Urban development dynamics and zoning

Chase Abram
University of Chicago

Jordan Rosenthal-Kay*
University of Chicago

March 20, 2025

Abstract

How do local housing markets interact across space and time? To answer this question, this paper develops a tractable dynamic spatial equilibrium model in continuous time. Our model features forward-looking housing developers, and inherits the features of many quantitative urban models. We are able to globally solve for the transition dynamics absent linearization techniques by relying on the model's spatial equilibrium condition, which dramatically reduces the dimensionality of the state space. We show numerically that following a demand shock, housing adjustment paths may be nonmonotonic, as short-run demand increases may induce some developers to overshoot their long-run housing supply. We apply our model to study how zoning restrictions affect the dynamics of housing development following local housing supply and demand shocks in San Francisco. We infer de facto zoning restrictions using bunching in the building height distribution over different zoning classifications, and use our estimates of latent zoning parameters to quantify the model.

*Authors' emails: abram@uchicago.edu, jrosenthalkay@uchicago.edu. We thank Thomas Bourany, Jonathan Dingel, Michael Galperin, Tom Hierons, and participants of the Trade Working Group and the Mansueto Urban Doctoral Program at the University of Chicago for comments. All errors are our own.

1 Introduction

Cities are dynamic. At a given instant, one neighborhood might be in the midst of a 'boom' while others find themselves in a 'bust.' For example, London's Canary Wharf began development as the housing stock in neighborhoods like Brixton and Hackney decayed. Now, those neighborhoods gentrify while major firms flee the riverside financial center. In Chicago, the withdrawal of Sears Robuck from the Englewood neighborhood resulted in a population decrease on par, in percent terms, with the bombing of Hiroshima, while skyscraper construction boomed in the Loop. Each city has its own examples. This paper studies how these dynamics can coexist in a spatial equilibrium. We ask, *what are the spatial transition dynamics associated with localized shocks?* To answer our question, we develop a simple urban framework in which location-specific capital (housing) is accumulated by forward-looking developers. Following a shock, changes in economic outcomes are not instantaneous due to adjustment costs.

In answering our question, our goal is to provide a framework portable enough to apply to various situations and yield quantitative and qualitative insights into how location and investment decisions interact over time in equilibrium, and what model parameters determine the speed and geography of the transition dynamics. Our dynamic framework is simple in structure, inheriting many of the features of quantitative urban models used in applied work, and its steady state is the static equilibrium of these models. We set up the problem so that we can globally solve for the transition dynamics, without relying on exact-hat techniques or first-order approximations, and do so in an easy-to-implement, transparent, and computationally scalable way.

To illustrate the use of our framework in practice, we calibrate a version of our model to San Francisco. We use the quantified model to study the transition dynamics associated with localized housing supply and demand shocks in a city with zoning restrictions. To do so, we embed our dynamic housing development framework into a standard model of urban commuting in the style of Ahlfeldt et al. (2015), which we augment with zoning regulations. In our model, zoning regulations act as a quota on local housing supply, generating spatially heterogeneous long-run supply curves. Through the lens of the model, zoning regulations are an unobserved set of constraints that emerge from developers complying with the complex amalgam of rules stipulated in city municipal code. We show that these *de facto* regulations can be recovered using data on the share of floorspace bound by zoning regulations. We operationalize this theory with an estimation strategy similar to Brueckner et al. (2024): we use data on all structures in San Francisco and their zoning classification to detect bunching in the building height distribution, and use these bunching estimates to infer latent zoning parameters in the model. We use the calibrated model to perform two counterfactual exercises: a housing demand shock driven by a rise in wages downtown, and a housing supply shock from upzoning parts of the supply-constrained Sunset District. In these counterfactuals, we find spatially heterogeneous adjustment paths of the housing stock, with some locations quick to adjust, and others displaying a more laggard response.

In our model, there are a continuum of locations that differ in their amenities, productivities, capital accumulation adjustment costs, and application of policy instruments. The type of capital we focus on is housing, which is consumed by freely mobile households and whose stock is constructed and maintained by competitive development firms. The model is a full general equilibrium model that features both labor and goods markets. Moreover, our setup is 'neoclassical' in that all agents have perfect foresight, are price takers, and all markets clear in equilibrium. While we allow for commuting, we do not incorporate the sort of urban externalities of Lucas and Rossi-Hansberg (2002), so that we further restrict our attention to competitive equilibria.[1] The equilibrium of our model is described by a mean-field game in which developers take as given the evolution of the spatial distributions of housing and population, which they use to form rational expectations forecasts of prices. Mathematically, while a static spatial equilibrium is simply a set of prices across locations which satisfies a no-spatial-arbitrage condition, our dynamic equilibrium is set of time paths of prices which satisfy a no-spatial-arbitrage condition as well as a boundary value problem for each location.

Dynamic spatial models like ours are difficult to work with because of the large state space. However, relying on the spatial equilibrium condition of our model allows us to reduce the dimensionality of the state space, without relying on assumptions that eliminate agents' forward-looking behavior or the interdependence of locations through the spatial equilibrium. The key is that in our context, the spatial equilibrium condition pins down the relationship between prices at different locations at each instant. Thus, we only need to keep track of the time path of household utility (which is equalized across space) instead of the distribution of endogenous housing stock and population over time and space. In doing so, we are able to globally solve for

[1] We additionally restrict our attention to how the spatial distribution of economic activity within the city evolves over time, rather than how the spatial extent or size of the city changes following shocks. In short, we differ from Alonso-Muth-Mills class of models reviewed in Brueckner et al. (1987) and Duranton and Puga (2015), in which the spatial extent or population size of the city is endogenous.

the transition dynamics, absent linearization or exact-hat techniques.

Early work in the spatial dynamics literature constructed frameworks that eliminated the need to compute continuation values, reducing the dynamic-spatial problem to a sequence of static equilibria with inter-temporal spillovers (see, e.g., Desmet and Rossi-Hansberg (2010), whose approach is extended in Desmet and Rossi-Hansberg (2014) and Desmet, Nagy, et al. (2018)). Early attempts that featured forward-looking behavior included spatial interactions through spillovers, absent general-equilibrium price effects (as in Rossi-Hansberg, 2004a). Other parts of the literature focus on forward-looking workers facing migration costs, and uses 'dynamic exact hat algebra' developed in Caliendo et al. (2019) to solve for model counterfactuals. A notable exception is Crews (2023), who uses computational techniques from the study of mean field games to model human capital investment in a dynamic spatial growth model with mobility frictions, but is restricted to studying a balanced growth path. Kleinman et al. (2023) linearize a multi-location dynamic spatial model to study the speed of transition dynamics, while Sun (2024) uses tools from deep learning to solve a dynamic spatial model.

Our model is similar to both Bilal and Rossi-Hansberg (2021) and Bilal and Rossi-Hansberg (2023), who develop frameworks that allow for dynamic mobility and capital accumulation choices in a dynamic spatial framework, with the latter leveraging the 'Master equation' approach to solving first- and second-order approximations of dynamic models. The master equation has also been used in the applied mathematics literature to study toy linear city models (Barilla et al., 2021). Most similar to our work is Greaney et al. (2025), who constructs a rich dynamic urban model that features individuals with idiosyncratic income risk that can borrow and save, subject to a constraint, and make location decisions.[2] They use a 'mixed time' approach where time is continuous between years and location preference shocks and the ability to move occur at fixed intervals. While their model features a dynamic construction sector, there are no adjustment costs to development and development is irreversible, so developers only respond to contemporaneous demand. We view our approach as complimentary, by focusing on heterogeneous local capital accumulation, but not individual heterogeneity. While our framework cannot capture many of the features of Greaney et al. (2025), it far simpler to implement in practice and features fewer dynamic parameters. We hope that the features of our framework make it attractive for applied researchers who seek to include transition dynamics in their

[2] They also quantify their model to the San Francisco area, a pure coincidence attributable to the zeitgeist: urban development in the Bay – or lack thereof – has captured the public imagination.

urban models without modeling rich agent heterogeneity.

The rest of this paper is organized as follows. First, in Section 2, we lay out our model environment and characterize its equilibrium. In Section 3 we provide a numerical exploration of our model in toy economies to impart intuition on model behavior. In Section 4, we develop a quantitative version of our model with zoning restrictions and quantify the model by measuring zoning restrictions in San Francisco. We explore several counterfactuals with the model in Section 5. Section 6 concludes.

2 Model

In this section, we lay out all the model elements, accompanied by some discussion and feasible extensions. We then proceed to characterize the equilibrium solution.

2.1 Model outline

In our environment, time is continuous and indexed by $t \geq 0$. Space is continuous and indexed by locations $x \in [0, 1]$.[3] We use a sequence-space representation of our model for exposition, in which calendar time t is a state variable. A unit measure of households choose where to live each instant. The density of households at any instant and time is denoted $n(x, t)$. Having chosen their residence, households then choose how to allocate their income between a consumption good, which is freely traded and the numéraire, and housing, which has a location-specific price at each instant. All households inelastically supply a unit of labor to the final good firm, which produces the final good with a constant returns technology that only uses labor. The final good is then consumed or used in the construction of housing. As is standard in the urban literature, each location is owned by a landlord who consumes only the final good. This setup allows us to concentrate our attention on the housing market, the focus of our paper. We abstract from both wage determination (wages are marginal products) and the goods market (there is only one good). However, we allow income, $w(x, t)$ to depend on time and space, allowing for, e.g., time-varying commuting costs to be paid for in dollars.

[3] This is an arbitrary ordering of space. In empirical applications, space is discrete, and this has no bearing on our results. With continuous space, the model is a mean-field game, while with N locations, it is an N-player game. The difference is whether developers internalize their effect on the aggregate state. Our assumption is that they do not.

There is a unit of land in each location x in fixed supply. Housing developers rent land from landlords, and use it in the construction of housing. At each location and time, there is a land rental price, $r(x,t)$. Taking the flow price of housing $p(x,t)$ as given, developers construct or demolish housing. Housing production incurs two costs, both of which are paid for using the final good. Housing construction is subject to an adjustment cost, which is convex in the rate of adjustment, while the maintenance of the housing stock requires operating costs that are increasing and convex in the stock of housing.

The equilibrium is then a path of housing prices and quantities, such that households optimally consume and locate, and developers optimally build housing, and the housing, goods, and labor markets clear.

Households Freely mobile households maximize their utility by choosing a location x, and a bundle of goods at each instant. Having chosen x, a household residing at location x chooses how to allocate their income of $w(x,t)$ between consumption goods, and housing, priced at $p(x,t)$ final goods per unit of floorspace per unit time. They also obtain utility from local amenities, $A(x,t)$. Households have Cobb-Douglas preferences, and so utility is given by,

$$u(x,t) = \max_{h,c}\ A(x,t)\left(\frac{h}{\beta}\right)^{\beta}\left(\frac{c}{1-\beta}\right)^{1-\beta} \quad \text{subject to} \quad c + p(x,t)h \leq w(x,t). \tag{1}$$

Importantly, households do not have access to savings technology.

Developers At each location x, there is a developer who chooses how to develop housing, subject to adjustment frictions. They maximize the present discounted value of flow profits, $\mathcal{V}(H,0)$, taking as given the path of housing prices $p(x,t)$, a convex operating cost $\kappa(H)$, and adjustment costs $\Phi(\dot{H},H)$, which satisfy $\Phi(0,H)=0$, $\Phi_{\dot{H}\dot{H}}>0$. They discount the future at rate ρ. Each developer's problem is decoupled from all other locations, except through the general equilibrium determination of $p(x,t)$. With this in mind, we drop the x-indexing of location in this section for notational convenience, with the understanding that all prices and quantities are implicitly still indexed by location x. The recursive form of the developer problem is given by the following Hamilton-Jacobi-Bellman (HJB) equation,

$$\rho\mathcal{V}(H,t) - \frac{\partial\mathcal{V}(H,t)}{\partial t} = \max_{\dot{H}}\left\{p(t)H - \kappa(H) - \Phi(\dot{H},H) + \frac{\partial\mathcal{V}(H,t)}{\partial H}\dot{H}\right\}. \tag{2}$$

Equilibrium Given $A(x,t)$ and $w(x,t)$, an equilibrium in the model is a set of paths $\{p(x,t),h(x,t),n(x,t),H(x,t)\}$

such that,

(i) given $w(x,t)$ and $p(x,t)$, households solve (1), so that optimally demanded housing is $h(x,t)$, and there is no opportunity for spatial arbitrage: $u(x,t) = u(x',t), \forall x, x' \in [0,1]$;

(ii) The labor market clears: $\int_0^1 n(x,t)dx = 1$ at all t;

(iii) given $p(x,t)$, developers solve (2), optimally choosing housing supply paths $H(x,t)$;

(iv) rental prices $r(x,t)$ exhaust all development profit;

(v) housing markets clear at all locations and times, $H(x,t) = n(x,t)h(x,t)$.

Discussion Several points about this setup are worth noting. First, we have excluded a forward-looking problem for households because instantaneous free mobility precludes this. The full knowledge of the time path of prices does not affect any household's location decision problem; households are never 'stuck,' so location choice does not constitute an type of investment decision as in Bilal and Rossi-Hansberg (2021).

The way we handle profits – via absentee landlords – is standard in the literature. These landlords are 'absentee' in that they do not consume housing or participate in the final goods production; they simply 'eat' their land rents. An alternative assumption would be to rebate land rent back to households. Local rebating would introduce a pecuniary externality, and aggregating profits to a citywide portfolio owned by the households would rescale income and have no effect on the distribution of economic activity. Thus, our assumption allows us to trace out the effects of policy on both renters and rentiers. We view this as an attractive feature of our model.

Developers maintain a stock of housing capital and make dynamic decisions. This is our model's key innovation. To capture the fact that there are substantial time-to-build costs to housing development, and once housing capital is installed, it is costly to remove, we introduce an adjustment cost à la Rotemberg (1982).[4]

[4] Our adjustment cost is levied on developers, and reflects time-to-build considerations, holdup due to regulatory matters, and so on. Thus, it is different from the adjustment costs for housing levied on households, which include real transaction costs and psychic frictions of changing, e.g., school districts, as in Stokey (2009). Sluggish capital response to price changes is not new in the literature; see, e.g., Iacoviello (2005). However, housing construction is not typically modeled as subject to adjustment frictions, and most of the macroeconomic literature on housing focuses on it as an asset class (Piazzesi and Schneider, 2016). Our main innovation is to embed this setup into a spatial equilibrium model that can account for granular heterogeneity.

A developer at x needs to only know the time path of prices *at* x. While prices at other x' affect developers at x, they only do so through housing market clearing. Thus, the model has a block-recursive structure.

Master equation representation of the economy To understand why the developer's problem is hard, we cast a particular version of our problem in the state-space following Bilal (2023).[5] To move from the sequence-space to the state-space, we eliminate time-indexing through a change of variables by writing the developer's problem as a function of the aggregate state,

$$\mathcal{S}(t) = \{A(x,t), w(x,t), H(x,t), n(x,t) \ : \ x \in [0,1]\},$$

which captures the path of the exogenous spatial distribution of $A(x,t)$ and $w(x,t)$ and endogenous distributions of $H(x,t)$ and $n(x,t)$. We can write floorspace prices $p(x,t) = \mathcal{P}(x,\mathcal{S})$. That is, location-specific floorspace prices are a functional of the full spatial distribution of amenities, wages, housing, and population. To forecast prices, developers must know the full evolution of the economy over time and space.

To construct a fully recursive (i.e., Markovian) representation of the economy requires that we derive Kolmogorov forward equations (KFEs) for $H(x)$ and $n(x)$ as a function of the aggregate state, $\mathcal{S}$. First we define the continuation value operator $\mathcal{H}(x,\mathcal{S})[\mathcal{V}] = \frac{\partial \mathcal{V}(x)}{\partial H}\dot{H}(x,\mathcal{S})$. Its adjoint operator, $\mathcal{H}^*(x)[\mathcal{S}] = \dot{H}(x,\mathcal{S})$ is the KFE for H. Our first simplification is to note that instantaneous floorspace market clearing implies that, $n(x,t) = H(x,t)/h(x,t)$, which means that the evolution of $n(x)$ depends entirely on the evolution of H and prices.[6] This reduces the dimensionality of the aggregate state $\mathcal{S}$ by reducing its dependence on $n(x,t)$, since the distribution and evolution of population can be written as a function of $\{A, w, H\}$ only.

