

Working Paper 2017-4-2
April 2017

An Approach for Empirical Work in Spatial Dynamics

Kyle Mangum
Georgia State University

This paper can be downloaded at: <http://uwrg.gsu.edu>

An Approach for Empirical Work in Spatial Dynamics

Kyle Mangum*

March 2017

Abstract

This paper illustrates how to incorporate forward-looking behavior into empirical spatial equilibrium models with locational heterogeneity. The main insight is that the standard spatial equilibrium already embeds a natural starting point for dynamics: the presence of spatial indifference conditions can dramatically simplify the state space of forward-looking agents. One needs to know how utility evolves with the variables in the system, but not every individual state, effectively side-stepping the curse of dimensionality. Standard numerical rational expectations methods can be applied to derive approximate solutions to the full dynamic specification. The paper uses an example model of landowners exercising real options in construction of housing to show how the approximation's errors are negligibly small, and importantly, much smaller than the differences between dynamic and myopic specifications of the same model.

Keywords. spatial equilibrium; spatial dynamics; rational expectations; housing supply

JEL codes: R13, C63, D58, R31

1 Introduction

Spatial indifference is the hallmark of the classical spatial equilibrium model. In the typical model, free mobility among a sufficiently large number of agents ensures equal utility at all points in space, else someone would move. This, when combined with other conditions specific to the context, pins down the endogenous variables of the system (e.g. wages, rents, travel costs, proximity to employment or amenities, etc.) because the particular source(s) of the utility flow can vary over space. The literature leveraging these models, whether within-city models like the monocentric city model and its variants (Alonso (1964), Mills (1967), Muth (1969), summarized in Brueckner (1987)), or the between-city regional model of Roback (1982), is voluminous.

The nature of these models is fundamentally static: indifference at a given moment. But cities are full of durable elements, from geographic features to evolving capital stocks to municipal

*Department of Economics, Georgia State University. Correspondence to: kmangum@gsu.edu. I thank the participants of the conference on Advances and Applications of Spatial Equilibrium, sponsored by Economic Research Initiatives at Duke (ERID). I also thank the Duke Urban-Environmental Working Group for access and hosting of data. Errors are my own.

boundaries and policy regimes, meaning a static model with myopic agents is a major simplification. The dynamic spatial equilibrium model has made limited inroads to the literature in urban and regional economics. While there are many conceptual reasons for forward-looking behavior, there are also many practical reasons to avoid it, as state spaces in spatial models, where location itself is a form of heterogeneity, quickly suffer from the curse of dimensionality. Yet, there are reasons to believe that dynamic and myopic models can deliver substantially different results, so eschewing a dynamic approach in the interest of convenience may be consequential.

This paper suggests a readily feasible method for incorporating forward-looking behavior into a standard class of spatial equilibrium models. The goal is to make dynamic versions of such models easy to implement empirically for simulation and estimation. The basic insight is that the classic spatial equilibrium model already contains a natural structure for incorporating dynamics. In models with conditions for spatial indifference, forward-looking agents need not keep track of all states in the economy—which can balloon out of control in a model of heterogeneous spaces—but rather track only how they expect the utility condition to evolve. The evolution of this utility condition could be endogenously subject to the actions of the other agents in the economy, meaning a rational expectations equilibrium will be obtained when agents correctly predict (in expectation) the evolution of the utility condition factoring in the responses of other agents, who in turn factor in their expectations of other agents’ reactions, and so on. This approach combines components of standard spatial equilibrium and dynamic optimization models: a free mobility condition pinning down utility, but durable features evolving under the control of forward looking agents. It then borrows intuition from the well-known literature on rational expectations models with heterogeneous agents, where certain variables evolve under endogenously determined laws of motion.¹ The paper illustrates the logical progression from a single agent dynamic model with exogenous variables describing the aggregate economy to a dynamic equilibrium version with an endogenous transition of those aggregate variables.

The approach is designed for empirical models (i.e. those to be estimated and simulated) and relies on a numerical implementation. It is an approximation to the fully specified (infeasible) equilibrium when agents are heterogeneous, because their reactions may depend on their own features and the sources of utility. That is, an agent may know the aggregate state and forecast utility to evolve to a certain point, but deviations from the rule could result from dependence on the distribution of shocks across agents.² For example, consider a set of local markets of varying supply elasticity for some consumption good. Conditioning on aggregate demand, if a local relative demand shock occurs in a market with more elastic supply, a higher quantity of the

¹For a textbook treatment of rational expectations models, see Ljungqvist and Sargent (2001). The numerical approach is based on Krusell and Smith (1998), although applications with finite numbers of agents, such as this paper’s, are also comparable to Weintraub et al. (2008).

²Identical agents are a special case in which this approach is not an approximation, but also one in which the curse of dimensionality may not be present.

good would result in equilibrium than if the relative demand shock occurred in an low supply elasticity local market.

It is effectively impossible to know exactly how much error is in the approximation relative to the fully specified equilibrium because solution of the full model is utterly infeasible, but features of the underlying model give insight as to how much heterogeneity is being smoothed away with aggregate laws of motion. The paper uses an example model to show that the perturbations to an agent's policy function are small, and importantly, much smaller than the differences between dynamic and myopic models. In general, the approach performs better as the number of agents in the economy grows, because laws of large numbers apply to the shocks across agents. Thus, in many cases, using this approach for a dynamic model is preferable to assuming away dynamics altogether, and it performs better just as the dynamics would be increasingly difficult to incorporate.

The relevant environment is one in which mobility of some factor(s) can enforce a no arbitrage/equal utility condition, but there exist agents that are heterogenous and fixed in place, live for more than one unit of time, and face constraints or frictions that introduce state variables to their problem. Obvious applications for the forward-looking agents are landowners and local or regional governments.³

The number of forward-looking agents is of practical importance. On one extreme, in a two-agent model (like a two region economy), each agent has one response function to consider, so it is feasible to track all agents state spaces, meaning no approximation is necessary.⁴ At the other extreme of an infinite number of agents, and the approach is essentially that of Krusell and Smith (1998) in a uniquely spatial context. This paper will illustrate the in-between cases, where the number of forward-looking agents is finite but there are too many to compute full solutions.⁵ These cases naturally correspond to many empirically relevant contexts, like cities in a national economy, states in a federation, or neighborhoods in a metropolitan area.⁶

The context illustrated here is one of capital construction in heterogenous spaces; specifically, housing construction across metropolitan areas in the U.S. In this context, the population of housing consumers enforces a spatial equilibrium utility condition across cities, but the land-controlling agents in the cities must decide how much durable and irreversible housing to add

³The equal utility condition is of course a common assumption in urban models, but some recent island economies models, discussed below, seek to relax it.

⁴Certain small spatial economy models can get complicated when considering multiple person types and sectors, but different solution approaches apply. See also footnote 1.

⁵An infinite number of a discrete type of agents resembles a weighted version of the example of this paper. Indeed, in that the evolution of \bar{U} in this model hinges on a continuous variable that is affected by agents of different sizes, it is essentially a share-weighted aggregation of discrete types.

⁶The focus here is on a fixed number of locations, not an endogenously determined number as encountered in systems of cities models (Henderson (1974), Black and Henderson (1999)). This is not consequential, since a suitably defined partitioning of land would deliver a fixed number of agents, and nothing about the method imposes size or density restrictions on land partitions. In fact, the method could be useful in analyzing the dynamics of lands at different levels of development.

in a given period, knowing that construction exercises a real option over its land inputs and can affect future prices and costs. The model features forward-looking agents in heterogeneous spaces with continuous choice variables (land and capital inputs). Housing construction is a natural place to exhibit the approach because it is emblematic of conditions that make dynamic models both desirable and difficult. The literature has shown renewed interest in understanding housing construction, but has yet to incorporate forward-looking behavior into an equilibrium setting.⁷ This is despite a history of theoretical and empirical research on housing developers exercising real options on land.⁸ The paper closes with examples of policy simulations for dynamic and myopic versions of the model, demonstrating where the specifications depart from each other. The differences between models can be complex and nuanced, dependent on situations in the economy and attributes of the heterogeneous agents. The lesson is that researchers may not know just how different the empirical results can be until both specifications are used. It is the goal of this paper to make feasible the dynamically-specified alternative.

As Desmet and Rossi-Hansberg (2010) describe, the literature on spatial dynamics is still relatively small, with few papers incorporating dynamics into environments with rich spatial heterogeneity. The current paper relates to three strands of work at the frontier of spatial dynamics. One strand uses evolving static equilibria, in which states of the economy change but agents solve static problems (e.g. Nieuwerburgh and Weill (2010), Glaeser et al. (2014), Morten and Bryan (2015), Morten and Oliveira (2016), Diamond (2016)). Sometimes these models are augmented with “adjustment costs” which impose frictions in any given period but do not alter the agent’s decision horizon. While these can be useful tools, they are not fully specified internally consistent models, since we must impose some naivete on the agents—they are always surprised to wake up tomorrow. As the paper shows below, the out-of-sample predictions of myopic and forward-looking agents can be materially different.

