1)$, calibrated from wealth moments, to their benchmark model values. We apply the calibrated value for the maximum owner-occupied house size ($\max \mathbb{H}$) to rental housing in the no-ownership model to ensure that our results are not driven by differences in feasible housing consumption. Then, to show the importance of saving and borrowing choices, we build another version of the model where not only all households are renters but they also live hand-to-mouth.

In both the no-ownership and the hand-to-mouth models, we allow changes in local real estate values to affect household utility. To make these models comparable to our main model, where 60.9% of the value of residential real estate in the Bay Area is owned by homeowners, we assume that 60.9% of residential real estate in each location is owned by location-specific REITs that redistribute all gains or losses equally to local households. This preserves the lack of spatial diversification that homeowners in the main model experience when shocks hit. This approach is similar to how several other spatial models have incorporated homeownership without explicitly modeling it.³⁹ The remaining residential

³⁸In the quantitative model, homeownership multiplies the size of our state space by the number of wealth grid points times the number of house size grid points. In our quantitative model, the wealth grid has 50 points and the house size grid has 7 points. Thus, our model has a state space that is 350 times larger than a model with hand-to-mouth renters.

³⁹See [Redding and Rossi-Hansberg \(2017\)](#) for a discussion of this approach.

Table 3: Distribution of Welfare Gains

Welfare gains, %	Upzoning			HSR		
	Main	No own.	HTM	Main	No own.	HTM
Mean	0.50	0.45	0.35	0.40	0.30	0.09
Variance	1.30	0.17	0.04	0.59	0.13	0.01
1st percentile	-3.85	0.02	0.04	-1.29	-0.44	-0.19
5th percentile	-1.42	0.05	0.07	-0.37	-0.12	-0.05
10th percentile	-0.74	0.09	0.10	-0.14	-0.05	-0.02
25th percentile	-0.16	0.20	0.19	0.02	0.07	0.02
Median	0.20	0.35	0.33	0.28	0.27	0.09
75th percentile	0.67	0.54	0.46	0.51	0.47	0.15
90th percentile	1.36	0.96	0.62	0.78	0.65	0.20
95th percentile	1.90	1.31	0.75	1.11	0.80	0.23
99th percentile	3.01	2.05	1.01	2.87	1.51	0.40
Fraction harmed, %	34.25	0.09	0.02	22.12	14.46	15.56

Note: The table shows the distribution of welfare gains in the two counterfactuals, using the main model, the model without homeownership, and the hand-to-mouth model. The last line reports the fraction of households who experience welfare losses. Welfare gain is measured as the percentage change in the consumption aggregate $c^{1-\eta}h^\eta$, applied in the initial steady state for the remainder of a household's lifetime, that would make it indifferent between remaining in the initial steady state or experiencing the counterfactual transition.

real estate and commercial real estate are held by spatially diversified REITs, as in the main model. Households own equal shares of these REITs, and we set the fractions that are owned by households in the model equal to that from the main model.⁴⁰

We re-estimate all other parameters of the no-ownership and hand-to-mouth models. We then re-run the HSR and upzoning counterfactuals. Relocations of residents and jobs are similar to the main model. Transitional dynamics are also broadly similar; see Appendix Sections B.3 and B.4 for details.

However, as Table 3 shows, welfare gains differ substantially between the main model and the two alternative models. Without homeownership, welfare gains from the HSR are smaller for half the population. A large source of welfare gains in the main model comes from house price appreciation experienced by homeowners close to train stations. Exposure to local REITs cannot fully capture this effect because households in our main model

⁴⁰Recall that in the main model, households invest fraction $f^S = 0.0845$ of their liquid assets in the residential REIT and fraction $f^C = 0.0574$ in the commercial REIT. In both sectors, this implies that not all real estate is owned by households in the model. We assume that the remainder is held by absentee owners.

select into homeownership and hold owner-occupied housing with leverage. Conversely, welfare gains from upzoning are higher for the majority of residents in the model without homeownership. In the main model, many homeowners, particularly old homeowners, lose from the relaxation of housing supply constraints and the ensuing fall in house prices. In the model without homeownership, these losses are more evenly distributed across local residents.

The average welfare gains in the hand-to-mouth model are much smaller than in our main model and in the no-ownership model. Not allowing households to insure themselves by saving and borrowing lowers average welfare gains from upzoning by one-quarter and from HSR by over two-thirds.


The variance of welfare gains is substantially larger in our main model than in the other two models (second row), and there are far more losers from the policy in our main model (last row). In the upzoning experiment, 34.25% of households lose in our main model compared to a mere 0.09% in the model without homeownership and 0.02% in the hand-to-mouth model, even though average gains are highest in the main model. In the alternative models, real estate losses experienced by REITs, which result in lower transfers to locals, are offset by lower rents so that nearly no one loses from the policy. In the HSR experiment, 22.12% of the population experiences losses in our main model compared to 14.46% in the no-ownership model and 15.56% in the hand-to-mouth model, even though average welfare gains are higher in our main model.

Simply put, homeownership increases the exposure of households to localized shocks, such as the introduction of the HSR or the relaxation of housing supply constraints. This is because owner-occupied housing is a spatially undiversified asset that many households own with substantial leverage. For many households, it is also the largest asset in their portfolios. Therefore, heterogeneity in tenure status substantially increases the dispersion in welfare changes that households experience. Approximating homeownership by a local REIT cannot capture the effects of homeownership.

The heterogeneity in welfare changes not only occurs in the population at large, it also manifests itself in nearly every location. Panel A of Figure 11 shows that the upzoning experiment creates a large number of losers in all locations, including in the locations where housing supply constraints are relaxed. Surely, many homeowners in those locations are not happy about lower house prices due to housing supply expansion. Strikingly, panel B shows that the majority of owners lose from the policy in all but 5 locations, including many locations not directly treated. However, in the model without homeownership (panel C) and in the hand-to-mouth model (panel D), there are virtually no losers.

In the HSR counterfactual, there are many locations far from the HSR stations in the

Figure 11: Spatial Distribution of Fraction Harmed, Upzoning




Note: The maps show the fraction of households who experience welfare losses in each location in the upzoning counterfactual. Panels A and B show the fractions of all households and homeowners, respectively. Panels C and D show the fractions in the model without homeownership and the hand-to-mouth model. Bordered areas represent the locations where upzoning takes place.

northeast of the CZ where a large number of residents lose from the policy (panel A of Figure 12). These losers are homeowners who see their property values fall when some residents move out to live closer to the HSR. Indeed, panel B shows that many owners outside of the four locations with stations lose from the policy, and in many locations the fraction of losers is greater than 50 percent. However, as shown in panels C and D, in the model without homeownership and in the hand-to-mouth model, losers are concentrated close to the HSR. This result is driven by older renters in those areas who suffer from higher rents but do not benefit from better connectivity brought about by the HSR.

These results suggest that a model with homeownership is well suited to rationalize why policies targeted at expanding housing supply are often met with stiff opposition from local residents. Our model generates that the majority of homeowners lose from the policy in many locations. Homeowners are more active in local politics than renters and are also more likely to vote (DiPasquale and Glaeser, 1999; Hall and Yoder, 2022). Moreover, as shown in panel C of Figure 10, wealthy residents lose from the upzoning

Figure 12: Spatial Distribution of Fraction Harmed, HSR



Note: The maps show the fraction of households who experience welfare losses in each location in the HSR counterfactual. Panels A and B show the fractions of all households and homeowners, respectively. Panels C and D show the fractions in the model without homeownership and the hand-to-mouth model. Bordered areas represent the locations where HSR stations are built.

policy. If higher wealth is associated with greater political influence, wealthy homeowners can use their economic power to thwart housing supply reforms.

Similarly, our model can account for difficulties in implementing transportation improvements. As panel B of Figure 12 demonstrates, the locations where the majority of homeowners lose tend to be the neighbors of the locations that receive a station, such as several locations in Silicon Valley and areas south of Downtown San Francisco. However, the HSR tracks would have to pass through these locations. Our model justifies opposition by the local residents on economic grounds.

5 Conclusion

We propose a new computational dynamic spatial equilibrium framework. Our method solves urban models with forward-looking, risk-averse agents who face idiosyncratic risk. They make dynamic consumption-savings decisions as well as periodic location and

homeownership decisions. We use the model to study the transitional dynamics following policy changes that affect some places within a city differently from others. Prominent examples are policies that increase residential density in the urban core and improve transportation infrastructure. The results indicate slow transitions and spatially heterogeneous responses due to the presence of moving costs and a legacy stock of immobile real estate. They also show that accounting for dynamics and homeownership is indispensable to establish both aggregate and distributional implications of local policies and shocks.