Next, we use the chain rule to replace the time derivative term in (2) with a term that depends only on the distributions of housing. Denote by $\frac{\partial \mathcal{V}}{\partial \mathcal{S}}$ the Fréchet derivative of $\mathcal{V}$ with respect to $\mathcal{S}$, so $\frac{\partial \mathcal{V}}{\partial \mathcal{S}}(x,x')$ is the effect on $\mathcal{V}(x,\ldots)$ of changes in $\mathcal{S}(x')$. Then,

$$\frac{\partial}{\partial t}\mathcal{V}(x) = \int \frac{\partial \mathcal{V}}{\partial \mathcal{S}}(x,x')\frac{\partial \mathcal{S}}{\partial t}(x')dx' \ .$$

[5] Doing so limits us to studying unanticipated, permanent shocks to $A(x,t)$ and $w(x,t)$. This assumption guarantees that the current state of exogenous variables is a sufficient statistic for forecasting the (distribution of) paths for those variables. To consider arbitrary time paths of exogenous variables, those time paths must be part of the aggregate state. In our main analysis, we do not rely on the master equation formulation and thus do not limit ourselves to shocks of this nature.

[6] To see this, use $h(x,t) = \beta w(x,t)/\mathcal{P}(x,\mathcal{S})$ from household optimization, and $h(x,t)n(x,t) = H(x,t)$. This implies $n(x,t) = \frac{H(x,t)\mathcal{P}(x,\mathcal{S})}{\beta w(x,t)}$.

The term $\frac{\partial \mathcal{S}}{\partial t}$ jointly encodes the KFE for housing (and therefore population). Since the evolution of $\mathcal{S}$ depends solely on the evolution of H, $\frac{\partial \mathcal{S}}{\partial t}(x) = \mathcal{H}^*(x,\mathcal{S})[\mathcal{S}]$. The master equation is,

$$\rho\mathcal{V}(x,H,\mathcal{S}) = \max_{\dot{H}} \mathcal{P}(x,\mathcal{S})H - \kappa(H) - \Phi(\dot{H},H) + \mathcal{H}(x,\mathcal{S})[\mathcal{V}] + \int \frac{\partial \mathcal{V}}{\partial \mathcal{S}}(x,x')\mathcal{H}^*(x',\mathcal{S})[\mathcal{S}]dx' \,.$$

This representation shows that solving for $\mathcal{P}(x,\mathcal{S})$ is key to solving the equilibrium of our model.

In other dynamic spatial models (e.g., Bilal and Rossi-Hansberg, 2023), one needs to keep track of the dependence of prices on $\mathcal{S}$ at all x. By imposing free mobility of households, and therefore spatial equilibrium, we dramatically reduce this complexity. The spatial equilibrium pins down the relationship between any pair of prices $\mathcal{P}(x,\mathcal{S})$ and $\mathcal{P}(x',\mathcal{S})$, meaning that the dependence of $\mathcal{P}$ on $\mathcal{S}$ is not x-specific. As we will see, we only need to keep track of a unidimensional object – the equilibrium level of household utility, which we call $\mathcal{U}(t)$ – which allows us to construct prices across x given $\mathcal{S}$, instead of tracking the full distribution H. In other words, $\mathcal{U}(t)$, is a sufficient statistic for the price distribution. It is a moment of $\mathcal{S}$ that depends on the full spatial distribution of H.

In economic terms, developers' choices affect the decision problems of developers elsewhere only through their effect on prices. This is a mean field game, in which these interactions between developers – 'coupling,' in the language of Lasry and Lions (2007) and Cardaliaguet (2010) – are summarized through $\mathcal{U}(t)$, upon which each developer exerts infinitesimal influence. Such mean field games are an active research area in macroeconomics, but our formulation appears novel to the literature (see Achdou et al., 2014, for an early review).

2.2 Equilibrium Characterization

To solve our model, we leverage the sequence-space representation of the economy to derive an analytic representation of the full spatial dynamics of our economy, relying on insights from the state-space approach.

Household optimality Given their income of $w(x,t)$ at each instant, their indirect utility at (x,t) is, $u(x,t) = A(x,t)\frac{w(x,t)}{p(x,t)^\beta}$. Free mobility ensures that there is a spatial equilibrium: for any x, $u(x,t) = \mathcal{U}(t)$. At any

(x,t), we then have that,

$$p(x,t) = \left(\frac{A(x,t)w(x,t)}{\mathcal{U}(t)}\right)^{\frac{1}{\beta}}.$$

Developer optimality The optimality condition for adjustment $\dot{H}$ is to increase (or decrease) housing up to the point where the marginal value of more housing equals the marginal adjustment cost, so that at any x,

$$\Phi_{\dot{H}}(\dot{H}, H) = \frac{\partial \mathcal{V}}{\partial H}$$

Combining the envelope condition with a time-differentiated version of the optimality condition, the solution to the developer's problem is a second-order ordinary differential equation (ODE) in H,

$$\rho\Phi_{\dot{H}} - \Phi_{\dot{H}\dot{H}}\ddot{H} - \Phi_{\dot{H}H}\dot{H} = p(t) - \kappa'(H) - \Phi_H. \tag{3}$$

To finish characterizing the solution, we require two boundary conditions. One is given by the initial housing $H(0) = H_0$, which the developer must take as given when deciding how to adjust. The other condition is given by assuming that, beyond some time T, $p(t)$ converges towards a steady-state p, and thus H must also converge towards its steady state, which solves the HJB when $p(t) = p$ and $\dot{H} = \ddot{H} = 0$, so the developer just adjusts H so that prices equal marginal operating costs,

$$p(T) = \kappa'(H).$$

In the absence of adjustment costs, developers can instantaneously adjust, and this is their profit-maximizing condition. We refer to this as a static equilibrium.

Theorem 2.1. A unique static equilibrium exists if κ' is invertible and $\kappa' \geq 0$.

Proof. A static equilibrium is a common utility value $\mathcal{U}$ consistent with all equilibrium conditions. The structure of the equilibrium allows us to show that it exists and is unique if,

$$\mathcal{U}^{1/\beta} = \frac{1}{\beta}\int_0^1 A(x)^{1/\beta} w(x)^{1/\beta - 1} (\kappa')^{-1}\left(w(x)^{1/\beta} A(x)^{1/\beta} \mathcal{U}^{-1/\beta}\right) \mathrm{d}x$$

has a unique solution for $\mathcal{U}$, which is true if $(\kappa')^{-1}$ exists and is nonnegative. This expression is derived by substituting prices for $\mathcal{U}$ and combining housing and labor market clearing. See Appendix A for details. $\square$

Theorem 2.2. The steady-state of the dynamic equilibrium is the static equilibrium.

Proof. In the steady state, $\Phi_H(0, H) = \Phi_{\dot{H}}(0, H) = 0$ which implies $\kappa'(H) = p$ at each x. This corresponds with a static equilibrium; see Appendix A. $\square$

We now seek to characterize the dynamic equilibrium. To do so, we make two functional form restrictions for analytical exposition. We retain these for the remainder of the paper. We restrict our attention in the main text to situations in which,

$$\kappa(H) = \frac{H^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}}, \qquad \Phi(\dot{H}, H) = \frac{\xi}{2}\dot{H}^2.$$

Relying on these functional forms only simplifies the representation of the ODEs which characterize dynamic developer optimality. To consider alternative dynamic models, a researcher needs only to substitute for κ and Φ or re-derive the developer ODEs in their context.

Our first assumption will makes the steady-state housing supply isoelastic with elasticity γ. This is isomorphic to a specification in which maintaining a stock H required combining land and the final good in a Cobb-Douglas fashion with expenditure share $1/(1+\gamma)$ on land. As land is a fixed factor, these adjustment costs are convex in the final good. The second assumption maintains that adjustment only depends on the rate of change of development, not on the current level or growth rate of development.

Under these assumptions, the equilibrium of our model is described by an integro-differential system,

$$\begin{aligned} \rho\xi\dot{H}(x,t) - \xi\ddot{H}(x,t) + H(x,t)^{1/\gamma} &= p(x,t) = w(x,t)^{1/\beta}A(x,t)^{1/\beta}\mathcal{U}(t)^{-1/\beta} \\ \mathcal{U}(t) &= \left(\int_0^1 \frac{1}{\beta}H(x,t)A(x,t)^{1/\beta}w(x,t)^{1/\beta-1}\mathrm{d}x\right)^\beta \\ H(x,0) &= H_0(x) \\ \lim_{T\to\infty} H(x,T) &= H_\infty(x) \end{aligned} \tag{4}$$

This equation describes each developer's boundary value problem (BVP), which are then coupled through an integral equation.

Theorem 2.3. For an arbitrary price path $p(x,t)$, the developer's boundary value problem has a unique solution.

Proof. We prove this by first proving that boundary value problems of the form $\ddot{y} = f(y, \dot{y}, t)$, $y(0) = y_0$, $y(T) = y_T$ have a unique solution $y \in \mathcal{C}^2$ provided $\partial f/\partial y > 0$ and $|\partial f/\partial \dot{y}| \leq M$ for some $M \geq 0$, using a proof technique from Keller (1966). We then show our problem fits into this class of problems, provided $\kappa'' > 0$ and $\left|\frac{\Phi'\Phi'''}{(\Phi'')^2}\right| \leq \tilde{M}$ for $\tilde{M} \geq 0$ (which is true for power functions). □

It is unknown to us at the time of writing whether it can be proven that a unique solution to the integro-differential system (4) exists. However, we have reason to believe it does. Recall that for any $u(t) \in \mathcal{C}^2$, and thus $p(x,t) \in \mathcal{C}^2$, a solution to the BVP exists. Call this solution $H(x,t)[u]$. A solution to the integro-differential system exists if there exists a unique $\mathcal{U}(t)$ such that,

$$\mathcal{U}(t) = \left(\int_0^1 \frac{1}{\beta} H(x,t)[\mathcal{U}] A(x,t)^{\frac{1}{\beta}} w(x,t)^{1/\beta - 1} dx\right)^{1/\beta}.$$

Provided $H(x,t)[\mathcal{U}] : \mathcal{U} \mapsto H(x,t)$ is a continuous and compact mapping on some closed, bounded, convex, and nonempty subset of $\mathcal{C}^2$, a solution would exist by Schauder's fixed point theorem. To argue uniqueness, consider the following contradiction argument: For two solutions to be different over a nonzero measure of time, it must the the case that their implied $H(x,t)$ also differ over a nonzero measure of time. Then use the fact that a higher level of $\mathcal{U}(t)$ means a uniformly (over x) lower level of $p(x,t)$, which is only consistent with uniformly higher $H(x,t)$. But over the interval where prices are lower, developers will also want to offer less housing, which contradicts the necessity of higher $H(x,t)$. It is worth again emphasizing that these arguments are not formal proofs, but intuitive reasoning. The economics of the argument are simple: if the utility path is higher for some time interval, that must mean households are facing a higher supply schedule, which is making equilibrium prices lower. But if developers knew that prices were going to be lower, they would not choose the higher supply.

2.3 Extensions

A number of extensions are possible in our framework.

(1) Commuting The model presented so far takes wages as given. Representing the equilibrium as a simple integro-differential system remains possible even with a rich spatial structure on wages, provided they are decoupled from the investment decisions of developers. For example, we can allow commuting as in Ahlfeldt et al. (2015) and the proceeding quantitative urban literature with the following structure.

First, we assume that at every instant, firms at x' produce the final good (which is the numéraire) with a linear-in-labor technology,

$$y(x', t) = Q(x', t)L(x', t),$$

where $L(x', t)$ is labor at the firm located at x' at time t and $Q(x', t)$ is its productivity. Under perfect competition for goods and labor, wages are marginal products, $w(x', t) = Q(x', t)$, and therefore do not depend on the size of the firm $L(x', t)$, which could otherwise be endogenous to $H(x, t)$.

To generate a rich spatial structure on wages, we assume a timing decision. In each instant t, households first choose where to live x based on *expected* wages. Upon choosing x, a vector of productivity shocks expressed in efficiency units of labor $\varepsilon(x')$ is realized iid from a Fréchet distribution with shape parameter $1/\epsilon$. Households choose where to work based on both these shocks and iceberg commuting costs $\tau(x, x') \geq 1$. Expected wages then take the following form,

$$\mathbb{E}[w(x') \mid x] \propto \left(\int_0^1 \left(w(x')/\tau(x, x') \right)^{\epsilon} dx' \right)^{1/\epsilon}.$$

This structure allows for productivity shocks at $Q(x')$ to affect housing demand across all locations in the city.

(2) Idiosyncratic location preference shocks Our framework also allows for households to have idiosyncratic location preferences that are iid Fréchet. In this formulation, each household solves,

$$\max_x u(x, t)\vartheta(x, t),$$

where $\vartheta(x, t)$ is Fréchet distributed with shape parameter $1/\theta$. Then, the analogous object to $\mathcal{U}(t)$ is ex-ante expected utility,

$$\mathcal{W}(t) \propto \left(\int_0^1 (u(x, t))^{\theta} dx \right)^{1/\theta}$$

and the pricing function is,

$$\mathcal{P}(x,t) = \left(\frac{\beta w(x,t)}{H(x,t)}\right)^{\frac{1}{1+\theta\beta}} \left(\frac{A(x,t)w(x,t)}{\mathcal{W}(t)}\right)^{\frac{\theta}{1+\theta\beta}},$$

where the term $\frac{\beta w(x,t)}{H(x,t)}$ substitutes for $n(x,t)$. This captures the fact that Fréchet shocks are isomorphic to introducing additional local congestion in an environment with free mobility (Allen and Arkolakis, 2014). Substituting for $p(x,t)$ in developers' BVPs, the optimal path of H at x is described by,

$$H(x,t)^{\frac{1}{1+\theta\beta}} \left(\rho\xi\dot{H}(x,t) + \xi\ddot{H}(x,t) + H(x,t)^{1/\gamma}\right) = (\beta)^{\frac{1}{1+\theta\beta}} \left(\frac{A(x,t)w(x,t)^{1+1/\theta}}{\mathcal{W}(t)}\right)^{\frac{\theta}{1+\theta\beta}}.$$

Thus this matches System (4) when $\theta \to \infty$. See Appendix A.4 for a full derivation.

(3) Neoclassical investment We now briefly lay out an alternative version of the developer problem which takes the more standard form of capital accumulation with depreciation subject to convex investment costs. Operating costs arise to offset depreciation.[7]

In this alternative formulation, we assume that a developer at x solves,

$$\rho\mathcal{V}(H,t) - \partial_t\mathcal{V}(H,t) = \max_I \{p(t)H - \Phi(I,H) + V_H \cdot (I - \delta H)\}$$

where $\Phi(I,H)$ is the cost to increase housing at rate I, given current H.[8] To obtain a differential equation that describes developers' investment path, we proceed as before. We time-differentiate the first order condition and combine with the envelope condition, assume $\Phi(I,H) = \frac{\xi}{2}I^2$, and use $\dot{H} = I - \delta H$ and $\ddot{H} = \dot{I} - \delta\dot{H}$, to derive,

$$-\xi\ddot{H}(x,t) + \rho\xi\dot{H}(x,t) + \xi\delta(\rho+\delta)H(x,t) = p(x,t). \tag{5}$$

[7]To be clear, our formulation separates κ and Φ, effectively assuming that developers have different operating costs and adjustment costs. The developer's problem then revolves around moving between operating scales subject to difficulty in adjustment. The neoclassical approach links the two through convex investment costs and implies a substantially more elastic long run supply of housing. Given that housing appears fairly inelastically supplied at granular levels in many U.S. cities (see, e.g., Baum-Snow and Han, 2024, who estimate very inelastic long-run housing supply at the tract level in many U.S. cities), we prefer our formulation that allows us to modulate the long-run supply elasticity through κ.

[8]That is, their problem is,

$$\max_{\{I(t)\}} \int_0^\infty e^{-\rho t}\left[p(t)H(t) - \Phi\big(I(t),H(t)\big)\right] \mathrm{d}t$$
$$\dot{H} = I - \delta H.$$

Equation (5) replaces the ODE in System (4), and otherwise the problem inherits the same housing demand structure. See Appendix A.5 for a full derivation.

(4) Endogenous wages Our framework additionally allows for wages to be endogenous to H, provided the coupling is local – i.e., there is no commuting or frictional trade.[9] To illustrate this in practice, we consider an example where each location produces a unique, freely traded intermediate good, whose variety is indexed by the location of production, x. There is a final goods firm that assembles these varieties into the final good with a constant elasticity of substitution technology with elasticity σ, and the final good remains the numéraire.