Another strand concerns endogenous growth and spatial development (e.g. Black and Henderson (1999), Quah (2002), Rossi-Hansberg and Wright (2007), Desmet and Rossi-Hansberg (2010), Desmet and Rossi-Hansberg (2014), Desmet et al. (2015)). These models tend to be very rich and complicated, but have features such that in equilibrium the spatial heterogeneities disappear and/or the dynamics reduce to simple time-invariant conditions, like balanced growth paths.⁹ This paper resembles these approaches in that it reduces the state space to be a function of aggregate state variables, but the rational expectations approach here need not be restricted

⁷Recent comparative studies of regional housing supply in spatial equilibrium include Saiz (2010), Glaeser et al. (2014), and Albouy and Ehrlich (2016). Housing supply models with forward-looking agents include Paciorek (2013) and Murphy (2013). These are partial equilibrium settings.

⁸For theoretical treatments, see Arnott and Lewis (1979), Titman (1985), Capozza and Helsley (1990), Capozza and Li (1994), Capozza and Li (2002), Novy-Marx (2007). For empirical evidence, see Cunningham (2006), Cunningham (2007), Bulan et al. (2009).

⁹The richest model of spatial dynamics to my knowledge appears in Desmet and Rossi-Hansberg (2014) and Desmet et al. (2015). There, particular assumptions on free entry of firms and diffusion of knowledge allow for the solution of static problems instead of dynamic programs.

to steady states or balanced growth paths, allowing for analysis of shocks or policies at any point in the state space, including transitional paths. However, this paper envisions applications of policy analysis near observed equilibria, and is not designed to study endogenous growth, and therefore these are complementary works for different types of questions.

A third related strand studies population dynamics in island economies models (see Rappaport (2004), Rappaport (2005) for neoclassical growth models with labor mobility, and see Davis et al. (2013), Yoon (2014), Nenov (2015), Mangum (2015) for dynamic island economies settings). These models incorporate sticky features like moving costs, housing, or physical and human capital accumulation, and in so doing abandon the standard spatial equilibrium framework of no arbitrage in utility. Different equilibrium concepts apply. Yoon (2014) in particular is related because it also uses a forecasting rule to approximate rational expectations equilibria, although in Yoon (2014), the state space complexity comes from multiple types of people and sectors in a two-region economy.¹⁰

Comparing to the existing literature, the approach suggested here rigorously incorporates dynamics and does not avoid forward-looking behavior, but by offering a modification of a common setting, provides a convenient starting point for practitioners. The analyses can be conducted using standard equilibrium concepts and dynamic optimization techniques.

The paper proceeds in six parts. Section 2 introduces a partial equilibrium model of a forward-looking agent deciding whether and when to develop housing on vacant land. Dynamic and myopic versions are compared. The following section expands the model to a general equilibrium setting of a collection of forward-looking agents and proposes an approximated solution method. Section 4 evaluates the performance of the approximation method in the context of the example model. Section 5 provides a brief overview of estimation of the housing construction model; details are relegated to appendix section A. Section 6 compares the results for dynamic and myopic specifications in fit, estimation, and counterfactual simulations. Section 7 concludes. Appendix B provides details on the computation of the dynamic model.

2 Models With Forward-Looking Agents

The paper will illustrate the intuition of the approach to solving dynamic spatial equilibrium models by using a running example of housing suppliers across heterogeneous locations in a closed economy. I introduce the essentials of the model in this section. Construction is a natural

¹⁰Given the different environment, Yoon (2014) is more related to Lee (2005) and Lee and Wolpin (2006). These estimate models of worker sectoral mobility (discrete choice) with iterative rational expectations routines to determine the endogenous sectoral skill prices. These also adopt forecasting rules to approximate rational expectations equilibria, though the aggregation is over a large number of individual agents in a small number of sectors. The current model has single agents in a large number of locations (parallel to sectors), which is more similar to a dynamic game with many players.

context for an example since the capital is durable and development involves the exercise of real options. With that said, the point here is not to argue whether all housing models should be dynamic, but rather admit the possibility that some models should be. Nor is this the only type of application envisioned for the techniques, but merely a relevant example. I illustrate the differences in behavioral predictions in myopic and dynamic models and offer suggestions for their tractable implementation when appropriate.

Before proceeding, notice that there are multiple senses of the word “dynamic.” One is that the economy has durable features (i.e. state variables) that transition over time. Another sense is that the agents of the economy are aware of the durability and live long enough to be impacted by it. The former contrasts a dynamic environment versus a static (or malleable) one, while the latter concerns forward-looking versus myopic agents. In terms of solving and implementing models, it is clearly the horizon of the agent that matters most. Dynamic optimization models can be substantially harder to solve, estimate, and simulate than myopic ones. As described in the introduction, dynamics have been introduced to spatial models via myopic agents solving static problems under evolving conditions. This is fundamentally not the same as modeling behavior of an agent who knows s/he lives for more than an instant of time and that future outcomes will depend in part on current decisions. The behavior of that agent may be substantially different. I will refer to the two models as “dynamic” and “myopic” to emphasize the difference is in the decision horizon, while both versions have evolution of state variables.

First I consider whether the myopic and dynamic versions of this model are materially different in partial equilibrium; that is, with the agent capable of altering his own states but not affecting prices or the utility level of the economy.

2.1 A Model of Housing Supply in Heterogeneous Locations

An economy consists of a finite number of locations J which offer unique bundles of income, housing, and amenities to consumers. Each location consists of one agent responsible for housing supply; consider one of these agents. The housing supplier is profit maximizing, lives infinitely over discrete periods, and discounts future profits at rate β . The agent is endowed with ownership of undeveloped land and decides whether and when to convert vacant land into housing, which can be sold (including its underlying land input) to consumers; hence, I will refer to this agent as “the builder.” There is a stock of housing H built atop a stock of land area A , and stocks depreciate at rate δ . The builder can add to the stock but cannot intentionally remove it. It is this irreversibility which makes the builder’s problem dynamic: vacant land has continuation value, so its development exercises a real option.¹¹

¹¹The assumptions of infinite life and complete irreversibility are not necessary for imposing dynamics. Any nonzero removal cost would also impose a dynamic externality over multiple periods of life. The point is that the builder lives for multiple periods and something about the stock is durable, i.e. “putty-clay,” not “putty-putty”

The builder’s problem is written recursively as

$$V(X) = \max_x [\pi(X, x) + \beta E(V(X'|X, x))] \quad (1)$$

where V is the value function for behaving optimally (“Bellman”), $\pi()$ represents flow profits, X are state variables, x are choice variables, and $E()$ is the expectation taken over future states. Note the expectation is conditional on (i) current states, so there may be persistence, and (ii) current period choices, so decisions may affect the evolution of states (in particular, the builder’s stock variables).

2.1.1 Production and Input Factors

The builder makes housing services by combining land and capital. The production function is given by

$$i = F(a, k, \phi, \alpha) \quad (2)$$

where i is housing services added to stock, a is land input, k is capital, and ϕ and α are parameters governing housing productivity and the elasticity of substitution between inputs, respectively. In the empirics below, I use the standard Cobb-Douglas production function, $i = \phi a^\alpha k^{1-\alpha}$.¹²

The builder decides how much of each input to use each period. He adds to the housing stock according to (2), and the relative quantities of k and a determine density of building. A unit of capital comes at cost κ . This represents materials (wood, brick, glass, etc.) and the labor and equipment costs of their installation. A unit of land comes at cost ρ . This represents the cost of converting virgin land to suitability for building.¹³ Additionally, the builder faces a cost convex $c(\cdot)$ in the amount of housing added in one period. This cost captures anything inelastically supplied in the city in a short time horizon (frictions in procuring new land, delays at the permit office, opposition to “excessive” new building, etc.). For exposition, I refer to it as the land assembly cost (i.e. the difficulty of getting all the housing together at the same point in time). Convexity is mathematically important as well because it ensures a finite solution. Also, convexity effectively allows for within-city heterogeneity in costs and returns obscured by a single builder, meaning that under reasonable parameterizations, there is nonzero construction flow in each period, which is always a feature of the data for wide enough geographies (e.g. U.S. counties or larger).

and rebuilt each period (see also Rossi-Hansberg and Wright (2007)).

¹²Recent work on the housing production function has found a constant returns to scale Cobb-Douglas function to be a reasonably good approximation. See Epple et al. (2010), Combes et al. (2012), Ahfeldt and McMillen (2014), Albouy and Ehrlich (2016).

