Dynamic Urban Economics*

Brian Greaney†
Andrii Parkhomenko‡
Stijn Van Nieuwerburgh§

December 18, 2024

Abstract

We build a dynamic model of an urban area which combines key features of quantitative spatial models and macro-housing models. It integrates a large number of locations, forward-looking households, commuting, costly migration, uninsurable income risk, and housing tenure choice. We cast the model in continuous time, while shocks arrive and some choices are taken at discrete time intervals. This "mixed time" approach allows us to efficiently compute both steady state equilibrium and transition dynamics, even when there are thousands of location pairs. Using a quantitative model of the San Francisco Bay Area, we show how accounting for forward-looking behavior, spatial frictions, and transition dynamics changes the estimated effects of spatially heterogeneous shocks and policies that have traditionally been studied with static models.

Key Words: urban, dynamics, spatial equilibrium, tenure choice, life cycle, house prices.

JEL Codes: C63, G11, J61, R10, R21, R23, R31, R52

*We thank our discussants David Albouy, Adrien Bilal, Carlos Hurtado, Tomoya Mori, and Miguel Zerezero, as well as Milena Almagro, Lucas Conwell, Diego Daruich, Fabian Eckert, Elisa Giannone, Selo Imrohoroglu, Paolo Martellini, Conor Walsh, and seminar and conference participants at USC, the University of Wisconsin at Madison, the West Coast Spatial Workshop at UCSD, the University of Washington, the AREUEA National Conference, the UEA European meeting in Copenhagen, the IEB Workshop in Urban Economics in Barcelona, the SED meeting in Barcelona, the WEAI conference in Seattle, and the NBER Summer Institute Urban Economics. First draft: July 16, 2024.

†University of Washington.
‡University of Southern California.
§Columbia University and NBER.

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) framework 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.

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 residence for some households and workplace 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. 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 forty 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 of much larger urban areas or even entire countries.

The second contribution of this paper is to show how accounting for forward-looking behavior, 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 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 second experiment consists of an increase in housing supply ("upzoning") in locations with below-the-median construction productivity.

We find that the HSR attracts more residents and jobs to most locations where stations are built, as well as to some adjacent areas. However, due to migration frictions and the durability of structures, spatial reallocation is gradual and the full transition takes over 50 years. 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. 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.

The upzoning experiment has similar effects on population and employment in treated locations as the HSR experiment, but for a different reason: upzoned locations are better places for residents, due to more abundant housing, and for firms, due to more abundant workers nearby. However, the welfare implications are different. The gains for renters, especially the young, are much larger than for homeowners, whose housing wealth falls as house prices decline. Older and less productive homeowners experience sizable welfare losses, as housing accounts for more of their lifetime wealth.

We show that sluggish adjustment to shocks, as well as reallocation costs, mean that

welfare gains that account for transition dynamics are lower than the gains calculated from comparing two steady states. There is also rich heterogeneity of impacts across locations, so our results cannot be summarized by a model with stylized geography. Our results suggest that urban policy evaluation with a model that abstracts from forward-looking agents, homeownership, transition dynamics, or realistic geography 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 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), 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 labour 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, Mabille, 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 include 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),

[1] See 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.

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) use models with two types: 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 and, as a result, obtains closed-form solutions for migration and saving decisions. The key to tractability 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) and Greaney (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.[3] 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.[4] In contrast to our approach, his method can handle aggregate risk.[5] On the other hand, we

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

[3]Examples include Sun (2024) who develops a deep learning method to solve a dynamic spatial equilibrium model with a large number of locations.

[4]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.

[5]We require that aggregate uncertainty be resolved within a finite time horizon. There *can* be uncertainty

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.[6]

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.

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.

[6]Irreversible construction is necessary to study situations in which the durability of real estate influences 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

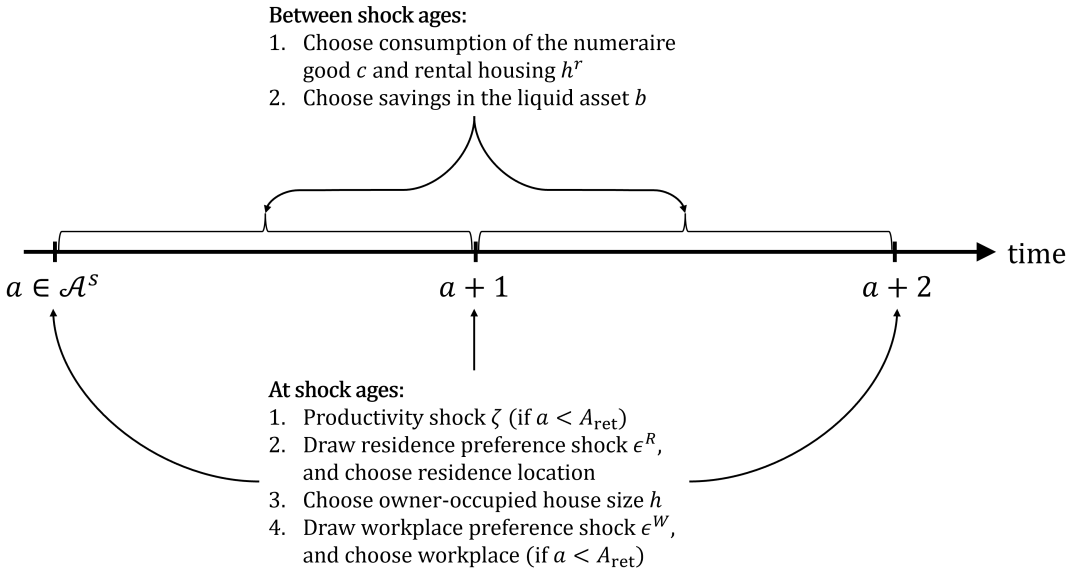

2.1.1 Timing

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 $\mathcal{A}^s \equiv \{0, 1, ..., A-1\}$, and the rest of the lifespan. As described below, households receive a series of shocks at shock ages, as well as make location and owner-occupied house size choices.[7] The timeline of household decisions is illustrated in Figure 1. In the quantitative model that we describe in Section 3, the interval between shock ages is one year.

We assume that the age distribution is uniform. As a result, at any moment in time there is a 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

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

[7]This timing assumption is made primarily for computational reasons, discussed in Section 2.5.

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

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

Households discount the future at rate ρ.

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 $\overline{z}(a)$:

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

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 $\mathfrak{b}(\zeta)$ that is a function of their labor productivity at the time of retirement. Earnings are subject to a payroll tax of rate τ_z. Payroll tax revenues go toward social security transfers and wasteful government spending. We denote the sum of after-tax earnings and pensions by

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

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 is

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

Housing services per unit vary between rented and owner-occupied housing, and the parameter χ captures the non-pecuniary benefits from owner-occupied housing. Renting and owning simultaneously is prohibited, i.e., an owner must have $h^r = 0$ and a renter must have $h = 0$.

Housing is non-tradable across space, so its price varies across locations. Rental prices are denoted by r_{it} and 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.

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 can save and borrow in a risk-free, liquid asset b at interest rate q. We assume that the interest rate q is set externally. Whenever a household purchases a house, it must satisfy a collateral constraint which prohibits borrowing more than a fraction ϕ of

its housing wealth:

$$b \geq -\phi p_{it} h. \tag{2.2}$$

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

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

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 renters are not allowed to borrow.[8]

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

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

At shock ages, liquid wealth changes discontinuously for homeowners who change residential location and/or households who adjust owner-occupied house size (recall that owner-occupied housing must be sold prior to changing residential location). Immediately after choosing residential location i', the liquid wealth and owner-occupied housing of a household with initial states (b, h, i) are:[9]

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

After choosing residential location, the household chooses owner-occupied housing.[10] Immediately after adjusting owner-occupied house size to h', the liquid wealth of a household with initial 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 shock-age household with state (b, h, i) is

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

[8] Our framework can accomodate unsecured borrowing without difficulty. We abstract from it in our baseline model for expositional clarity.

[9] In our calibration, $\psi < 1 - \phi$, so a sale is feasible for every residential location choice.

[10] Unlike Greaney (2023), we only allow adjustments of owner-occupied housing at shock ages.

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:

$$\begin{aligned} \rho V_t(\Omega) = \max_{c, h^r} & \, u(c, \mathbf{h}(h, h^r)) + \partial_b V_t(\Omega)\dot{b}_t(\Omega, c, h^r) + \partial_a V_t(\Omega) + \partial_t V_t(\Omega) \\ & \text{s.t. } \dot{b}_t(\Omega, c, h^r) \geq 0 \text{ if } b \leq -\phi p_{it}h, \\ & h^r = 0 \text{ if } h > 0. \end{aligned} \tag{2.4}$$

The first term of the right-hand side of (2.4) 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

$$V_t^W(\Omega) = \lim_{\iota \downarrow 0} \begin{cases} \nu^W \log\left(\sum_j \exp([V_{t+\iota}(b, h, i, \zeta, j, a + \iota) + E_j^W - d_{ij}]/\nu^W)\right) & \text{if } a < A_{\textbf{ret}}, \\ V_{t+\iota}(b, h, i, \zeta, a + \iota) & \text{otherwise.} \end{cases} \tag{2.5}$$

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

$$\begin{aligned} V_t^H(\Omega) =& V_t^W(\tilde{b}_t^H(\Omega, \tilde{h}_t(\Omega)), \tilde{h}_t(\Omega), i, \zeta, a), \\ \tilde{h}_t(\Omega) =& \underset{h' \in \mathbb{H}_t(\Omega)}{\text{argmax}} V_t^W(\tilde{b}_t^H(\Omega, h'), h', i, \zeta, a). \end{aligned} \tag{2.6}$$

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

$$V_t^R(\Omega) = \nu^R \log\left(\sum_{i'} \exp\left([V_t^H(\tilde{b}_t^R(\Omega, i'), \tilde{h}^R(\Omega, i'), i', \zeta, a) + E_{i'}^R - \mu_{ii'}(a))]/\nu^R\right)\right). \tag{2.7}$$

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

$$V_t(\Omega) = \begin{cases} \int V_t^R(b, h, i, \zeta', a)\, f(\zeta'|\zeta)\, \mathrm{d}\zeta' & \text{if } a < A_{\mathbf{ret}}, \\ V_t^R(\Omega) & \text{otherwise} \end{cases} \tag{2.8}$$

where $f(\zeta'|\zeta)$ is the conditional probability density function of individual productivity ζ' (that is, the normal p.d.f. with mean $\theta_z\zeta$ and variance σ_z^2).

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 instead leave the economy. Finally, new-born 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:

$$\partial_t g_t(\Omega) = -\partial_b[\dot{b}_t(\Omega) g_t(\Omega)] - \partial_a g_t(\Omega), \tag{2.9}$$

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:

$$g_t^z(\Omega) = \begin{cases} \sum_j \int \lim_{\iota \downarrow 0} g_{t-\iota}(b, h, i, \zeta', j, a - \iota)\, f(\zeta|\zeta')\, \mathrm{d}\zeta' & \text{if } a < A_{\mathbf{ret}}, \\ \lim_{\iota \downarrow 0} g_{t-\iota}(b, h, i, \zeta, a - \iota) & \text{if } a \geq A_{\mathbf{ret}}. \end{cases} \tag{2.10}$$

After productivity shocks are observed, households choose their residential location.