Production technology remains linear-in-labor, $y(x,t) = Q(x,t)L(x,t)$, and there is no commuting so that $L(x,t) = n(x,t)$ in equilibrium. Goods prices are simply marginal costs, $w(x,t)/Q(x,t)$. Market clearing for the x good at time t, expressed in wages, is,

$$w(x,t)n(x,t) = (1-\beta)\int_0^1 \left(w(x,t)/Q(x,t)\right)^{1-\sigma} w(s,t)n(s,t)ds\ .$$

Using the fact that the wage bill $w(x,t)n(x,t) = \frac{1}{\beta}p(x,t)H(x,t)$, and $n(x,t) = H(x,t)/h(x,t)$, we can solve for wages as a function of $p(x,t)$ and $H(x,t)$. To do so, we define city GDP as,

$$\mathcal{Y}(t) = \frac{1}{\beta}\int_0^1 p(x,t)H(x,t)dx.$$

Then wages can be expressed as,

$$w(x,t) = Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)\mathcal{U}(t)^{1/\beta}}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}$$

Using the definition of prices, we can write $\mathcal{Y}(t)$ as a function only of $A(x,t)$, $H(x,t)$, parameters, and $\mathcal{U}(t)$. Similarly, we can replace wages in the definition of $\mathcal{U}(t)$ as a function of $\mathcal{Y}(t)$. This means we can again write the pricing function $p(x,t)$ as a function of $A(x,t)$, $\mathcal{U}(t)$, and $\mathcal{Y}(t)$. In the state-space approach, this means $\mathcal{P}(x,\mathcal{S}) \equiv \mathcal{P}(A(x),\mathcal{U},\mathcal{Y})$. When $\sigma \to \infty$, wages are no longer endogenous to the spatial distribution

[9] When labor supply to any firm depends on the full spatial distribution of labor (e.g., through commuting), or goods prices depend on the spatial distribution of households because of trade costs that engender heterogeneous shifters of 'market access', then we can not use $\mathcal{U}(t)$ as a sufficient statistic for $\mathcal{S}(t)$ everywhere. It is in these situations, when the state space is larger than the number of locations, that dynamic spatial models become computationally difficult and researchers rely on approximations, like the using first order approximations to the master equation (the FAME) as in Bilal and Rossi-Hansberg (2023).

of H, and $\mathcal{Y}(t)$ is no longer needed to forecast prices. All derivations are available in Appendix A.6.

The key insight here is that allowing wages to respond to the distribution of H introduces another time-dependent variable. Now, there are 'coupled' interactions that operate through the labor market, as each developer's choice of H exerts infinitesimal influence on city GDP, $\mathcal{Y}(t)$. This is because in equilibrium, GDP responds to the spatial distribution of labor, as productivity $Q(x,t)$ is spatially heterogeneous. In short, when wages are endogenous, both moments $\mathcal{U}(t)$ and $\mathcal{Y}(t)$ of the aggregate state $\mathcal{S}(t)$ are required to forecast prices.

3 Numerical exploration of model behavior

Before we turn towards a quantitative application, we first demonstrate how our model behaves in toy economies. Our goal is to communicate intuition on how the model behaves.

3.1 The dynamic propagation of housing demand shocks across space

We consider a city composed of a large, discrete number of locations in the unit square.[10] We consider a city with commuting, following Extension 1, Section 2.3. We study how the paths of housing respond in equilibrium to a productivity shock to $Q(x')$ at a single location x'. This shock affects housing demand at all locations, but heterogeneously so: spillovers, which operate through adjusting expected wages, decay with distance to x'.

In Figure 1, we display the results of an exercise in which productivity is positively shocked at $(0.25, 0.5)$. Given the commuting structure of the economy, this raises expected wages everywhere, with the strength of the spillover decaying in distance to the treated site. In the left panel, we show how $H(x)$ responds just after the impact of the shock, before developers have fully adjusted. The solid black line on the figure displays the boundary at which there are zero changes to the housing stock, while dashed black line displays changes that are equal to a 1.5% increase in H from baseline. In the right panel, we show the steady state, retaining the

[10] This may seem contrary to our description of space as $x \in [0,1]$. The labeling $x \in [0,1]$ is an arbitrary ordering of space. As we are working with a large but finite number of locations, we can map them onto $[0,1]$. Moreover $[0,1] \times [0,1]$ has the same cardinality as $[0,1]$.

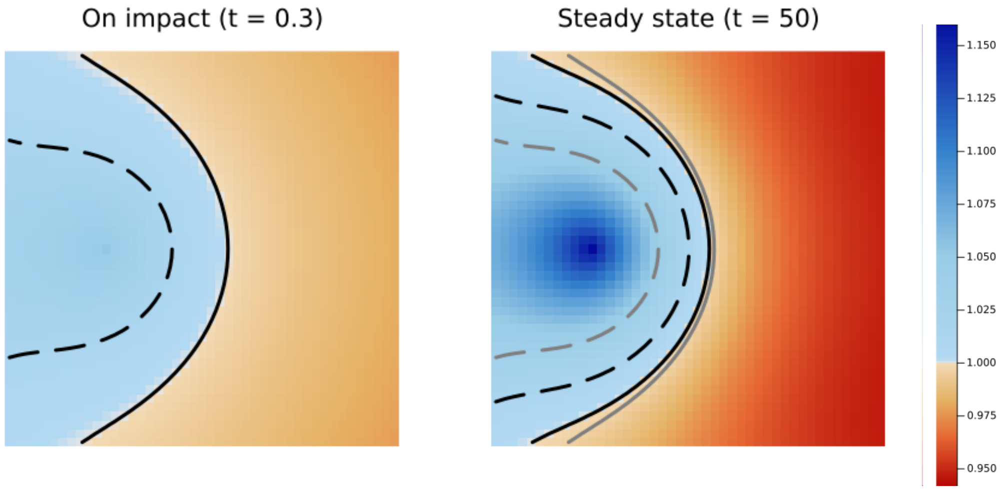

Figure 1: The dynamics of housing adjustment following productivity shock to Q at (0.25,0.5). Left: the spatial distribution of $H(x,t)$ just after impact. Right: The spatial distribution of $H(x,t)$ in the steady state. The dashed black line indicates the contour for 1.5% increases in H, while the solid black line indicates the 0% change contour. Grey lines in the right panel are the same contours as in the left panel. We set $A(x,t) = Q(x,t) = 1$ for all x and t, and let iceberg commuting costs be $\tau(x,x') = \exp(-\kappa\|x-x'\|)$. We additionally set $\kappa = 1.5, \beta = 0.25, \gamma = 1.85, \xi = 200$ and $\epsilon = 2.1$.

original contour lines (in grey) and plotting new ones (in black). The spatial dynamics are non-monotonic: between the solid black and grey lines are regions for which housing supply initially grew, but shrink in at the steady state. The reason this region exists is that initially, housing supply cannot meet housing demand, so workers spread over space. As housing supply near the treated location adjusts, nearby locations absorb workers. In short, the initial 'blast radius' of the shock is larger than its steady state level, because short-run housing supply is more inelastic.

In this example, we have $N = 2,025$ locations and time runs for 50 periods before settling to a steady state. Despite this large space, we are able to solve for the full equilibrium transition dynamics at every location in less than a minute with fewer than 185 lines of Julia code.

3.2 Zooming in: non-monotonic transition dynamics and the importance of forward-looking behavior

We now examine the time-paths of housing and zoom in on the non-monotonic nature of the transition dynamics in an economy with 3 locations labeled $x = 1, 2, 3$. We consider a permanent change in the distribution of amenities $A(x)$. In this example, there are 3 locations, and they are initially identical in every

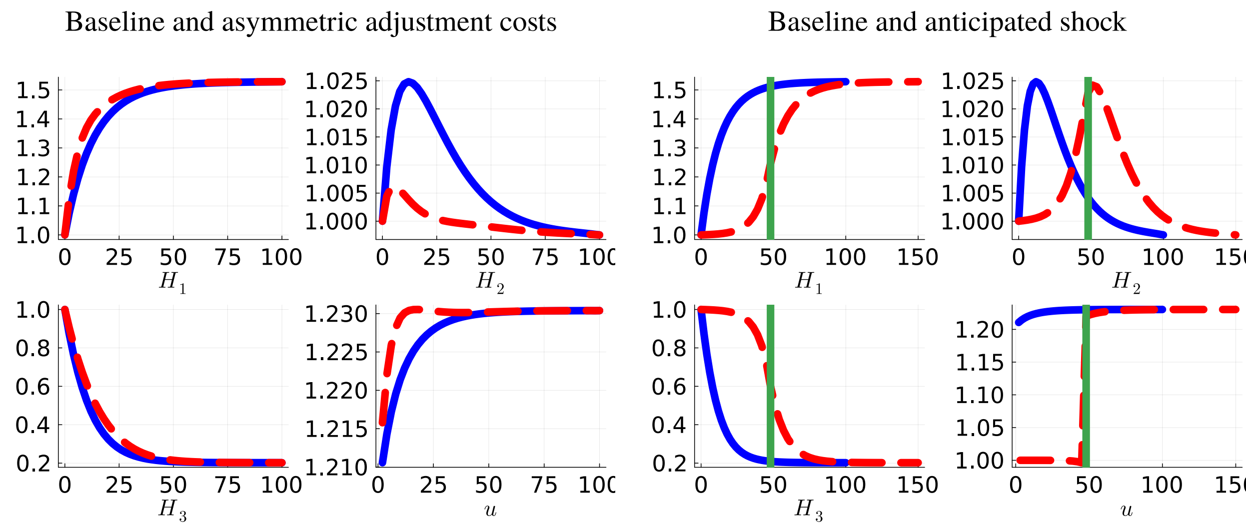

Figure 2: The transition dynamics of housing in response to a demand shock at locations 1 and 2. Left panel: housing paths with symmetric adjustment (blue), and asymmetric adjustment costs (red). Right panel: anticipation effects of a shock announced at $t = 0$ that arrives at $t = 50$ (in green). Blue paths are the same as those in the left panel, while red paths display the transition dynamics when the shock is announced.

dimension. At time $t = 0$, a permanent shock occurs, which immediately increases $A(1)$ dramatically, increases $A(2)$ moderately, and does nothing to $A(3)$.

The left panel of Figure 2 displays the time paths of housing in each location, as well as the household utility level in blue. The paths for H_1 and H_3 are standard: housing demand shifts towards the location with the better amenity, so long-run equilibrium housing increases in 1 and decreases in 3, and the adjustment is not immediate due to the developer cost of adjusting housing. Similarly, the path for household utility increases in the long-run, but is sluggish because the sluggish response of developers amounts to a short-run supply constraint, maintaining high short-run prices. In location 2, housing increases in the short-run, but decreases in the long-run. This non-monotonicity is due to the convexity in housing adjustment costs. In the short-run, location 1 developers are slow to adjust the housing stock, so developers in location 2 find it optimal to increase housing supplied and absorb excess demand at location 1. After some time, however, location 1 developers increase housing supply enough that it is no longer optimal for location 2 developers to continue building, and begin to decrease the housing stock to its long run level.

In short, developers smooth investment in a way that generates a nonmonotonic transition path at location 2, which we view as a counterintuitive reversals in housing and population growth over time. Obtaining this result is only possible in a framework where housing adjustment dynamics are frictional, as any static model will only be able to speak to long-run outcomes, or perhaps immediate short-run outcomes, but never the shape of the time path of housing in the medium-run.

Asymmetric adjustment costs Our baseline experiment above assumed symmetric adjustment costs, so that developers increasing total housing at a rate of 1 unit/t pay the same adjustment cost as those decreasing housing 1 unit/t. This assumption is unrealistic; housing is durable and not demolished piecemeal. This means that housing supply appears particularly inelastic in response to negative demand shocks (Glaeser and Gyourko, 2005). To account for this feature, we simply modify the adjust cost Φ to include asymmetry as follows,[11]

$$\Phi(\dot{H}) = \begin{cases} \frac{\xi_U}{2}(\dot{H})^2 & \dot{H} \geq 0 \\ \frac{\xi_D}{2}(\dot{H})^2 & \dot{H} < 0. \end{cases}$$

We assume $\xi_U < \xi_D$, so that downward adjustment is more costly than upward adjustment, and fix $\frac{\xi_U + \xi_D}{2} = \xi$. In the left panel of Figure 2, we plot the transition paths for housing when adjustment costs are asymmetric with dashed red lines. As developers at location 2 seek to avoid the costly 'building down' of the housing stock, they do not react as strongly to the excess demand for location 1. To smooth adjustment costs over time, developers at location 3 build down less rapidly. In this example, the swift decline in prices at location three are not large enough to discourage migration to location 1. Faced with more pressure to build, developers at location 1 build faster.

Anticipated demand shock As our model features developers that are forward looking, we can consider how the announcement of shocks influences dynamic equilibrium behavior. In the right panel of Figure 2, we consider the case when the shock to $A(x)$ is announced at $t = 0$ and arrives at $t = 50$. The long-run outcome of this shock is the same as above, as all locations face the same long-run fundamentals. However, in anticipation of increased future housing demand, developers slowly build up or build down the stock of housing to smooth costly adjustment. As no fundamentals have yet shifted for households, the population distribution only reacts to changes in prices engendered by changing housing supply, but jumps on impact of the shock.

[11] In practice, we use a smooth approximation to this Φ so that Φ'' is twice-differentiable at 0. This approximation respects that Φ'' is the same as in the symmetric case to which we are comparing, and the asymmetry only comes through when $\dot{H} \neq 0$.

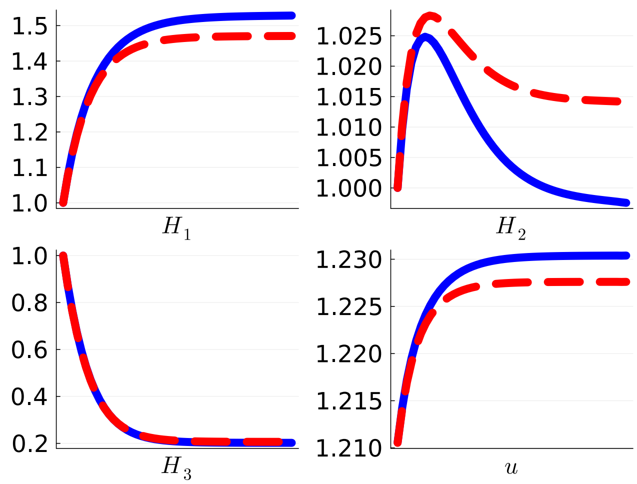

Figure 3: Transition dynamics of H when Z_1 is smaller than the long-run H_1 absent the policy.

Zoning We now incorporate a particular policy that is relevant for our quantitative exercise. For a variety of reasons, cities may restrict building height or floorspace allocations in some parts or all of a city. In the context of our model, this is simply a policy restriction $H(x) \leq Z(x)$, where $Z(x)$ is set of a central authority, and we refer to any policy Z under the umbrella term *zoning*.

Unlike a tax or subsidy, zoning policies directly affect feasible building choices for developers, and thus cannot be mapped into simple parameter changes in the model above, so are not nested in the framework we have discussed thus far. We incorporate zoning by replacing $\kappa(H)$ with a function $\tilde{\kappa}(H; Z)$ that approximates $\kappa(H)$, except near Z, where it approaches infinity.[12] We now show how binding zoning affects the baseline experiment from above.

Consider the same $t = 0$ permanent change in A as above, except we set Z_1 to be smaller than than the long-run H_1 were zoning not present. Figure 3 shows the resulting housing paths without zoning (blue) and with zoning (dashed red). Location 1 is no longer able to increase its housing to the full extent from the baseline, and this allows housing demand to permanently shift towards location 2 increasing long-run H_2 (this occurs for H_3 as well, the magnitude is just small in this particular parameterization). Additionally, the

[12]That is, we approximate the combination of the maintenance cost $\kappa(H)$ and the constraint $H \leq Z$ with an altered $\tilde{\kappa}(d, H)$ that satisfies: (i) $\tilde{\kappa}(d, 0) = 0$, for all d; (ii) $\lim_{H\to Z} \tilde{\kappa}(d, H) = \infty$, for all d; (iii) $\lim_{d\to\infty} \tilde{\kappa}(d, H) = \kappa(H)$ if $H < Z$ and ∞ otherwise. The d parameter proxies for how close the function $\tilde{\kappa}$ function is to the true constraint. The advantage of using this formulation is that the above machinery can now be used by simply replaced κ with $\tilde{\kappa}$, and avoid having to directly deal with the additional constraint function imposed by zoning in each developer's problem. Attacking this problem directly would require each developer to keep track of a continuum of constraints $H(x, t) \leq Z(x, t)$ when solving for the optimal housing path, and in general any measure of the continuum of constraints may be slack. Our proxy approach yields quantitatively similar results to the direct solution of the problem, but sidesteps the additional analytic difficulties.

increase in long-run H_2 will mean that any short-run building up is less costly, because less of it will need to be demolished, so developers in location 2 increase short-run housing even more. This example shows that zoning constraints affect the transition dynamics by creating spatial spillovers that persist over time.