¹³Recall that the builder already owns the land and resells it in post-production form. Any land costs are therefore in the acquisition, holding, and conversion of land, not its market value.

The housing services output is sold at a price R to be explained below. The builder's flow profits are

$$\pi = Ri - \kappa k - \rho a - \frac{c(\cdot)}{1 + \nu} i^{1+\nu}$$

where $\nu \geq 0$ governs the extent of the convexity. Within a small area like a city, land is in limited supply. To reflect the exhaustibility of land,¹⁴ I specify the land assembly cost function to be increasing with the amount of land already employed, $c = c(A)$, $c'(A) \geq 0$.¹⁵ This means that the cost of adding more housing within a period is increasing in the amount of land already used up; or in other words, the short run elasticity is decreasing as land disappears. Intuitively, this could reflect increasing difficulty in land assembly, or greater likelihood of facing opposition from existing residents. Empirically, many large U.S. cities have seen housing prices increase faster than construction flows, indicative of decreasing price elasticity.

Assuming differentiability of the value function allows for the derivation of analytical first order conditions for optimization.¹⁶ Taking first order conditions yields, after some manipulation,

$$\frac{F_a}{F_k} = \frac{\rho - \beta V_A}{\kappa} \quad (3)$$

For example, if F is Cobb-Douglas,

$$\frac{F_a}{F_k} = \frac{\phi \alpha a^{\alpha-1} k^{1-\alpha}}{\phi (1-\alpha) a^\alpha k^{-\alpha}} \implies \frac{k}{a} = \frac{1-\alpha}{\alpha} \frac{\rho - \beta V_A}{\kappa} \quad (4)$$

which provides a condition for the demand for capital as a function of land (or vice versa). Using $\Delta = \frac{\rho - \beta V_A}{\kappa}$, the production function becomes

$$i = \phi (\Delta a)^{1-\alpha} a^\alpha = \phi \left(\frac{1-\alpha}{\alpha} \Delta \right)^{1-\alpha} a \quad (5)$$

which shows that the density of construction depends on the production function parameters and the relative costs of land and capital, Δ . The important point to note is that the cost of land includes its shadow price, V_A , the effect it has on continuation value, which will be negative since land is exhaustible (so $-V_A$ is positive, like a cost). Otherwise (3) and (4) are typical expressions of the rate of factor substitution. When the shadow value is larger in magnitude, because, say, land is near its exhaustion or future demand is expected to be high, builders will

¹⁴I use the term "exhaustible" for exposition, though technically I am not imposing a maximum amount of land available. This is consistent with standard models of urban areas that often endogenize the city boundary.

¹⁵This is not the only way to capture exhaustibility, but one specification which I found to be consistent with the data.

¹⁶The flow profit function is smooth and continuously differentiable, so the value function should inherit these properties (Stokey et al. (1989), Thm. 9.10). The policy function may be discontinuous at a threshold; at low enough output prices, the builder can decide to construct zero.

choose denser construction (more capital per unit of land) in order to conserve land.¹⁷

The factor demand condition allows for the re-expression of (1) as a continuous choice of a single input. Using $\Phi = \phi(\frac{1-\alpha}{\alpha}\Delta)^{1-\alpha}$,

$$V(X) = \max_a [R\Phi a - (\Phi\kappa + \rho)a - \frac{c(A)}{1+\nu}(\Phi a)^{1+\nu} + \beta E(V(X'|X, a))]$$

Maintaining the assumption of the differentiability of V , the above yields an analytical expression for the policy functions.

$$a = \left(\frac{R\Phi - (\Phi\kappa + \rho) + \beta(V_H\Phi + V_A)}{c(A)\Phi^{1+\nu}} \right)^{\frac{1}{\nu}} \quad (6a)$$

$$i = \Phi a \quad (6b)$$

From the upper expression, (6a), we can see that the amount of land exercised depends positively on the price of the output, and negatively on current costs, densities, and shadow costs. The shadow price of additional housing stock itself, V_H , independent of the land, is present in this expression. This will become important in the general equilibrium version of the model, but will be zero under the assumption that the solitary builder has no effects on prices.

Consider how two locations with heterogeneity in primitives might have different policy functions. Higher land or capital costs will cause one to build less than another. Higher land costs relative to capital will result in greater density through factor substitution. Higher assembly costs have the direct effect of less building, mitigated by the fact that these also increase density through the shadow value of land. Similarly, a greater shadow value of land results in less land exercised, partly through increased density and partly through greater reluctance to build at all.

2.1.2 Output Prices

Local utility, from consumption of numeraire goods n , housing h , and amenities μ , equals a reservation utility offered by the economy. The between-city condition is

$$u(n, h, \mu) = \bar{U}$$

I am treating \bar{U} as exogenous in this section, but later consider it determined by spatial equilibrium conditions throughout the economy.

Within cities, utility maximizing consumers tradeoff between the consumption of numeraire goods and housing services, subject to a budget constraint determined by the city's income. Housing services are rented for one period at rate r . Consumers are freely mobile and do not

¹⁷Additions to the housing stock may also have a shadow cost, V_H , which can affect the total amount of building, but this drops out of the relative first order conditions determining density.

have assets, so they solve a static problem. For example, if consumers have a standard log Cobb-Douglas utility function, $u = \ln(n) + \gamma \ln(h) + \mu$, the intra-location first order condition for optimization combined with the budget constraint $y = n + rh$ yields

$$r = \frac{\gamma}{1 + \gamma} \frac{y}{h} \quad (7)$$

for a given stock of housing, H , in the location. If the stock of housing is divided equally among the population of residents in the location, p , then $h = \frac{H}{p}$. The population—and hence, the rental rate and housing services per person—is pinned down by the reservation utility condition

$$\ln \frac{H}{p} = \bar{U} - \ln\left(\frac{1}{1 + \gamma} y\right) - \mu \quad (8)$$

Intuitively, (8) shows that consumers are willing to suffer less housing when local income and/or amenities are greater (when rents are higher). Moreover, combining (7) and (8) shows that rents will be lower, *ceteris paribus*, when the outside option is higher (when other places are relatively more attractive to consumers).

Consumers rent housing each period but builders sell it after construction. I want the model to focus on the builder’s real options problem, not an asset pricing problem. Hence, I model the construction decision only, not a landlord’s problem of renting existing stock to consumers. A modeling device for this is to assume the builder sells the stock to risk neutral middlemen who in turn rent it to residents. The price of this transaction should be determined as the present value of the future stream of rents. For simplicity, I assume risk neutrality and that incomes are a random walk, so the price is the discounted sum of current rents, inclusive of depreciation of the stock.¹⁸

$$R = \frac{1}{1 - \delta\beta} r \quad (9)$$

2.1.3 The Builder’s Solution

The state variables for the partial equilibrium version of the model are local income, y , land already used, A , and the reservation utility. I will assume for simplicity that local amenities do not evolve, though this is not necessary for any results. Housing stock itself is not a state variable in partial equilibrium (seen by combining (7) and (8)), but will be in general equilibrium. The partial equilibrium builder’s problem is written as

$$V(y, A, \bar{U}) = \max_a [\pi(y, A, \bar{U}, a) + \beta E(y', A', \bar{U}' | y, A, \bar{U}; a)] \quad (10)$$

¹⁸A practical issue for empirical implementation is that the builders either need to cover their cash cost of physical materials or have access to a financing market, and the latter would greatly complicate the model without much value added.

This is solved by (6a) and (6b) given laws of motions for the exogenous states y and \bar{U} . The endogenous state evolves according to $A' = (1 - \delta)A + a$.

2.2 Dynamic and Myopic Versions of the Model

Consider two versions of the above model: one in which the builder is forward looking, and another in which he is not, though the states evolve as specified in the dynamic version. The latter case is a myopic model, in that the builder is unaware of effects his actions have on future states.¹⁹ The above derivations would apply, except that any shadow values would no longer be relevant—the V_A, V_H terms drop out. How different might these be?

The next exercise gives an empirical comparison of these specifications for reasonable calibrations some (imaginary) U.S. metro areas.

2.2.1 Comparing Policy Functions

I consider a generic location with attributes of a typical large U.S. metro area: a population of 2.5 million residents, and housing services per person of 511 sqf. at a floor area ratio of 0.053 (which determines the land employed). The mean per capita income set to the national average of \$28,500. The builder in the location may face high (+1*sd*), medium (± 0 *sd*), and low (−1*sd*) demand states of income and national utility, for a matrix of nine potential demand states. These yield output prices of (9) via (7), (8), and then choices according to (6a),(6b) after solving (10).²⁰ I calculate the elasticity of each of these policy functions relative to the medium-medium demand state. I do this at high and low values of parameters on assembly cost (using $c(A) = cA$, $c = 100 \pm 50$, $\nu = 1$) and housing productivity (using $\phi = .05 \pm .025$), for each specification of the model, to see how the elasticities might vary with local attributes. Using the elasticity relative to the midpoint abstracts from level differences between cities with different attributes. I emphasize that all other variables of the model are the same, so that this exercise illustrates the impact of time horizon in a ceteris paribus setting.