Our approach is well-suited to study the welfare effects of place-based policies that aim to provide social insurance to people in left-behind places, adjustment dynamics to work-from-home shocks with local fiscal policy consequences that risk triggering an urban doom loop (Gupta et al., 2025), and to quantitatively assess the general equilibrium impact of neighborhood-scale urban investment projects.

Bibliography

- Achdou, Yves, Han, Jiequn, Lasry, Jean-Michel, Lions, Pierre-Louis and Moll, Benjamin. (2022). ‘Income and wealth distribution in macroeconomics: A continuous-time approach’, *The Review of Economic Studies* 89(1), 45–86.
- Acosta, Camilo. (2022), The Incidence of Land Use Regulations. Working Paper.
- Ahlfeldt, Gabriel M. and Pietrostefani, Elisabetta. (2019). ‘The economic effects of density: A synthesis’, *Journal of Urban Economics* 111, 93–107.
- Ahlfeldt, Gabriel M., Redding, Stephen J., Sturm, Daniel M. and Wolf, Nikolaus. (2015). ‘The Economics of Density: Evidence From the Berlin Wall’, *Econometrica* 83(6), 2127–2189.
- Allen, Treb and Arkolakis, Costas. (2022). ‘The Welfare Effects of Transportation Infrastructure Improvements’, *The Review of Economic Studies* 89(6), 2911–2957.
- Allen, Treb, Arkolakis, Costas and Li, Xiangliang. (2016), Optimal City Structure. Working Paper.
- Allen, Treb and Donaldson, Dave. (2022), Persistence and path dependence in the spatial economy, Technical report, National Bureau of Economic Research.
- Almagro, Milena and Domínguez-Iino, Tomás. (2024), Location Sorting and Endogenous Amenities: Evidence from Amsterdam, Technical report, National Bureau of Economic Research.
- Artuç, Erhan, Chaudhuri, Shubham and McLaren, John. (2010). ‘Trade shocks and labor adjustment: A structural empirical approach’, *American economic review* 100(3), 1008–1045.
- Baum-Snow, Nathaniel and Han, Lu. (2023), The Microgeography of Housing Supply. Working Paper.

- Berger, David, Guerrieri, Veronica, Lorenzoni, Guido and Vavra, Joseph.** (2018). 'House prices and consumer spending', *The Review of Economic Studies* 85(3), 1502–1542.
- Bilal, Adrien.** (2023), Solving Heterogeneous Agent Models with the Master Equation, Working Paper 31103, National Bureau of Economic Research.
- Bilal, Adrien and Rossi-Hansberg, Esteban.** (2021). 'Location as an Asset', *Econometrica* 89(5), 2459–2495.
- Bilal, Adrien and Rossi-Hansberg, Esteban.** (2023), Anticipating Climate Change Across the United States, Working Paper 31323, National Bureau of Economic Research.
- Cai, Sheng, Caliendo, Lorenzo, Parro, Fernando and Xiang, Wei.** (2022), Mechanics of spatial growth, Technical report, National Bureau of Economic Research.
- Caliendo, Lorenzo, Dvorkin, Maximiliano and Parro, Fernando.** (2019). 'Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock', *Econometrica* 87(3), 741–835.
- Campbell, John Y and Cocco, Joao F.** (2007). 'How do house prices affect consumption? Evidence from micro data', *Journal of monetary Economics* 54(3), 591–621.
- Carroll, Christopher D.** (2006). 'The method of endogenous gridpoints for solving dynamic stochastic optimization problems', *Economics letters* 91(3), 312–320.
- Chen, Liming, Hasan, Rana, Jiang, Yi and Parkhomenko, Andrii.** (2024). 'Faster, Taller, Better: Transit Improvements and Land Use Policies', *Journal of Development Economics* p. 103322.
- Crews, Levi.** (2023), A Dynamic Spatial Knowledge Economy. Working Paper.
- Davis, Morris A., Larson, William D., Oliner, Stephen D. and Shui, Jessica.** (2021). 'The price of residential land for counties, ZIP codes, and census tracts in the United States', *Journal of Monetary Economics* 118, 413–431.
- Davis, Morris A. and Van Nieuwerburgh, Stijn.** (2015). 'Housing, Finance, and the Macroeconomy', *Handbook of Regional and Urban Economics* 5, 753–811.
- Delventhal, Matthew J. and Parkhomenko, Andrii.** (2024), Spatial Implications of Telecommuting. Working Paper.
- De Nardi, Mariacristina.** (2004). 'Wealth inequality and intergenerational links', *The Review of Economic Studies* 71(3), 743–768.
- Desmet, Klaus, Nagy, Dávid Krisztián and Rossi-Hansberg, Esteban.** (2018). 'The Geography of Development', *Journal of Political Economy* 126(3), 903–983.
- DiPasquale, Denise and Glaeser, Edward L.** (1999). 'Incentives and social capital: Are homeowners better citizens?', *Journal of urban Economics* 45(2), 354–384.
- Dvorkin, Maximiliano.** (2023), Heterogeneous Agents Dynamic Spatial General Equilibrium. Working Paper.
- Eckert, Fabian and Kleineberg, Tatjana.** (2021), Saving the American Dream? Education Policies in Spatial General Equilibrium. Working Paper.
- Fajgelbaum, Pablo D, Gaubert, Cecile, Gorton, Nicole, Morales, Eduardo and Schaal, Edouard.** (2023), Political Preferences and Transport Infrastructure: Evidence from California's High-Speed Rail, Working Paper 31438, National Bureau of Economic Research.
- Favilukis, Jack, Ludvigson, Sydney C. and Van Nieuwerburgh, Stijn.** (2017). 'The

- Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium', *Journal of Political Economy* 125(1), 140 – 223.
- Favilukis, Jack, Mabilie, Pierre and Van Nieuwerburgh, Stijn.** (2022). 'Affordable Housing and City Welfare', *The Review of Economic Studies* . rdac024.
- Floden, Martin and Lindé, Jesper.** (2001). 'Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?', *Review of Economic Dynamics* 4(2), 406–437.
- Giannone, Elisa.** (2019), Skill-Biased Technical Change and Regional Convergence. Working Paper.
- Giannone, Elisa, Li, Qi, Paixao, Nuno and Pang, Xinle.** (2023), Unpacking Moving, Working paper.
- Glaeser, Edward L and Gyourko, Joseph.** (2005). 'Urban decline and durable housing', *Journal of political economy* 113(2), 345–375.
- Graham, Matthew R., Kutzbach, Mark J. and McKenzie, Brian.** (2014), Design Comparison of LODES and ACS Commuting Data Products, Technical Report 14-38. U.S. Census Bureau.
- Greaney, Brian.** (2023), Homeownership and the Distributional Effects of Uneven Regional Growth. Working Paper.
- Greaney, Brian.** (2024), Housing constraints and spatial misallocation: Comment.
- Gupta, Arpit, Mittal, Vrinda and Van Nieuwerburgh, Stijn.** (2025). 'Work from home and the office real estate apocalypse', *American Economic Review* forthcoming.
- Gyourko, Joseph and Molloy, Raven.** (2015), Chapter 19 - Regulation and Housing Supply, in Gilles Duranton, J. Vernon Henderson and William C. Strange., eds, 'Handbook of Regional and Urban Economics', Vol. 5 of *Handbook of Regional and Urban Economics*, Elsevier, pp. 1289–1337.
- Hall, Andrew B and Yoder, Jesse.** (2022). 'Does homeownership influence political behavior? Evidence from administrative data', *The Journal of Politics* 84(1), 351–366.
- Hansen, Gary D.** (1993). 'The cyclical and secular behaviour of the labour input: Comparing efficiency units and hours worked', *Journal of Applied Econometrics* 8(1), 71–80.
- Heblich, Stephan, Redding, Stephen and Sturm, Daniel M.** (2020). 'The Making of the Modern Metropolis: Evidence from London', *The Quarterly Journal of Economics* 135(4), 2059–2133.
- Hsieh, Chang-Tai and Moretti, Enrico.** (2019). 'Housing Constraints and Spatial Misallocation', *American Economic Journal: Macroeconomics* 11(2), 1–39.
- Kaplan, Greg, Mitman, Kurt and Violante, Giovanni L.** (2020). 'The Housing Boom and Bust: Model Meets Evidence', *Journal of Political Economy* 128(9), 3285 – 3345.
- Kleinman, Benny, Liu, Ernest and Redding, Stephen.** (2023). 'Dynamic Spatial General Equilibrium', *Econometrica* 91(2), 385–424.
- Kopecky, Karen A and Suen, Richard MH.** (2010). 'Finite state Markov-chain approximations to highly persistent processes', *Review of Economic Dynamics* 13(3), 701–714.
- Landvoigt, Tim, Piazzesi, Monika and Schneider, Martin.** (2015). 'The Housing Market(s) of San Diego', *American Economic Review* 105(4), 1371–1407.