The density of state variables after residential locations are chosen is:

$$g_t^R(\Omega) = \sum_{i'} \int \int \mathbf{1}_t^R(b', h', i'; b, h, i)\, \pi_t^R(b', h', i', \zeta, a; i)\, g_t^z(b', h', i', \zeta, a)\, \mathrm{d}b' \mathrm{d}h', \tag{2.11}$$

where

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

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

$$\pi_t^R(\Omega; i') = \frac{\exp([V_t^H(\tilde{b}_t^R(\Omega, i'), \tilde{h}^R(\Omega, i'), i', \zeta, a) + E_{i'}^R - \mu_{ii'}(a))]/\nu^R)}{\sum_{i''} \exp([V_t^H(\tilde{b}_t^R(\Omega, i''), \tilde{h}^R(\Omega, i''), i'', \zeta, a) + E_{i''}^R - \mu_{ii''}(a)]/\nu^R)} \tag{2.12}$$

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:

$$g_t^H(\Omega) = \int \int \mathbf{1}_t^H(b', h', i, \zeta, a; b, h)\, g_t^R(b', h', i, \zeta, a)\, \mathrm{d}b' \mathrm{d}h', \tag{2.13}$$

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:

$$g_t(\Omega) = \begin{cases} \pi_t^W(b, h, i, \zeta, a; j)\, g_t^H(b, h, i, \zeta, a) & \text{if } a < A_{\mathbf{ret}}, \\ g_t^H(\Omega) & \text{otherwise.} \end{cases} \tag{2.14}$$

where

$$\pi_t^W(\Omega; j) = \lim_{\iota \downarrow 0} \frac{\exp([V_{t+\iota}(b, h, i, \zeta, j, a) + E_j^W - d_{ij}]/\nu^W)}{\sum_{j'} \exp([V_{t+\iota}(b, h, i, \zeta, j', a) + E_{j'}^W - d_{ij'}]/\nu^W)} \tag{2.15}$$

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, Redding, Sturm, and Wolf (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.[11]

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:

$$w_{jt} = \alpha Z_j L_{jt}^{\alpha-1} H_{Cjt}^{1-\alpha}, \tag{2.16}$$

$$r_{Cjt} = (1-\alpha) Z_j L_{jt}^{\alpha} H_{Cjt}^{-\alpha}. \tag{2.17}$$

Rearranging equation (2.16), labor demand can be written as:

$$L_{jt} = \left(\frac{\alpha Z_j}{w_{jt}} \right)^{\frac{1}{1-\alpha}} H_{Cjt} \tag{2.18}$$

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

$$L_{jt} = \sum_i \int \int \int \int e^{\bar{z}(a)+\zeta} g_t(b, h, i, \zeta, j, a) \, \mathrm{d}b \, \mathrm{d}h \, \mathrm{d}\zeta \, \mathrm{d}a. \tag{2.19}$$

2.2.2 Developers

There are two types of developers that specialize in building one of the two types of floorspace, indexed by m: residential ($m = S$) and commercial ($m = C$). Developers

[11] Here, Z_j is an exogenous parameter. In Appendix Section B.1, we endogenize Z_j as a function of local employment density and show that this has a minor effect on the quantitative results presented in Section 4.

construct new floorspace according to the production function

$$Y^h_{mit} = Z^h_{mit} K_{mit},$$

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

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

$$H_{Sit} = \sum_j \int \int \int \int [h^r_t(b, h, i, \zeta, j, a) + h] \, g_t(b, h, i, \zeta, j, a) \, \mathrm{d}b \, \mathrm{d}h \, \mathrm{d}\zeta \, \mathrm{d}a. \tag{2.20}$$

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

$$N_{it} = \sum_j \int \int \int \int g_t(b, h, i, \zeta, j, a) \, \mathrm{d}b \, \mathrm{d}h \, \mathrm{d}\zeta \, \mathrm{d}a. \tag{2.21}$$

Rearranging equation (2.17), commercial floorspace demand can be written as

$$H_{Cjt} = \left[\frac{(1 - \alpha) Z_j}{r_{Cjt}} \right]^{\frac{1}{\alpha}} L_{jt} \tag{2.22}$$

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

$$\dot{H}_{mit} = Y^h_{mit} - \delta H_{mit}, \tag{2.23}$$

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

Denote floorspace construction costs by

$$\hat{p}_{mit} = \frac{1}{Z^h_{mit}}.$$

If floorspace price p_{mit} equals construction cost $\hat{p}_{mit}$, 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:

$$Y^h_{mit} = \begin{cases} 0 & \text{if } p_{mit} < \hat{p}_{mit}, \\ \max\{\dot{H}_{mit} + \delta H_{mit}, 0\} & \text{if } p_{mit} = \hat{p}_{mit}, \\ \infty & \text{if } p_{mit} > \hat{p}_{mit}. \end{cases} \tag{2.24}$$

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

$$\begin{aligned} p_{mit} = \hat{p}_{mit} \text{ and } Y^h_{mit} > 0 \\ \text{or } p_{mit} < \hat{p}_{mit} \text{ and } Y^h_{mit} = 0. \end{aligned} \tag{2.25}$$

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: $Y^h_{mit} = \delta H_{mit} > 0$. 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). 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:

$$r_{mit} = (\delta + q + \tau_h)p_{mit} - \dot{p}_{mit}. \tag{2.26}$$

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 (2.4)–(2.8) are satisfied.
2. The density of state variables is consistent with household optimization: equations (2.9)–(2.15) are satisfied.

3. Firms optimize: equations (2.16)–(2.17) hold.
4. Labor markets clear in all locations.
5. Residential and commercial rental space markets clear: equation (2.26) holds for $m = S, C$ in all locations.
6. Residential and commercial ownership space markets clear: equation (2.25) 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.[12] Finally, fixed housing adjustment costs imply that households face a stopping time problem of when to adjust housing. The 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 using standard 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).[13] It allows us to make use of both efficient discrete- and continuous-time numerical methods. As is common in quantitative urban models, the fact

[12]Our 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 $2I$ local amenities (residential and workplace in each location).

[13]In contrast to this paper, Greaney (2023) allows housing adjustment at any time. This, however, implies that the number of locations that the model can handle is limited. Greaney (2023)'s model does not have commuting and has 50 locations.

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 2.5, 2.7, 2.12, and 2.15). Between shock ages, location and owner-occupied housing are fixed, and the agent's problem reduces to a simple consumption-saving choice (equations 2.4 and 2.9). 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.[14] 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.[15] 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 some cases, orders of magnitude faster) in continuous time than in discrete time. The speed gains are greatest in models with nonconvexities, such as ours.

Limiting idiosyncratic shocks and discrete choices to shock ages has two main advantages. First, it allows us to represent the HJB and KF equations (2.4 and 2.9, respectively), as sparse matrix equations. This matters because the efficiency of Achdou, Han, Lasry, Lions, and Moll (2022)'s algorithm depends on sparsity. Suppose instead that the opportunity to change residential location arrived according to a Poisson process, so that residential moves could occur at any time.[16] The HJB and KF equations would then contain I additional terms, one for each possible location choice. This significantly reduces the

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

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

[16] This assumption is made by Crews (2023), Bilal (2023), and Bilal and Rossi-Hansberg (2023) in their continuous-time spatial models.

sparsity when the number of locations is large, making the model intractable.[17] Second, it allows us to replace workplace with a lower-dimensional state variable. Since workplace is freely chosen and workplace amenities are enjoyed and commuting costs are paid at shock ages, the only workplace characteristic that is relevant between shock ages is the wage. As a result, we can replace the discrete state variable workplace with the continuous state variable wage. Wage can be discretized on a grid whose cardinality is independent of the number of locations. This is the key insight which keeps the state space size manageable even when there are many location pairs.

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 typically 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. Importantly, this implies that homeowners who receive adverse shocks can avoid default by moving to a cheaper location and/or changing tenure status.[18]

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).[19] Figure 2 shows the relationship between the number of locations and time to solve equilibrium.[20] Solving the model for a small CZ such as Portland takes just 12 seconds, while New York takes 252 seconds. Solving it for the San Francisco Bay Area, the CZ that we use for our quantitative analysis in Sections 3 and 4, takes 40 seconds. 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,

[17] Theoretically, this issue can be avoided by employing the "explicit" updating method described by Achdou, Han, Lasry, Lions, and Moll (2022). However, this method is less numerically stable and requires finer age and time grids than the "implicit" updating method. Since our shock age assumption ensures sparsity, the implicit method is both more accurate and efficient for solving our model.

[18] This statement is conditional on the size of the shocks, but correct for the shock sizes we consider.

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

[20] For this demonstration, we used a MacBook Pro laptop computer with a 1.4 GHz Intel core 1.5 processor.

Figure 2: Number of locations and computational time

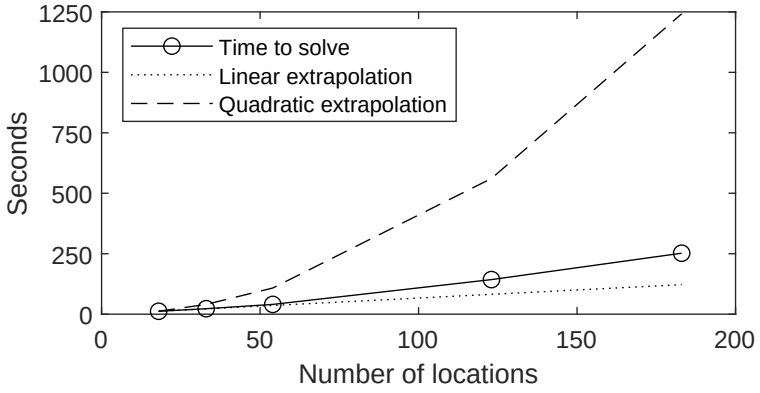

Note: The solid line shows the relationship between computational time required to solve the stationary equilibrium and the number of locations in an urban area. The lower dotted line shows extrapolated (using the two smallest CZs) computational time if it were a linear function of the number of locations.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

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.

3 Application

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

3.1 Data

Model locations. Following Delventhal and Parkhomenko (2023), we define a model location as the intersection of a Census Public Use Microdata Area (PUMA) and a county.[21] 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.[22] This

[21]PUMA 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.

[22]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

number of locations implies 3,025 residence-workplace location pairs.

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 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 price-rent ratio in the entire Bay Area. We calculate the price-rent ratio for the Bay Area as follows. We build hedonic price and rent indices from the 2012–2016 ACS data of individual observations, controlling for property characteristics. Then we divide the price index by the rent index evaluated at a house of median size, and obtain the value of 26.67.

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 dataprovider Compstak. The correlation between

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. For more information, see the website of the Association of Bay Area Governments: https://abag.ca.gov/.

commercial and residential rents across model locations is 90%. Appendix A contains more details.

Commute times. We use 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 summarized 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.

Preferences. We calibrate the weight of housing in utility η to match the ratio of housing expenditure to earnings in the Bay Area. This share is 0.24 and includes both expenditures of renters and imputed rents of homeowners. The calibrated η is 0.3146. We calibrate the preference for homeownership χ to match the homeownership rate. In the San Francisco Bay Area, the homeownership rate is 54.2% and the calibrated value of χ is 1.0514.

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.[23] Then we calibrate ν^R to match the variance across locations of log average earnings (0.01043) and obtain $\nu^R = 2.4559$.

Parameter ν^W determines how far workers are willing to commute. If idiosyncratic preferences are strong, then workers should be relatively insensitive to commute times

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

Table 1: Model parameters

Parameter	Description	Value	Target or source	Value
		Internally calibrated		
ρ	discount factor	0.0198	median wealth-earnings ratio	1.647
η	weight of housing in utility	0.3146	median rent-earnings	0.24
χ	preference for homeown.	1.0514	homeownership rate	0.540
ϑ_0	bequest motive	0.3669	homeown., ages 80–84 vs 20–24	0.643
ϑ_1	bequest motive	0.4679	p75/p25 wealth ratio, age 65	3.321
ν^R	Gumbel scale, resid. shocks	2.4559	100× variance of log earnings	1.013
ν^W	Gumbel scale, work. shocks	0.2554	p90 commute time, minutes	67.0
κ	cost of commuting	0.0104	commuting gravity coeff.	-0.0408
μ^0	moving cost, intercept	14.8961	cross-county migration rate	0.0224
μ^a	moving cost, age coeff.	0.1734	migration, ages 20–24 vs 80–84	0.0827
$\max \mathbb{H}^r$	largest rental size	4.5291	p90 rental size / p10 own. size	1.32
$\min \mathbb{H}$	smallest own. size	3.4256	p90 own. size / p10 own. size	2.64
$\max \mathbb{H}$	largest own. size	9.0582	frac. owners in <p50 own. size	0.5
		Externally calibrated		
q	interest rate	0.02	data	
γ	risk aversion	2	standard value	
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	0.0239	BEA	
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	

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

when choosing their workplace. On the other hand, if the preferences are weak, 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 $\nu^W = 0.2554$.[24]

[24]Out of 139 mln commuters we observe in the nationwide 2012–2016 LODES data, 9.8 mln travel between locations that are over 3 hours apart. Due to reasons outlined in Graham, Kutzbach, and McKenzie (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.