Table 1 reports the percentage changes in policy functions for each type of model. (I will loosely use the term “elasticity” to describe the percentage changes.) The upper panel reports the percentage difference for land employed, while the lower panel reports housing services. The table is read in three-by-three blocks; for example, in the first block of numbers in the table, corresponding to the myopic model of the low cost, low productivity city, it shows that land developed is 10 percent lower in the low income, low national utility state than in medium

¹⁹This could be naivete or because the builder lives for only one period. A short decision horizon is unappealing here because the “builder” is really the landowner/custodian—not a fly-by-night construction firm—which is meant to capture real options for profit maximizing entities as well as constraints imposed by local governments who care about future externalities.

²⁰I take the observed AR1 transition of local incomes y and \bar{U} from the data described in appendix A.

income, medium utility state. (The middle entry of each block is zero by construction.) Recall that national utility, i.e. the value of the outside option, has an inverse effect on local demand, so that the lower left corner is the highest demand scenario for each city (high local income and low outside option), and the upper right the lowest (vice versa). For example, land added is 59 percent higher in high income, low national utility, and 49 percent lower for low income, high utility in the myopic, low cost, low productivity city.

The point of this table is to compare the myopic and forward looking agents. They are facing exactly the same states and parameters and differ only in the ability to project the future and internalize dynamic externalities. As the table shows, this difference is not trivial. The forward-looking agent is considerably more responsive to the various demand scenarios. For example, in the low cost, low productivity city, the forward-looking agent adds 86 percent more land in high income, low utility state instead of the 59 percent difference by the myopic agent. Why? Intuitively, the forward-looking agent expects the especially high demand states to be temporary and uses more of his exhaustible resource in the current opportunity. Likewise, he withholds more land in the low income, high utility demand state than does the myopic agent (-67 percent compared to -49). This is essentially intertemporal substitution in the timing development of the exhaustible resource, which the myopic agent cannot do.

Furthermore, heterogeneity in the builder's primitives matters in the comparison of myopic and forward looking agents. For a high cost, high productivity city, the gaps between the agents are smaller (52.6 compared to 43.9, and -41.9 compared to -36.6), because with higher costs, (i) the agent is more restricted within a single period, and (ii) his continuation value is lower, diminishing incentives for intertemporal substitution. Thus, the degree of misspecification could vary by attributes of the agent. Notice also that the myopic agent does not distinguish between the high and low cost regimes: the upper two blocks and lower two blocks of the panel are the same for the myopic agent. (There are level differences, but recall the table reports elasticities abstracting from levels.) Without his internalizing the rising cost, he makes the same proportional response to changing demand conditions. In contrast, costs affect the forward-looking agent's ability to substitute over time, and hence he is more sensitive in the low cost regime.

The lower panel, exhibiting the elasticities of housing services added (i.e. land plus capital), shows a similar pattern, but another dimension of comparison arises between myopic and forward-looking agents. The forward-looking agent is even more sensitive to changing demand conditions, choosing not only to add more land in good states and withhold in bad, but also to build more capital-intensively in the good states. This mitigates the burden on future costs while allowing him to reap the current period's profits. The myopic agent, however, is too naive to make the substitution and chooses inputs according to current prices alone; hence, the upper and lower panels are the same for the myopic agent.

Table 1: Comparison of Models: Elasticity of Policy Function With Respect to States

Location Parameter		States \ Model	Policy Function: Land					
c	ϕ		Myopic			Dynamic		
		Res. Utility: Income	Low	Med	High	Low	Med	High
Low	Low	Low	-10.02	-30.79	-49.52	-2.89	-36.79	-67.07
		Med	24.48	0.00	-22.08	40.65	0.00	-36.41
		High	59.29	31.06	5.61	86.25	38.87	-3.65
Low	High	Low	-7.42	-22.79	-36.66	-0.64	-27.80	-52.01
		Med	18.13	0.00	-16.35	32.75	0.00	-29.30
		High	43.90	23.00	4.15	68.40	30.12	-4.20
High	Low	Low	-10.02	-30.79	-49.52	-7.03	-32.36	-55.12
		Med	24.48	0.00	-22.08	30.06	0.00	-27.03
		High	59.29	31.06	5.61	67.94	33.13	1.80
High	High	Low	-7.42	-22.79	-36.66	-4.23	-24.11	-41.95
		Med	18.13	0.00	-16.35	23.68	0.00	-21.28
		High	43.90	23.00	4.15	52.61	25.13	0.41

Location Parameter		States \ Model	Policy Function: Housing					
c	ϕ		Myopic			Dynamic		
		Res. Utility: Income	Low	Med	High	Low	Med	High
Low	Low	Low	-10.02	-30.79	-49.52	-11.40	-43.26	-70.93
		Med	24.48	0.00	-22.08	42.50	0.00	-37.26
		High	59.29	31.06	5.61	105.79	51.75	4.11
Low	High	Low	-7.42	-22.79	-36.66	-12.21	-37.65	-59.53
		Med	18.13	0.00	-16.35	35.04	0.00	-30.54
		High	43.90	23.00	4.15	91.18	45.72	5.79
High	Low	Low	-10.02	-30.79	-49.52	-12.66	-37.04	-58.62
		Med	24.48	0.00	-22.08	31.10	0.00	-27.62
		High	59.29	31.06	5.61	80.08	41.74	7.62
High	High	Low	-7.42	-22.79	-36.66	-14.82	-33.67	-50.18
		Med	18.13	0.00	-16.35	25.36	0.00	-22.36
		High	43.90	23.00	4.15	71.53	39.12	10.40

NOTES: The table reports the percent change in the land (upper panel) and housing (lower panel) policy functions (see (6a), (6b)) with respect to changes in demand conditions for each time-horizon specification of the model. The location/agent is four versions of a generic city, as described in the main text. There are two demand states, local income and the economy's reservation utility level, structured in a 3x3 matrix of low/medium/high states; the differences are taken with respect to the center cells (medium/medium).

In summary, whether an agent is forward looking or not has consequences for his state-dependent actions, and in heterogeneous ways depending on the agent’s characteristics. This serves as motivation for a dynamic model with forward-looking agents. So far, I have regarded only a partial equilibrium setting. The next section introduces a spatial equilibrium version of the model of forward-looking agents and discusses associated empirical challenges.

3 The Dynamic Equilibrium Model

I now extend the model to a dynamic spatial equilibrium by relaxing the assumption of an exogenously fixed reservation utility. In the canonical Roback (1982) model, the level of utility in the economy, \bar{U} , is enforced by the free mobility of a sufficiently large set of agents—here, the consumers/residents—though the source of utility can vary by location (the relative amounts of consumer goods, housing goods, amenities, etc. in the bundle). The Roback model and many recent applications of it²¹ occur in the cross section, in single or repeated snapshots of time. Yet I suggest that this spatial equilibrium condition is a natural place to incorporate dynamics. The model above featured a dynamic optimization problem solved by an agent taking \bar{U} as transitioning by some exogenous law of motion. An equilibrium version of this model is one in which \bar{U} is determined endogenously, perhaps with the single agent having non-negligible effects on its evolution, but otherwise with a very similar structure. Of course, one needs to know how the reservation utility evolves over time with states of the economy. I now define the type of equilibrium the model readily accommodates, and discuss the challenges of empirical implementation.

3.1 The Equilibrium Definition

Consider the individual agent’s optimization problem rewritten as

$$\begin{aligned}
 V_{\tilde{v}}(\bar{U}, X) &= \max_x [\pi(\bar{U}, X, x) + \beta E(V_{\tilde{v}}(\bar{U}', X' | \bar{U}, X, x))] \\
 \bar{U}' &= v(\bar{U}, x(\tilde{v}), x_{-i}(\tilde{v}), \theta) \\
 X' &= F(X, x(\tilde{v}), \theta)
 \end{aligned}
 \tag{11}$$

where I have split out the notation on the equilibrium object \bar{U} from the agent’s own state variables X and defined generic laws of motion for these states. The reservation utility transition may depend on the agent’s action x and the actions of other agents, x_{-i} , and possibly other features of the environment collected in θ . Notice that the agents’ actions depend on their belief

²¹Notably, the Roback model has been extended in various ways by David Albouy and coauthors in Albouy (2009), Albouy (2012), Albouy and Stuart (2014), Albouy and Ehrlich (2016), Albouy (2016).

of the law of motion for \bar{U} . For a rational expectations equilibrium to hold, the actions of the individual agents must deliver the law motion for \bar{U} that the agents expect. The rational expectations equilibrium is

$$\begin{aligned} V_v(\bar{U}, X) &= \max_x [\pi(\bar{U}, X, x) + \beta E(V_v(\bar{U}', X' | \bar{U}, X, x))] \\ \bar{U}' &= v(\bar{U}, x(v), x_{-i}(v), \theta) \\ X' &= F(X, x(v), \theta) \end{aligned} \tag{12}$$

The difference between (11) and (12) is that in the latter, the rule the agents use corresponds to the actual rule, $\tilde{v} = v$. Hereafter, I refer to this as the rational expectations rule, “RER.”