- Luccioletti, Claudio.** (2023), Should Governments Subsidize Homeownership? Working Paper.
- Martellini, Paolo.** (2022), Local labor markets and aggregate productivity. Working Paper.
- Merton, Robert C.** (1969). 'Lifetime portfolio selection under uncertainty: The continuous-time case', *The review of Economics and Statistics* pp. 247–257.
- Ortalo-Magné, François and Prat, Andrea.** (2016). 'Spatial asset pricing: A first step', *Economica* 83(329), 130–171.
- Piazzesi, M. and Schneider, Martin.** (2016), Housing and Macroeconomics, Vol. 2, Elsevier, chapter Chapter 19, pp. 1547–1640.
- Redding, Stephen J and Rossi-Hansberg, Esteban.** (2017). 'Quantitative spatial economics', *Annual Review of Economics* 9(1), 21–58.
- Rouwenhorst, K Geert.** (1995), Asset Pricing Implications of Equilibrium Business Cycle Models, in 'Frontiers of business cycle research', Princeton University Press, pp. 294–330.
- Saiz, Albert.** (2010). 'The Geographic Determinants of Housing Supply', *The Quarterly Journal of Economics* 125(3), 1253–1296.
- Severen, Christopher.** (2021). 'Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification', *The Review of Economics and Statistics* pp. 1–99.
- Steinberg, Joseph B.** (2019). 'Brexit and the macroeconomic impact of trade policy uncertainty', *Journal of International Economics* 117, 175–195.
- Sun, Jeffrey E.** (2024), The Distributional Consequences of Climate Change: The Role of Housing Wealth, Expectations, and Uncertainty. Working Paper.
- Takeda, Kohei and Yamagishi, Atsushi.** (2023), History versus Expectations in the Spatial Economy: Lessons from Hiroshima. Working Paper.
- Tolbert, Charles M. and Sizer, Molly.** (1996), U.S. Commuting Zones and Labor Market Areas: A 1990 Update, Staff Reports 278812, United States Department of Agriculture, Economic Research Service.
- Tsivanidis, Nick.** (2023), The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio. Working Paper.
- Valentinyi, Akos and Herrendorf, Berthold.** (2008). 'Measuring Factor Income Shares at the Sectoral Level', *Review of Economic Dynamics* 11(4), 820–835.
- Vanhapelto, Tuuli.** (2022), House Prices and Rents in a Dynamic Spatial Equilibrium. Working Paper.
- Warnes, Pablo.** (2024), Transport Infrastructure Improvements and Spatial Sorting: Evidence from Buenos Aires. Working Paper.
- Zerecero, Miguel.** (2021), The Birthplace Premium. Working Paper.

Online Appendix

A Data

Wages. We use the Census Transportation Planning Products (CTPP) database and the ACS data for 2012–2016 to obtain estimates of average wage for each location. We use the data reported for the period from 2012 to 2016. We use the variable “earnings in the past 12 months (2016 \$), for the workers 16-year-old and over,” which is based on the respondents’ workplace locations. The variable provides the estimates of the number of people in each of the several earning bins in each workplace tract.⁴¹

We calculate mean labor earnings for tract k as $\bar{w}_k = (\sum_b N_{b,k} \bar{w}_b) / \sum_b N_{b,k}$, where $N_{b,k}$ is the number of workers in bin b in tract k , and $\bar{w}_b$ is mean earnings in bin b for each PUMA, calculated from the ACS microdata. Next, to control for possible effects of workers’ heterogeneity on tract-level averages, we estimate

$$\bar{w}_k = \alpha + \beta_1 age_k + \beta_2 sexratio_k + \sum_r \beta_{2,r} race_{r,k} + \sum_d \beta_{3,d} ind_{d,k} + \sum_o \beta_{4,o} occ_{o,k} + \epsilon_k, \quad (A.1)$$

where age_k is the average age; $sexratio_k$ is the proportion of males to females in local labor force; $race_{r,k}$ is the share of race $r \in \{Asian, Black, Hispanic, White\}$; $ind_{d,k}$ is the share of jobs in industry d ; and $occ_{o,k}$ is share of jobs in occupation o in tract k .⁴² The estimated tract-level wage index is the sum of the estimated constant and the tract fixed effect: $\hat{w}_k^0 \equiv \hat{\alpha} + \hat{\epsilon}_k$. We then construct wage indices for each location j , $\hat{w}_j^0$, as the employment-weighted average of the values of $\hat{w}_k^0$ for each tract k that pertains to model location j .

Residential floorspace rents and prices. To estimate the citywide price-rent ratio, we use the 2012–2016 ACS data. We keep only household heads to ensure that the analysis is at the level of a residential unit. We exclude observations who live in group quarters;

⁴¹The bins are $\leq \$9,999$; $\$10,000$ – $\$14,999$; $\$15,000$ – $\$24,999$; $\$25,000$ – $\$34,999$; $\$35,000$ – $\$49,999$; $\$50,000$ – $\$64,999$; $\$65,000$ – $\$74,999$; $\$75,000$ – $\$99,999$; and $\geq \$100,000$.

⁴²We use the following *industry* categories: Agricultural; Armed force; Art, entertainment, recreation, accommodation; Construction; Education, health, and social services; Finance, insurance, real estate; Information; Manufacturing; Other services; Professional scientific management; Public administration, Retail. We use the following *occupation* categories: Architecture and engineering; Armed Forces; Arts, design, entertainment, sports, and media; Building and grounds cleaning and maintenance; Business and financial operations specialists; Community and social service; Computer and mathematical; Construction and extraction; Education, training, and library; Farmers and farm managers; Farming, fishing, and forestry; Food preparation and serving related; Healthcare practitioners and technicians; Healthcare support; Installation, maintenance, and repair; Legal; Life, physical, and social science; Management; Office and administrative support; Personal care and service; Production; Protective service; Sales and related.

live in farm houses, mobile homes, trailers, boats, tents, etc.; are younger than 20 years old; and live in a dwelling that has no information on the year of construction. Then we estimate the following hedonic rent and price indices for each PUMA using self-reported housing rents and prices:

$$\ln \mathbf{q}_{i,t} = \beta_0 + \beta_1 \mathbf{X}_{i,t} + \varphi_t + \varepsilon_{i,t}. \quad (\text{A.2})$$

Here, $\mathbf{q}_{i,t}$ is the rent or the price reported by household i in year t , while $\mathbf{X}_{i,t}$ is a vector of controls that includes the number of rooms in the dwelling, the number of units in the structure (e.g., single-family detached, 2-family building), and the year of construction. Parameter φ_t is the year fixed effect, and $\varepsilon_{i,t}$ is the error term. We then use the price and rent indices evaluated at a median house type (single-family, detached, three bedroom), and divide them to obtain the price-rent ratio of 26.7.

Commercial floorspace rents and prices. We use lease transaction data for office, retail, and industrial properties located in California from the data provider CompStak, spanning the period 2000 until 2023. We then estimate a regression of the log real rent annual net effective per square foot on a set of ZIP code fixed effects and a series of control variables. Net effective rent adjusts the contract rent schedule over the life of the lease for landlord concessions (tenant improvements and free rent). The control variables include space type (office, retail, or industrial), lease length, building age, building size (log square feet), and building quality (A, B, or C). We then create the ZIP-code rent index as the sum of the ZIP code FE, the average of the 2015 and 2016 time FE, for a class-A office of average age, size, and with average lease-length structure. In a final step we aggregate up from the ZIP code to our model locations in the Bay Area using population weights.

Commute times. The CTPP data divides commuting times into 10 bins: less than 5 minutes, 5 to 14 minutes, 15 to 19 minutes, 20 to 29 minutes, 30 to 44 minutes, 45 to 59 minutes, 60 to 74 minutes, 75 to 89 minutes, 90 or more minutes, and work from home.

First, we calculate travel time between each pair of locations as the average of all tract-to-tract times with an origin inside one location and a destination in the other. We discard the calculation for any pair for which fewer than 10% of all possible tract-to-tract times are reported by CTPP. We also exclude times that imply a speed of more than 100 km/hour or less than 5 km/hour. We perform this same calculation for the average distance of each location *from itself*, obtaining data-based estimates of internal travel times.

Second, we take the primitive connections and the travel times between them, detailed above, as the first-order connections in a transport network. We use Dijkstra's algorithm to find the smallest possible travel times through this network between any pair of model locations for which travel times cannot be calculated directly.