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

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

We assume fixed migration costs because distance is unlikely to be an important consideration within cities. 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.[25] We calibrate the age coefficient μ^a to match the percentage point difference in migration rates between 20–24 year olds and 80–84 year olds, which is equal to 8.46 percentage points. We obtain $\mu^0 = 14.8961$ and $\mu^a = 0.1734$.

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

$$d_{ij} = \kappa \times \text{time}_{ij}, \tag{3.2}$$

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,"

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

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 $\kappa = 0.0104$.

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 labor productivity, the persistence term θ_z and the variance term σ_z, are taken from Floden and Lindé (2001). The support of the distribution of $\bar{z}(a)$ for a given age is discretized to a five-point grid.

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 64.3 percentage points and our calibrated ϑ_0 is equal to 0.3669. We calibrate the curvature parameter of the bequest function ϑ_1 to match the ratio

[25] Migration across county borders is relatively uncommon. The migration rate for any residential move within the Bay Area is much higher: 9.7 percent.

of the 75th to the 25th percentiles of the wealth distribution at age 65 from the 2016 Survey of Consumer Finances (SCF), and obtain $\vartheta_1 = 0.4679$.

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 residential fixed assets, and obtain the annual depreciation rate of 0.0239.

Next, we describe how we build the sets of possible sizes of owner-occupied and rental units, $\mathbb{H}$ and $\mathbb{H}^r$. We let $\mathbb{H}$ be discrete and $\mathbb{H}^r$ continuous.[26] 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 to match the following two moments. First, 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. Second, we ensure that the fraction of homeowners who live in owner-occupied houses of median size or smaller is 50%. Our calibrated smallest and largest owner-occupied house sizes are 3.43 and 9.06.[27]

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 AHS. Our calibrated largest rental unit size is 4.53.

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 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 (2.26), this interest rate results in the price-rent ratio of 19.61.[30] We calibrate the discount factor ρ to match

[26]For 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.

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

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

[29]See 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).

[30]The equivalent moment in the data for 2012–16 is 26.67, as noted above. Taking a longer-term perspective results in a lower price-rent ratio. Using the Zillow data on levels of prices and rents, and Case-Shiller and BLS data on changes, we calculated that the average price-rent ratio in the San Francisco MSA from

the median wealth-income ratio of 1.647 from the 2016 SCF, and obtain the value of 0.018.

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.

Construction. The construction productivity functions Z_{mit}^h are chosen to match empirical estimates of floorspace supply elasticities. Denote the rental price that obtains when construction is positive (so that price p_{mit} equals construction cost $\hat{p}_{mit}$) by

$$\hat{r}_{mit} = (\delta + q + \tau_h)\hat{p}_{mit} - \partial_t \hat{p}_{mit}.$$

Since $\hat{p}_{mit} = 1/Z_{mit}^h$, the function Z_{mit}^h determines $\hat{r}_{mit}$.

In the residential segment, we assume that Z_{Sit}^h is a function of the number of local housing units (i.e., local residential population N_{it}). Specifically, we set Z_{Sit}^h so that, when construction is positive, there is a log-linear relationship between residential rents and the number of housing units:

$$\hat{r}_{Sit} = \bar{r}_{Si} N_{it}^{1/\xi_i^S}. \tag{3.4}$$

In the previous expression, ξ_i^S is the housing unit supply elasticity in the residential market and $\bar{r}_{Si}$ is the exogenous component of residential construction productivity. We obtain ξ_i^S by combining 30-year MSA-level elasticities from Saiz (2010) and 10-year tract-level elasticities from Baum-Snow and Han (2023). This allows us to have long-run elasticities at the level of model locations. Appendix A provides more details on how we construct supply elasticities for residential floorspace. The productivity shifter $\bar{r}_{Si}$ is chosen to match observed rents r_{Si}^{data} in the initial steady state:[31]

$$\bar{r}_{Si} = r_{Si}^{\text{data}} N_{it}^{-1/\xi_i^S}.$$

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

1987–2023 was 23.1.

[31] As discussed in Section 2.2.2, construction is always positive in a steady state, so $\hat{r}_{Si} = r_{Si}$.

H_{Cit}:

$$\hat{r}_{Cit} = \bar{r}_{Ci} H_{Cit}^{1/\xi_i^C}. \tag{3.5}$$

Due to the absence of local estimates of commercial floorspace supply elasticities, we assume that these elasticities are the same for residential and commercial real estate and infer $\bar{r}_{Ci}$ in the same way as for the residential sector but using the Compstak commercial rent data.[32] Appendix A provides details on how we construct supply elasticities for commercial floorspace.

3.3 Non-targeted Moments

This section describes the baseline model's implications for commuting, 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. 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 0.88. Panel B shows that our model also produces a distribution of commuting times that is similar to the one in the data.

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 C of Figure 3 shows that our model is fairly 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 D of Figure 3, our model matches the entire age profile of migration remarkably well. Importantly, our model generates the flattening of the migration profile with age.

4 Policy Experiments

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

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

Figure 3: Non-targeted moments

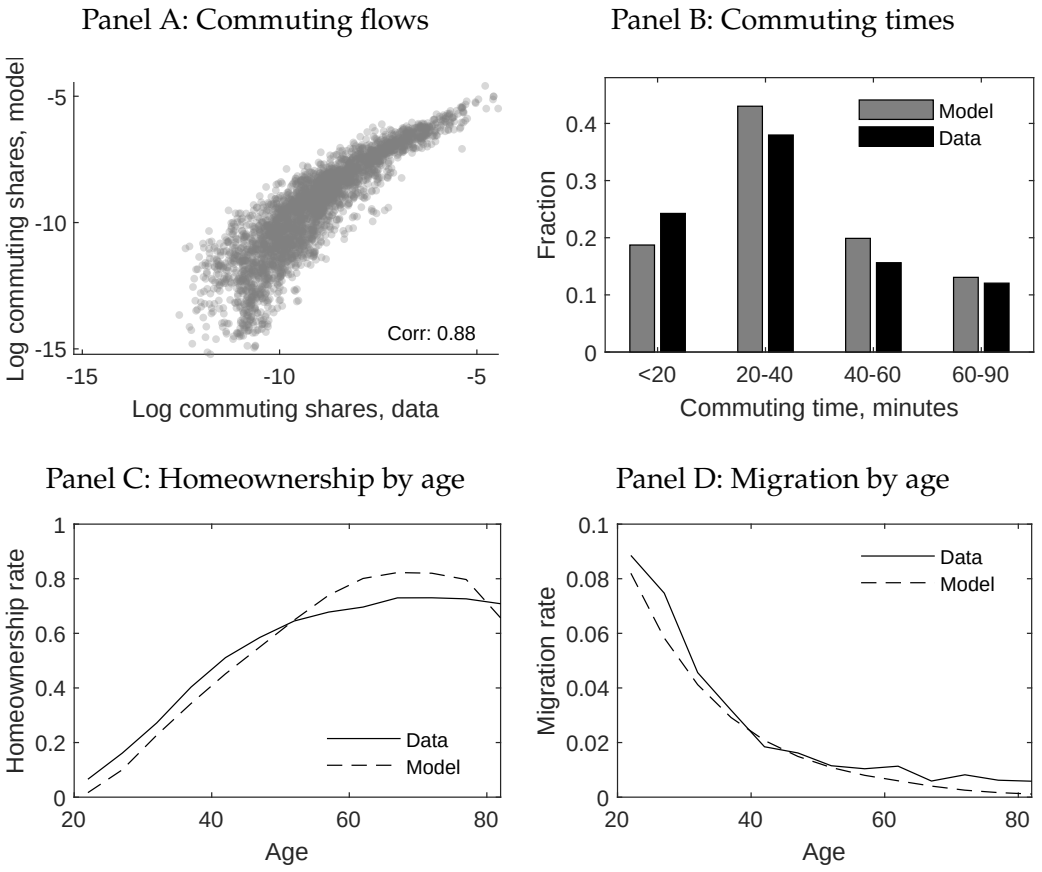

Note: The maps figure shows several non-targeted moments, in the model and in the data. See the text for more details.

4.1 Effects of Transportation and Housing Policies

We use the calibrated model of the Bay Area to run two counterfactual experiments: (1) improve the transportation network and (2) increase housing supply. In both experiments, we assume that the population of the Bay Area is fixed, i.e., it is a closed city.

4.1.1 Counterfactual Experiments

High-Speed Rail. In our first 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).[33] It envisages four stations in the Bay Area: San Francisco, Millbrae/SFO, San Jose, and Gilroy. The projected travel time between San

[33] Fajgelbaum, Gaubert, Gorton, Morales, and Schaal (2023) also study the impact of the California HSR using a static quantitative spatial model of California.

Figure 4: Counterfactual changes in travel times and housing productivity

Panel A: HSR, outbound time decrease (min.)

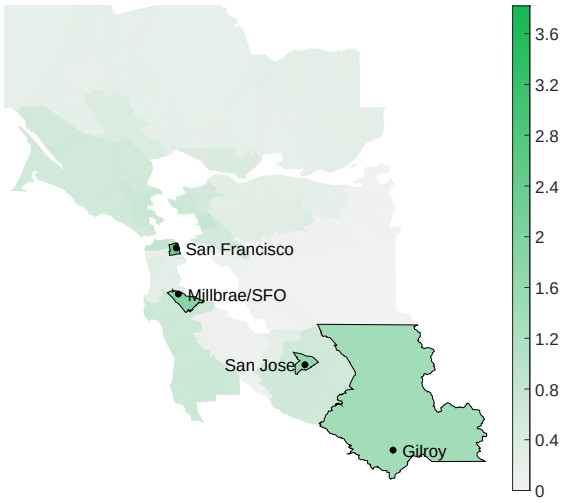

Panel B: HSR, inbound time decrease (min.)

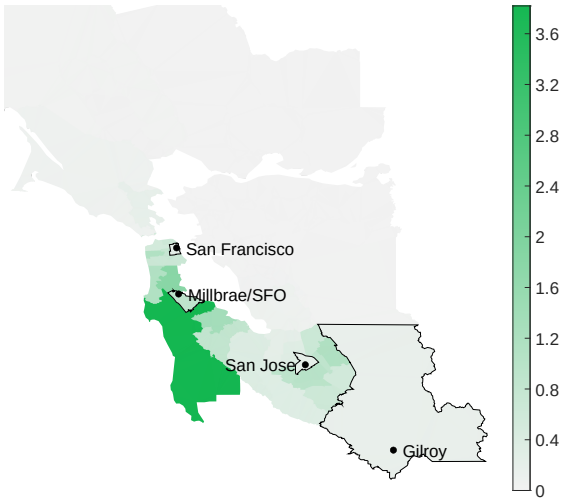

Panel C: Upzoning, productivity increase

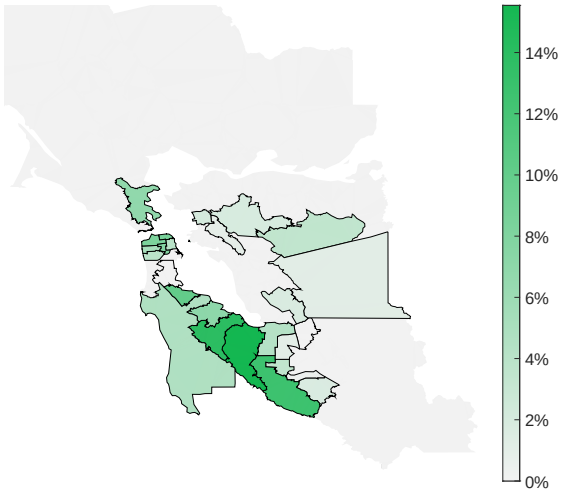

Note: Panel A shows the locations of HSR stations in the HSR counterfactual, as well as the reduction in outgoing travel times in minutes for each location, weighted by pre-HSR commuting flows. Panel B shows the reduction in incoming travel times, weighted by pre-HSR commuting flows. Panel C shows the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

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

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 A and B of Figure 4 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. While time savings are the largest in locations with stations, they are also substantial for many neighboring areas. Moreover, 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.