3.2 The Challenge in Practice

The principle is simple enough, but in practice it can be very hard to derive conditions for the RER, much less the rule itself, especially when agents are heterogeneous and representative agent techniques would not apply. In spatial contexts, location itself is a form of heterogeneity, so the difficult cases are hard to avoid.

On the other hand, the spatial equilibrium model is already far simpler than a dynamic game with the same numbers of agents. Condensing the equilibrium object to a single reservation utility helps to avoid the intractable curse of dimensionality in a dynamic game in which every state of every agent enters the value function of each agent. For a large number of agents, such would be intractable to solve even once for one agent, much less jointly many times for all agents. Thus, the first insight this paper offers is that the classic spatial equilibrium model is a special case of dynamic heterogeneous agent models where empirics are made feasible by inherent model conditions.

Still, the model can be troublesome to apply empirically. Consider an economy consisting of J heterogeneous agents, like the model above. The transition of \bar{U} will depend on the individual actions of all the J agents in the economy, which means that J value functions have to be solved jointly, and the RER v updated accordingly.²² How does one do this?

Here, I turn to the suggestion of Krusell and Smith (1998): guess a RER, see if it fits, update that rule until it produces the joint conditions of (12). That is, iterate over functional forms and parameters until agents acting according to the RER generate the RER they condition upon. The approximated RER is written as

$$S(X, \theta) \approx v(X, \theta) \tag{13}$$

²²Even if estimation may be able to avoid full solution of the model, simulation for the purpose of policy evaluation would require model solution.

This is a numerical approximation to the “true” solution, much as a discretized or projected value function iteration is an approximation to the underlying functional equation. But note that the approximation is not the normalization itself, but its transition. The spatial equilibrium’s provision of reservation utility is the *exact condition* for agents in the economy *at a point in time*. But, the *transition* of \bar{U} could depend on precisely how economic activity is distributed over the locations—i.e. which agents are facing which states. Hence, using some S in (13) is an approximation to the true v , while \bar{U} is actually the true reservation utility of the model.

In summary, there are two insights. First, by conditioning on no-arbitrage in utility, the spatial equilibrium model naturally embeds a simplified approach for dynamics. One needs to be able to solve for the transition of the economy’s utility condition, not a high-dimensional set of best response functions. The second insight is that well-established numerical techniques allow one to find the necessary transition rules, to an approximation. The remainder of the paper will demonstrate an application of this method to the housing builder problem and show that it is an attractive and reasonable way to incorporate dynamics into empirical applications.

3.3 Application to the Builder’s Problem

I return to the builder’s problem of section 2, but now allowing the reservation utility to be determined endogenously. The environment is as before, but I must specify how a reservation utility across locations is obtained.

3.3.1 Defining Reservation Utility

The closed economy of J heterogeneous locations has an aggregate population of P , which for simplicity will evolve exogenously. In the model’s environment, consumers get utility from consumption goods, housing and amenities, with budget determined by local income. The stocks of land and housing and the budget constraints of the various locations evolve over time, and the suppliers of housing are forward looking about their durable stock. Housing evolves endogenously and income exogenously, and amenities are assumed fixed.²³

The first step in turning the model into a general equilibrium setup is to define how the reservation utility is determined. The interest of this paper is that the utility comes from within the closed economy of J locations, and is not determined by, say, some unchanging hinterland (Glaeser et al. (2014)). This is the more difficult case, as the latter could reduce to a static condition. I will specify the level of utility offered by the economy to be at the national average housing services per capita and consumption occurring at the intratemporal optimum for a consumer at this level of housing, with national amenities are normalized to be zero. That is,

²³Amenities need not be fixed, but the model already contains the features essential for the point of this paper.

$$\bar{U} = \ln\left(\frac{1}{1+\gamma}Y\right) + \gamma\ln\left(\frac{H_N}{P}\right)$$

where Y is national average income per capita, H_N is the level of housing stock nationally, and P is national population. The reservation utility contains features evolving both exogenously (Y, P) and endogenously (H_N). To be clear, this is not the only way to impose no spatial arbitrage, merely one that is empirically convenient.²⁴ Some normalization is necessary, but this particular one is not.

National utility rises in income and housing stock and falls with population. For an individual builder in some location j , the price of his housing services output has the following properties.

$$\frac{\partial R}{\partial y} > 0, \quad \frac{\partial R}{\partial Y} < 0, \quad \frac{\partial R}{\partial P} > 0, \quad \frac{\partial R}{\partial H_N} < 0$$

Thus, rents are affected by demand, rising with local income and population but falling in national average income (because the outside option is higher), and supply, falling with existing aggregate housing stock. The local housing stock is a component of housing supply, so that adding housing locally lowers future prices (hence, local stock has a shadow value), but at the same time, housing added by suppliers in other locations also suppress prices, *ceteris paribus*.

Individual state variables are the local income and land employed in housing stock. The economy's utility evolves according to the transition of national income, population, and the stock of housing. I assume these are known by the builders of each location, with income and population following exogenous rules. It will be necessary to find a RER because housing stock

²⁴One obvious alternative would be some reference location, perhaps some amalgamation of regions like rural lands, small cities collectively, and so on. These and others could serve as suitable normalizations, so long as one can characterize their evolution. An advantage of using national averages in the closed economy is that local population shares and rents are defined in closed form as ratios relative to national averages. These are

$$\begin{aligned} \ln r_j &= \ln\left(\frac{\gamma}{1+\gamma}\right) + \frac{1+\gamma}{\gamma}\ln(y_j) - \frac{1}{\gamma}\ln(Y) + \frac{1}{\gamma}\mu_j + \ln\frac{P}{H_N} \\ \ln\frac{p_j}{P} &= \frac{1}{\gamma}\ln\frac{y_j}{Y} + \ln\frac{H_j}{H_N} + \frac{1}{\gamma}\mu_j \end{aligned}$$

evolves endogenously. With state variables now specified, the builder’s problem is written as

$$\begin{aligned}
V(y, A, Y, P, H_N) &= \max_a [\pi(y, A, Y, P, H_N, a) + \beta E(y', A', Y', P', H'_N | y, A, Y, P, H_N; a)] \\
A' &= (1 - \delta)A + a \\
y' &= f_y(y, \theta_y) \\
Y' &= f_Y(\vec{y}, \vec{\theta}_y) \\
P' &= f_P(P, \theta_P) \\
H'_N &= (1 - \delta)H_N + S(Y, H_N, P, \theta_H)
\end{aligned} \tag{14}$$

where f_y, f_Y, f_P are functions denoting the exogenous laws of motion for incomes and population. (National income is a function of the vector of local states and laws of motion.) In (14), the transition functions are general, and functions appropriate to the context can be applied in practice. The RER is then found using repeated simulation of the model.

3.3.2 Finding the RER

The algorithm for finding the approximated RER S is

1. Start with an initial guess for S_0
2. Solve the value functions for each location’s agent conditioning on this guess, obtaining policy functions $i_j = \Phi_j(S_0) \cdot a(X, S_0)$
3. Use these policy functions to generate a new law of motion for the endogenous aggregate state, $H'_N - (1 - \delta)H_N = \sum_j i_j(X, S_0)$
4. Update the RER to $S_1 \approx \sum_j i_j(X, S_0)$
5. Repeat until convergence in S , i.e. when $\|S_{r+1} - S_r\| < \epsilon$ for some sufficiently small ϵ

Using this algorithm involves first selecting a suitable function for S . Like many numerical methods, this process takes some trial and error, involving iteration over functional form in addition to parameters. There are two general guidelines. One is that the enlisted RER have “small” errors, and the other is simply that the simulated equilibrium matches the data well. Another possible guideline, of course, is that simplicity is preferred.

The simplest starting point is a linear-in-parameters function which can be easily updated within a simulation routine. In this environment, the relevant variables are national demand states (population and income), and national supply (the stock of housing services).²⁵ After

²⁵More variables—say, to account individually for demand states in certain particular locations—can in principle be included, although these would come at the cost of additional variables in the agent’s state space.

some specification testing, I found a very simple functional form, $S \equiv v_0 + v_1 Y + v_2 \frac{H_N}{P}$ satisfies each of the guidelines. The predictive power is high, with R^2 values in excess of 0.98, and the model fits well, as will be shown below.