Floorspace supply elasticities. We obtain housing supply elasticities ξ_i^S from [Saiz \(2010\)](#) and [Baum-Snow and Han \(2023\)](#). The benefit of the former is that it provides long-run 30-year supply elasticities, compared to 10-year elasticities in [Baum-Snow and Han \(2023\)](#). The benefit of the latter is that it provides elasticities at the Census tract level, compared to at the MSA level in [Saiz \(2010\)](#).

First, we take the elasticities from [Saiz \(2010\)](#) using the value of 0.66 for model locations that belong to the San Francisco MSA and 0.76 for locations that belong to the San Jose MSA. These elasticities were estimated using population changes, which means that these are elasticities of the supply of housing units, not total floorspace. Second, we take the 2011 *housing units* elasticities estimated with the FMM-IV model at the tract level from [Baum-Snow and Han \(2023\)](#) (variable *gamma11b_units_FMM*), then aggregate them to the level of model locations and restrict the elasticities to be at least 0.05, and obtain ξ_i^{BSH} for each model location. Third, we calculate

$$\xi_i = \frac{\xi_i^{\text{BSH}}}{\bar{\xi}^{\text{BSH}}} \xi_i^{\text{Saiz}},$$

where ξ_i^{Saiz} is the [Saiz \(2010\)](#) elasticity and $\bar{\xi}^{\text{BSH}}$ is the population-weighted average [Baum-Snow and Han \(2023\)](#) elasticity. This approach ensures that the average elasticity in our model is the same as in [Saiz \(2010\)](#) and the variance across locations is the same as in [Baum-Snow and Han \(2023\)](#)

To obtain supply elasticities in the commercial real estate sector, we use the same procedure with the only exception that we use the *total floorspace* elasticity from [Baum-Snow and Han \(2023\)](#) (variable *gamma11b_space_FMM*).

Table A.1: Location Characteristics

Location	N_i	N_j^W	E_i^R	E_j^W	Z_j	ξ_i^S	ξ_i^C
Alameda (North)–Berkeley and Albany Cities	72.7	73.8	0.00	0.00	1.00	0.64	0.53
Alameda (Northwest)–Oakland (Northwest) and Emeryville Cities	110.0	141.1	-0.09	-0.02	1.06	0.89	0.70
Alameda (Northeast)–Oakland (East) and Piedmont Cities	90.2	35.2	-0.02	-0.07	1.03	1.15	0.75
Alameda (North Central)–Oakland City (South Central)	63.1	18.2	-0.16	-0.09	0.94	0.65	0.59
Alameda (West)–San Leandro, Alameda and Oakland (Southwest) Cities	117.1	90.7	-0.05	-0.04	1.01	0.97	0.79
Alameda (North Central)–Castro Valley, San Lorenzo and Ashland	90.7	24.1	-0.09	-0.09	0.96	0.83	0.66
Alameda (Central)–Hayward City	96.9	62.1	-0.10	-0.05	0.96	0.94	0.75
Alameda (Southwest)–Union City, Newark and Fremont (West) Cities	96.5	96.2	-0.05	0.01	0.92	0.65	0.65
Alameda (South Central)–Fremont City (East)	131.0	39.3	0.00	-0.08	1.01	0.63	0.62
Alameda (East)–Livermore, Pleasanton and Dublin Cities	146.0	105.1	0.01	-0.02	0.99	1.01	0.77
Contra Costa (Far Southwest)–Richmond (Southwest) and San Pablo Cities	64.4	31.3	-0.18	-0.10	0.95	0.73	0.64
Contra Costa (Far Northwest)–Richmond (North), Hercules and El Cerrito Cities	80.1	24.6	-0.11	-0.09	0.96	1.01	0.69
Contra Costa (Northwest)–Concord (West), Martinez and Pleasant Hill Cities	80.8	69.0	-0.10	-0.07	1.03	1.01	0.74
Contra Costa–Walnut Creek (West), Lafayette, Orinda Cities and Moraga Town	68.9	56.3	-0.01	-0.08	1.10	2.37	1.31
Contra Costa (South)–San Ramon City and Danville Town	81.9	52.4	0.02	-0.09	1.09	2.15	1.22
Contra Costa (Central)–Concord (South), Walnut Creek (East) and Clayton Cities	71.9	40.4	-0.04	-0.10	1.06	1.13	0.78
Contra Costa (North Central)–Pittsburg and Concord (North and East) Cities	69.5	27.3	-0.15	-0.08	0.92	0.84	0.67
Contra Costa (Northeast)–Antioch City	59.7	21.3	-0.13	-0.11	0.96	0.67	0.59
Contra Costa (East)–Brentwood and Oakley Cities	59.2	14.2	-0.08	-0.09	0.93	2.06	1.28
Marin (North and West)–Novato and San Rafael (North) Cities	68.5	46.4	-0.02	-0.08	1.05	2.16	1.22
Marin (Southeast)–San Rafael (South), Mill Valley and Sausalito Cities	79.6	64.7	0.07	-0.08	1.14	1.54	0.93
Napa–Napa City	91.5	68.6	-0.03	-0.04	1.02	1.65	1.06
San Francisco (North and West)–Richmond District	103.5	67.6	0.09	-0.12	1.24	0.87	0.64
San Francisco (North and East)–North Beach and Chinatown	85.5	99.9	0.08	-0.09	1.27	1.10	0.76
San Francisco (Central)–South of Market and Potrero	94.6	354.9	0.03	-0.03	1.26	1.47	0.95
San Francisco (Central)–Inner Mission and Castro	89.4	47.6	0.07	-0.10	1.19	0.82	0.61
San Francisco (Central)–Sunset District (North)	80.9	32.6	0.05	-0.14	1.17	0.73	0.57
San Francisco (South Central)–Sunset District (South)	86.5	26.8	0.03	-0.11	1.10	0.62	0.51
San Francisco (South Central)–Bayview and Hunters Point	82.6	31.9	-0.02	-0.14	1.09	0.73	0.57
San Mateo (North Central)–Daly City, Pacifica Cities and Colma Town	105.5	28.1	-0.03	-0.12	1.09	1.01	0.69
San Mateo (North Central)–South San Francisco, San Bruno and Brisbane Cities	91.9	78.8	-0.03	-0.12	1.20	0.72	0.56
San Mateo (Central)–San Mateo (North), Burlingame and Millbrae Cities	67.2	59.9	0.08	-0.11	1.20	0.88	0.68
San Mateo (South and West)–San Mateo (South) and Half Moon Bay Cities	93.3	66.1	0.04	-0.09	1.19	1.45	0.96
San Mateo (East Central)–Redwood City, San Carlos and Belmont Cities	87.2	78.1	0.06	-0.10	1.21	0.81	0.67
San Mateo (Southeast)–Menlo Park, East Palo Alto Cities and Atherton Town	66.0	54.5	0.10	-0.14	1.30	1.19	0.87
Santa Clara (Northwest)–Mountain View, Palo Alto and Los Altos Cities	126.6	205.4	0.14	-0.05	1.26	0.84	0.77
Santa Clara (Northwest)–Sunnyvale and San Jose (North) Cities	104.6	102.2	0.03	-0.04	1.10	0.43	0.63
Santa Clara (Northwest)–San Jose (Northwest) and Santa Clara Cities	97.0	180.4	-0.01	-0.03	1.07	0.38	0.62
Santa Clara (North Central)–Milpitas and San Jose (Northeast) Cities	95.0	56.3	-0.03	-0.09	1.05	0.36	0.57
Santa Clara (North Central)–San Jose City (East Central) and Alum Rock	70.0	12.6	-0.06	-0.13	1.00	0.31	0.49
Santa Clara (East)–Gilroy, Morgan Hill and San Jose (South) Cities	66.8	29.4	-0.02	-0.10	1.00	1.10	0.89
Santa Clara (Southwest)–Cupertino, Saratoga Cities and Los Gatos Town	84.9	67.6	0.12	-0.11	1.20	0.81	0.71
Santa Clara (Central)–San Jose (West Central) and Campbell Cities	90.9	47.5	0.02	-0.09	1.09	0.31	0.55
Santa Clara (Central)–San Jose City (Northwest)	77.4	108.3	-0.06	-0.12	1.16	0.31	0.59
Santa Clara (Central)–San Jose City (Central)	99.3	45.3	-0.02	-0.13	1.11	0.31	0.50
Santa Clara (Central)–San Jose City (South Central/Branham) and Cambrian Park	74.1	24.6	-0.02	-0.13	1.06	0.31	0.46
Santa Clara (Central)–San Jose City (Southwest/Almaden Valley)	72.2	25.9	0.00	-0.16	1.12	0.31	0.46
Santa Clara (Central)–San Jose City (Southeast/Evergreen)	69.8	14.5	-0.03	-0.14	1.02	0.45	0.55
Santa Clara (Central)–San Jose City (East Central/East Valley)	61.4	20.5	-0.10	-0.14	1.05	0.31	0.53
Solano (Southwest)–Vallejo and Benicia Cities	88.3	40.9	-0.15	-0.04	0.86	0.52	0.50
Solano (Central)–Fairfield and Suisun City Cities	75.9	37.6	-0.13	-0.04	0.90	1.51	1.01
Solano (Northeast)–Vacaville and Dixon Cities	62.8	31.9	-0.12	-0.04	0.90	1.41	0.97
Sonoma (North)–Windsor Town, Healdsburg and Sonoma Cities	99.3	61.3	-0.02	0.02	0.92	2.58	1.46
Sonoma (South)–Petaluma, Rohnert Park and Cotati Cities	77.1	43.7	-0.06	-0.02	0.91	1.30	0.91
Sonoma (Central)–Santa Rosa City	109.1	74.1	-0.07	0.01	0.92	0.97	0.78