Upzoning. In our second 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.[34] In particular, we first calculate the median housing productivity $\bar{r}_{Si}$ in the Bay Area and then, for each location below the median, we increase its $\bar{r}_{Si}$ 10% toward the Bay Area median. Locations with low calibrated $\bar{r}_{Si}$ are those where prices are relatively high, which may indicate insufficient supply of housing. Panel C of Figure 4 shows that upzoning is largely concentrated around the Silicon Valley and San Francisco–areas that are notorious for the difficulty to build and for high prices.

In these counterfactuals, the introduction of the HSR or the Upzoning are unanticipated shocks. As soon as they occur, agents correctly anticipate the entire future path of prices (wages and prices and rents in both the residential and commercial property sectors) in each location. The computational appendix 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 does not lead to meaningful differences in results.

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

4.1.2 Transitional and Long-Run Spatial Effects

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 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 5. However, as we can see in panel A of Figure 6, these changes only happen gradually due to moving costs, housing transaction costs, and durability of structures. The full transition unfolds over nearly 60 years. The magnitude of long-run population growth in locations with stations ranges from 2.5% in San Jose to 11% in San Francisco. One important reason for this heterogeneity is relatively high estimated housing supply elasticity in San Francisco. This allows it to have a lower price growth (see panel D) and over 4 times larger population growth than San Jose where the supply elasticity is much lower.[35]

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 6. However, as time passes and new residents move into locations with stations, jobs partly recover. For example, Gilroy loses 2.6% jobs on impact but the long run decline in jobs is only 1.9%. Still, it 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 15% long-run increase in jobs (see panel C of Figure 5). Even though it doesn't have an HSR station of its own, its workers experience the largest fall in travel times of all locations (see panel B of Figure 4). 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 project has a sizable effect on real estate markets in locations with HSR stations. House prices jump immediately by between 1.9 and 6 percent in anticipation of future growth of rents (panel D of Figure 6). In the long run, price appreciation reaches 2.8–7.6 percent, and the variation in magnitudes reflects different population growth rates and differences in 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

[35] In our model, 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 5: Long-run spatial effects of HSR

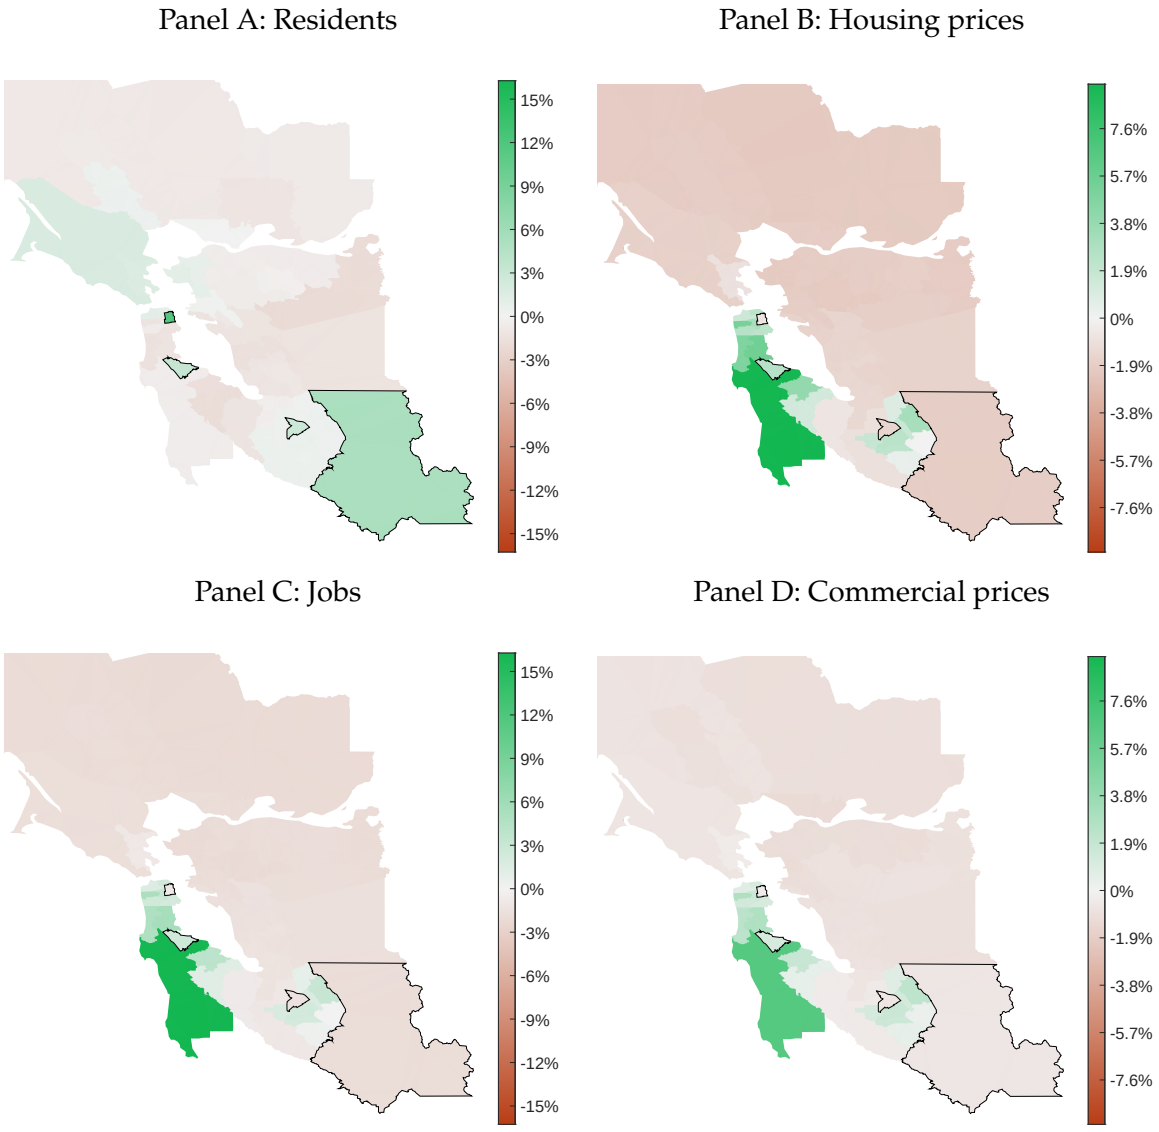

Note: The maps show the long-run percentage changes in residents, jobs, and real estate prices in the HSR counterfactual. Bordered areas represent the locations where HSR stations are built.

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 3 percentage points in San Jose and nearly 2 percentage points in San Francisco in the first 20 years before recovering in the long run. At the same time, outflow of residents from many other locations lead to lower price-rent ratios and higher homeownership.

Panels G to I of Figure 6 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

Figure 6: Transitional effects of HSR

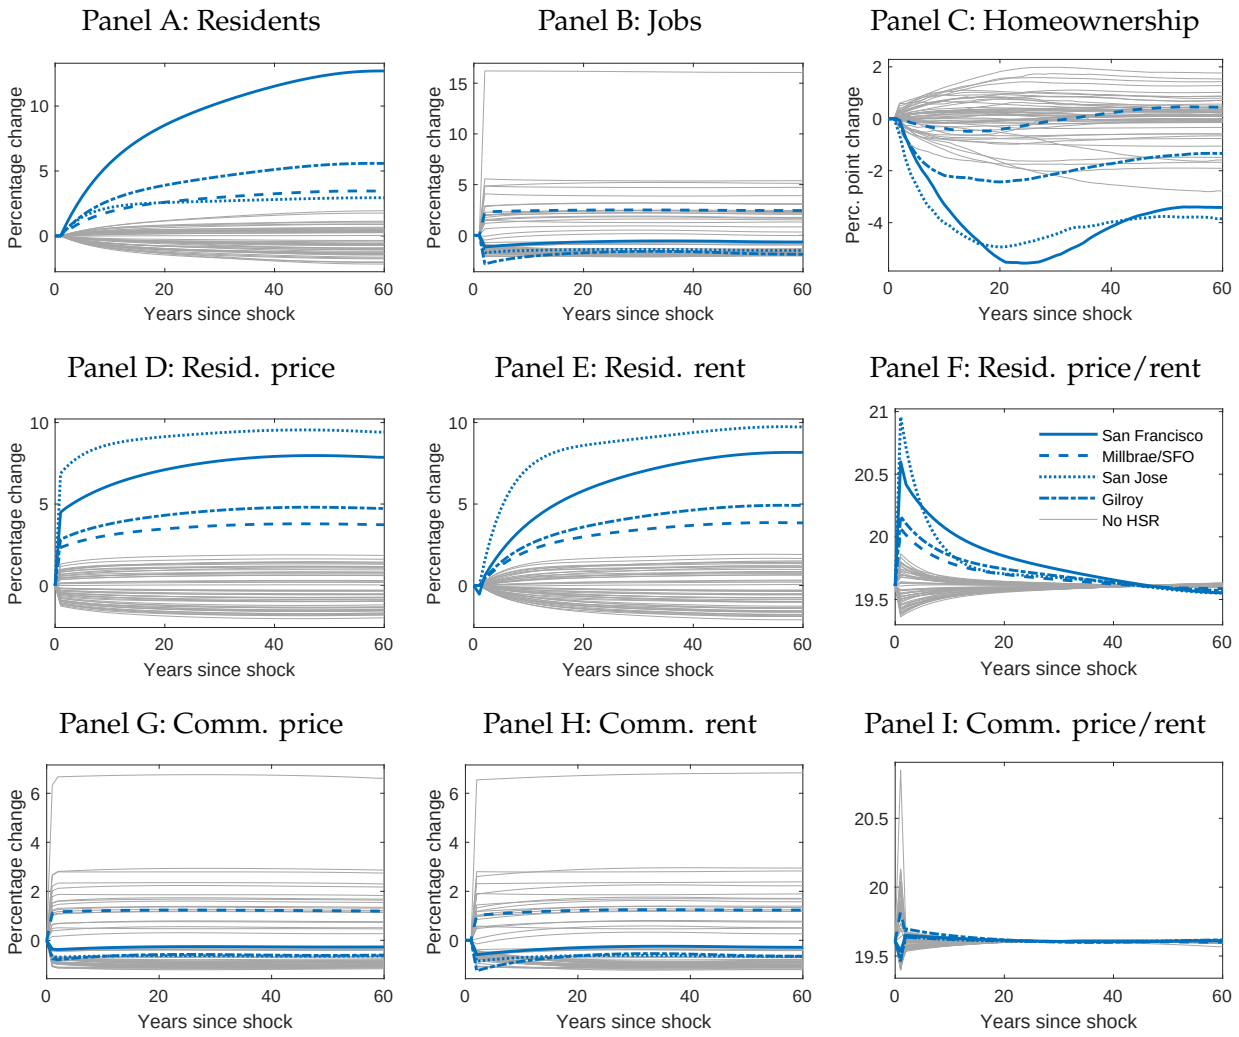

Note: The figure shows the evolution of residents, jobs, homeownership rates, real estate prices, rents, and price-rent ratios over the transition in the HSR counterfactual. Thick blue lines represents the four model locations that contain the HSR stations. Thin gray lines represent other locations.

long-run levels in just a few years. Panels B and D of Figure 5 show that long-run changes in residential and commercial prices largely mimic the changes in residents and jobs.