4 Quantitative model

In this section we present a quantitative model that we use to study the dynamic, spatial propagation of shocks in a city with zoning restrictions. In the model inherits a similar structure to the model presented in Section 2. It differs on two dimensions: wage determination and housing supply. We work with discrete space (census tracts in the San Francisco-Oakland-Fremont MSA), and so we index census tracts by i or j.

In the quantitative model, a continuum of measure $\bar{N}$ agents are freely mobile across locations indexed by i and have Cobb-Douglas utility over housing floorspace, on which they spend a fraction β of their expenditure, and the freely traded numéraire final good produced everywhere. At each location, their utility is multiplicatively shifted by the amenity value of that location, A_i.

Goods technology and the labor market In the quantitative model, we allow for commuting, following Extension (1), Section 2.3. All workplaces operate a constant returns technology to produce the freely traded final good, $y_j = Q_j L_j$, where L_j is measured in efficiency units of labor supplied to the firm.

Once each household ν chooses their location i, they then choose their workplace, j. Their choice of workplace depends on a productivity shock $\varepsilon_j(\nu)$ that scales the number of efficiency units they supply that is iid Fréchet with shape parameter $1/\epsilon$. However, commuting is costly, and so conditional on choosing i, they can supply $\varepsilon_j(\nu)/\tau_{ij}$ efficiency units of labor to each location. They choose their workplace j to maximize their factor reward, $w_j \varepsilon_j(\nu)/\tau_{ij}$. Subsequently, conditional on living in i, their expected wage,

$$\mathbb{E}[w_j \mid i] \propto \left(\sum_j \tau_{ij}^{-\epsilon} w_j^{\epsilon} \right)^{1/\epsilon} .$$

Housing services We incorporate zoning into the quantitative model by employing ideas developed in Martynov (2022) and Baum-Snow and Han (2024). Housing at each location H_i is composed of floorspace built atop parcels ω, $f(\omega)$. The measure of parcels per location is T_i. Floorspace per parcel is maintained using the same technology described in Section 2, shifted by a parcel-specific operating efficiency $e(\omega)$, so that when developers profit-maximize, at a stationary equilibrium, and zoning constraints do not bind, floorspace per parcel is given by,

$$f_i(\omega) = (e_i(\omega)p_i)^\gamma.$$

Zoning, Z_i, is a restriction on the amount of floorspace that can be built per parcel; $f_i(\omega) \leq Z_i$. When choosing over which parcels to develop, developers use the most efficient parcels first. Parcel efficiency is drawn from a distribution, $e_i(\omega) \sim G_i$. The efficient parcels are built up until the zoning constraint binds. We denote the marginal efficiency that equates the free-market level of floorspace with the zoning restriction $\tilde{e}_i = Z_i^{1/\gamma}/p_i$. Thus, total floorspace can be denoted,

$$H_i = T_i \left[\int_{E_i}^{\tilde{e}} p_i^\gamma e^\gamma dG(e) + \int_{\tilde{e}}^{\infty} Z_i dG(e) \right],$$

which decomposes the measure of parcels into those for which zoning binds $e > \tilde{e}$ and those for which zoning does not bind. For tractability, we assume parcel efficiency is Pareto distributed with location-specific scale parameter E_i and shape parameter φ. In each instant, the Pareto efficiency shocks are redrawn. Households are free to adjust the distribution of floorspace across parcels in each period without cost, and adjustment costs only depend on changes to the total quantity of floorspace H_i at the tract level.

This structure on housing supply allows us to infer the shape of a tract-level housing supply curve. Overall tract-level housing supply integrates over parcel-level floorspace services supply curves. However, whenever parcel floorspace is provided to the point which the zoning regulation binds, we are on the vertical component of the supply curve and cannot infer parameters that govern housing supply on the non-vertical component of supply. Our provided structure allows us to use aggregate data on the share of parcels for which zoning binds, total parcel area, and housing prices, to construct tract-level supply curves in the presence of zoning. Our technique also allows us to infer the appropriate average zoning restriction, Z_i, at the tract level with the same data.

Stationary equilibrium Given model parameters $\{\beta, \gamma, \varphi\}$ and location specific fundamentals, $\{A_i, E_i, Q_j, Z_i\}$, a stationary spatial equilibrium is a set of floorspace prices p_i and wages w_j, such that

1. taking prices as given, households optimally choose h, and households choose locations i to maximize ex-ante utility u_i, and choose workplaces j to maximize their received wage, $w_j L_j(\nu)/\tau_{ij}$;

2. goods firms profit maximize, so that wages are marginal products $w_j = Q_j$;

3. developers optimally choose how much floorspace to provide;

4. and a spatial equilibrium holds, $u_i = \mathcal{U}$.

Stationary equilibrium characterization The utility of living in any location i is,

$$u_i = \frac{A_i \left(\sum_j \tau_{ij}^{-\epsilon} \bar{Q}_j^{\epsilon}\right)^{1/\epsilon}}{p_i^{\beta}} \implies p_i = \left(\frac{A_i \left(\sum_j \tau_{ij}^{-\epsilon} Q_j^{\epsilon}\right)}{\mathcal{U}}\right)^{1/\beta}$$

as u_i must be everywhere equalized at a spatial equilibrium. When households optimally pick workplaces j, commuting flows are given by,

$$N_{ij} = \frac{\tau_{ij}^{-\epsilon}(Q_j)^{\epsilon}}{\sum_k \tau_{ik}^{-\epsilon}(Q_k)^{\epsilon}} N_i.$$

When developers optimally choose floorspace, total long-run tract-level floorspace supply is,

$$H_i = T_i(E_i p_i)^{\gamma} \left(\frac{\gamma}{\gamma - \varphi} Z_i^{1-\varphi/\gamma} E_i^{\varphi-\gamma} p_i^{\varphi-\gamma} - \frac{\varphi}{\gamma - \varphi}\right), \tag{6}$$

which is displayed in Appendix Figure A1.

Dynamics As $\kappa'(H_i)$ is the long-run supply curve, in our case it is implicitly defined through Equation (6). As adjustment costs depend on H_i and not the distribution of floorspace across parcels, use $\Phi(\dot{H}_i) = \frac{\xi}{2}(\dot{H}_i)^2$.

Then the dynamic equilibrium is, given paths of $A_i(t), w_i(t)$ a set of paths $\{n_i(t), H_i(t), p_i(t)\}$ such that conditions (i)-(v) of the dynamic equilibrium in Section 2 are satisfied.

4.1 Quantifying the model

We quantify the model by first assuming that the observed allocations in San Francisco in 2018 reflect a stationary equilibrium of the model. To rationalize data as a steady-state equilibrium of the model, we require that the data identify parameters A_i, E_i, Q_i, Z_i as well as commuting costs τ_{ij}, given spatially invariant parameters β, γ, ϵ and φ. To do this we use data on housing rents, floorspace, zoning, and commuting patterns in San Francisco in 2018. To do this, we use data on floorspace prices p_i, residency n_i, commuting flows n_{ij}, parcel area T_i, and the share of parcels for which the zoning constraint is binding, which we call S_i.

Our choice of parameter values is standard in the literature. We choose $\beta = 0.24$, following Davis and Ortalo-Magné (2011), who use U.S. Census data to estimate that 24% of income on average is spent on housing. We set $\gamma = 1.85$, following Combes et al. (2021), who estimate land's share in housing production as $(1+\gamma)^{-1} = 0.35$ using French data. We set $\epsilon = 2.18$, based on the estimate in Severen (2023), who estimates this parameter using labor demand instruments and data from Los Angeles. We set $\varphi = 12$, implying little dispersion in parcel productivity, so long-run floorspace supply curves closely resemble those used in Section 3.[13]

4.1.1 Data

Our data comes from several sources. We aggregate all data to Census tracts in San Francisco county (i.e., the city of San Francisco) and aggregate all data to the level of the county for all other counties in the San Francisco commuting zone (e.g., Alameda county, which includes Oakland and Berkeley).

Rents Prices p_i come from Pennington (2021), who collects geocoded rental data from Craigslist ads for apartment rentals in the Bay Area from 2000-2018. We residualize these observed rents on year and quarter fixed effects, as well as observed covariates, like number of bedrooms, bathrooms, and squarefootage using a standard hedonic regression. We then scale the regression residual so rents on average match the price of a 1 bedroom apartment in the Bay Area in 2018. We then aggregate these averages to the Census tract level and

[13] We do not estimate φ and do not know of good estimates from the literature. In principle, the relationship that $\log p_i = \zeta_i + \varphi^{-1} \log S_i$ could identify φ, provided that zoning stringency, S_i, was randomly assigned. In areas with tougher zoning, prices would be higher. In reality, zoning is far from randomly assigned: developers petition more strongly for zoning relaxation in areas where demand is highest, and zoning goes unchallenged when it puts little upward pressure on rents. In effect, the relationship between zoning stringency and prices is *negative*. To estimate φ, a researcher would need either an instrument for zoning uncorrelated with unobserved shifters of price, or time-varying information on S_i and p_i that would allow for location fixed effects, neither of which can implement. We view the use of structural models of zoning combined with quasi-experimental variation to discipline the relationship between observed prices and regulations as an exciting avenue for future research – the U.S. historical record is replete with natural experiments in housing development.

use a distance-based kernel to smoothly interpolate rents for the handful of Census tracts with no available rental data.

Commuting data To measure commutes, we use counts of primary job flows from home to work tracts, N_{ij} using Census LEHD-LODES commuting flows for 2017 and 2018. We use the model's implied gravity equation to recover τ_{ij} and Q_j. We estimate the regression,

$$N_{ij} = \tau_{ij}^{-\epsilon} \xi_i \xi_j,$$

with a Poisson estimator with origin and destination fixed effects (Silva and Tenreyro, 2006; Correia et al., 2020). To do so, we assume that $\tau_{ij}^{-\epsilon} = \exp(-\epsilon X_{ij}' \delta) t_{ij}^{-\epsilon}$ where X_{ij} are observable bilateral components of distance and for t_{ij}, the unobservable component of distance, $\mathbb{E}[t_{ij}^{-\epsilon} \mid X_{ij}, \xi_i, \xi_j] = 1$. We then recover estimates of $\tau_{ij}^{-\epsilon}$ using the predicted values from the gravity regression, $\hat{\tau}_{ij}^{-\epsilon} = \exp(-\epsilon X_{ij}' \hat{\delta})$. This procedure effectively smooths the data to account for the sparsity of the commuting matrix; that is, we account for granular uncertainty, which can lead to poor performance of counterfactuals in granular spatial models (Dingel and Tintelnot, 2020). We include in X_{ij} drivetime, transit, and walktime durations from querying Google Maps API for all pairwise tracts in San Francisco and its surrounding counties. We use Google's standard traffic model and query durations for an anticipated Monday, 9am arrival. Using the destination fixed effects, $\hat{\xi}_j$, we estimate Q_j^{ϵ}, and combine these estimates to recover the expected wage at each location, $E[w_j \mid i] = (\sum_j \hat{\tau}_{ij}^{-\epsilon} \hat{\xi}_j)^{1/\epsilon}$.

Zoning and floorspace data We use building footprint and height data San Francisco Enterprise Geographic Information Systems Program. The data reflect a 3D topographic map of the city for the year 2010. Additionally, we use the Zoning Districts shapefile from DataSF. Building footprint squarefootage gives us T_i. In Section 4.1.2, we show how we use this data to identify S_i, the share of parcels in a lot which up against the zoning constraint. With S_i in hand, we can use the model's structure to recover Z_i and E_i.

Relying on the assumption that parcel efficiency is Pareto distributed and developers optimally select which parcels to develop until the zoning constraint binds, we recover,

$$S_i = (E_i / Z_i^{1/\gamma})^{\varphi} p_i^{\varphi}.$$

When this relationship is inserted into the floorspace supply equation, we can invert the equation analytically to recover,

$$Z_i = \left[\frac{\gamma}{\gamma - \varphi} S_i - \frac{\varphi}{\gamma - \varphi} S_i^{\gamma/\varphi}\right]^{-1} \frac{H_i}{T_i}$$

which solves for Z_i when the floorspace market clears, $H_i = \beta w_i n_i / p_i$. With Z_i in hand, we can directly invert the S_i equation to recover E_i for each parcel.[14]

Recovering amenities From the spatial equilibrium condition, amenities can be recovered up to scale since,

$$A_i = \mathcal{U} \times \frac{p_i^\beta}{\left(\sum_j \tau_{ij}^{-\epsilon} Q_j^\epsilon\right)^{1/\epsilon}}.$$

As $\mathcal{U}$ is a function of some average of the A_is, it cannot be recovered from data if the A_is are not identified up to scale. Thus, we set it to 1 at baseline.

4.1.2 Identification of zoning cutoffs

Zoning, as prescribed by San Francisco Municipal Code, is multifaceted. However, we view the economic *consequences* of zoning as twofold: *de facto* restrictions on the amount of habitable floorspace per parcel, and changes in local amenity value. The economic benefits of zoning, by creating harmonious neighborhoods, or simply realigning the spatial distribution of economic activity so as to internalize urban externalities (as in Rossi-Hansberg, 2004b) are not the focus of our paper. The amenity benefits are captured cross section by location-specific amenities, A_i, and we hold land-use patterns (i.e., where housing and production occur in space) fixed. We focus our attention on the restrictions on available floorspace per unit land (captured by building height). For example, city code describes that parcels zoned "RH-1,"

> *"...are occupied almost entirely by single-family houses on lots 25 feet in width, without side yards. Floor sizes and building styles vary, but tend to be uniform within tracts developed in distinct time periods. Though built on separate lots, the structures have the appearance of small-scale row housing, rarely exceeding 35 feet in height. Front setbacks are common, and ground level open space is generous. In most cases the single-family character of these Districts has been maintained for a considerable time. "* §209.1, San Francisco Municipal Code (2024)

[14] For locations outside of San Francisco county, we assume there is no zoning, $S_i = 0$, and that total floorspace is comparable to that of the city itself.

Approximately 49% of residential or mixed use structures and 26% of residential or mixed use built volume in San Francisco are zoned RH-1. While code occasionally specifies exact building height limits,[15] these *de jure* limits do not align with *de facto* restrictions on floorspace, which come from the assemblage of other regulations, including mandated front setbacks, rear yards ("30% of lot depth, but in no case less than 15 feet"), side years, residential design guidelines, street frontage and parking regulations, and so on. These regulations limit the types of structures developers can build, and in effect limit the available floorspace per parcel.

To measure *de facto* floorspace restrictions, we look for bunching in the distribution of building height within in residential zoning class. Our strategy is similar to that of Brueckner et al. (2024), who use a bunching estimator and building height data in New York City to assess the costs of height regulation. In contrast, our strategy is designed to recover the structural primitives Z_i for use in counterfactual analysis. In line with our model of development under zoning, the basic idea of our empirical strategy is that the distribution of building height will be right-truncated at the zoning limit, Z_i. In the data, not all parcels are built until the zoning restriction (and others are built beyond it, due to 'grandfather clauses,' zoning variances, illegal construction, and so on), but there is bunching in the height distribution within each zoning class, which we take as indicative of an unobserved Z_i.

We detect Z_i by looking for discontinuities in the building height density by residential zoning class. Inferring cutoffs by looking for bunching in the distribution of an endogenous variable is an approach used in the education literature (see, e.g., Mountjoy (2024) using ACT cutoffs for public universities), and is in spirit similar to the 'structural break' identification for determining realized zoning, as in Song (2021) and Cui (2024). We implement our strategy by looking for changes in the derivative of the empirical cumulative distribution function (CDF), where we can employ a regression 'kink' design with unknown threshold (similar to Hansen (2017), but without an outcome variable). Of course, looking for changes in the derivative of a random variable's CDF is equivalent for looking for discontinuities in its density. However, smoothly approximating the density around a discontinuity encounters empirical problems, like specifying parameters like binwidth when implementing tests like that of McCrary (2008) for manipulation testing at a cutoff. Instead, we follow Cattaneo et al. (2020) in using the empirical CDF, which mitigates these concerns.

[15]e.g., for RH-1, code stipulates that "[no] portion of a Dwelling may be taller than 35 feet. Structures with uses other than Dwellings may be constructed to the prescribed height limit, which is generally 40 feet. Per §261 the height limit may be decreased or increased based on the slope of the lot."