Several methods are available for solving the value functions. Projection methods were convenient in this case. Computation time was not burdensome. More details are available in appendix section B.

4 Approximation Quality

An obvious question is then, how bad is the approximation? This question is not easy to answer completely. The point of the approximation is not merely to speed up computation, but to make it feasible in the first place. The “true” dynamic equilibrium, where the transition rule depends on every state of every agent in the economy, is effectively unknowable for all but small collections of locations. So one cannot simply compare the full and approximate solutions once and proceed with empirics when the approximation is sufficiently close. However, the approximation quality can be evaluated to some degree.

There is, first, a conceptual argument for an approximation. Like Krusell and Smith (1998) and Weintraub et al. (2008), I consider S under a notion of bounded rationality. The full solution of the model—i.e. solving it like a dynamic game with each agent considering the individual strategies of all other agents—is not necessarily the most realistic characterization of the way agents behave. It is not a stretch to believe agents are forward looking, trying to predict future prices and the impact these have on current actions; however, tracking the states of each player and trying to anticipate their individual actions is a complicated information problem, even for sophisticated agents—one they probably not actually doing in the real world. Thus, the full solution could be more of a red herring than gold standard. The approximation may be closer to the “truth” in terms of modeling economic behavior.

With that said, I proceed to an quantitative exercise to evaluate how serious is the error imposed by the approximation. Recall the manner in which condensing of the information in a spatial equilibrium model is an approximation. In the current period, the aggregate states are actually what determine prices, so there is no loss of information in their summary. However, the *evolution* of these states will depend on how the components are distributed over space. For example, if there is income growth in a low cost location, the housing stock would grow larger than if the income shock had occurred in a high cost location. Thus, it is the laws of motion that contain approximation error.

4.1 In a Generic Economy

To get a sense of how much error this approximation imposes on an agent's current period policy function, I run a series of simulations which hold fixed the aggregate states but perturb their distribution over space. Agents will choose policies based on their rational expectations law of motion, and then I compare what their optimal policy would have been had they known exactly what the next period states would be. The idea is to measure how ignorance of the particular spatial distribution of the states imposes variance to policy functions.

I first do this in the generic model in order to work in a controlled environment. The agent of interest is the builder in the generic location, with population of 2.5 million residents and mean housing services per person of 511 sqf., mean per capita income set to the national average of \$ 28,500, and parameters of $c(A) = 2000A$, $\mu = 0$, and $\phi = 0.0385$. The rest of the economy is drawn to have high and low income ($\pm 25\%$ of mean), large and small size (50% smaller and 100% larger), and elastic or inelastic housing supply ($c(A) = (2000 \pm 100)A$). The interaction of these attributes yields $2^3 = 8$ locations. Another residual location is included to absorb perturbations to states in order to obtain the same national average; this residual otherwise has the attributes of the first generic location, and brings the total locations in the economy to ten. Another economy adds a high and low amenity dimension ($\mu = \pm 0.025$), yielding a total of $1 + 2^4 + 1 = 18$ locations in that scenario.

The perturbation exercise varies which of the other locations gets an income shock, but maintains the same level of national income; that is, it fixes the information that is known within the approximation, but varies the detail that is lost by not tracking each agent's states individually. Measuring the variance in the policy functions is one way to quantify how greatly the approximation could mischaracterize economic behavior relative to the full model.²⁶

The perturbation exercise proceeds as follows. First I simulate the economy to yield a baseline rational expectations law of motion for the aggregate endogenous state H . Next, I find the policy functions for all locations at the mean states. Then, I impose a positive income shock of size $q \in \{5\%, 10\%, 25\%\}$ to the income of one of the locations, maintaining the national average by subtracting an appropriately weighted amount of income from the last (10th or 18th) location, and leaving all other locations at their means. I then evaluate the policy functions at the perturbation, which may yield an evolution of H different from that expected using the RER. Finally, I return to the first location to evaluate its policy function had it known the aggregate state with perfect foresight, having not thrown away information on which location had received the income shock. The comparison of the two policy functions, $i(E(Y_{t+1}), E(H_{t+1}(S)))$ versus

²⁶Note that this is a first order perturbation only. That is, the other locations are still operating under the RER, S . This seems the most appropriate test, as to go further in allowing other locations to behave according to different rules would require updating the law of motion, which would return the simulation to another slightly different RER; this would be a comparison of two RERs, not a RER to a non-approximated law of motion.

$i(E(Y_{t+1}), H_{t+1})$, helps to quantify how great is the change in behavior as a function of the lost information.

Table 2 reports the results of the exercise for the economies of ten and 18 locations for income shocks of 5%, 10%, and 25%. For brevity, I summarize the differences by groupings of the locations receiving the income shocks. For example, the first row of the table reports the mean absolute percentage difference in the policy function when the small, low income locations (of any elasticity or amenity value) receive the income shock. The lower panel groups by elasticity and amenity, summarizing over size and income. In the first column of 5% income shocks, the table shows that the approximated rational expectations policy function and the perfect foresight policy vary by 0.04 to 0.14 percent, with the larger differences sensibly coming when the large and elastic locations are recipients of the shock. The differences grow roughly proportionally with the size of the shocks. When the economy expands to 18 locations, the differences drop considerably—falling more than 50 percent—because each individual location contributes less to the aggregate law of motion. This highlights that the approximation should perform better as the number of locations grows.

Comparing this table with Table 1, one can see that the potential error introduced by the approximation is orders of magnitude smaller than the differences between the static and dynamic models. The static and dynamic models in Table 1 are in many cases 20-40 percentage points apart, while in Table 2, the differences are less than a percentage point. This suggests there is far more potential for model bias in assuming myopia than forward-looking bounded rationality. Moreover, the issue of forward-looking behavior versus myopia is present in a single agent model and does not dissipate with more agents, whereas the errors of the dynamic approximation do dissipate as the number of locations grows.

4.2 Using U.S. Data

Continuing in this line of thinking, I also conduct perturbation exercises in a more realistic setting, using the economy on which the model is estimated. The data and estimation routine are described in the next section and appendix A. For now, it suffices to state that the model is estimated using the 49 largest metro areas in the U.S., with all other smaller metros aggregated into a residual fiftieth location.

As in Table 2, I evaluate the policy function at the mean of the states in the data as well as at perturbations of those states, constrained so that the national aggregates remain fixed, and then compare the policy functions. Instead of using the controlled environment of one income shock to one city at a time, I draw a full set of income shocks, positive and negative in the bounds $[-q, q]$, to $J - 2$ cities. The left out cities are the location of interest (i.e. the one whose policy function I am comparing) and the residual location, which absorbs the net change to the national aggregate. This more directly corresponds to the performance of the approximation

Table 2: Approximation Error in a Generic Model: Mean Absolute Percentage Differences to Policy Functions, by Location Receiving the Shock (%)

Economy:		Ten Locations			18 Locations		
Perturb. Size:		0.05	0.10	0.25	0.05	0.10	0.25
Attributes of Perturbed Location							
Size	Income						
Small	Low	0.098	0.203	0.562	0.017	0.036	0.098
Small	High	0.044	0.093	0.262	0.020	0.040	0.103
Large	Low	0.110	0.227	0.608	0.023	0.047	0.123
Large	High	0.140	0.105	0.268	0.017	0.041	0.103
Supply Amenity							
Elastic	Low	0.109	0.137	0.390	0.023	0.053	0.139
Elastic	High	-	-	-	0.023	0.046	0.119
Inelastic	Low	0.087	0.177	0.460	0.011	0.023	0.061
Inelastic	High	-	-	-	0.021	0.042	0.109

NOTES: The table reports the percent deviation in the policy functions between that using the approximated rational expectations equilibrium and the perfect foresight outcome. The economies are organized of generic locations as described in the main text, with two size economies present. The experiment changes which underlying receives a positive demand shock, holding fixed the aggregate states, and the table reports means over collections of recipient locations.

in an actual empirical setting, where there would be various shocks to many locations at once, while the agent continues to track only the aggregates.

The perturbation performed many times in loops: 100 perturbations for each of the $J - 1$ locations of interest (the residual locale is omitted from study), for both income shocks and changes to the distribution of the housing stock. I report the mean absolute percentage difference over the 100 perturbation draws for each location and state variable. The complete list of locations is reported in appendix Table C1, but for brevity, I summarize by location type in Table 3.