Upzoning. The increase in construction productivity allows developers to build more housing in upzoned locations. As can be seen in panel A of Figure 7, this attracts more residents from locations where upzoning did not take place. However, not all upzoned locations gain residents. For example, because upzoning is more agressive in the Silicon Valley (see panel C of Figure 4), it draws some of its new residents from East Bay where upzoning was more moderate. Panel C of Figure 7 shows the long-run reallocation of

Figure 7: Long-run spatial effects of upzoning

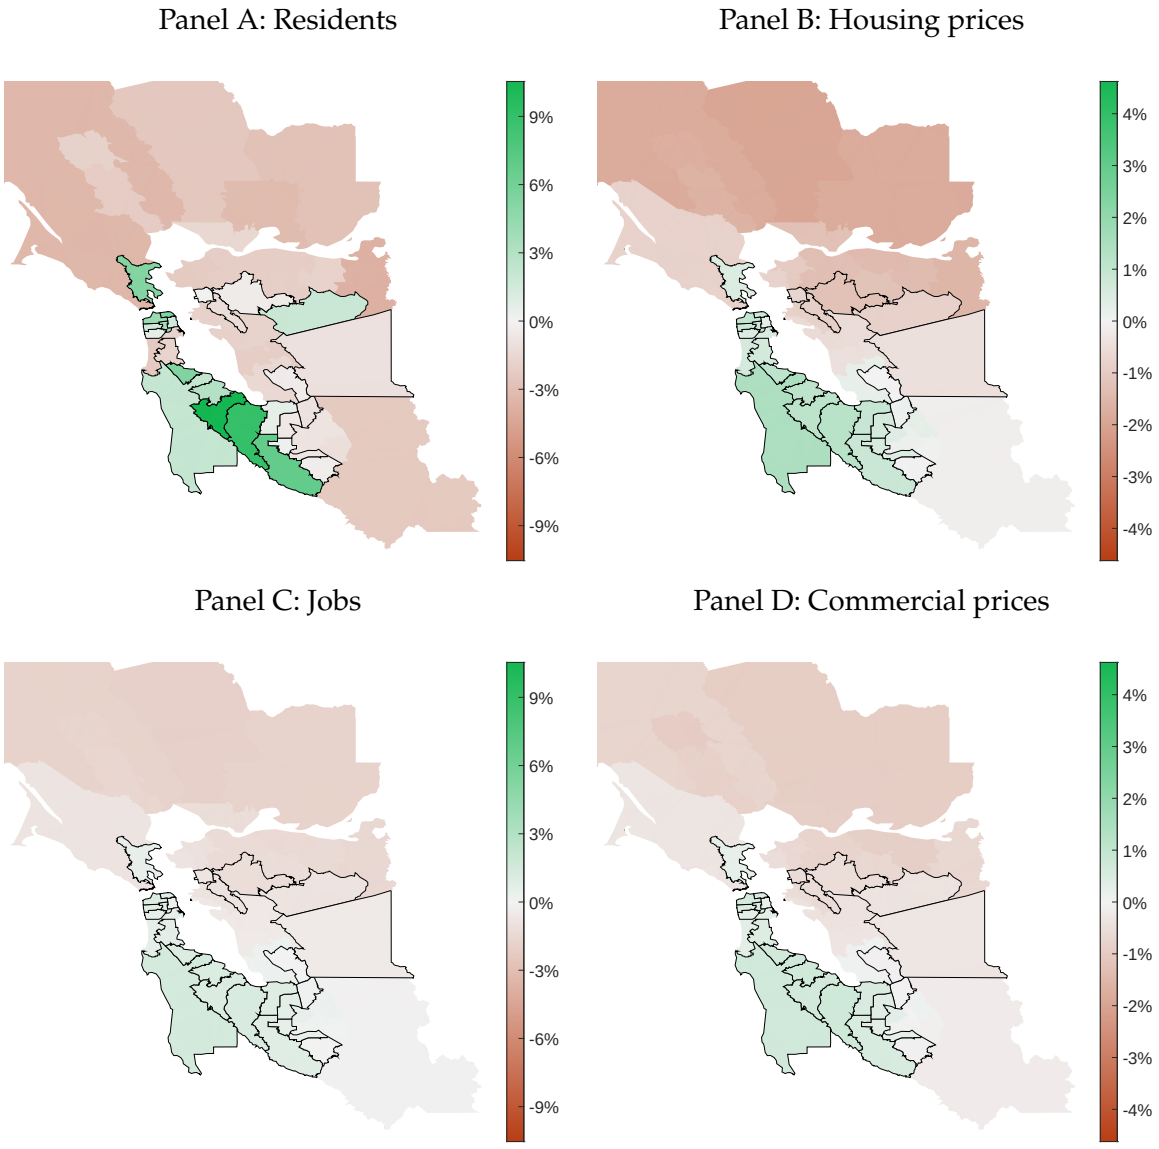

Note: The maps show the long-run percentage changes in residents, jobs, and real estate prices in the upzoning counterfactual. Bordered areas represent the locations where upzoning takes place.

jobs to the areas around the 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. Panels A and B in Figure 8 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 nearly 60 years for the number of residents and jobs to reach the new steady state.

Residential prices and rents fall in upzoned areas on impact and remain lower in the

Figure 8: Transitional effects of upzoning

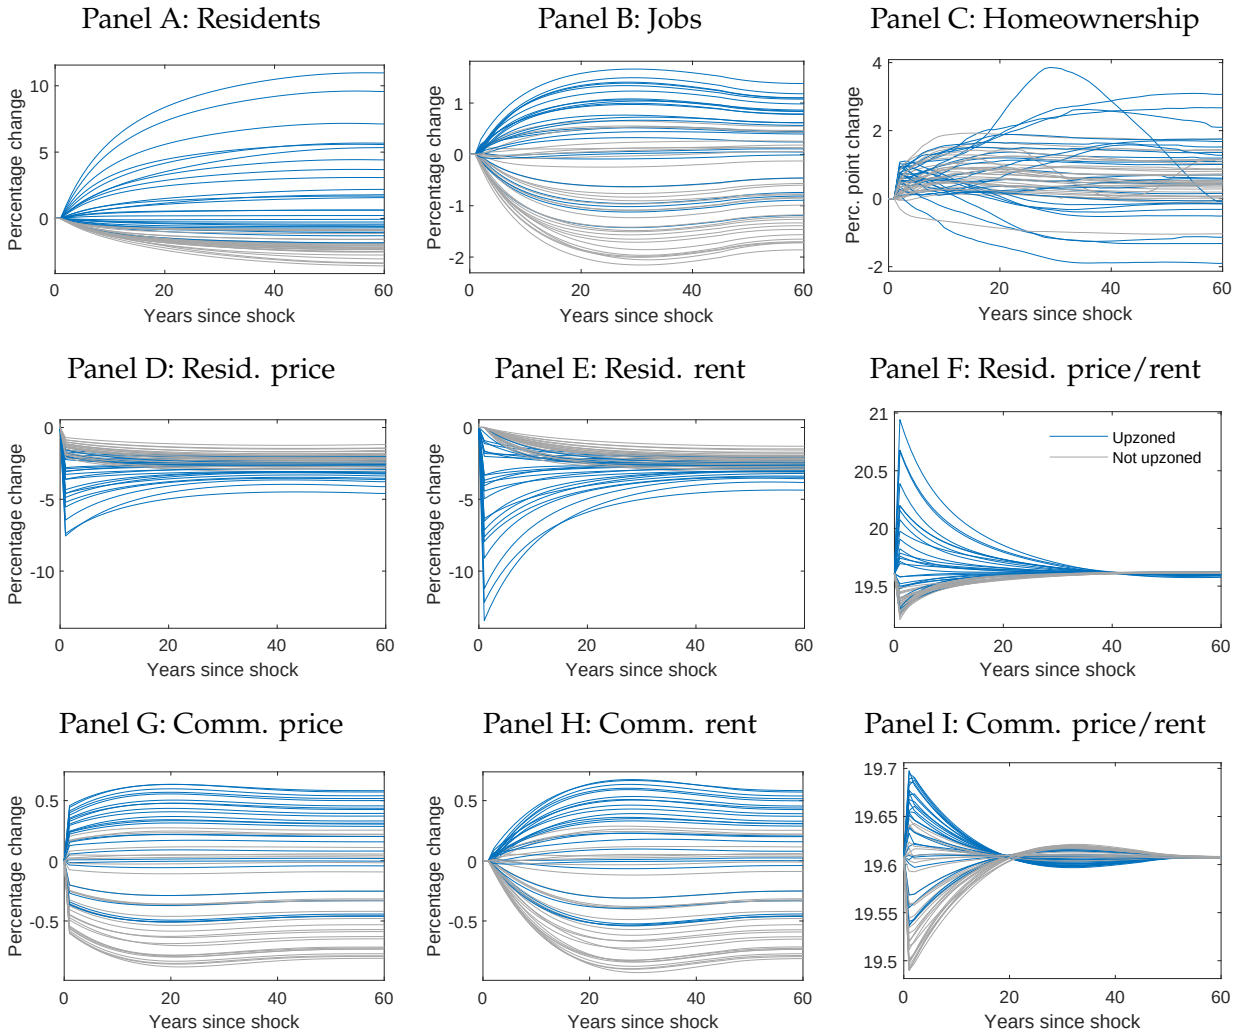

Note: The figure shows the evolution of residents, jobs, homeownership rates, real estate prices, rents, and price-rent ratios over the transition in the upzoning counterfactual. Blue lines represent the model locations where housing supply productivity was increased. Gray lines represent other locations.

long run, as demonstrated in panels D and E of Figure 8. As residents move away from non-treated locations, prices and rents there fall as well, and every single location in the Bay Area experiences long run decline in house prices (panel C of Figure 7). 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 8 and panel D 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 8). This is because, as new housing comes onto the market each year, residential

Table 2: Aggregate Welfare

Welfare gain, % median earnings	HSR	Upzoning
Accounting for transition	6.46	7.98
Comparing two steady states	6.36	10.76

Note: The table describes aggregate welfare gains in the two counterfactual experiments. Welfare is measured as the increase in wealth, measured as a percentage of initial median earning, that all households would need to receive in the initial equilibrium to equalize average utility pre- and post-counterfactual.

rents increase after an initial drop, while prices are forward-looking and immediately reflect the entire expected future rent path. Moreover, 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 8).

4.1.3 Welfare Analysis

Next, we compute welfare gains from each of the policies. We calculate the welfare gain as a one-time increase in wealth, measured as percentage of median annual earnings, that a household would need to receive in the pre-counterfactual steady state to be indifferent between remaining in this initial steady state and experiencing the counterfactual transition. Table 2 shows that the HSR results in an aggregate welfare gain that is equivalent to an increase in wealth of 4.43% of median earnings. The upzoning leads to a welfare gain in the amount of 5.97% of median earnings. We also compute welfare gains by ignoring the transition and simply comparing the pre-policy and the long-run steady states. In case of the HSR, the difference between the two measures is small. However, in the upzoning counterfactual, accounting for the transition lowers welfare gains by over one-quarter. This is because the expansion of housing supply leads to sizable wealth losses among homeowners in the upzoned locations, which lowers their utility in the beginning of the transition due to legacy real estate exposure. Another reason why accounting for transition lowers welfare gains in both counterfactuals is that it is costly to relocate to locations that offer higher utility due to moving and housing transaction costs.

Our dynamic setting with rich agent heterogeneity allows us to dissect welfare gains along several dimensions. Figure 9 shows welfare effects separately by age, worker pro-

ductivity level, tenure status, and residential location before the shock took place. Panel A shows that in the HSR experiment, welfare gains are higher for younger individuals and for homeowners. This is because younger individuals can enjoy the benefits of the HSR for a longer period of time, while homeowners benefit from real estate appreciation due to greater connectivity in many locations. At the same time, gains are larger for 30–40 year olds than for very young workers. One reason is that young workers are more geographically mobile and are more likely to move away from neighborhoods with stations. Proximity to the HSR matters too. While welfare gains are concentrated among residents who live in locations where HSR stations are built, 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. Even though the vast majority of residents experience gains, old renters who live in locations where HSR stations are built experience losses. This is because they do not benefit from greater connectivity brought by the HSR but their rents still go up. Moreover, Panel B shows that more productive individuals benefit more. This is because individual and workplace-specific productivity are complements, so access to the HSR disproportionately benefits more productive workers.