In particular, we estimate a 'regression kink' (RK) design to test for discontinuities in the derivative of the empirical CDF of building height. Given the building height distribution of residential buildings in zoning class C (whose elements we index by indexed by i), we estimate the change in slope of the empirical CDF ($\text{Percentile}_{i,C}$) at each potential kink K by estimating by weighted least squares,

$$\text{Percentile}_{i,C} = \alpha + \beta_C(K)1(\text{Height}_{i,C} - K) \times 1(\text{Height}_{i,C} < K) + f(\text{Height}_{i,C}, K) + u_{i,C}, \quad (7)$$

where f is a second-order polynomial in the running variable on each side of the kink. Around each kink K, we estimate this regression over a bandwidth of 12 feet using quadratic kernel weights.[16] We estimate $\beta_C(K)$ over the distribution of $\text{Height}_{i,C}$ in each zoning class, and use,

$$\hat{Z}_C = \text{argmax}_K\, \hat{\beta}_C(K)$$

as the estimate of the building height restriction. We estimate Z_C for 11 different zoning classes. In Appendix Figure A2, we map our residential zoning maps, which capture over 98% of residential zoning classifications in San Francisco county. Appendix Table A1 reports each class' share of structures, built volume, and the measured height restriction.

To compute tract-level analogues of Z_C, and to moreover convert Z_C to model-relevant units, we compute S_i as the share of area that is built to a height at or above Z_C for each zoning class C. This is because in practice, there are often several different zoning classifications within a tract. Thus S_i is an average measure of the extent to which zoning restrictions bind, weighted by the shares of area within a parcel that belong to each zoning class.

The outcome of this empirical strategy is illustrated in Figure 4 for RH-1. In the left panel, the empirical CDF of the height distribution is plotted, as is the estimate of the local change between the left and right derivatives at each point, from estimating (7). The dashed line corresponds to the height cutoff K which maximizes $\beta(K)$ estimate. The same dashed line appears in the histogram of building height (right panel), where bunching near the estimated cutoff is visibly apparent. The cutoff, about 21 feet in height (approximately two storeys) is well below the legal height limit of 35 feet, and reflects the fact that it is difficult to build taller structures

[16]Experimentation with the choice of bandwidth or weighting kernel had no noticeable effect on the results and we do not report robustness to this choice in the paper.

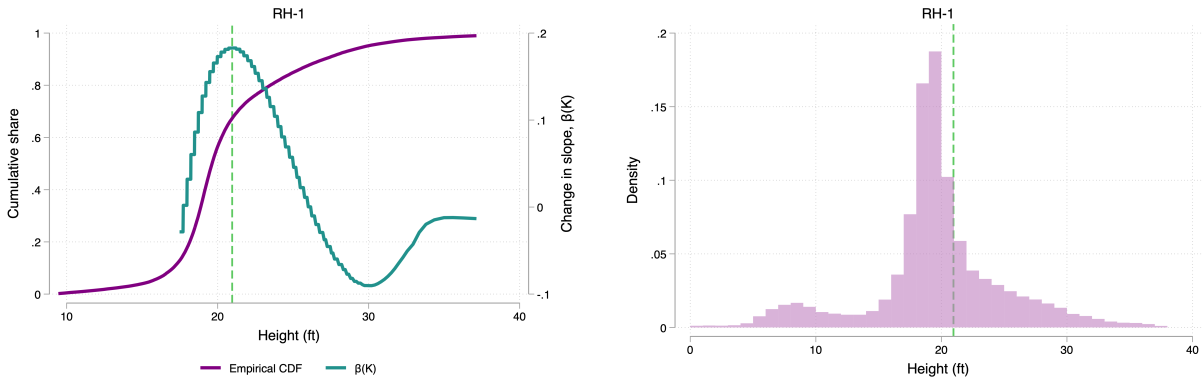

Figure 4: Illustration of the results of the RK design for structures zoned RH-1. Left panel: empirical CDF and point estimates of β across the height distribution. Right: histogram of building heights. The estimated Z_C is displayed by a light green dashed line in both figures.

that adhere to RH-1 requirements.

Inversion results Our parameters are correlated with related data. In Appendix Figure A3, we display observed built volume and recovered H_i, as well as observed average height and Z_i. Both show a positive correlation. In Appendix Figure A4, we display the results of the inversion for p_i, w_i, A_i, and Z_i. Amenities roughly track floorspace prices, though are markedly lower in the Tenderloin – potential scarring from the fentanyl epidemic. Zoning appears most restrictive in the Marina District, other neighborhoods around the Presidio, and the Sunset District.

5 Results

We now use the quantified model to explore two housing demand and supply shock experiments. To shock housing demand, we consider a counterfactual in which wages rise in the downtown tract in San Francisco. This raises wages everywhere through the commuting network. We view this 'return to the downtown' as a counterfactual in which the loses to downtown employment stemming from the 'tech' downturn and Covid-19 pandemic, are reversed. As a housing supply shock, we consider upzoning the tracts along the light rail corridor which abuts Golden Gate Park from the south.

We calibrate our dynamic parameters by setting $\rho = 0.05$, for a 5% annual discount rate, and $\xi = 1e - 4$. The choice of ξ allows our exercises to be illustrative in a qualitative sense and relative sense, but is not meant

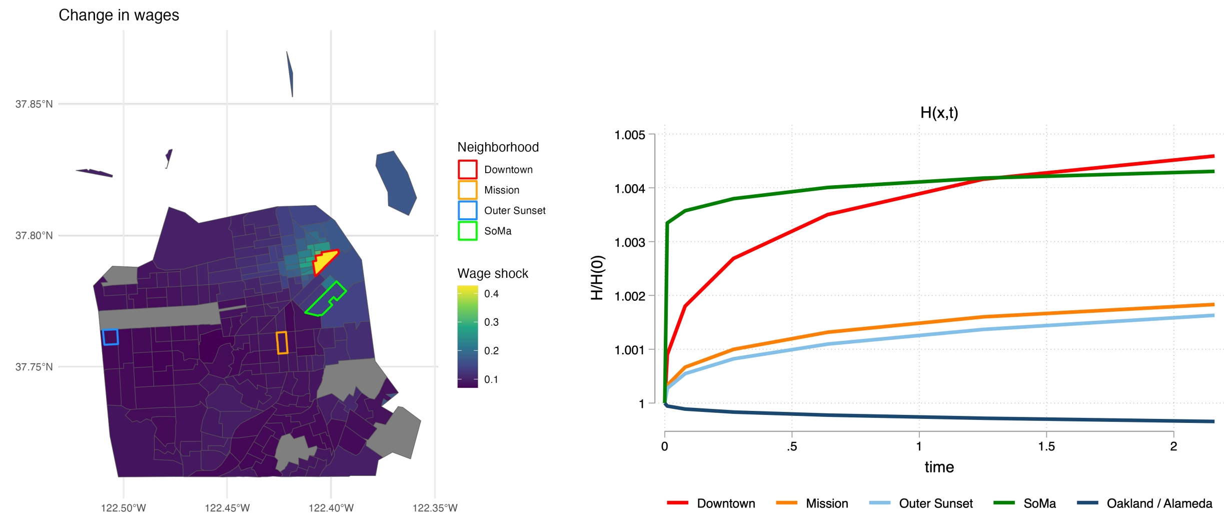

Figure 5: Left: geography of the wage shock in San Francisco. Right: time paths of the evolution of the housing stock in some select neighborhoods.

to be taken as serious calibration of adjustment costs.

5.1 Return to the downtown

We assume the baseline equilibrium is a stationary equilibrium of our model, and we shock wages in the downtown tract in San Francisco by 2.5%. This is the tract that contains the headquarters of, e.g., Wells Fargo Bank, the Salesforce Tower, and the Transamerica Pyramid.

Figure 5 displays the results of this exercise. The left panel shows the wage shock we feed into the model across tracts in San Francisco, and highlights several tracts. We highlight the treated downtown tract (red); the Mission (orange), which is well connected to the downtown via the heavy-rail BART transit system; SoMa (South of Market, green); and a tract in Outer Sunset (blue), which is far from the downtown but connected on the Muni light rail system.

The right panel shows the time paths of housing development at each location, as well as total development in one of the commuter counties, Alameda, which contains the city of Oakland. SoMa, which has lax zoning at baseline, develops rapidly but hits its zoning constraint, while the very dense downtown tract, develops more slowly, but ultimately sees a rise in its housing stock (in percent terms) that exceeds that of SoMa.

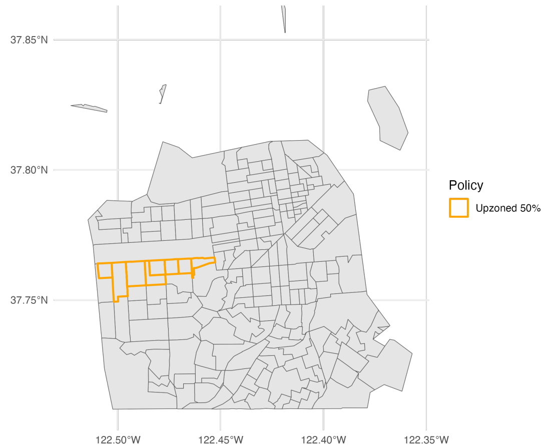

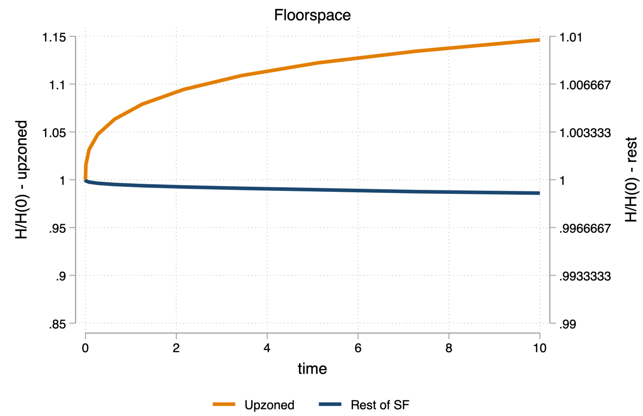

Figure 6: Left: the upzoned tracts. Right: time paths of the evolution of the housing stock in the upzoned neighborhoods and the rest of San Francisco.

Commuter neighborhoods like the Mission and Outer Sunset grow slowly but steadily, while the housing stock in Oakland slowly builds down.

5.2 Upzoning the Sunset District

We now consider a counterfactual in which some of the tracts in the Sunset District – those which abut Golden Gate Park and through which the Muni light rail runs – are upzoned, meaning that the housing supply constraints are relaxed. To do so, we raise Z_i in these tracts by 50%. This counterfactual aligns with the common policy recommendation that housing supply ought to be increased around transit hubs.

Figure 6 displays the results of this exercise. The left panel shows the geography of the treated tracts, while the right panel shows the time paths of the housing stock in the upzoned tracts and the rest of the city.

While a demand shock incentives rapid development, this supply shock generates a much more protracted response as the city slowly responds to this change in the housing supply elasticity at these locations. Despite allowing 50% more development at these tracts, in the long run, the housing stock only grows by 15%. This is because despite zoning binding at baseline, demand pressures are not too strong at baseline.

6 Conclusion

This paper develops a dynamic spatial equilibrium model of housing development in cities. In our model, competitive developers make forward-looking decisions on how to build the housing stock and face adjustment costs. While the elements of our model are simple, the spatial dynamic equilibrium is, in general, hard to solve because of the large state space. By casting the problem as a mean-field game, we show that the equalized value of utility across sites in the city is a sufficient statistic for the aggregate state – the spatial distributions of population, housing, and amenities – that allows developers to forecast prices. This insight allows us to dramatically reduce the dimensionality of the state space and globally solve the transition dynamics. Our approach stands in contrast to existing models that rely on first-order approximations.

In our model, we find that developers' forward-looking behavior can generate non-monotonic transition dynamics, demonstrating how expectations, adjustment costs, and zoning constraints shape the evolution of cities in ways that static or myopic models fail to capture. Moreover, our model is easy to implement and solve on the computer, making it accessible for researchers seeking to model the transition dynamics when studying policy counterfactuals in quantitative urban models.

To illustrate our framework in practice, we calibrated our model to match the San Francisco MSA, and augmented our model with zoning regulations. These zoning regulations generated spatially heterogeneous long-run housing supply curves. We estimated latent zoning parameters using data on building heights and zoning classifications within the city, and used the model to study the dynamics associated with housing supply and demand shocks. We found that the equilibrium rate of housing adjustment was spatially heterogeneous, suggesting that the short-run responses to shocks in this urban environment are not sufficient to infer the long-run equilibrium response.

In short, we provide a tractable yet flexible model of dynamic urban development, offering new tools for researchers interested in studying the dynamic response of the spatial distribution of economic activity within cities to policy interventions or economic shocks.

References

Achdou, Yves et al. (2014). “Partial differential equation models in macroeconomics”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372.2028, p. 20130397.

Ahlfeldt, Gabriel M. et al. (2015). “The Economics of Density: Evidence from the Berlin Wall”. In: *Econometrica* 83.6, pp. 2127–2189.

Allen, Treb and Costas Arkolakis (2014). “Trade and the Topography of the Spatial Economy”. In: *The Quarterly Journal of Economics*, pp. 1085–1139.

Barilla, Cesar, Guillaume Carlier, and Jean-Michel Lasry (2021). “A Mean Field Game Model for the Evolution of Cities”. In: *Journal of Dynamics and Games*.

Baum-Snow, Nathaniel and Lu Han (2024). “The microgeography of housing supply”. In: *Journal of Political Economy* 132.6, pp. 1897–1946.

Bilal, Adrien (2023). *Solving heterogeneous agent models with the master equation*. Tech. rep. National Bureau of Economic Research.

Bilal, Adrien and Esteban Rossi-Hansberg (2021). “Location as an Asset”. In: *Econometrica* 89.5, pp. 2459–2495.

— (2023). *Anticipating Climate Change Across the United States*. Tech. rep. National Bureau of Economic Research.

Brueckner, Jan K et al. (1987). “The structure of urban equilibria: A unified treatment of the Muth-Mills model”. In: *Handbook of regional and urban economics* 2.20, pp. 821–845.

Brueckner, Jan K, David Leather, and Miguel Zerecero (2024). “Bunching in real-estate markets: Regulated building heights in New York City”. In: *Journal of Urban Economics* 143, p. 103683.

Caliendo, Lorenzo, Maximilliano Dvorkin, and Fernando Parro (2019). “Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock”. In: *Econometrica* 87.3, pp. 741–835.

Cardaliaguet, Pierre (2010). *Notes on mean field games*. Tech. rep. Technical report.

Cattaneo, Matias D, Michael Jansson, and Xinwei Ma (2020). “Simple local polynomial density estimators”. In: *Journal of the American Statistical Association* 115.531, pp. 1449–1455.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2021). “The production function for housing: Evidence from France”. In: *Journal of Political Economy* 129.10, pp. 2766–2816.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin (2020). "Fast Poisson estimation with high-dimensional fixed effects". In: *The Stata Journal* 20.1, pp. 95–115.

Crews, Levi (2023). "A Dynamic Spatial Knowledge Economy". In: *Working Paper*.

Cui, Tianfang (2024). "Did Race Fence Off The American City? The Great Migration and the Evolution of Exclusionary Zoning". In: *Working Paper*.

Davis, Morris A and François Ortalo-Magné (2011). "Household expenditures, wages, rents". In: *Review of Economic Dynamics* 14.2, pp. 248–261.

Desmet, Klaus, David Krisztian Nagy, and Esteban Rossi-Hansberg (2018). "The Geography of Development". In: *Journal of Political Economy* 126.3.

Desmet, Klaus and Esteban Rossi-Hansberg (2010). "On spatial dynamics". In: *Journal of regional science* 50.1, pp. 43–63.

— (2014). "Spatial Development". In: *American Economic Review* 104.4, pp. 1211–1243.

Dingel, Jonathan I and Felix Tintelnot (2020). *Spatial economics for granular settings*. Tech. rep. National Bureau of Economic Research.

Duranton, Gilles and Diego Puga (2015). "Urban land use". In: *Handbook of regional and urban economics*. Vol. 5. Elsevier, pp. 467–560.

Glaeser, Edward L and Joseph Gyourko (2005). "Urban decline and durable housing". In: *Journal of political economy* 113.2, pp. 345–375.

Greaney, Brian, Andrii Parkhomenko, and Stijn Van Nieuwerburgh (2025). *Dynamic Urban Economics*. Tech. rep. National Bureau of Economic Research, Inc.

Hansen, Bruce E (2017). "Regression kink with an unknown threshold". In: *Journal of Business & Economic Statistics* 35.2, pp. 228–240.

Iacoviello, Matteo (2005). "House prices, borrowing constraints, and monetary policy in the business cycle". In: *American economic review* 95.3, pp. 739–764.