Note: The table shows the number of employed and retired residents (N_i) and workers (N_j^W) in thousands, residential and workplace amenities (E_i^R and E_j^W), total factor productivity in the traded good sector (Z_j) and floorspace supply elasticities of residential and commercial construction (ξ_i^S and ξ_i^C) in each model location.

B Extensions

B.1 Agglomeration Externalities

In the model presented in Section 2, the productivity of traded-good firms is an exogenous parameter Z_j . However, extensive empirical evidence suggests that local productivity depends on employment density and it has been common in quantitative urban models to model productivity as such. Following Ahlfeldt et al. (2015), we specify the productivity as

$$Z_{jt} = \bar{Z}_j L_{jt}^\lambda,$$

where $\bar{Z}_j$ is the exogenous component of productivity and λ is the elasticity of productivity with respect to the efficiency units of labor in a given location.

To examine the sensitivity of our results to endogenizing productivity, we follow existing empirical evidence (Ahlfeldt and Pietrostefani, 2019) and set $\lambda = 0.05$. Then we re-estimate the model. Model parameters are listed in Table B.1. After that, we perform the HSR and the upzoning counterfactuals. An attractive feature of a model with endogenous productivity is that it makes policy counterfactuals less reliant on exogenous productivity differences across locations. The variance of $\ln \bar{Z}_j$ goes down from 0.0109 in the model without agglomeration to 0.0088 in the model with agglomeration.


Table B.1: Internally calibrated parameters, model with agglomeration

Parameter	Description	Value	Target or source	Value
ρ	discount factor	0.0215	median wealth-earnings ratio	1.647
η	weight of housing in utility	0.3282	median rent-to-earnings	0.240
χ	preference for homeown.	1.0524	homeownership rate	0.542
ϑ_0	bequest motive	0.4940	homeown., ages 80–84 vs 20–24	0.642
ν^R	Gumbel scale, resid. shocks	2.0797	100× variance of log earnings	1.043
ν^W	Gumbel scale, work. shocks	0.2270	p90 commute time, minutes	67.0
κ	cost of commuting	0.0093	commuting gravity coeff.	-0.0408
μ^0	moving cost, intercept	12.5636	cross-county migration rate	0.0224
μ^a	moving cost, age coeff.	0.1523	migration, ages 20–24 vs 80–84	0.0853

Note: The table describes internally calibrated parameters in the model with agglomeration externalities. See the text for more details.

Figure B.1 compares local long-run changes in residents and jobs in the baseline version of the model that we described in Section 4 and the changes in the version with endogenous productivity. We can see that changes in jobs are somewhat larger in the version of the model with agglomeration externalities but the ranking of locations by job gains is

Figure B.1: Effect of agglomeration externalities on spatial reallocation



Note: The figure compares local changes in residents and jobs between the baseline version of the model and the version with endogenous productivity in the upzoning (panel A) and the HSR (panel B) counterfactuals. The solid diagonal line is the 45-degree line.

preserved. That is, including agglomeration externalities slightly amplifies the effects of the HSR and upzoning on job changes but does not lead to any qualitative differences. At the same time, changes in residents are only indirectly affected by agglomeration forces and, therefore, are nearly identical regardless of whether productivity is endogenous or not. We also checked that the difference between welfare gains in both types of counterfactuals is negligible.

B.2 Frequency of Shock Ages

In our baseline calibration, we set the frequency of shock ages to 1 year. In this section, we examine how sensitive our results are to the frequency of shock ages. To do this, we repeat our main analysis when shock ages are 1/2 years apart instead of 1 year. We first re-calibrate the model. Table B.2 shows the values of internally calibrated parameters. The calibration of the model with 1/2-year shock ages requires different moving costs to obtain the same implied annual migration rates.


We then run the same upzoning and HSR counterfactuals as in Section 4. Figure B.2 compares long-run reallocations of residents and jobs in both counterfactuals using each model. Our overall conclusion is that the results are very similar. The transition dynamics are slightly slower than in the benchmark model. The longer transitions occur because the “directedness” of moves is lower. One could recalibrate ν_R and ν_W to avoid this, but the change is small.

Table B.2: Internally calibrated parameters, shock ages 1/2 years apart

Parameter	Description	Value	Target or source	Value
ρ	discount factor	0.0222	median wealth-earnings ratio	1.647
η	weight of housing in utility	0.3229	median rent-to-earnings	0.240
χ	preference for homeown.	1.0551	homeownership rate	0.542
ϑ_0	bequest motive	0.3693	homeown., ages 80–84 vs 20–24	0.642
ν^R	Gumbel scale, resid. shocks	2.0061	100× variance of log earnings	1.043
ν^W	Gumbel scale, work. shocks	0.1827	p90 commute time, minutes	67.0
κ	cost of commuting	0.0075	commuting gravity coeff.	-0.0408
μ^0	moving cost, intercept	13.5481	cross-county migration rate	0.0112
μ^a	moving cost, age coeff.	0.1453	migration, ages 20–24 vs 80–84	0.0427

Note: The table describes internally calibrated parameters in the model where shock ages happen every half a year. See the text for more details.

Figure B.2: Spatial reallocation, main model vs model with shock ages 1/2 years apart



Note: The figure compares local changes in residents and jobs between the baseline version of the model and the version with shock ages 1/2 years apart in the upzoning (panel A) and the HSR (panel B) counterfactuals. The solid diagonal line is the 45-degree line.

B.3 No Homeownership Model

Existing quantitative urban models sometimes attempt to account for homeownership by assuming that housing is owned by local REITs, which redistribute rents to local residents (Redding and Rossi-Hansberg, 2017). Following this approach, we develop a version of our model that proxies for homeownership in this way but does not have tenure choice.

The household problem is similar to that in the main model, except that (i) households rent housing, and (ii) residents of location i receive a transfer $\mathcal{T}_{it}$ from REITs. There are three kinds of REITs: local residential REITs, a spatially diversified residential REIT, and

a commercial REIT. Local residential REITs own fraction $\mathcal{F}^{LS} = 0.6087$ of the residential real estate in location i , the same as the fraction of residential real estate value that is owner-occupied in the main model. The remaining residential real estate is owned by the spatially diversified residential REIT. As in the main model, the commercial REIT is also diversified across space.

As in the main model, the rent charged by REITs is such that the (ex-ante) return on real estate investments is q :

$$r_{it}^m = (q + \delta + \tau^h)p_{it}^m - \dot{p}_{it}^m$$

Note that this implies that, in the initial steady state, the present value of rents less taxes on a unit of floorspace equals the cost of constructing it:

$$p_i^{m,ss} = \int_0^\infty e^{-(q+\delta)s} (q + \delta) p_i^{m,ss} ds \quad (\text{B.1})$$

$$= \int_0^\infty e^{-(q+\delta)s} (r_i^{m,ss} - \tau^h p_i^{m,ss}) ds \quad (\text{B.2})$$

After an unexpected shock, the present value of pre-shock floorspace will in general not equal the cost paid to construct it:

$$p_i^{m,ss} H_i^{m,ss} = H_i^{m,ss} \int_0^\infty e^{-(q+\delta)s} (q + \delta) p_i^{m,ss} ds \quad (\text{B.3})$$

$$\neq H_i^{m,ss} \int_0^\infty e^{-(q+\delta)s} (r_{is}^m - \tau^h p_{is}^m) ds \quad (\text{B.4})$$

$$= H_i^{m,ss} \int_0^\infty e^{-(q+\delta)s} [(q + \delta) p_{is}^m - \dot{p}_{is}^m] ds \quad (\text{B.5})$$

Denote the excess return earned on sector- m floorspace in i after an unexpected shock by

$$\Pi_{it}^m = H_i^{m,ss} e^{-(q+\delta)t} [(q + \delta)(p_{it}^m - p_i^{m,ss}) - \dot{p}_{it}^m] \quad (\text{B.6})$$

Note that the present value of excess returns equals the capital gain caused by the shock:

$$\int_0^\infty \Pi_{it}^m dt = H_i^{m,ss} (p_{i0}^m - p_i^{m,ss}). \quad (\text{B.7})$$

As is standard in quantitative urban models, we assume that local residents own equal shares of their local residential REIT. The per-household transfer in location i from the

local residential REIT is therefore

$$\mathcal{T}_{it}^{LS} = \frac{\mathcal{F}^{LS} \Pi_{it}^S}{N_{it}^R} \quad (\text{B.8})$$

where N_{it}^R is the population of location i .