In the upzoning counterfactual, welfare gains are larger for younger individuals and for renters, as shown in panel C in Figure 9. This is because younger individuals can enjoy the benefits of the upzoning for a longer period of time, while many homeowners, especially the older ones, lose housing wealth as a result of housing supply expansion. Homeowners' welfare gains depend little on whether their location was upzoned or not because, due to spatial equilibrium effects, non-upzoned areas lose residents and 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. However, the gains for renters peak around retirement age and are much larger in upzoned locations. One important reason for that is that older renters are less mobile and are more likely to remain in their current locations and reap the benefits of housing supply expansion. Panel D shows that more productive renters benefit more from the upzoning policy.

4.2 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, Hasan, Jiang, and Parkhomenko, 2024; Fajgel-

Figure 9: Welfare effects

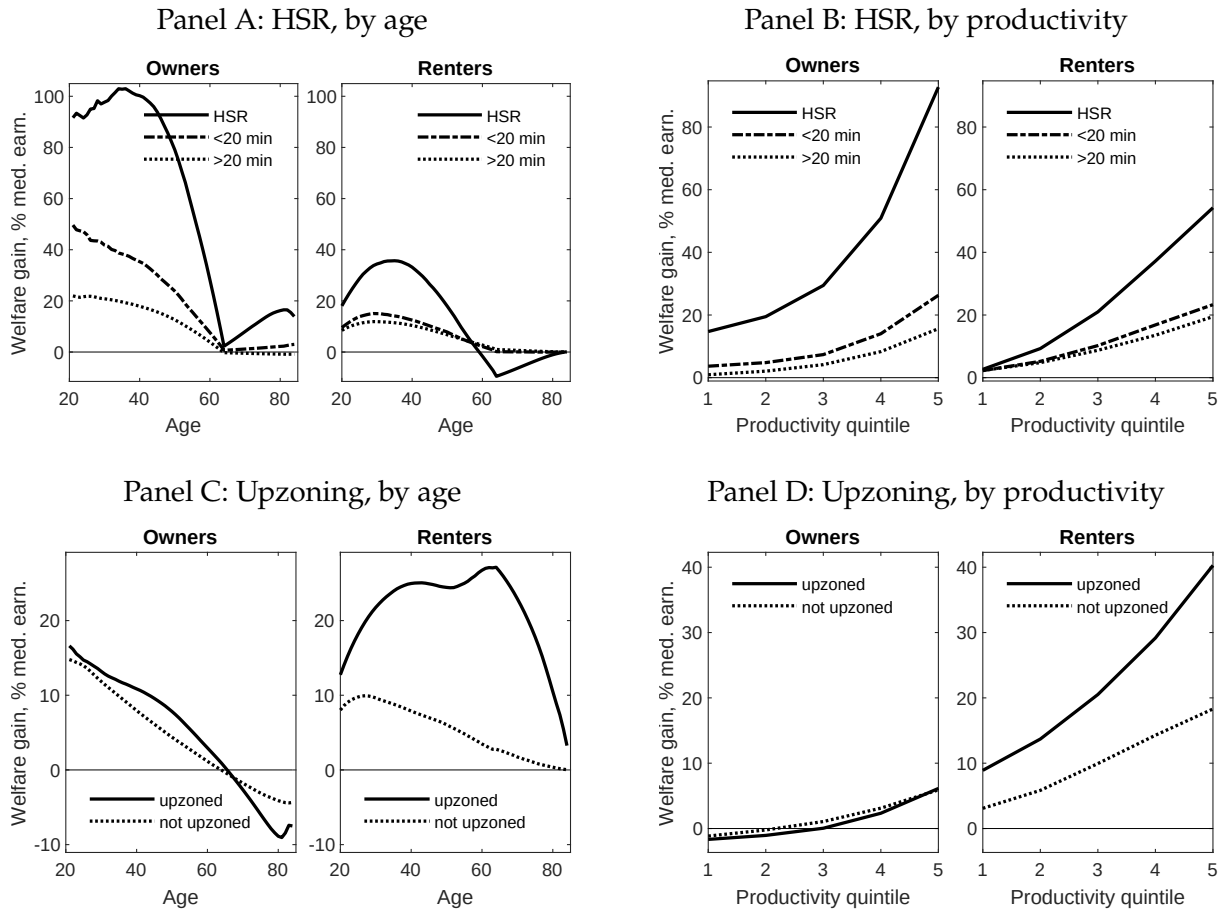

Note: The figure shows counterfactual welfare gains by age (panels A and C) and individual labor productivity (panel B and D) in the HSR and upzoning counterfactuals. Welfare gains are shown separately for owners (left side of each panel) and renters (right side), as well as separately for residents of treated locations (solid line), locations less than 20 minutes away from treated locations (dash-dotted line), and other locations (dotted line).

baum, Gaubert, Gorton, Morales, and Schaal, 2023). Policy counterfactuals with increasing housing supply within a city have been studied either using static frameworks (Allen, Arkolakis, and Li, 2016; Acosta, 2022) or using dynamic frameworks with a small number of within-city locations (Favilukis, Mabille, and Van Nieuwerburgh, 2022).

Our policy experiments demonstrate how a dynamic model, such as ours, produces important results that cannot be obtained from a static model. First, a static model cannot produce transitional dynamics that arise from individual choices.[36] As we show in Figures

[36]It is possible to obtain transitional dynamics from a sequence of interdependent static equilibria. However, these transitions would be determined by path dependence and not by forward-looking decisions, which would limit the scope of welfare analysis. Moreover, modeling dynamics as a sequence of static equilibria requires assuming that model periods are sufficiently long and prevents short-run policy analysis.

6 and 8, transitions can be long and sometimes non-monotonic. This allows us to conduct policy impact analysis for an arbitrary time horizon and renders our model more suitable for policymaking. A 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 gradually over several decades and, as prices and rents adjust at different paces, the policy has different effects on renters and owners.

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. That is, workers choose where to live and work not only based on wages, housing costs, commuting costs, and amenities, as in the static model. They also take into account which locations will allow them to make optimal tenure choice, optimal saving decisions, and insure 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 9, 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 and that one cannot simply collapse the geography into a small number of locations, e.g., treated and 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 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.

5 Conclusion

We propose a new computational dynamic spatial equilbrium 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 trans-

portation 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.

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, and to quantitatively assess the general equilibrium impact of neighborhood-scale urban investment projects.

Bibliography

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022): "Income and wealth distribution in macroeconomics: A continuous-time approach," *The Review of Economic Studies*, 89(1), 45–86.

Acosta, C. (2022): "The Incidence of Land Use Regulations," Working Paper.

Ahlfeldt, G. M., and E. Pietrostefani (2019): "The economic effects of density: A synthesis," *Journal of Urban Economics*, 111, 93–107.

Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015): "The Economics of Density: Evidence From the Berlin Wall," *Econometrica*, 83(6), 2127–2189.

Allen, T., and C. Arkolakis (2022): "The Welfare Effects of Transportation Infrastructure Improvements," *The Review of Economic Studies*, 89(6), 2911–2957.

Allen, T., C. Arkolakis, and X. Li (2016): "Optimal City Structure," Working Paper.

Allen, T., and D. Donaldson (2022): "Persistence and path dependence in the spatial economy," Discussion paper, National Bureau of Economic Research.

Almagro, M., and T. Domínguez-Iino (2024): "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," Discussion paper, National Bureau of Economic Research.

Artuç, E., S. Chaudhuri, and J. McLaren (2010): "Trade shocks and labor adjustment: A structural empirical approach," *American economic review*, 100(3), 1008–1045.

Baum-Snow, N., and L. Han (2023): "The Microgeography of Housing Supply," Working Paper.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018): "House prices and consumer spending," *The Review of Economic Studies*, 85(3), 1502–1542.

Bilal, A. (2023): "Solving Heterogeneous Agent Models with the Master Equation," Working Paper 31103, National Bureau of Economic Research.

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

——— (2023): "Anticipating Climate Change Across the United States," Working Paper 31323, National Bureau of Economic Research.

Cai, S., L. Caliendo, F. Parro, and W. Xiang (2022): "Mechanics of spatial growth," Discussion paper, National Bureau of Economic Research.

Caliendo, L., M. Dvorkin, and F. Parro (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87(3), 741–835.

Campbell, J. Y., and J. F. Cocco (2007): "How do house prices affect consumption? Evidence from micro data," *Journal of monetary Economics*, 54(3), 591–621.

Carroll, C. D. (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics letters*, 91(3), 312–320.

Chen, L., R. Hasan, Y. Jiang, and A. Parkhomenko (2024): "Faster, Taller, Better: Transit Improvements and Land Use Policies," *Journal of Development Economics*, p. 103322.

Crews, L. (2023): "A Dynamic Spatial Knowledge Economy," Working Paper.

Davis, M. A., and S. Van Nieuwerburgh (2015): "Housing, Finance, and the Macroeconomy," *Handbook of Regional and Urban Economics*, 5, 753–811.

De Nardi, M. (2004): "Wealth inequality and intergenerational links," *The Review of Economic Studies*, 71(3), 743–768.

Delventhal, M. J., and A. Parkhomenko (2023): "Spatial Implications of Telecommuting," Working Paper.

Desmet, K., D. K. Nagy, and E. Rossi-Hansberg (2018): "The Geography of Development," *Journal of Political Economy*, 126(3), 903–983.

Dvorkin, M. (2023): "Heterogeneous Agents Dynamic Spatial General Equilibrium," Working Paper.

Eckert, F., and T. Kleineberg (2021): "Saving the American Dream? Education Policies in Spatial General Equilibrium," Working Paper.

Fajgelbaum, P. D., C. Gaubert, N. Gorton, E. Morales, and E. Schaal (2023): "Political Preferences and Transport Infrastructure: Evidence from California's High-Speed Rail," Working Paper 31438, National Bureau of Economic Research.

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (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, J., P. Mabille, and S. Van Nieuwerburgh (2022): "Affordable Housing and City Welfare," *The Review of Economic Studies*, rdac024.

Floden, M., and J. Lindé (2001): "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," *Review of Economic Dynamics*, 4(2), 406–437.

Giannone, E. (2019): "Skill-Biased Technical Change and Regional Convergence," Working Paper.

Giannone, E., Q. Li, N. Paixao, and X. Pang (2023): "Unpacking Moving," Working paper.

Glaeser, E. L., and J. Gyourko (2005): "Urban decline and durable housing," *Journal of political economy*, 113(2), 345–375.

Graham, M. R., M. J. Kutzbach, and B. McKenzie (2014): "Design Comparison of LODES and ACS Commuting Data Products," Discussion Paper 14-38, U.S. Census Bureau.

Greaney, B. (2023): "Homeownership and the Distributional Effects of Uneven Regional Growth," Working Paper.

Gyourko, J., and R. Molloy (2015): "Chapter 19 - Regulation and Housing Supply," in *Handbook of Regional and Urban Economics*, ed. by G. Duranton, J. V. Henderson, and W. C. Strange, vol. 5 of *Handbook of Regional and Urban Economics*, pp. 1289–1337. Elsevier.

Hansen, G. 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, S., S. Redding, and D. M. Sturm (2020): "The Making of the Modern Metropolis: Evidence from London," *The Quarterly Journal of Economics*, 135(4), 2059–2133.

Kaplan, G., K. Mitman, and G. L. Violante (2020): "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 128(9), 3285 – 3345.

Kleinman, B., E. Liu, and S. Redding (2023): "Dynamic Spatial General Equilibrium," *Econometrica*, 91(2), 385–424.

Kopecky, K. A., and R. M. Suen (2010): "Finite state Markov-chain approximations to highly persistent processes," *Review of Economic Dynamics*, 13(3), 701–714.

Martellini, P. (2022): "Local labor markets and aggregate productivity," Working Paper.

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

Ortalo-Magné, F., and A. Prat (2016): "Spatial asset pricing: A first step," *Economica*, 83(329), 130–171.

Piazzesi, M., and M. Schneider (2016): "Housing and Macroeconomics," vol. 2, chap. Chapter 19, pp. 1547–1640. Elsevier.

Redding, S. J., and E. Rossi-Hansberg (2017): "Quantitative spatial economics," *Annual Review of Economics*, 9(1), 21–58.

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

Saiz, A. (2010): "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125(3), 1253–1296.

Severen, C. (2021): "Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification," *The Review of Economics and Statistics*, pp. 1–99.