Keller, Herbert B. (1966). "Existence theory for two-point boundary value problems". In: *Bulletin of the American Mathematical Society* 72.5, pp. 728–731.

Kleinman, Benny, Ernest Liu, and Stephen J Redding (2023). "Dynamic spatial general equilibrium". In: *Econometrica* 91.2, pp. 385–424.

Lasry, Jean-Michel and Pierre-Louis Lions (2007). "Mean field games". In: *Japanese journal of mathematics* 2.1, pp. 229–260.

Lucas, Jr., Robert E. and Esteban Rossi-Hansberg (2002). “On The Internal Structure of Cities”. In: *Econometrica* 70.4, pp. 1445–1476.

Martynov, Petr (2022). *Welfare effects of zoning: Density restrictions and heterogeneous spillovers*. Tech. rep. Working paper, University of California, Berkeley.

McCrary, Justin (2008). “Manipulation of the running variable in the regression discontinuity design: A density test”. In: *Journal of econometrics* 142.2, pp. 698–714.

Mountjoy, Jack (2024). “Marginal Returns to Public Universities”. In: *Working Paper*.

Pennington, Kate (2021). “Does building new housing cause displacement?: the supply and demand effects of construction in San Francisco”. In: *The Supply and Demand Effects of Construction in San Francisco (June 15, 2021)*.

Piazzesi, Monika and Martin Schneider (2016). “Housing and macroeconomics”. In: *Handbook of macroeconomics* 2, pp. 1547–1640.

Rossi-Hansberg, Esteban (2004a). “Cities Under Stress”. In: *Journal of Monetary Economics* 51, pp. 903–927.

— (2004b). “Optimal urban land use and zoning”. In: *Review of Economic Dynamics* 7.1, pp. 69–106.

Rotemberg, Julio J (1982). “Sticky prices in the United States”. In: *Journal of political economy* 90.6, pp. 1187–1211.

San Francisco Municipal Code (2024). *Planning Code*. URL: http://www.amlegal.com.

Severen, Christopher (2023). “Commuting, labor, and housing market effects of mass transportation: Welfare and identification”. In: *Review of Economics and Statistics* 105.5, pp. 1073–1091.

Silva, JMC Santos and Silvana Tenreyro (2006). “The log of gravity”. In: *The Review of Economics and statistics* 88.4, pp. 641–658.

Song, Jaehee (2021). “The effects of residential zoning in US housing markets”. In: *Available at SSRN 3996483*.

Stokey, Nancy L (2009). “Moving costs, nondurable consumption and portfolio choice”. In: *Journal of Economic Theory* 144.6, pp. 2419–2439.

Sun, JE (2024). *The Distributional Consequences of Climate Change: The Role of Housing Wealth, Expectations, and Uncertainty*. Tech. rep. Working Paper.

A Proofs

A.1 Proof of Thrm 2.1

The static equilibrium must satisfy conditions for optimal housing demand by households, spatial equilibrium for households, optimal housing supply by firms, the housing market clearing at all locations, and the total population summing to unity. Household preferences yield constant housing expenditure β, so housing per household at x is

$$h(x) = \beta \frac{w(x)}{p(x)}$$

Then the indirect utility at location x is

$$u(x) = A(x) \frac{w(x)}{p(x)^{\beta}}$$

Spatial equilibrium imposes $u(x) = u(x') = \bar{u}$ for all x, x' with positive population, therefore prices satisfy

$$\bar{u} = \frac{w(x)A(x)}{p(x)^{\beta}} \implies p(x) = w(x)^{1/\beta} A(x)^{1/\beta} \mathcal{U}^{-1/\beta}$$

Optimal housing supply is

$$\kappa'\big(H(x)\big) = p(x)$$

and housing market clearing is

$$H(x) = n(x)h(x)$$

Inverting profit maximization yields, $H(x) = (\kappa')^{-1}(p(x))$. We combine this with the housing market clearing expression, and substitute for housing demand, and use the spatial equilibrium condition to replace prices, so that at each x,

$$n(x) = \frac{H(x)}{h(x)} = \frac{(\kappa')^{-1}\left(A(x)^{1/\beta}\mathcal{U}^{-1/\beta}\right)}{\frac{1}{\beta}A(x)^{1/\beta}\mathcal{U}^{-1/\beta}}.$$

By labor market clearing, the total labor sums to unity. That is,

$$\begin{aligned} 1 &= \int_0^1 n(x)\mathrm{d}x \\ &= \int_0^1 \frac{(\kappa')^{-1}\left(w(x)^{1/\beta}A(x)^{1/\beta}\mathcal{U}^{-1/\beta}\right)}{\frac{1}{\beta}w(x)^{1-1/\beta}A(x)^{-1/\beta}\mathcal{U}^{1/\beta}}\mathrm{d}x \\ \implies \mathcal{U}^{1/\beta} &= \frac{1}{\beta}\int_0^1 A(x)^{1/\beta}w(x)^{1/\beta-1}(\kappa')^{-1}\left(w(x)^{1/\beta}A(x)^{1/\beta}\mathcal{U}^{-1/\beta}\right)\mathrm{d}x \end{aligned}$$

The equation above has a unique solution when κ is differentiable and has a weakly positive derivative. Therefore, the equilibrium exists and is unique.

Furthermore, if we restrict our attention to the case in which $\kappa'(H) = H^{1/\gamma}$, the closed-form solution is

$$\overline{u} = \beta^{-\frac{\beta}{1+\gamma}} \left(\int_0^1 A(x)^{\frac{1+\gamma}{\beta}} \mathrm{d}x \right)^{\frac{\beta}{1+\gamma}}$$

A.2 Proof of Thrm 2.2

Recall the optimality equation

$$\rho \Phi_{\dot{H}} - \Phi_{\dot{H}\dot{H}} \ddot{H} - \Phi_{\dot{H}H} \dot{H} + \Phi_H + \kappa'(H) = p(t)$$

Impose steady state, and use that $\Phi_H(0, H) = \Phi_{\dot{H}}(0, H) = 0$, to find

$$\kappa'(H) = p$$

This condition holds at every location, and is the same optimality condition in the static problem. The remainder of the equilibrium conditions for market clearing are also identical, so the equilibrium is the same as the static equilibrium.

A.3 BVP proof

Theorem A.1. Let $f(t, y, y')$ be continuously differentiable in in y and y', and satisfy,

$$\frac{\partial f}{\partial y} > 0$$
$$\left| \frac{\partial f}{\partial y'} \right| \leq M,$$

for some $M \geq 0$ for all $y \in C^2$. Then there exists a unique solution to the boundary-value problem (BVP),

$$\begin{aligned} y'' &= f(t, y, y'), \quad y \in [0, T] \\ y(0) &= y_0 \\ y(T) &= y_T \end{aligned}$$

for $(y_0, y_T) \in \mathbb{R}^2$.

Proof. Consider the initial value problem (IVP)

$$\begin{aligned} y'' &= f(t, y, y'), \quad y \in [0, T] \\ y(0) &= y_0 \\ y'(0) &= s \end{aligned}$$

This problem has a unique solution, call it $y(t; s)$, and the solution is continuously differentiable in s (see any basic text on ODEs). Then y solves the BVP iff $y(T; s) = y_T$. Let $\xi(t; s) \equiv \frac{\partial y(t;s)}{\partial s}$, so that a sufficient condition for existence and uniqueness of a solution to the BVP is that there exists some $m > 0$ such that $\xi(T; s) \geq m$ for all s ($\xi > 0$ is necessary for uniqueness, and $\xi \geq m > 0$ is sufficient for existence). Now use $y'' = f(t, y, y')$ to note the following

$$\begin{aligned} \frac{\partial y''}{\partial s} &= \frac{\partial f}{\partial y}\frac{\partial y}{\partial s} + \frac{\partial f}{\partial y'}\frac{\partial y'}{\partial s} \\ \xi'' &= \frac{\partial f}{\partial y}\xi + \frac{\partial f}{\partial y'}\xi' \end{aligned}$$

The boundary conditions of the IVP imply

$$\begin{aligned} \xi(0; s) &= 0 \\ \xi'(0, s) &= 1 \end{aligned}$$

Thus there exists a δ such that there exists a $t > 0$ in the ball of radius δ around 0 such that $\xi(t; s) > 0$. Suppose there exists $\hat{t}$ such that $\xi(\hat{t}; s) < 0$. Then $\xi(t; s) > \xi(0; s) > \xi(\hat{t}, s)$, and there exists a local maxima in $[0, \hat{t}]$ (it may or may not be at t), which is necessarily greater than zero. But at that maxima it must be that $\xi' = 0$ and $\xi'' \leq 0$, so

$$\begin{aligned} \xi &= \frac{\xi'' - \frac{\partial f}{\partial y'}\xi'}{\frac{\partial f}{\partial y}} \\ &= \frac{\xi''}{\frac{\partial f}{\partial y}} && \left(\tfrac{\partial f}{\partial y'} \text{ bounded}\right) \\ &\leq 0 && \left(\tfrac{\partial f}{\partial y} > 0, \xi'' \leq 0\right) \end{aligned}$$

Thus the local maxima is ≤ 0, contradicting that $\xi(\hat{t}; s) > 0$, so there cannot be a local maxima, and our supposition that there exists $\hat{t}$ such that $\xi(\hat{t}; s) < 0$ was incorrect. Then for all $t > 0$ and all $s \in \mathbb{R}$ we have $\xi(t; s) > 0$, thus $\xi(T; s) > 0$. Therefore, $\frac{\partial f}{\partial y}\xi > 0$, and $\xi'' \geq \frac{\partial f}{\partial y'}\xi'$. Let $p(t) \equiv \frac{\partial f(t, y(t;s), y'(t;s))}{\partial y'}$, and directly note the following from $\xi'' \geq p(t)\xi'$,

$$\int_0^t e^{-\int_0^x p(x)\mathrm{d}x}\xi''(x)\mathrm{d}x \geq \int_0^t e^{-\int_0^x p(\tau)\mathrm{d}\tau}p(x)\xi'(x)\mathrm{d}x$$

Additionally, integration by parts yields,

$$\int_0^t e^{-\int_0^x p(\tau)\mathrm{d}\tau}\xi''(x)\mathrm{d}x = e^{-\int_0^t p(\tau)\mathrm{d}\tau}\xi'(t) - \xi'(0) + \int_0^t e^{-\int_0^x p(\tau)\mathrm{d}\tau}p(x)\xi'(x)\mathrm{d}x$$

Thus,

$$e^{-\int_0^t p(\tau)\mathrm{d}\tau}\xi'(t) - \xi'(0) \geq 0$$

Thus,

$$\begin{aligned}\xi'(t) &\geq \xi'(0)e^{\int_0^t p(\tau)\mathrm{d}\tau} \\ &\geq 1e^{-Mt}\end{aligned}$$

Thus,

$$\begin{aligned}\xi(t) &= \xi(0) + \int_0^t \xi'(x)\mathrm{d}x \\ &\geq 0 + \int_0^t e^{-Mx}\mathrm{d}x \\ &= \frac{1}{M}\left[1 - e^{-Mt}\right] \\ &> 0.\end{aligned}$$

Let $m = \frac{1}{M}\left[1 - e^{-MT}\right]$, and the proof is complete.

□

Theorem A.2. If $\kappa(H) = \frac{1}{1+\frac{1}{\gamma}}H^{1+\frac{1}{\gamma}}$, $\Phi(\dot{H}, H) = \xi\dot{H}^2$, then for any continuous price path $p(t)$ the BVP has a solution.

Proof. The BVP is

$$\begin{aligned}\rho\xi\dot{H} - \xi\ddot{H} + H^{\frac{1}{\gamma}} &= p(t) \\ H(0) &= H_0 \\ H(T) &= H_T\end{aligned}$$

Rewriting the equation in the format of the above theorem yields

$$f(t, H, \dot{H}) = \ddot{H} = \frac{\rho\xi\dot{H} + H^{\frac{1}{\gamma}} - p(t)}{\xi}$$

Then we verify,

$$\begin{aligned}\frac{\partial f}{\partial H} &= \frac{1}{\gamma\xi}H^{\frac{1}{\gamma}-1} > 0 \\ \left|\frac{\partial f}{\partial \dot{H}}\right| &= \rho\end{aligned}$$

The conditions are satisfied, thus there exists a unique solution. □

A.4 Model with idiosyncratic location preferences

Recall baseline dynamic model,

$$\rho\xi\dot{H}(x,t) - \xi\ddot{H}(x,t) + H(x,t)^{1/\gamma} = p(x,t)$$

where,

$$p(x,t) = \left(\frac{w(x,t)A(x,t)}{\mathcal{U}(t)}\right)^{1/\beta}$$

and,

$$\mathcal{U}(t) = \left(\frac{1}{\beta}\int_0^1 H(x,t)A(x,t)^{1/\beta}w(x,t)^{1/\beta-1}dx\right)^{\beta}$$

with boundary conditions on H.

Solving the model with idiosyncratic preference shock requires obtaining $p(x,t)$ as function of endogenous t-variables and exogenous (x,t) variables.

Start with Fréchet household choice,

$$n(x,t) = \frac{(A(x,t)w(x,t)p(x,t)^{-\beta})^{\theta}}{\int_0^1 (A(s,t)w(s,t)p(s,t)^{-\beta})^{\theta}ds} = \left(\frac{A(x,t)w(x,t)p(x,t)^{-\beta}}{\mathcal{W}(t)}\right)^{\theta}$$

Note that this isn't quite 'Fréchet' but rather because we have a continuum of draws. Regardless, we rely on its max stable analog for drawing in this space.

Now we could easily put p(x,t) on the LHS but we still have $n(x,t)$ which is endogenous, peskily floating around. However, as housing markets clear instantaneously,

$$p(x,t)H(x,t) = \beta w(x,t)n(x,t) \implies L(x,t) = \frac{1}{\beta}\frac{p(x,t)H(x,t)}{w(x,t)}$$

So therefore,

$$\begin{aligned}
\frac{1}{\beta}\frac{p(x,t)H(x,t)}{w(x,t)} &= \left(\frac{A(x,t)w(x,t)p(x,t)^{-\beta}}{\mathcal{W}(t)}\right)^{\theta} \\
\implies \frac{1}{\beta}\frac{H(x,t)}{w(x,t)}p(x,t)^{1+\beta\theta} &= \left(\frac{A(x,t)w(x,t)}{\mathcal{W}(t)}\right)^{\theta} \\
\implies p(x,t)^{1+\beta\theta} &= \beta\frac{w(x,t)}{H(x,t)}\left(\frac{A(x,t)w(x,t)}{\mathcal{W}(t)}\right)^{\theta} \\
\implies p(x,t) &= \left(\beta\frac{w(x,t)}{H(x,t)}\right)^{\frac{1}{1+\beta\theta}}\left(\frac{A(x,t)w(x,t)}{\mathcal{W}(t)}\right)^{\frac{\theta}{1+\theta\beta}}
\end{aligned}$$

returning to the ODE,

$$\rho\xi\dot{H}(x,t) - \xi\ddot{H}(x,t) + H(x,t)^{1/\gamma} = p(x,t)$$
$$\rho\xi\dot{H}(x,t) - \xi\ddot{H}(x,t) + H(x,t)^{1/\gamma} = \left(\beta\frac{w(x,t)}{H(x,t)}\right)^{\frac{1}{1+\beta\theta}}\left(\frac{A(x,t)w(x,t)}{\mathcal{W}(t)}\right)^{\frac{\theta}{1+\theta\beta}}$$
$$H(x,t)^{\frac{1}{1+\theta\beta}}\left(\rho\xi\dot{H}(x,t) - \xi\ddot{H}(x,t) + H(x,t)^{1/\gamma}\right) = (\beta)^{\frac{1}{1+\beta\theta}}\left(\frac{A(x,t)w(x,t)^{\frac{1+\theta}{\theta}}}{\mathcal{W}(t)}\right)^{\frac{\theta}{1+\theta\beta}}$$

Note that as $\theta \to \infty$ this takes the form of the original ODE in the System (4).