Households in the model own fraction $\mathcal{F}^{DS} = 0.0529$ of the spatially diversified residential REIT and fraction $\mathcal{F}^C = 0.0199$ of the commercial REIT. These are set to the same fractions of real estate value that are held by households in the main model. Households own equal shares of the spatially diversified REITs. The per-household transfer from the spatially diversified residential REIT is

$$\mathcal{T}_t^{DS} = \mathcal{F}^{DS} (1 - \mathcal{F}^{LS}) \sum_i \Pi_{it}^S. \quad (\text{B.9})$$

The per-household transfer from the commercial REIT is

$$\mathcal{T}_t^C = \mathcal{F}^C \sum_i \Pi_{it}^C. \quad (\text{B.10})$$

The total transfer from REITs in location i is

$$\mathcal{T}_{it} = \mathcal{T}_{it}^{LS} + \mathcal{T}_t^{DS} + \mathcal{T}_t^C \quad (\text{B.11})$$

We impose a preference for homeownership $\chi = 0$ so that it is never optimal to own a house of any size. We continue to allow individuals to save in the risk-free asset. We re-estimate the parameters of the quantitative model using the same approach as in Section 3, with two exceptions. First, we fix the parameters $(\rho, \vartheta_0, \vartheta_1)$, calibrated from wealth moments, to their benchmark model values. Second, we apply the calibrated value for the maximum owner-occupied house size (max IH) to rental housing in the no-ownership model to ensure that our results are not driven by differences in feasible housing consumption. All other parameters are re-estimated to give the no-homeownership model the best shot. Table B.3 reports the resulting parameters for the no-homeownership model.


Then we rerun the upzoning and the HSR counterfactuals. Figure B.3 shows that the magnitudes of local changes in residents and jobs between the old and the new steady states are nearly the same as in the main model. We also compare transition speeds in both models. In the main model, the transition of residential population halfway to the counterfactual steady state takes 9.5 years for an average location in the upzoning counterfactual and 10.4 years in the HSR counterfactual. In the no-homeownership model,

Table B.3: Internally calibrated parameters, model without homeownership

Parameter	Description	Value	Target or source	Value
η	weight of housing in utility	0.3091	median rent-to-earnings	0.240
ν^R	Gumbel scale, resid. shocks	2.2748	100× variance of log earnings	1.043
ν^W	Gumbel scale, work. shocks	0.2257	p90 commute time, minutes	67.0
κ	cost of commuting	0.0092	commuting gravity coeff.	-0.0408
μ^0	moving cost, intercept	13.7158	cross-county migration rate	0.0224
μ^a	moving cost, age coeff.	0.1708	migration, ages 20–24 vs 80–84	0.0853

Note: The table describes internally calibrated parameters of the model without homeownership. See the text for more details.

Figure B.3: Spatial reallocation in the main model and the model without homeownership



Note: The figure compares local changes in residents and jobs between the baseline version of the model and the version without homeownership in the upzoning (panel A) and the HSR (panel B) counterfactuals. The solid diagonal line is the 45-degree line.

these numbers are 9.7 and 10.1 years, respectively. That is, transitions have similar speeds in both types of models despite the fact that the model without homeownership is missing an important impediment to migration—the housing transaction cost. The reason is because the parameters of the no-ownership model are re-estimated to match the same empirical data moments. In particular, the parameters that describe moving costs (μ^0 and μ^a) and residential location preferences (ν^R) are higher than in the main model. The higher moving costs in the no-ownership model naturally slow down transitions in that model.

B.4 Hand-to-Mouth Model

Finally, we develop a version of our model where households are not allowed to save and borrow. The model is the same as the no-homeownership described above, except that households consume their income every period; there are no savings or bequests.

Table B.4: Internally calibrated parameters, hand-to-mouth model


Parameter	Description	Value	Target or source	Value
η	weight of housing in utility	0.2924	median rent-to-earnings	0.240
ν^R	Gumbel scale, resid. shocks	0.8746	100× variance of log earnings	1.043
ν^W	Gumbel scale, work. shocks	0.0543	p90 commute time, minutes	67.0
κ	cost of commuting	0.0022	commuting gravity coeff.	-0.0408
μ^0	moving cost, intercept	5.2742	cross-county migration rate	0.0224
μ^a	moving cost, age coeff.	0.0657	migration, ages 20–24 vs 80–84	0.0853

Note: The table describes internally calibrated parameters of the hand-to-mouth model. See the text for more details.

As Table B.4 demonstrates, not allowing for saving and borrowing behavior leads to substantial differences in parameter values. In particular, we calibrate much lower scale parameters for the distributions of residence and workplace preferences. When consumption smoothing is not possible, fundamental features of locations are more important than idiosyncratic preferences of households. We also calibrate much lower commuting and migration costs. When insurance against shocks is not available, it must be less costly for households to commute and relocate in order to rationalize the migration rates that we observe in the data.

Despite the differences in parameter values, the magnitudes of resident and job movements are reasonably similar in the upzoning counterfactual, as can be seen in panel A of Figure B.4. However, as shown in panel B, the model with hand-to-mouth agents strongly underpredicts the extent of spatial relocation in the HSR experiment.

Figure B.4: Spatial reallocation in the main model and the hand-to-mouth model



Note: The figure compares local changes in residents and jobs between the baseline version of the model and the the hand-to-mouth model in the upzoning (panel A) and the HSR (panel B) counterfactuals. The solid diagonal line is the 45-degree line.

C Computational Appendix

C.1 Solving the Household's Problem

We solve the household's problem using the finite differences method described by Achdou, Han, Lasry, Lions and Moll (2022). For each house size $h \in \mathbb{H}$, we make a grid for liquid wealth b with lower bound $\underline{b}(h)$ and upper bound $\bar{b}(h)$. We verify ex-post that the lower bounds are always less than or equal to the collateral constraint (that is, $\underline{b}(h) \leq \min_{i,t}\{-\phi p_{it}h\}$) and that the upper bounds do not bind. We use unevenly spaced grids for liquid wealth, with nodes concentrated near the lower bound. We discretize idiosyncratic individual labor productivity ζ using Rouwenhorst's method (Rouwenhorst, 1995; Kopecky and Suen, 2010). We discretize wages using an evenly spaced grid with lower bound $\underline{w}$ and upper bound $\bar{w}$. We verify ex-post that $\underline{w} \leq \min_{j,t} w_{jt}$ and $\bar{w} \geq \max_{j,t} w_{jt}$. Finally, we discretize age using an evenly spaced grid with intervals Δ .⁴³

Given the value function at the maximum age, $V_t(b, h, i, w, \zeta, A) = v(b + (1 - \psi)p_{it}h)$, we use the finite differences algorithm to solve the HJB equation (2.5) for $V_{t-\Delta}(b, h, i, w, \zeta, A - \Delta)$. Repeating this n times yields the value function immediately after the last shock age, $\lim_{t \downarrow 0} V_{t-1+t}(b, h, i, w, \zeta, A - 1 + t)$. Evaluating equations (2.6) - (2.9) yields the value function at the last shock age, $V_{t-1}(b, h, i, w, \zeta, A - 1)$.

We use linear interpolation to evaluate wage values between wage gridpoints for equation (2.6) and liquid wealth values between liquid wealth gridpoints for equations (2.7) and (2.8). The conditional expectation in equation (2.9) is calculated using the transition matrix given by Rouwenhorst's method. We iterate backwards in this fashion from age A to 0 to obtain the full discretized value and policy functions.