Steinberg, J. B. (2019): "Brexit and the macroeconomic impact of trade policy uncertainty," *Journal of International Economics*, 117, 175–195.

Sun, J. E. (2024): "The Distributional Consequences of Climate Change: The Role of Housing Wealth, Expectations, and Uncertainty," Working Paper.

Takeda, K., and A. Yamagishi (2023): "History versus Expectations in the Spatial Economy: Lessons from Hiroshima," Working Paper.

Tolbert, C. M., and M. Sizer (1996): "U.S. Commuting Zones and Labor Market Areas: A 1990 Update," Staff Reports 278812, United States Department of Agriculture, Economic

Research Service.

Tsivanidis, N. (2023): "The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio," Working Paper.

Valentinyi, A., and B. Herrendorf (2008): "Measuring Factor Income Shares at the Sectoral Level," *Review of Economic Dynamics*, 11(4), 820–835.

Vanhapelto, T. (2022): "House Prices and Rents in a Dynamic Spatial Equilibrium," Working Paper.

Warnes, P. (2024): "Transport Infrastructure Improvements and Spatial Sorting: Evidence from Buenos Aires," Working Paper.

Zerecero, M. (2021): "The Birthplace Premium," Working Paper.

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.[37]

We calculate mean labor earnings for tract k as $\bar{w}_k = \left(\sum_b N_{b,k}\bar{w}_b\right) / \sum_b N_{b,k}$, 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

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

where age_k is the average age; $sexratio_k$ is the proportion of males to females in local labor force; $race_{r,j}$ 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.[38] 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; 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

[37] The 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 $\geq \$100,000$.

[38] We 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.

housing rents and prices:

$$\ln \mathbf{q}_{\iota,t} = \beta_0 + \boldsymbol{\beta}_1 \boldsymbol{\mathcal{X}}_{\iota,t} + \varphi_t + \varepsilon_{\iota,it}. \tag{A.2}$$

Here, $\mathbf{q}_{\iota,it}$ is the rent or the price reported by household ι in year t, while $\boldsymbol{\mathcal{X}}_{\iota,t}$ 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 $\varepsilon_{\iota,t}$ 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 26.67.

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 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 MSAs in Saiz (2010). We construct our elasticities by combining the benefits of both measures. 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*), and then aggregate them to the level of model locations and restrict the elasticites to be at least 0.05. Finally, our elasticity for each model location i is

$$\xi_i = \xi_i^{\text{Saiz}} + \overline{\xi}\tilde{\xi}_i^{\text{BSH}},$$

where ξ_i^{Saiz} is the Saiz (2010) elasticity, $\tilde{\xi}_i^{\text{BSH}}$ is the Baum-Snow and Han (2023) elasticity, and the normalization constant $\overline{\xi}$ ensures that the variance of elasticities across locations is the same in our model 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*).

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, Redding, Sturm, and Wolf (2015), we specify the productivity as

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

where $\overline{Z}_j$ 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 and 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.[39]

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 preserved. That is, including agglomeration externalities amplifies the effects of the HSR and upzoning on job changes but does not lead to any qualitative changes. 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 countefactuals is negligible.

Figure B.1: Welfare effects

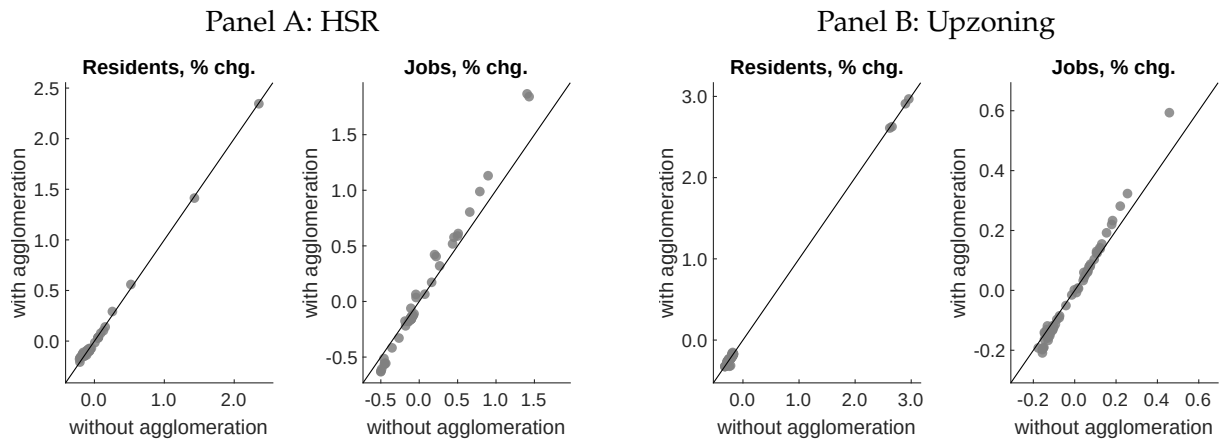

Note: The figure compares local changes in residents and jobs between the baseline version of the model and the version with endogenous productivity in the HSR (panel A) and the upzoning (panel B) counterfactuals. The solid diagonal line is the 45-degree line.

[39]The variance of $\ln \overline{Z}_j$ goes down from 0.0113 in the model without agglomeration to 0.0095 in the model with agglomeration.

B.2 Frequency of Shock Ages

We examine a version of the model where shock ages are 1/2 years apart instead of 1 year. We first re-calibrate the model. Figure B.2 shows the non-targeted moments. The fit is very good, including for migration rates. The calibration of the model with 1/2-year shock ages requires different moving costs to obtain the same implied annual migration rates. Then we run the same zoning counterfactual as in Section 4. Figure B.3 shows long-run spatial counterfactual effects. Figure B.4 shows transitional effects. Our overall conclusion is that the results are very similar than in the benchmark. 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.

Figure B.2: Non-targeted moments, dt = 0.5

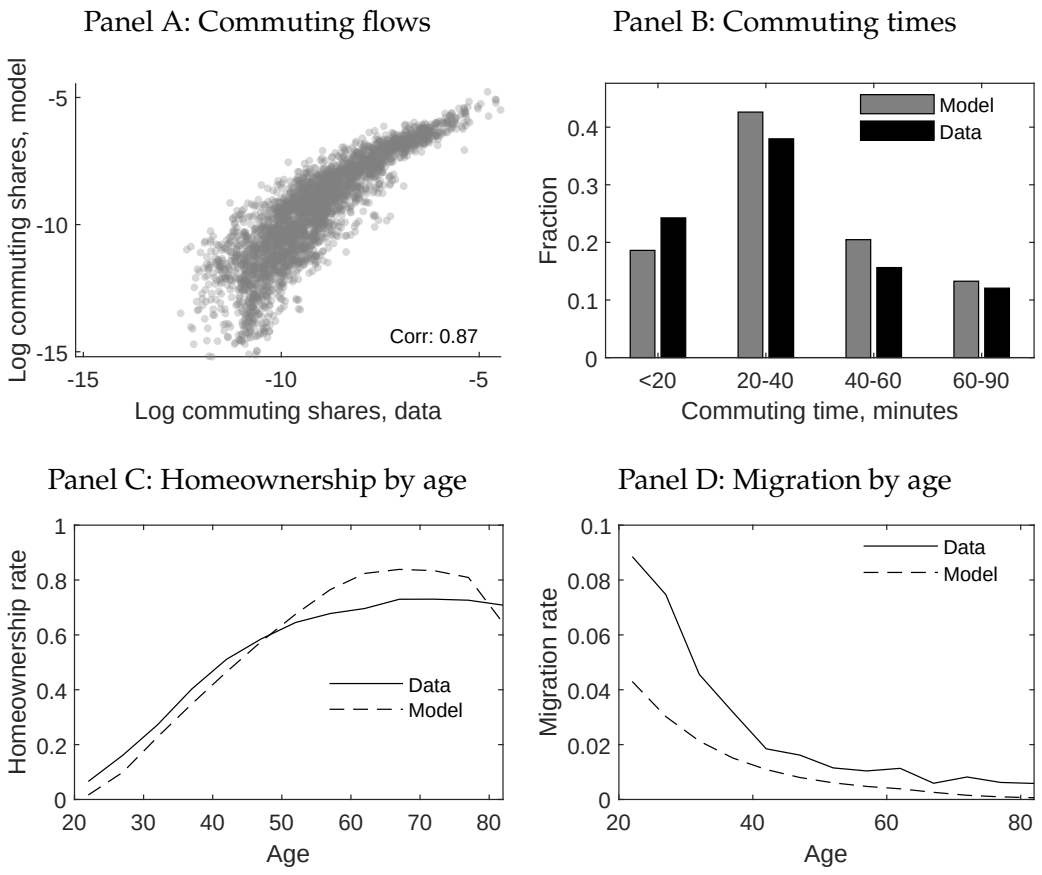

Note: The maps figure shows several non-targeted moments, in the model and in the data. See the text for more details.

Figure B.3: Long-run spatial effects of upzoning, dt = 0.5

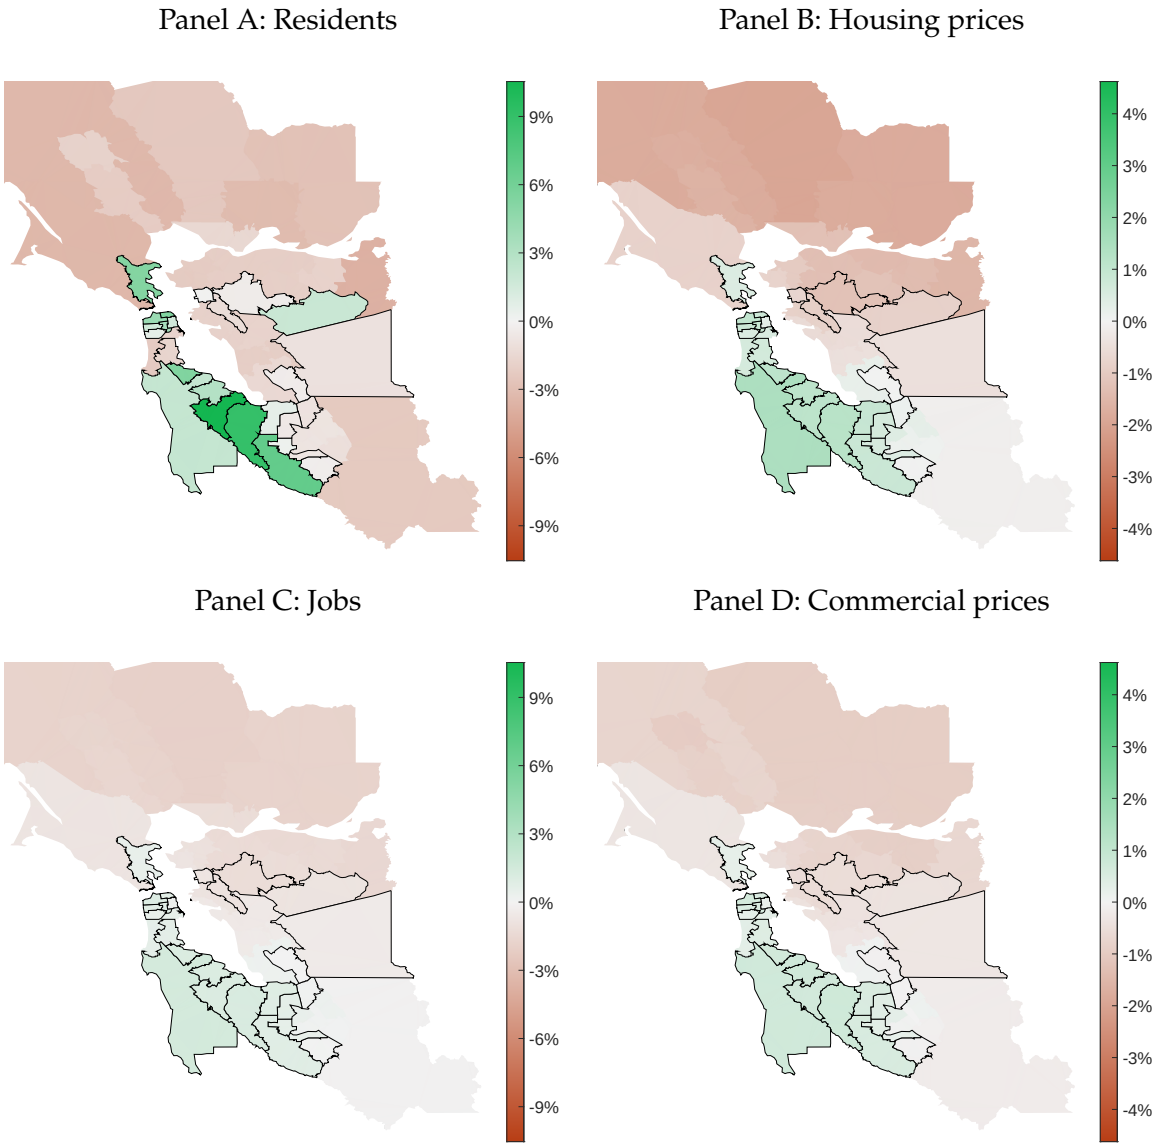

Note: The maps show the long-run percentage changes in residents, jobs, and real estate prices in the upzoning counterfactual. Bordered areas represent the locations where upzoning takes place.