Now, to write welfare,

$$\mathcal{W}(t) = \left(\int_0^1 [A(x,t)w(x,t)p(x,t)^{-\beta}]^{\theta}dx\right)^{1/\theta}$$
$$\mathcal{W}(t) = \left(\int_0^1 \left[A(x,t)w(x,t)\left(\beta\frac{w(x,t)}{H(x,t)}\right)^{\frac{-\beta}{1+\beta\theta}}\left(\frac{A(x,t)w(x,t)}{\mathcal{W}(t)}\right)^{\frac{-\beta\theta}{1+\theta\beta}}\right]^{\theta}dx\right)^{1/\theta}$$
$$\mathcal{W}(t)^{1-\frac{\beta\theta}{1+\beta\theta}} = \left(\int_0^1 \left[A(x,t)w(x,t)\left(\beta\frac{w(x,t)}{H(x,t)}\right)^{\frac{-\beta}{1+\beta\theta}}(A(x,t)w(x,t))^{\frac{-\beta\theta}{1+\theta\beta}}\right]^{\theta}dx\right)^{1/\theta}$$
$$\mathcal{W}(t)^{\frac{1}{1+\beta\theta}} = \left(\int_0^1 \left[A(x,t)w(x,t)\left(\beta\frac{w(x,t)}{H(x,t)}\right)^{\frac{-\beta}{1+\beta\theta}}(A(x,t)w(x,t))^{\frac{-\beta\theta}{1+\theta\beta}}\right]^{\theta}dx\right)^{1/\theta}$$
$$\mathcal{W}(t) = \left(\int_0^1 [A(x,t)w(x,t)]^{\frac{\theta}{1+\beta\theta}}\left(\beta\frac{w(x,t)}{H(x,t)}\right)^{\frac{-\theta\beta}{1+\beta\theta}}dx\right)^{\frac{1+\beta\theta}{\theta}}$$
$$\mathcal{W}(t) = \beta^{-\beta}\left(\int_0^1 \left(\frac{H(x,t)}{w(x,t)}\right)^{\frac{\theta\beta}{1+\beta\theta}}[A(x,t)w(x,t)]^{\frac{\theta}{1+\beta\theta}}dx\right)^{\frac{1+\beta\theta}{\theta}}$$

Note that this is $\mathcal{W}(t) \to \mathcal{U}(t)$ as $\theta \to \infty$.

A.5 Neoclassical capital accumulation

We begin with the development problem,

$$\max_{\{I(t)\}} \int_0^{\infty} e^{-\rho t}\left[p(t)H(t) - \Phi(I(t),H(t))\right]\mathrm{d}t$$
$$\dot{H} = I - \delta H$$

where $\Phi(I,H)$ is the cost to increase housing at rate I, given current H. The recursive form of this problem (the Hamilton-Jacobi-Bellman) is,

$$\rho\mathcal{V}(H,t) - \partial_t\mathcal{V} = \max_I \left\{ p(t)H - \Phi(I,H) + \mathcal{V}_H\dot{H} \right\}$$
$$\dot{H} = I - \delta H.$$

To solve the HJB equation, we first take the first order condition,

$$\Phi_I = \mathcal{V}_H.$$

The envelope condition is,

$$\rho\mathcal{V}_H - \partial_t\mathcal{V}_H = p(t) - \Phi_H + \mathcal{V}_{HH}[I - \delta H] - \delta\mathcal{V}_H.$$

We differentiate the first order condition with respect to time and combine with the envelope condition to yield,

$$\rho\Phi_I = p(t) - \Phi_H + \Phi_{II}\dot{I} + \Phi_{IH}\dot{H} - \delta\Phi_I$$

We then specialize to quadratic investment costs (this will make solving the ODE easier below): $\Phi(I,H) = \frac{\xi}{2}I^2$. The equation becomes,

$$\rho\xi I = p(t) + \xi\dot{I} - \delta\xi I$$

Now we use $\dot{H} = I - \delta H$ and $\ddot{H} = \dot{I} - \delta\dot{H}$ to write the equation as a second-order ODE in H. Combining like terms yields,

$$-\xi\ddot{H} + \rho\xi\dot{H} + \xi\delta(\rho+\delta)H = p(t).$$

This is almost an inhomogeneous equation, and if we knew $p(t)$ in terms of elementary functions there would even be a closed-form. Alas, $p(t)$ is endogenous, and trying to guess a form for $p(t)$ appears futile. Nonetheless, we can do pretty well and make our lives computationally easier by going as far as we can with analytics.

The solution to the ODE above the sum of the solution to the relevant homogeneous equation and the particular solution for the nonhomogeneous part: $H(t) = H_h(t) + H_p(t)$. Recall the relevant homogeneous equation is,

$$-\xi\ddot{H} + \rho\xi\dot{H} + \xi\delta(\rho+\delta)H = 0$$

The general solution to such an equation is,

$$H_h(t) = c_1 e^{\lambda_t t} + c_2 e^{\lambda_2 t}$$

where λ_i are the roots of the characteristic equation. The particular solution is trickier, but the general form is,

$$H_p(t) = -H_1(t)\int \frac{H_2(s)g(s)}{W(H_1(s),H_2(s))}\mathrm{d}s + H_2(t)\int \frac{H_1(s)g(s)}{W(H_1(s),H_2(s))}\mathrm{d}s$$

where W is the Wronskian,

$$W(H_1(t), H_2(t)) = H_1(t)H_2'(t) - H_1'(t)H_2(t)$$

and $H_i(t) = c_i e^{\lambda_i t}$.

While this may seem unnecessary, this is good news for computation. To solve the BVP in the general case, we guess a path of prices, then solved the BVP for each location, then checked if markets cleared. Solving the BVP step is expensive, and now it is gone, because for any $p(x,t)$ guess we can just compute the path of H directly via the integrals above. In other words, it is (generally) easier to compute an integral than to solve a nonlinear equation, so this solution method is faster (but of course we are solving a different problem).

A.6 Endogenous wages with differentiated goods

To illustrate this in practice, we consider an example where each location produces a unique, freely traded intermediate good, whose variety is indexed by the location of production, x. There is a final goods firm that assembles these varieties into the final good with a constant elasticity of substitution technology with elasticity σ, and the final good remains the numéraire.

Production technology remains linear-in-labor, $y(x,t) = Q(x,t)L(x,t)$, and there is no commuting so that $L(x,t) = n(x,t)$ in equilibrium. Goods prices are simply marginal costs, $w(x,t)/Q(x,t)$. Market clearing for the x good at time t, expressed in wages, is,

$$w(x,t)n(x,t) = (1-\beta)\int_0^1 \left(w(x,t)/Q(x,t)\right)^{1-\sigma} w(s,t)n(s,t)ds$$

Using the fact that the wage bill $w(x,t)n(x,t) = \frac{1}{\beta}p(x,t)H(x,t)$, and $n(x,t) = H(x,t)/h(x,t)$, we can solve for wages as a function of $p(x,t)$ and $H(x,t)$, so that,

$$w(x,t) = Q(x,t)\left((1-\beta)\frac{p(x,t)H(x,t)}{\int_0^1 p(s,t)H(s,t)ds}\right)^{\frac{1}{1-\sigma}}$$

Since a fraction of all income β is spent on housing, we can define city GDP as,

$$\mathcal{Y}(t) = \frac{1}{\beta}\int_0^1 p(x,t)H(x,t)dx$$

So that wages are,

$$w(x,t) = Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{p(x,t)H(x,t)}{\mathcal{Y}(t)}\right)^{\frac{1}{1-\sigma}}$$

Recall that prices are, $p(x,t) = (A(x,t)w(x,t)/\mathcal{U}(t))^{1/\beta}$, so we can write this,

$$\begin{aligned} w(x,t) &= Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{(A(x,t)w(x,t)/\mathcal{U}(t))^{1/\beta}H(x,t)}{\mathcal{Y}(t)}\right)^{\frac{1}{1-\sigma}} \\ &= Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}w(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)\mathcal{U}(t)^{1/\beta}}\right)^{\frac{1}{1-\sigma}} \\ &= Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)\mathcal{U}(t)^{1/\beta}}\right)^{\frac{1}{1-\sigma}} w(x,t)^{\frac{1}{\beta(1-\sigma)}} \\ \implies w(x,t)^{\frac{\beta(\sigma-1)+1}{\beta(\sigma-1)}} &= Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)\mathcal{U}(t)^{1/\beta}}\right)^{\frac{1}{1-\sigma}} \end{aligned}$$

Which implies,

$$\implies w(x,t) = Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)\mathcal{U}(t)^{1/\beta}}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}$$

Using prices we can solve for GDP,

$$\begin{aligned} \mathcal{Y}(t) &= \frac{1}{\beta}\int_0^1 (A(x,t)w(x,t)/\mathcal{U}(t))^{1/\beta} H(x,t)dx \\ &= \frac{1}{\beta}\int_0^1 \left(A(x,t)\left[Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)\mathcal{U}(t)^{1/\beta}}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]/\mathcal{U}(t)\right)^{1/\beta} H(x,t)dx \\ &= \frac{1}{\beta}\int_0^1 \left(A(x,t)\left[Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]\mathcal{U}(t)^{\frac{1/\beta}{1/\beta+(\sigma-1)}-1}\right)^{1/\beta} H(x,t)dx \\ &= \frac{1}{\beta}\int_0^1 \left(A(x,t)\left[Q(x,t)\left(\frac{(1-\beta)}{\beta}A(x,t)^{1/\beta}H(x,t)\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]\mathcal{Y}(t)^{\frac{1}{1/\beta+(\sigma-1)}}\mathcal{U}(t)^{\frac{-(\sigma-1)}{1/\beta+(\sigma-1)}}\right)^{1/\beta} H(x,t)dx \\ &= \frac{1}{\beta}\int_0^1 \left(A(x,t)\left[Q(x,t)\left(\frac{(1-\beta)}{\beta}A(x,t)^{1/\beta}H(x,t)\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]\right)^{1/\beta} H(x,t)dx\mathcal{Y}(t)^{\frac{1/\beta}{1/\beta+(\sigma-1)}}\mathcal{U}(t)^{\frac{-(\sigma-1)}{1+\beta(\sigma-1)}} \end{aligned}$$

which implies,

$$\mathcal{Y}(t)^{\frac{(\sigma-1)}{1/\beta+(\sigma-1)}} = \mathcal{U}(t)^{\frac{-1/\beta(\sigma-1)}{1/\beta+(\sigma-1)}}\frac{1}{\beta}\int_0^1 A(x,t)^{1/\beta}Q(x,t)^{1/\beta}\left(\frac{(1-\beta)}{\beta}A(x,t)^{1/\beta}H(x,t)\right)^{\frac{-1/\beta}{1/\beta+(\sigma-1)}} H(x,t)dx$$

Which implies,

$$\mathcal{Y}(t) = \mathcal{U}(t)^{-1/\beta} \left(\frac{\beta}{1-\beta}\right)^{\frac{1/\beta}{\sigma-1}} \left(\frac{1}{\beta}\int_0^1 Q(x,t)^{1/\beta} A(x,t)^{1/\beta\left(\frac{(\sigma-1)}{1/\beta+(\sigma-1)}\right)} H(x,t)^{\frac{(\sigma-1)}{1/\beta+(\sigma-1)}} dx\right)^{\frac{1/\beta+(\sigma-1)}{\sigma-1}}$$

Plugging in for wages in utility,

$$\begin{aligned}
\mathcal{U}(t) &= \left(\int_0^1 \frac{1}{\beta} H(x,t) A(x,t)^{1/\beta} \left[Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)\mathcal{U}(t)^{1/\beta}}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]^{1/\beta-1} \mathrm{d}x\right)^{\beta} \\
&= \left(\int_0^1 \frac{1}{\beta} H(x,t) A(x,t)^{1/\beta} \left[Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)}\right)^{\frac{-1}{1/\beta+(\sigma-1)}} \mathcal{U}(t)^{\frac{1/\beta}{1/\beta+(\sigma-1)}}\right]^{1/\beta-1}\right. \\
&= \mathcal{U}(t)^{\frac{1/\beta-1}{1/\beta+(\sigma-1)}} \left(\int_0^1 \frac{1}{\beta} H(x,t) A(x,t)^{1/\beta} \left[Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]^{1/\beta-1}\right. \\
\implies \mathcal{U}(t)^{1-\frac{1/\beta-1}{1/\beta+(\sigma-1)}} &= \left(\int_0^1 \frac{1}{\beta} H(x,t) A(x,t)^{1/\beta} \left[Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]^{1/\beta-1} \mathrm{d}x\right)^{\beta}
\end{aligned}$$

Finally, we derive,

$$\mathcal{U}(t) = \left(\int_0^1 \frac{1}{\beta} H(x,t) A(x,t)^{1/\beta} \left[Q(x,t)\left(\frac{(1-\beta)}{\beta}\frac{A(x,t)^{1/\beta}H(x,t)}{\mathcal{Y}(t)}\right)^{\frac{-1}{1/\beta+(\sigma-1)}}\right]^{1/\beta-1} \mathrm{d}x\right)^{1/\sigma+\beta(1-1/\sigma)}$$

A.7 Zoning

Zoning is a restriction,

$$H(x) \leq \bar{Z}(x)$$

Recall that with no preference shocks, the equilibrium can boil down to finding a level of utility that clears all markets. This is done by solving for labor market clearing,

$$\int_0^1 n(x) dx = 1$$

and using housing market clearing to solve for $n(x)$,

$$h(x)n(x) = \text{housing supply}$$

where the price gradient $p(x)$ can be replaced with a unidimensional object, the common utility level,

$$p(x) = (A(x)/\mathcal{U})^{1/\beta}$$

Effective Housing supply $\tilde{H}(x)$ with zoning is,

$$\tilde{H}(x) = \min\{H(x), \bar{Z}(x)\} = \min\{p(x)^{\gamma}, \bar{Z}(x)\}$$

As housing demand is Cobb-Douglas and $w = 1$,

$$n(x) = \frac{1}{\beta}(A(x)/\mathcal{U})^{1/\beta} \min\{(A(x)/\mathcal{U})^{\gamma/\beta}, \bar{Z}(x)\}$$

and therefore $\mathcal{U}$ solves,

$$\frac{1}{\beta}\int_0^1 (A(x)/\mathcal{U})^{1/\beta} \min\{(A(x)/\mathcal{U})^{\gamma/\beta}, \bar{Z}(x)\}dx = 1$$

to confirm this has a unique solution, note,

$$\left(\int_0^1 (A(x))^{1/\beta} \min\{(A(x)/\mathcal{U})^{\gamma/\beta}, \bar{Z}(x)\}dx\right)^{\beta} = \beta^{\beta}\mathcal{U}$$

the RHS is upward sloping and the LHS is weakly downward sloping.

With $\mathcal{U}$ in hand, one can use housing market clearing to recover the population gradient $n(x)$.