There are some mildly interesting differences between types of locations—for instance, large and elastic locations are more sensitive to the approximation error—but the punchline of the table is in the units themselves. The average differences between the approximation and perfect foresight policy functions are of the order 10^{-6} percentage points. While the errors from Table 2 were small, these are many times smaller yet because of the greater number of locations (50, instead of 18 or ten), and the fact that the simultaneous local shocks cancel out each other. In other words, for every elastically supplied city with a positive income shock, there is another with a negative shock, and the RER performs well on average.²⁷

In summary, while using the RER approximation introduces some potential for error, the differences are much smaller than those between dynamic and myopic specifications. Moreover, the approximation performs better and better as the number of locations grows, just as the

²⁷To this point, notice that the variance of the shock draws is nearly irrelevant.

Table 3: Approximation Error in The Estimated Model: Mean Absolute Percentage Differences to Policy Functions, by Location of Interest ($\% 10^{-6}$)

Perturb. State:		Income			Housing Stock		
Perturb. Size:		0.05	0.10	0.25	0.05	0.10	0.25
Attributes of Interest Location							
All	All	0.523	0.530	0.554	0.276	0.276	0.277
Size							
Small	Low	0.338	0.342	0.358	0.178	0.178	0.179
Small	High	0.473	0.479	0.502	0.250	0.250	0.250
Large	Low	0.240	0.244	0.255	0.127	0.127	0.127
Large	High	0.812	0.823	0.859	0.429	0.429	0.429
Supply							
Elastic	Low	0.317	0.355	0.370	0.185	0.185	0.185
Elastic	High	0.543	0.530	0.553	0.276	0.276	0.276
Inelastic	Low	0.618	0.698	0.730	0.363	0.364	0.364
Inelastic	High	0.530	0.512	0.536	0.267	0.267	0.267

NOTES: The table reports the average percent deviation in the policy functions between that using the approximated rational expectations equilibrium and the perfect foresight outcome. The economy is the 50 location U.S. metro area model deriving from actual data. The experiment holds fixed the aggregate states and then randomly draws shocks to the distribution of the states across locations. The simulation takes means over 100 independent draws for each city, and the table reports means over groups of city types. The full results for each city are available in appendix Table C1.

full solution would be getting harder. In contrast, the concerns over model specification would be present even in a one-agent partial equilibrium model. Thus, a dynamic model with an approximate numerical solution is not only feasible, but desirable.

5 Estimation Overview

This section briefly describes the input data and how the example model is solved and estimated. While the details are important the estimation of this particular model, they are not essential to the general point of how to implement dynamic spatial equilibrium models of this class, so the detailed discussion is relegated to appendix section A. Here I provide a brief overview to fix ideas, and explain where the routine for solving the rational expectations equilibrium enters.

5.1 Data

The level of geography is the metropolitan area as given by the U.S. Census definition of Core-Based Statistical Area (CBSA). I focus on the 49 largest CBSAs by 2000 Census population and aggregate all others into a single residual location.²⁸ Identified cities comprise approximately

²⁸Fifty is an arbitrary number of locations that seemed a reasonably large cross section without beginning to confront data availability issues in smaller cities. Estimation need not be constrained to 50. The smallest

two-thirds the national urban population, and the residual the remaining third. Key pieces of data are unavailable for rural areas, so these are excluded.

Data on housing stock, construction, and population comes from the U.S. Census, focusing on the period 1980-2011. Information on local incomes derive from the Bureau of Economic Analysis. Materials costs for home construction are obtained from the RS Means Company, a construction data analytics firm.

An important component of the data concerns the relative factor intensities for land and capital. These derive from detailed microdata on the stock of housing in U.S. metro areas as of 2012. County tax assessor records list the living area, lot size, and year of construction of every single family home in metropolitan counties, so I can measure densities by location and vintage.

5.2 Estimation

Some parameters can be calibrated outside the model, including the transitions of income and national population, and the utility parameters γ and μ which can be derived from the model's spatial equilibrium conditions. The discount rate β is set to 0.95 and the depreciation parameter δ comes from an autoregression using the time series of housing stocks in the U.S.

The primitive cost parameters are estimated using simulation of the model, targeting moment conditions for each city derived from the living area/lot size data. Intuitively, the assembly cost parameters are identified from the level of land exercised in a period conditional on output prices (i.e. observed supply elasticity), and parameters governing input factor intensities from the density of construction, which is observed to vary with prices and land stocks. The most flexible, well-identified function for cost parameters set the convexity parameter $\nu = 1$, and estimated two parameters of $c(A) = c_1 + c_2 A_{jt}$, which are constrained to be nonnegative. The housing production TFP parameter ϕ is estimated, setting ρ and κ a priori (more details available in appendix A).

The estimation begins with the RER as measured from the time series in the data, and then updates internally. In the first iteration, I use S_0 from the data, estimate the parameters for each location, and then simulate the model to general a new aggregate series. I then run a regression to update the parameters of the RER to S_1 , re-estimate the model for each city, and repeat until the parameters of the locations' primitives and the RER can no longer be updated. Since this is a numerical technique, there is no general proof of convergence or uniqueness, but the model has properties indicating it will indeed converge, namely that more housing supply reduces continuation value ($V_H < 0$). In words, the routine converges because if the RER predicts more housing services added, the individual locations add less, which will decrease the next iteration's

specified cities in my data are Salt Lake City, UT and Rochester, NY; the largest to be aggregated to the residual location are Bridgeport, CT and Tulsa, OK. New Orleans, LA is excluded because of the disruption to housing stock and construction caused by Hurricane Katrina.

predicted housing services in RER, which will increase the individual agents' construction, and so on until these balance in equilibrium.

5.3 The Myopic Model

Estimation of the myopic model proceeds much like the dynamic, of course without the steps related to solving the value functions. There is still an equilibrium that needs to be solved for, in that the myopic builders must agree on the stock of housing services at any point in time. The simulated model must deliver the predicted sequence of housing services, period by period.

6 Comparing Model Specifications

6.1 Model Fit


One immediate way to compare two models is to “horserace” them and evaluate their ability to fit the observed data. The following subsection briefly compares the ability of each model to match data. Fit is certainly an important evaluation of any model, but it turns out not to be the major distinction between the myopic and dynamic model.

Figure 1 presents a series of plots measuring the fit of each model. The top panel contains scatter plots (*a* and *b*) for the choice variables of interest, land inputs and housing services (living area). On the horizontal axis are the actual values of land and housing added for each city in the estimation period, and the vertical plots the models' predicted values, with dynamic model values denoted with a circle and the myopic an “×.” All values line up very nearly along the diagonal, indicating that both models are capable of matching the ample cross-sectional variation in the level of construction. Note that this happens independently for housing services and land. In the middle panel, plot *c* shows their ratio (the density of construction) is also fit well within each model. The ability of the models to fit these moments is not surprising given the large degree of flexibility and locational heterogeneity present in both.

A difference between models emerges, however, when considering the *change* in density within each city over time. In most locations, construction density increased from the 1980s to 2000s. The dynamic model naturally embeds a mechanism for this: as the option value of land increases, builders use less of it for a given amount of housing. Hence, densities rise when demand is anticipated to be high (e.g. income growth) or when land costs rise (e.g. land is exhausted). Plot *d* shows the dynamic model's land/capital tradeoff within a given city is consistent with the data, if a bit understated on average. The myopic model contains no such mechanism, and therefore has constant density throughout, missing this feature of the data.

The bottom panel plots the distribution of the estimator's objective function, mean squared error in housing construction and density, across locations. As the top panel suggested, the two

Figure 1: Comparison of Models' Fit to Data



NOTES: The figure displays several plots evaluating the fit of the myopic and dynamic specifications of the housing construction model. Plots *a – d* plot actual versus model-predicted values of the land and housing services policies. Plots *e, f* show the distribution of the objective functions across locations.

models are similar in terms of ability to fit total housing construction, and also their respective “misses” (instances of higher error) are similar. The dynamic model does better in fitting the density distribution for the reasons just mentioned.

In summary, each model is flexible enough to match data. Though the myopic model misses a key feature of the evolution of density, this difference need not be too alarming on its own. If a researcher were concerned about this moment but nothing else about forward-looking behavior, some ad hoc adjustment to the myopic model could be made by, say, adding a density parameter as a function of states or time, $\phi(A, t, Y)$ etc. Interpretation of parameters and counterfactual simulation, discussed next, are more fundamental concerns.

6.2 Estimates

Difference in fit is likely a surmountable concern, especially in a flexible model that can target several moments in the data. A deeper question is whether the parameters are the “right” ones, which obviously is especially important if one takes the parameters themselves seriously—to do, say, interpretation of the coefficients and conduct counterfactual exercises. Recent literature comparing dynamic and static/myopic models has focused on proper interpretation of coefficients (see Murphy and Bishop (2016), Bishop (2008), Ma (2016)).