Given the density of state variables at age 0, $g_t(b, h, i, w, \zeta, 0)$ (described in Section 2.1.9), we use the finite differences algorithm to solve the Kolmogorov Forward equation (2.10) for $g_{t+\Delta}(b, h, i, w, \zeta, \Delta)$. Repeating this n times yields the density of state variables immediately before the first shock age, $\lim_{t \downarrow 0} V_{t+1-t}(b, h, i, w, \zeta, 1 - t)$. Evaluating equations (2.11) - (2.16) yields the density of state variables at the first shock age, $g_{t+1}(b, h, i, w, \zeta, 1)$. We use linear interpolation to assign mass that falls between wage and liquid wealth gridpoints to the adjoining nodes. We iterate forward in this fashion from age 0 to A to obtain the full discretized density of state variables.

⁴³We require $\Delta = 1/m$ for some integer m , so that shock ages occur every m steps on the age grid. In practice, we set $\Delta = 1$ in the main model and $\Delta = 1/2$ in the version of the model with shock ages every half a year.

C.2 Computing Stationary Equilibria

Given parameters, we compute stationary equilibria using the following algorithm:

1. Guess labor allocations L_i^0 and residential rents r_{Si}^0 . Total labor supply is normalized to 1, so this guess must satisfy $\sum_{i=1}^I L_i^0 = 1$ and $L_i^0 > 0$.
2. In a stationary equilibrium, construction is strictly positive and so $r_{Cj} = \bar{r}_{Cj}(H_{Cj}/H_{Cj0})^{1/\xi_j^C}$. Rearranging yields a supply equation for commercial floorspace: $H_{Cj} = (r_{Cj}/\bar{r}_{Cj})^{\xi_j^C} H_{Cj0}$. Subtract this from the commercial floorspace demand equation (2.23) to get the excess demand function

$$f(r_{Cj}) = [(1 - \alpha)Z_j/r_{Cj}]^{1/\alpha} L_j^0 - (r_{Cj}/\bar{r}_{Cj})^{\xi_j^C} H_{Cj0}$$

Solve these (independent) nonlinear equations for market-clearing commercial rents r_{Cj} .⁴⁴ Given H_{Cj} , compute wages using equation (2.17).

3. Use the algorithm described in Section C.1 to solve the household problem and compute the density of state variables. Then use equations (2.20) and (2.22) to compute labor and population allocations L_j and N_i^R . Since construction is strictly positive, residential rents are $r_{Si} = \bar{r}_{Si}(N_i^R/N_{i0}^R)^{1/\xi_i^S}$ (see equation 3.4).
4. If $\max_j |L_j^0 - L_j| < \epsilon$ and $\max_i |r_{Si}^0 - r_{Si}| < \epsilon$ for the numerical tolerance parameter $\epsilon > 0$, stop. Otherwise, update the guesses for L_j and r_{Si} using

$$\begin{aligned} L_j^0 &= L_j + \nabla(L_j - L_j^0), \\ r_{Si}^0 &= r_{Si} + \nabla(r_{Si} - r_{Si}^0). \end{aligned}$$

where $\nabla \in (0, 1]$ is a dampening parameter, and return to step (2). Note that by construction, L_i^0 always satisfies the criteria mentioned in step (1).

C.3 Computing Transitions

In this section, we describe our algorithm for computing transition dynamics after an unexpected shock. In all of our exercises, the economy is in an initial steady state at $t = 0$, and eventually converges to a new steady state after the shock. We discretize calendar time using an evenly spaced grid with lower bound 0, upper bound T , and interval length

⁴⁴Even when the number of locations is large, this can be done quickly using a standard nonlinear solver. We use Matlab's `fsolve` function.

equal to the one used to discretize age (Δ). We verify ex-post that T is sufficiently large that the economy has approximately converged to its new steady state by time T .

Our algorithm for computing transition dynamics is closely related to the one used to compute stationary equilibria. There are two main differences. The first is that we need to compute *paths* of prices, instead of a single price, for each market. The second is that, on a transition path, the construction irreversibility constraint may bind for a period of time in some locations. In this case, floorspace prices are not pinned down by construction costs, and we have to find prices at which demand equals the non-depreciated floorspace stock. For what follows, define the indicator functions $\mathbf{1}_{it}^m$, which is 1 if sector- m construction is strictly positive in i at time t , and 0 otherwise. Our algorithm for computing transition dynamics is as follows:

1. Compute the initial and final stationary equilibria using the algorithm described in Section C.2.
2. Guess labor allocations L_{it}^0 , residential rents r_{Sit}^0 , and the residential construction indicator $\mathbf{1}_{it}^{S0}$. Total labor supply is normalized to 1, so this guess must satisfy $\sum_{i=1}^I L_{it}^0 = 1$ and $L_{it}^0 > 0$.⁴⁵
3. Guess the commercial floorspace construction indicator $\mathbf{1}_{it}^{C0}$.
 - (a) Compute market-clearing commercial rents conditional on $\mathbf{1}_{it}^{C0}$. Commercial floorspace demand, given by equation (2.23), is $H_{Cit}(r_{Cjt}) = [(1 - \alpha)Z_j/r_{Cjt}]L_{jt}^0$. Commercial floorspace supply satisfies equation (3.5) when construction is positive, and falls at the rate of depreciation when construction is 0:

$$H_{Cit}^{\text{sup}}(r_{Cit}) = \begin{cases} (r_{Cit}/\bar{r}_{Ci})^{\xi_i^C} H_{Ci0}^{\text{ss}} & \text{if } \mathbf{1}_{it}^{C0} = 1, \\ H_{Ci0}^{\text{ss}} & \text{if } \mathbf{1}_{it}^{C0} = 0 \text{ and } t = 0, \\ (1 - \delta\Delta)H_{Cit-\Delta}^{\text{sup}}(r_{Cit}) & \text{if } \mathbf{1}_{it}^{C0} = 0 \text{ and } t > 0. \end{cases}$$

The excess demand function implied by $\mathbf{1}_{it}^{C0}$ is:

$$f(r_{Cjt}) = H_{Cit}(r_{Cjt}) - H_{Cit}^{\text{sup}}(r_{Cit})$$

Solve this system of equations for market-clearing commercial rents r_{Cjt} . Even

⁴⁵We use the initial guess $\mathbf{1}_{it}^{S0} = 1$ and $x_{jt}^0 = e^{-\iota t}x_{j0}^{\text{ss}} + (1 - e^{-\iota t})x_{jT}^{\text{ss}}$ for $x \in \{L, r_S\}$, where x_0^{ss} and x_T^{ss} indicate initial- and long-run steady state values, respectively. This ensures that our guess converges to the long-run equilibrium. We find that $\iota = 0.2$ yields a good initial guess.

when the number of locations and time periods is large, this can be done quickly using a standard nonlinear solver. We use Matlab's fsolve function.

- (b) Given H_{Cit} and L_{it}^0 , compute wages using equation (2.17).
- (c) Update the commercial construction indicator. First compute commercial construction demand:

$$Y_{Cit}^h = \begin{cases} H_{Ci0} - H_{Ci0}^{ss} & \text{if } t = 0, \\ (H_{Cit} - H_{Cit-\Delta})/\Delta + \delta\Delta H_{Cit-\Delta} & \text{if } t > 0. \end{cases}$$

Then update the commercial construction indicator using

$$\mathbf{1}_{it}^C = \begin{cases} 1 & \text{if } \mathbf{1}_{it}^{C0} = 0 \text{ and } p_{Cit} > 1/Z_{Cit}^h, \\ 0 & \text{if } \mathbf{1}_{it}^{C0} = 1 \text{ and } Y_{Cit}^h < 0, \\ \mathbf{1}_{it}^{C0} & \text{otherwise.} \end{cases}$$

In the first case, the price at which demand equals supply with no construction exceeds the construction cost $1/Z_{Cit}^h$. This is not an equilibrium outcome, as commercial construction firms would strictly prefer to construct more floorspace. In the second case, when commercial rents are determined by construction costs ($r_{Cit} = r_{Cit}(1/Z_{Cit}^h)$) negative construction is demanded. This is also not an equilibrium outcome because it violates the irreversible construction constraint. In these cases, and only these cases, the construction indicator needs to be changed.

If $\mathbf{1}_{it}^C = \mathbf{1}_{it}^{C0}$ for all (i, t) , proceed to step (4). Otherwise, set $\mathbf{1}_{it}^{C0} = \mathbf{1}_{it}^C$ and return to step (3a).

4. Given the value function at time T from the final steady state, use the finite differences algorithm to solve the HJB equation (2.5) for the value function at time $T - \Delta$. After computing the HJB equation, evaluate equations (2.6) - (2.9) at all shock ages. Iterate backward from $t = T$ to $t = 0$ to obtain the value and policy functions over the entire transition.