B Quantitative model details

B.1 Deriving housing supply

We assume developers operate housing services on a continuum of parcels ω of measure T_i in a neighborhood i. Each parcel has an efficiency $e(\omega) \sim G_i$. We assume developers provide housing services until the marginal cost of floorspace equals its price p, so that the floorspace in parcel ω in neighborhood i, under the zoning restriction Z_i is,

$$f_i(\omega) = \min\{\big(e(\omega)p\big)^\gamma, Z_i\}.$$

Developers optimally provide housing services by cost-minimizing when allocating resources to the production of housing services at each parcel. They do so by developing floorspace at the least cost (height $e(\omega)$) until the constraint binds. As there are a continuum of parcels, this cost minimization choice is summarized by a cutoff parcel efficiency $\tilde{e}$ so that all parcels with $e > \tilde{e}$ are developed until $f_i(\omega) = Z_i$. The remainder of parcels are developed so that $f_i(\omega) = \big(e(\omega)p_i\big)^\gamma$. Thus, total housing services can be written,

$$H_i = T_i \left(Z_i\big(1 - G(\tilde{e})\big) + \mathbb{E}[(ep)^\gamma \mid e < \tilde{e}]G(\tilde{e})\right)$$

or equivalently,

$$H_i = T_i \left(\int_{\tilde{e}_i}^{\infty} Z_i dG_i(e) + \int_{-\infty}^{\tilde{e}_i} (pe)^\gamma dG_i(e)\right).$$

Under the assumption that parcel efficiency is Pareto-distributed, so that $G_i(e) = 1 - (E_i/e)^\varphi$, we can solve the above integrals in closed form. First,

$$Z_i(1 - G_i(\tilde{e}_i)) = E_i^\varphi Z_i^{1-\varphi/\gamma} p_i^\varphi$$

then,

$$\begin{aligned}
\mathbb{E}[e^\gamma \mid e < \tilde{e}]G(\tilde{e}) &= \varphi E_i^\varphi \int_{E_i}^{\tilde{e}_i} e^{\gamma-\varphi-1} de \\
&= \varphi E_i^\varphi \left[\frac{1}{\gamma-\varphi} e^{\gamma-\varphi}\right]_{E_i}^{\tilde{e}_i} \\
&= \frac{\varphi}{\gamma-\varphi} E_i^\varphi \left[\tilde{e}_i^{\gamma-\varphi} - E_i^{\gamma-\varphi}\right] \\
&= \frac{\varphi}{\gamma-\varphi} E_i^\varphi \left[Z_i^{1-\varphi/\gamma} p_i^{\varphi-\gamma} - E_i^{\gamma-\varphi}\right]
\end{aligned}$$

Thus, combined,

$$H_i = T_i \left({}^\varphi_i Z_i^{1-\varphi/\gamma} p_i^\varphi + p_i^\gamma \frac{\varphi}{\gamma-\varphi} E_i^\varphi \left[Z_i^{1-\varphi/\gamma} p_i^{\varphi-\gamma} - E_i^{\gamma-\varphi}\right]\right)$$

Rearranged,

$$
\begin{aligned}
H_i =& T_i E_i^{\gamma} p_i^{\gamma} \left(Z_i^{1-\varphi/\gamma} E_i^{\varphi-\gamma} p_i^{\varphi-\gamma} + p_i^{\varphi-\gamma} \frac{\varphi}{\gamma-\varphi} E_i^{\varphi-\gamma} \left[Z_i^{1-i/\gamma} E_i^{\varphi-\gamma} - E_i^{\gamma-\varphi} \right] \right) \\
&= T_i E_i^{\gamma} p_i^{\gamma} \left(Z_i^{1-\varphi/\gamma} E_i^{\varphi-\gamma} p^{\varphi-\gamma} + \frac{\varphi}{\gamma-\varphi} \left[Z_i^{1-\varphi/\gamma} E_i^{\varphi-\gamma} p_i^{\varphi-\gamma} - 1 \right] \right) \\
&= T_i E_i^{\gamma} p_i^{\gamma} \left(\frac{\gamma}{\gamma-\varphi} Z_i^{1-\varphi/\gamma} E_i^{\varphi-\gamma} p_i^{\varphi-\gamma} - \frac{\varphi}{\gamma-\varphi} \right)
\end{aligned}
$$

Inversion to recover Z_i given information on S_i, H_i, T_i and p_i, we note,

$$
S_i = 1 - G_i(\tilde{e}) = E_i^{\varphi} Z_i^{-\varphi/\gamma} p_i^{\varphi}
$$

Rearranging this reveals,

$$
Z_i S_i^{\frac{\gamma}{\varphi}} = (E_i p_i)^{\gamma}
$$

which means the H_i equation can be inverted for Z_i analytically,

$$
Z_i = \left(\frac{\gamma}{\gamma-\varphi} S_i - \frac{\varphi}{\gamma-\varphi} S_i^{\gamma/\varphi} \right)^{-1} \frac{H_i}{T_i}.
$$

To recover $\kappa'(H)$, which is needed for the dynamic solution, we note that at the steady state $\kappa'(H) = p$, thus,

$$
H_i = T_i E_i^{\gamma} (\kappa'(H_i))^{\gamma} \left(\frac{\gamma}{\gamma-\varphi} Z_i^{1-\varphi/\gamma} E_i^{\varphi-\gamma} (\kappa'(H_i))^{\varphi-\gamma} - \frac{\varphi}{\gamma-\varphi} \right)
$$

defines $\kappa'(H_i)$ implicitly along the development trajectory.

B.2 Additional figures and tables

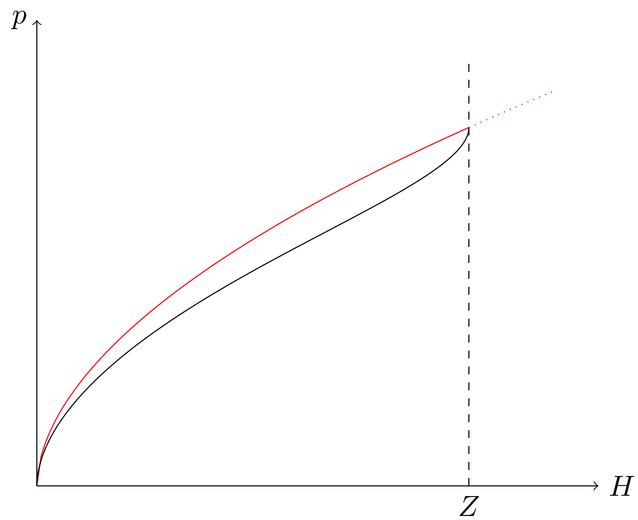

Figure A1: Visual depiction of supply curves with zoning. An isoelastic supply curve, $H = p^{\gamma}$, is shown in red, and supply curve with zoning and development on heterogeneous plots, is shown in black.

Zoning class, C	Share of structures	Share of volume	Estimated Z_C (ft)
Residential			
RH-1	48.6	25.6	21.0
RH-2	23.8	18.4	32.9
RH-3	7.0	8.2	32.9
RM-1	5.6	7.1	32.2
RM-2	1.4	3.2	44.1
RM-3	1.0	2.9	45.2
RTO	1.7	2.3	43.3
Other Residential	1.4	4.9	23.7
Mixed use			
NC	6.8	12.5	42.5
RC	0.8	4.8	42.5
Other Mixed	1.8	10.0	24.6

Table A1: Estimated Z_C and share of structures and floorspace bound by Z_C by zoning class in San Francisco.

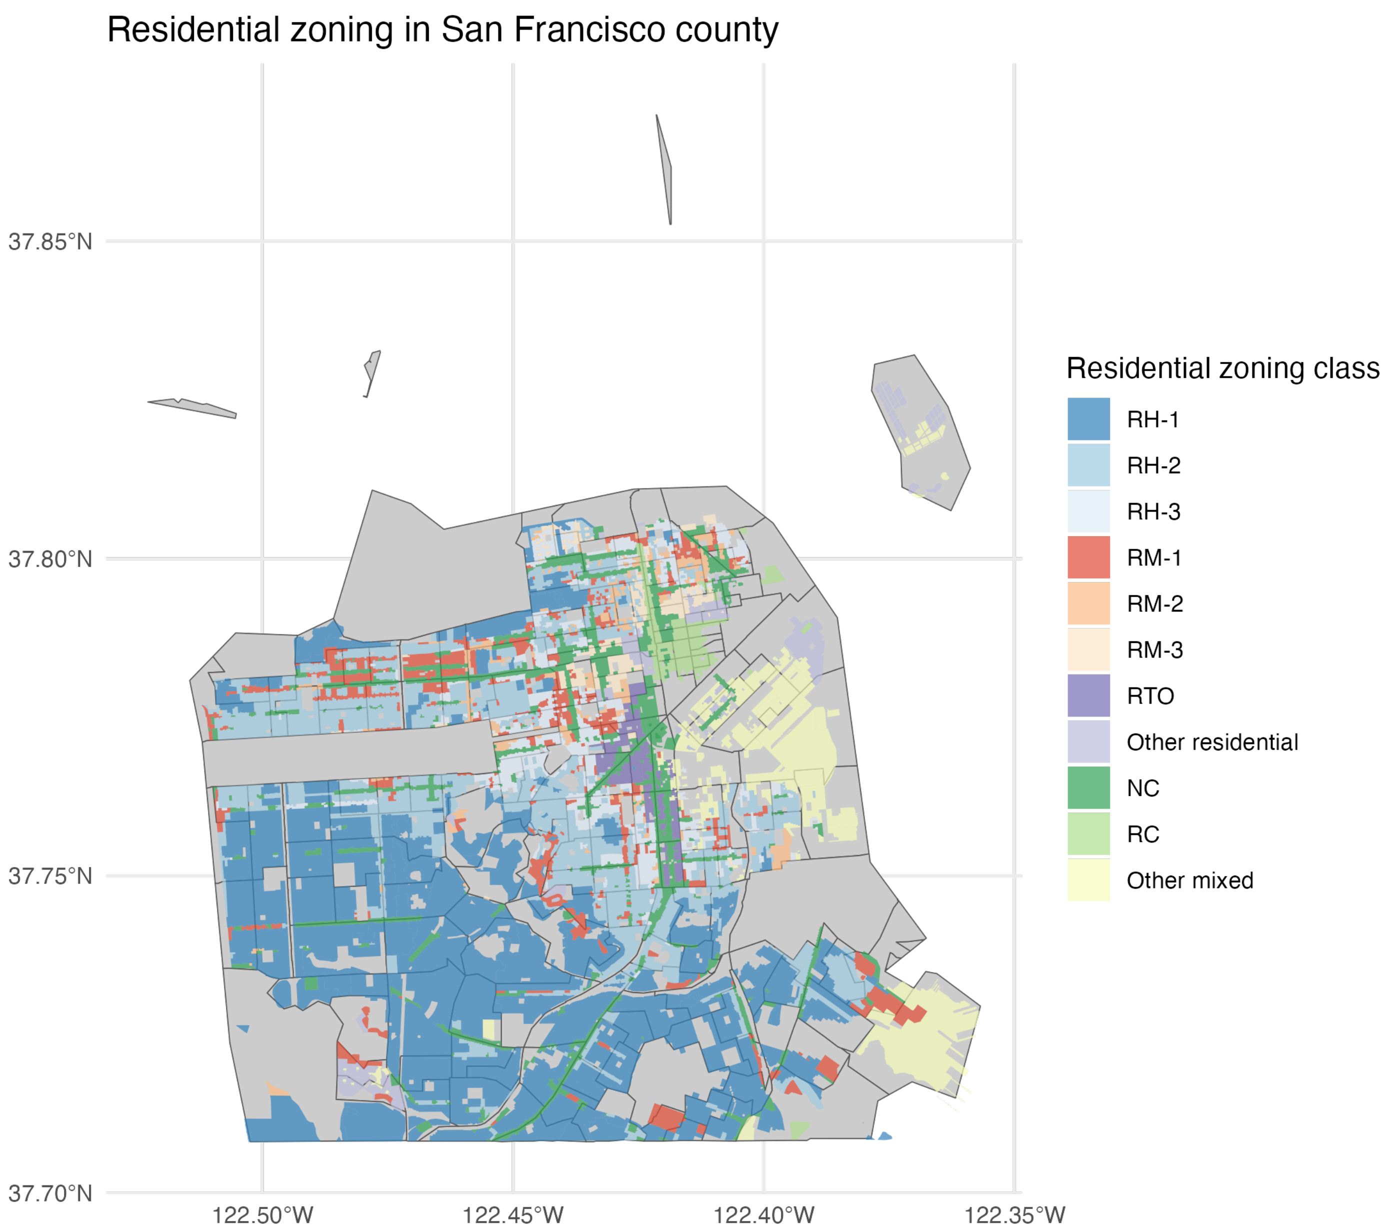

Figure A2: Map of residential zoning in San Francisco county

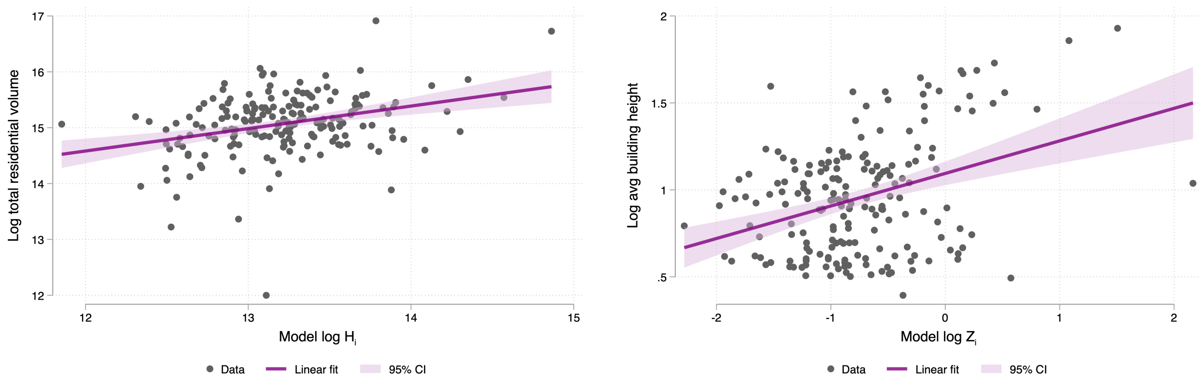

Figure A3: Tract-level model objects and aggregated data. Left: Real built volume vs model-implied H_i. Right: Observed avg. building height vs model-consistent estimates of Z_i.

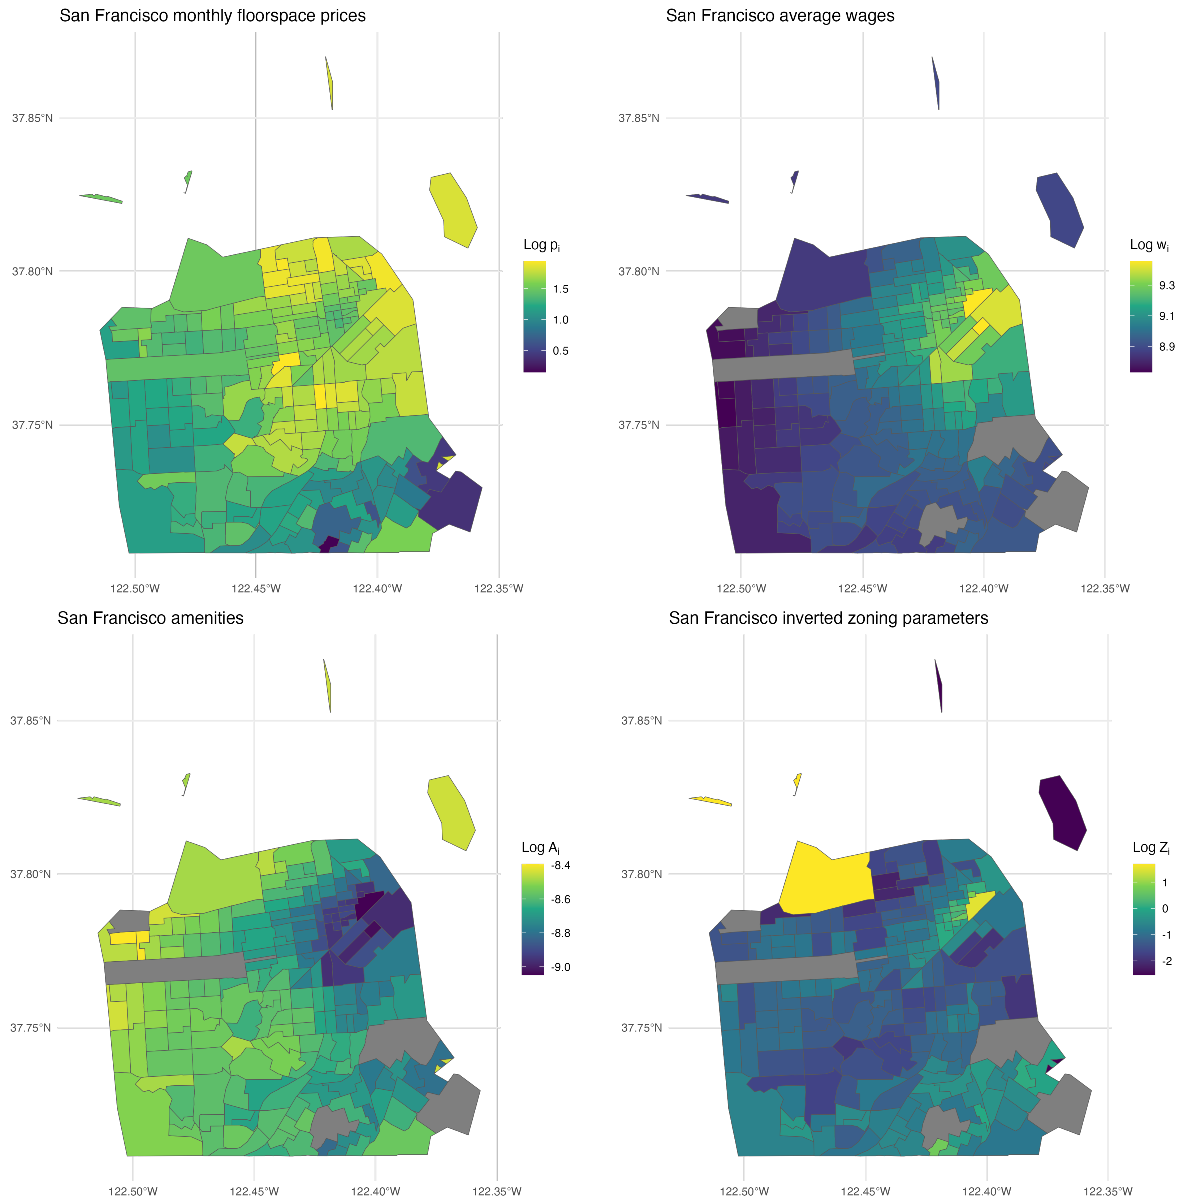

Figure A4: Clockwise: recovered p_i, w_i, A_i and Z_i.