Figure B.4: Transitional effects of upzoning, dt = 0.5

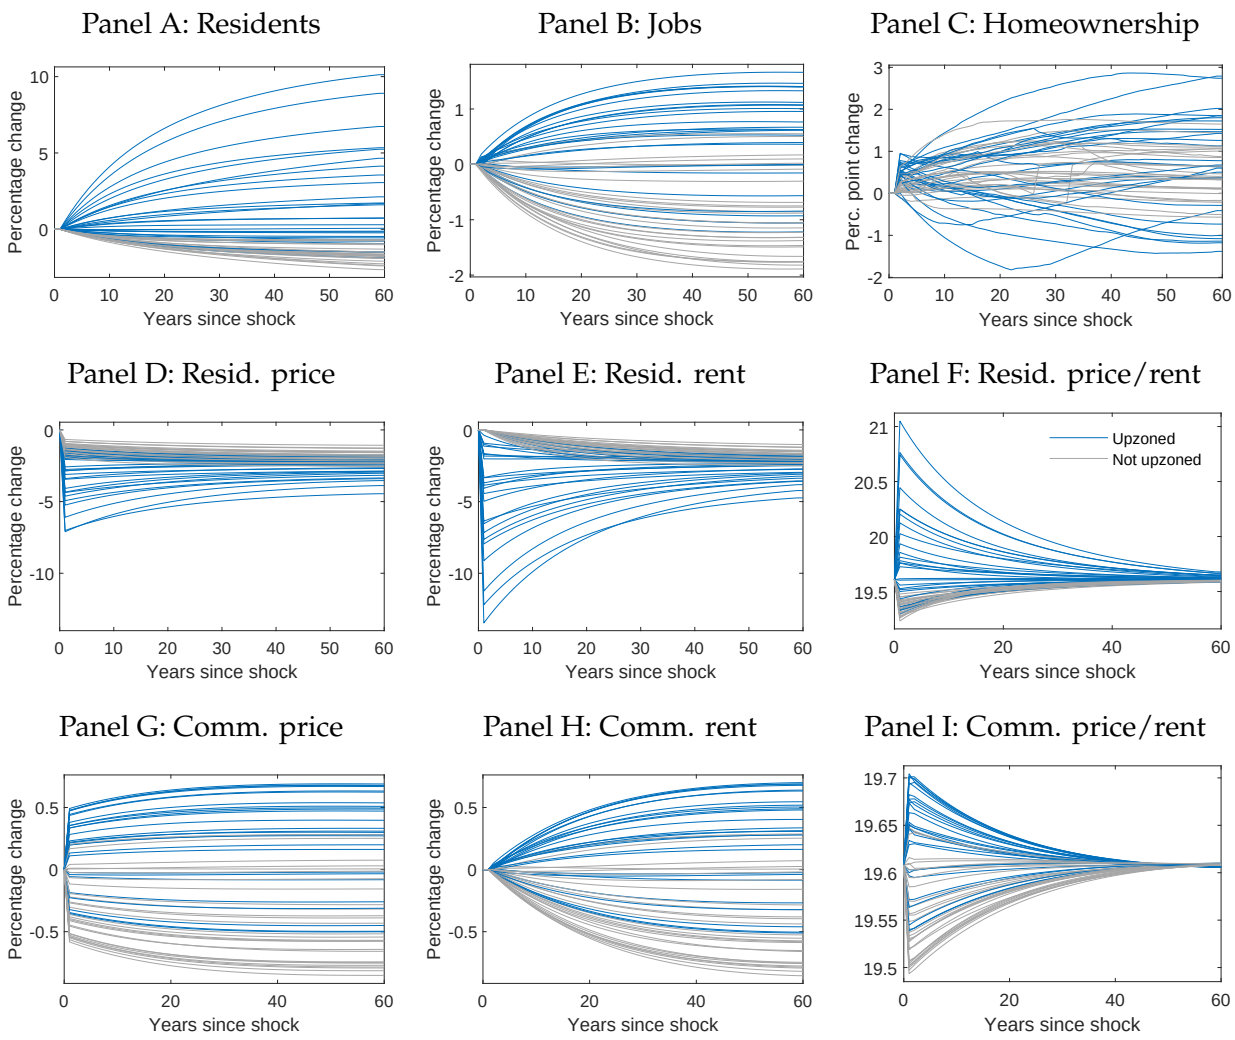

Note: The figure shows the evolution of residents, jobs, homeownership rates, real estate prices, rents, and price-rent ratios over the transition in the upzoning counterfactual. Blue lines represent the model locations where housing supply productivity was increased. Gray lines represent other locations.

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 $\overline{b}(h)$. 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 $\overline{w}$. We verify ex-post that $\underline{w} \leq \min_{j,t} w_{jt}$ and $\overline{w} \geq \max_{j,t} w_{jt}$. Finally, we discretize age using an evenly spaced grid with intervals Δ. We set $\Delta = 1/m$, $m \in \{1, 2, 3, ...\}$ so that there exists an integer n such that shock ages occur at every n steps on the age grid.

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 Hamilton-Jacobi-Bellman equation (2.4) 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, $\lim_{\iota \downarrow 0} V_{t-1+\iota}(b, h, i, w, \zeta, A - 1 + \iota)$. Evaluating equations (2.5) - (2.8) 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 (2.5) and liquid wealth values between liquid wealth gridpoints for equations (2.6) and (2.7). The conditional expectation in equation (2.8) 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 (2.9) 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, $\lim_{\iota \downarrow 0} V_{t+1-\iota}(b, h, i, w, \zeta, 1 - \iota)$. Evaluating equations (2.10) - (2.15) 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.

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 $r_{Cj} = \hat{r}_{Cj}$. Substituting into equation (3.5) and rearranging yields a supply equation for commercial floorspace: $H_{Cj} = (r_{Cj}/\bar{r}_{Cj})^{\xi_j}$. Subtract this from the commercial floorspace demand equation (2.22) to get the excess demand function

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

   Solve these (independent) nonlinear equations for market-clearing commercial rents r_{Cj}.[40] Given H_{Cj}, compute wages using equation (2.16).

3. Use the algorithm described in Section C.1 to solve the household problem and compute the density of state variables. Then use equations (2.19) and (2.21) to compute labor and population allocations L_j and N_i. Since construction is strictly positive, residential rents are $r_{Si} = \bar{r}_{Si} N_i^{1/\xi_i}$ (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

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

   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

[40] Even 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 function $\mathbf{1}^S_{it}$, which is 1 if residential 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^0_{jt}, residential rents r^0_{Sit}, and the residential construction indicator $\mathbf{1}^{S0}_{it}$. Total labor supply is normalized to 1, so this guess must satisfy $\sum_{i=1}^{I} L^0_{it} = 1$ and $L^0_{it} > 0$.

3. By equation (2.17), the supply of commercial floorspace when construction is positive is $H_{Cjt} = (r_{Cjt}/\bar{r}_{Cj})^{\xi_j}$. Combine with the floorspace law of motion (2.23) and construction irreversibility constraint to get the commercial floorspace supply equation on the discretized time grid:

$$H^{\text{sup}}_{Cjt}(r_{Cjt}) = \begin{cases} \max\{(r_{Cj0}/\bar{r}_{Cj})^{\xi_j}, H^{\text{ss}}_{Cj0}\} & \text{if } t = 0, \\ \max\{(r_{Cjt}/\bar{r}_{Cj})^{\xi_j}, (1-\delta\Delta)H^{\text{sup}}_{Cjt-\Delta}\} & \text{if } t > 0, \end{cases}$$

   where H^{ss}_{Cj0} is commercial floorspace in the initial steady state. Subtract this supply from the commercial floorspace demand equation (2.22) to get the excess demand function

$$f(r_{Cjt}) = [(1-\alpha)Z_j/r_{Cjt}]^{1/\alpha} L^0_{jt} - H^{\text{sup}}_{Cjt}(r_{Cjt})$$

   Solve this system of nonlinear equations for market-clearing commercial rents r_{Cjt}.[41] Given H_{Cjt} and L_{jt}, compute wages using equation (2.16).

4. Given the value function at time T from the final steady state, use the finite differences algorithm to solve the HJB equation (2.4) for the value function at time $T - \Delta$. After

[41] Even 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.

computing the HJB equation, evaluate equations (2.5) - (2.8) 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 (2.9) for the density of state variables at time Δ. Evaluate equations (2.10) - (2.15) 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 (2.19) - (2.21) to compute labor allocations L_{jt}, population allocations N_{it}, 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 $r_{Sit} = \bar{r}_{Si} N_{it}^{1/\xi_i}$ if construction is positive. Compute the market-clearing residential rent implied by $\mathbf{1}_{it}^{S0}$ and r_{Sit}^0:

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

If $\max_{j,t} |L_{jt}^0 - L_{jt}| < \epsilon$ 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_{jt} and r_{Sit} using

$$L_{jt}^0 = L_{jt}^0 + \nabla(L_{jt} - L_{jt}^0),$$
$$r_{Sit}^0 = r_{Sit}^0 + \nabla(r_{Sit} - r_{Sit}^0).$$

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

$$\mathbf{1}_{it}^{S} = \begin{cases} 1 & \text{if } \mathbf{1}_{it}^{S0} = 0 \text{ and } r_{Sit} > \bar{r}_{Si} N_{it}^{1/\xi_i}, \\ 0 & \text{if } \mathbf{1}_{it}^{S0} = 1 \text{ and } Y_{Sit}^{h} < 0, \\ \mathbf{1}_{it}^{S0} & \text{otherwise.} \end{cases}$$

In the first case, the rent at which demand equals supply with no construction exceeds $\hat{r}_{Sit}$. This is not an equilibrium outcome, as residential construction firms would strictly prefer to construct more housing. In the second case, when residential rents equal $\hat{r}_{Sit}$, 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 $\mathbf{1}_{it}^{S} = \mathbf{1}_{it}^{S0}$ 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 $\Theta = (\rho, \eta, \chi, \vartheta, \nu^R, \nu^W, \kappa, \mu^0, \mu^a, \min \mathbb{H}^r, \min \mathbb{H}, \max \mathbb{H}^r)$, 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 $12 + 5I = 287$ 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 $12 + 2(I - 1) = 120$ 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.[42] 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 derivate-free method. Our estimation algorithm is as follows:

[42]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.

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

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

   for $i, j > 1$ where $\nabla > 0$ is an ad-hoc updating parameter, and repeat this step.

3. Use a Newton-based method to find a Θ that matches the estimation targets listed in Table 3.2 to within ϵ.[43] 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 (2.16), (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.[44] 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.

[43] We use Matlab's fsolve function.

[44] 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.

Table C.1: Computation Time

	Steady State: P.E.	Estimate	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.

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 less than 55 locations, this task takes less than 30 seconds. 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	7.12 sec.
Seattle	33	1089	15.13 sec.
San Francisco	55	3025	27.21 sec.
Los Angeles	123	15129	1.47 min.
New York	183	33489	2.72 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.