Given the density of state variables at time 0 from the initial steady state, use the finite differences algorithm to solve the Kolmogorov Forward equation (2.10) for the density of state variables at time Δ . Evaluate equations (2.11) - (2.16) at all shock ages. Iterate forward from $t = 0$ to $t = T$ to obtain the density of state variables over the entire transition. Then use equations (2.20) - (2.22) to compute labor allocations

L_{it} , population allocations N_{it}^R , and residential floorspace demands H_{Sit} .

5. Residential floorspace supply equals demand when construction is positive, and falls at the rate of depreciation when construction is 0. Compute the residential floorspace supply implied by $\mathbf{1}_{it}^{S0}$:

$$H_{Sit}^{\text{sup}} = \begin{cases} H_{Sit} & \text{if } \mathbf{1}_{it}^{S0} = 1, \\ H_{Si0}^{\text{ss}} & \text{if } \mathbf{1}_{it}^{S0} = 0 \text{ and } t = 0, \\ (1 - \delta\Delta)H_{Sit-\Delta}^{\text{sup}} & \text{if } \mathbf{1}_{it}^{S0} = 0 \text{ and } t > 0. \end{cases}$$

Residential floorspace market clearing requires that supply equals demand and $r_{Sit} = \bar{r}_{Si}(N_{it}/N_{i0}^{\text{ss}})^{1/\xi_i^S}$ if construction is positive. Compute the market-clearing residential rent implied by $\mathbf{1}_{it}^{S0}$ and r_{Sit}^0 :

$$r_{Sit} = \begin{cases} r_{Sit}^0 H_{Sit}/H_{Sit}^{\text{sup}} & \text{if } \mathbf{1}_{it}^{S0} = 0, \\ \bar{r}_{Si}(N_{it}/N_{i0}^{\text{ss}})^{1/\xi_i^S} & \text{if } \mathbf{1}_{it}^{S0} = 1. \end{cases}$$

If $\max_{i,t} |L_{it}^0 - L_{it}| < \epsilon$ and $\max_{i,t} |r_{Sit}^0 - r_{Sit}| < \epsilon$ for the numerical tolerance parameter $\epsilon > 0$, proceed to step (6). Otherwise, update the guesses for L_{it} and r_{Sit} using

$$\begin{aligned} L_{it}^0 &= L_{it}^0 + \nabla(L_{it} - L_{it}^0), \\ r_{Sit}^0 &= r_{Sit}^0 + \nabla(r_{Sit} - r_{Sit}^0). \end{aligned}$$

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

6. Compute residential construction demand:

$$Y_{Sit}^h = \begin{cases} H_{Si0} - H_{Si0}^{\text{ss}} & \text{if } t = 0, \\ (H_{Sit} - H_{Sit-\Delta})/\Delta + \delta\Delta H_{Sit-\Delta} & \text{if } t > 0. \end{cases}$$

Then update the residential construction indicator using the same logic described in step 3(c):

$$\mathbf{1}_{it}^S = \begin{cases} 1 & \text{if } \mathbf{1}_{it}^{S0} = 0 \text{ and } p_{Sit} > 1/Z_{Sit}^h, \\ 0 & \text{if } \mathbf{1}_{it}^{S0} = 1 \text{ and } Y_{Sit}^h < 0, \\ \mathbf{1}_{it}^{S0} & \text{otherwise.} \end{cases}$$

If $\mathbf{1}_{it}^S = \mathbf{1}_{it}^{S0}$ for all (i, t) , stop. Otherwise, set $\mathbf{1}_{it}^{S0} = \mathbf{1}_{it}^S$ and return to step (3).

C.4 Estimation

In this section, we describe our algorithm for estimating the model. The estimated parameters can be grouped into three categories: the vector of economy-wide parameters $\Theta = (\rho, \eta, \chi, \vartheta_0, \nu^R, \nu^W, \kappa, \mu^0, \mu^a)$, the vectors of location-specific amenities (E_i^R, E_j^W) , and the vectors of location-specific productivities (Z_j, Z_i^S, Z_j^C) . In total, there are $9 + 5I = 284$ parameters to estimate. However, as discussed below, productivities can be read directly from data. In addition, without loss of generality, we normalize each amenity in location 1 to 0. As a result, there are $9 + 2(I - 1) = 117$ parameters that must be numerically estimated.

As discussed in Section 3.2, there is one target for each estimated parameter. In principle, we could use a derivative-based method to find parameter values that match these targets. However, given the large number of parameters to estimate, this would be prohibitively expensive. Instead, we take advantage of two facts to develop a more efficient algorithm.

First, it turns out that the economy-wide parameters Θ and location-specific amenities (E_i^R, E_j^W) have virtually independent effects on the objective function. Specifically, Θ affects the moments listed in Table 1 but has little effect on labor and population allocations. In contrast, amenities affect labor and population allocations but have little effect on the moments listed in Table 1.

Second, labor and population allocations respond smoothly to changes in amenities. As a result, we can estimate the relatively small vector Θ using a derivative-based algorithm, and the much larger vectors of amenities using a derivative-free method. Our estimation algorithm is as follows:

1. Guess (Θ, E_i^R, E_j^W) .
2. Use the algorithm described in Section C.1 to solve steady state conditional on parameters and observed prices. If labor and population allocations match the data to within a numerical tolerance of $\epsilon > 0$, proceed to step (3). Otherwise, update amenities using

$$\begin{aligned} E_i^R &= E_i^R + \nabla(N_i^{R\text{data}} - N_i^R) \\ E_j^W &= E_j^W + \nabla(N_j^{W\text{data}} - N_j^W) \end{aligned}$$

for $i, j > 1$ where $\nabla > 0$ is an ad-hoc updating parameter, and repeat this step. In the previous expression, N_i^R is residential population in location i and N_j^W is employment in location j .

3. Use a Newton-based method to find a Θ that matches the estimation targets listed in Table 3.2 to within ϵ . We use Matlab's `fsolve` function. For increased efficiency, use parallelized code to compute the Jacobian at each updating step.
4. If labor and population allocations match the data to within ϵ for the estimated Θ , proceed to step (5). Otherwise, return to step (2).
5. Use equation (2.17), (3.4), and (3.5) to back out productivities (Z_j, Z_i^S, Z_j^C) that match observed wages, residential rents, and commercial rents given model-implied labor allocations and floorspace quantities.

C.5 Computing Time

Table C.1 shows the amount of time required to complete each of the steps of our quantitative exercises. These include estimating the model, computing the long-run steady state after an unexpected shock, and computing the transition path between the initial and long-run steady states. In the first column, we also report the amount of time it takes to compute the steady state of the model given parameters and prices (that is, a partial equilibrium steady state). The model only needs to be estimated once, as the initial steady state is the same for both of our main counterfactuals. The time required to compute the long-run steady state and transition dynamics is similar for the transportation improvement and upzoning counterfactuals. It takes just over 45 minutes to estimate the model, and a little less than 3 hours to solve the counterfactuals we consider.

Table C.1: Computation Time

	Steady State: P.E.	Estimation	Steady State: GE	Transition
Transportation	27.21 sec.	46.60 min.	7.18 min.	2.85 hours
Upzoning	" "	" "	8.56 min.	2.50 hours

Note: The table shows the time required to complete the tasks required for our quantitative exercises. Each of these steps was performed using an Amazon Elastic Compute Cloud (EC2) m5.16xlarge instance. This instance has 64 virtual CPUs and 256 GB of RAM.

In Table C.2, we compare the computational time required to solve a partial equilibrium steady state of the model in commuting zones with varying numbers of locations. For commuting zones with fewer than 55 locations, this task takes less than 20 seconds on a personal computer. For the largest commuting zone (New York, which has 183 locations and 33,489 locations pairs), solving a partial equilibrium steady state requires a little less than 3 minutes. Figure 2 shows the relationship between computation time for the partial

equilibrium steady-state and number of locations. Importantly, even though the size of the state space grows quadratically with the number of locations, computational time grows nearly linearly. This is due to the fact that wage is a sufficient state variable for workplace between shock ages. As a result, it is possible to use our model to study cities with tens of thousands of location pairs.

Table C.2: Time to Compute Steady State: Partial Equilibrium

Commuting zone	Locations	Location pairs	Time to solve
Portland	18	324	7.63 sec.
Seattle	33	1,089	11.44 sec.
San Francisco	55	3,025	20.29 sec.
Los Angeles	123	15,129	1.13 min.
New York	183	33,489	2.91 min.

Note: The table shows the time required to compute a steady state of the model given prices and parameters for cities with varying numbers of locations. These times were obtained using a MacBook Pro laptop with 1.4 GHz Intel core i5 processor.