In this case, though the models may both target the same moments and succeed similarly in fitting the data, they do so at different parameters values. Table 4 reports coefficient estimates, city-by-city, in the dynamic and myopic specifications, with means and medians across cities in the bottom rows. (Appendix Table C2 reports the coefficient estimates with their standard errors.) There are too many parameters to discuss everything, but some patterns are worth highlighting. First, the myopic model is far more likely to arrive at the zero constraint for the c_2 parameters that governs how assembly costs change with used land stock. This is sensible because it is that parameter is most sensitive to decision horizon. In a dynamic model, expectations of future costs affect land developed in all states of the world, affecting construction densities and levels and their variance with local prices, so that many locations are permitted to exhibit fluctuations in land developed (through density or withholding of options). Only characteristically elastic markets in the central U.S. and some low demand Rust Belt cities arrive at the zero constraint for this parameter. In the myopic model, in contrast, the parameter only has bite in locations where housing service supply elasticities appear to be declining with urban growth (some characteristically constrained coastal markets). Second, this difference in finding c_2 has effects on the estimated levels of TFP (ϕ), as average productivity estimates must be higher to rationalize density without an option value to land. Third, the two sets of estimates differ in their average levels. A simple way to see this is to compare the sum of the mean c_1, c_2 parameters. The myopic model on average needs higher cost parameters to arrive at the appropriate level of construction activity because it lacks an option value component that rationalizes retention of land; with myopia, all reluctance is ascribed to current period costs. If one were trying to interpret these coefficients as indicative of underlying cost mechanisms—for instance, land use or building regulations or geographic constraints—these differences in model specification could alter the conclusions about the import of these mechanisms on costs.

Thus, one’s assumptions about model horizon could impact interpretation of parameters estimates, which can matter because it can affect prediction within and outside the model’s environment. I proceed to a comparison of the model’s elasticity estimates within the regime on which the models are estimated, and then close the section by looking at counterfactual policy scenarios.

Table 4: Parameter Estimates, Dynamic and Myopic Model

Location	Dynamic			Myopic		
	c_1	c_2	ϕ	c_1	c_2	ϕ
Residual	50.65	0.00	0.036	61.74	0.00	0.035
New York, NY	391.86	272.71	0.051	384.60	594.43	0.102
Los Angeles, CA	791.63	404.24	0.119	1,302.84	0.00	0.207
Chicago, IL	1,210.51	66.39	0.090	1,323.06	0.00	0.103
Philadelphia, PA	1,504.18	1,080.85	0.048	2,422.37	823.78	0.099
Dallas, TX	867.75	0.00	0.109	897.40	0.00	0.120
Miami, FL	1,938.75	379.63	0.131	2,464.88	0.00	0.205
Washington, DC	1,908.93	574.73	0.043	2,871.82	0.00	0.069
Houston, TX	1,030.19	0.00	0.129	1,068.06	0.00	0.148
Detroit, MI	1,679.75	195.16	0.055	1,946.13	0.00	0.060
Boston, MA	820.90	532.67	0.030	680.14	1,560.40	0.047
Atlanta, GA	569.81	118.65	0.045	790.05	0.00	0.054
San Francisco, CA	3,327.91	767.47	0.120	3,103.85	1,425.82	0.262
Riverside, CA	1,119.79	99.62	0.137	1,250.38	0.00	0.156
Phoenix, AZ	1,069.66	54.71	0.157	1,164.79	0.00	0.188
Seattle, WA	1,427.67	684.99	0.053	2,633.32	0.00	0.091
Minneapolis, MN	1,985.25	514.14	0.028	2,365.39	545.44	0.034
San Diego, CA	3,434.07	0.00	0.061	3,352.63	0.00	0.055
St Louis, MO	1,849.55	494.53	0.040	2,595.67	0.00	0.047
Baltimore, MD	5,439.27	1,540.71	0.046	6,163.50	1,959.02	0.083
Pittsburgh, PA	1,696.63	617.02	0.039	1,974.76	925.10	0.051
Tampa, FL	1,810.58	409.94	0.104	2,447.20	0.00	0.153
Denver, CO	3,063.77	252.02	0.071	3,586.74	0.00	0.094
Cleveland, OH	3,086.83	0.00	0.061	3,175.75	0.00	0.065
Cincinnati, OH	2,042.03	270.32	0.043	2,580.14	0.00	0.050
Portland, OR	1,634.19	388.27	0.074	2,499.14	0.00	0.111
Kansas City, MO	1,827.31	306.40	0.082	2,287.92	0.00	0.097
Sacramento, CA	1,574.20	808.76	0.079	2,675.79	0.00	0.116
San Jose, CA	3,619.89	6,497.68	0.073	0.00	13,457.61	0.240
San Antonio, TX	2,331.34	0.00	0.077	2,380.11	0.00	0.082
Orlando, FL	2,053.09	160.92	0.103	2,358.31	0.00	0.129
Columbus, OH	3,567.07	0.00	0.052	3,661.03	0.00	0.055
Providence, RI	1,505.23	518.74	0.034	756.20	2,255.66	0.042
Norfolk, VA	3,662.13	0.00	0.086	3,752.58	0.00	0.096
Indianapolis, IN	1,961.37	517.60	0.055	2,844.66	0.00	0.069
Milwaukee, WI	8,425.35	0.00	0.102	8,642.09	0.00	0.118
Las Vegas, NV	2,463.65	198.66	0.199	2,762.52	0.00	0.281
Charlotte, NC	1,546.42	826.09	0.054	2,784.16	0.00	0.074
Nashville, TN	1,504.63	508.62	0.028	2,658.72	0.00	0.034
Austin, TX	2,752.60	0.00	0.098	2,768.05	0.00	0.100
Memphis, TN	3,304.04	0.00	0.042	3,454.39	0.00	0.044
Buffalo, NY	4,933.48	0.00	0.035	4,995.04	0.00	0.035
Louisville, KY	2,276.97	0.00	0.017	2,183.32	0.00	0.017
Hartford, CT	3,657.74	716.66	0.035	4,992.54	0.00	0.040
Jacksonville, FL	3,683.06	831.22	0.087	4,962.18	0.00	0.125
Richmond, VA	3,456.80	0.00	0.034	3,386.32	0.00	0.033
Oklahoma City, OK	3,712.27	0.00	0.043	3,761.60	0.00	0.043
Birmingham, AL	1,367.04	518.63	0.020	3,498.89	0.00	0.023
Rochester, NY	7,866.75	0.00	0.027	7,924.39	0.00	0.027
Salt Lake City, UT	10,271.21	48.11	0.138	10,289.71	0.00	0.138
Mean	2,581.51	443.54	0.070	2,897.74	470.95	0.095
Median	1,950.06	261.17	0.055	2,614.49	0.00	0.082

NOTES: The table reports the parameter estimates for each location under the dynamic and myopic specifications of the model. The full description of the data and estimation routine are in appendix A. Appendix Table C2 reports the coefficients with standard errors.

6.3 Predicted Elasticities

Table 5 repeats the exercise of Table 1, comparing the policy functions of myopic and dynamic models in response to changes in demand. The three-by-three block structure of the table is as before, reporting elasticities to the demand shocks, which here are one standard deviation changes to local and national income. All other states held at their means, and the center cell is zero by construction. Note there are a few differences in context from the initial table. The model is now in general equilibrium, and the parameters are the estimates for actual cities as reported in Table 4. For brevity, I highlight just a few example cities, San Francisco, New York, Boston, and Austin, TX.

In both types of models, construction and land exercised respond to the changes in demand, with the high local income, low national income case being the greatest increase to local relative demand. Naturally, elasticities are larger in more elastically supplied (lower c_1, c_2) locations. However, the models differ in their predictions on the size of the response—and in which dimension—depending on the conditions of the city in question. San Francisco, with high assembly costs but relatively high productivity (as governed by ϕ), shows large differences between myopic and dynamic models in predictions on the amount of land used; in the dynamic model, the response to higher demand comes at greater density, as scarce land is withheld for future use. In Boston, with lower housing productivity but lower assembly costs, the gaps in land use are small but differences in construction amounts are large. In New York, there are some of both effects. The sizes can be as large as 10 percentage points despite differing only in the model’s specified horizon, which I emphasize again is many orders of magnitude larger than approximation errors found in Table 3. Now, in Austin, the two model’s predictions are quite similar. There, the two models have very similar parameters, including the estimated change in land costs parameter, c_2 , approaching the zero constraint. Thus, the dynamic model is nearly the same as the static, and the models deliver similar predictions for that city.

These are a few examples, and differences across cities run the gamut.²⁹ As before, the takeaway is that the differences in prediction and the potential misspecification can vary with attributes of the agents. Researchers must be especially cautious of time horizon assumptions in models with rich spatial heterogeneity, where the biases imposed in one place may not apply to the same extent in another.

6.4 Simulations

Often, models of this class are used to evaluate hypothetical policy regimes and states of the world. The potentially most serious difference between dynamic and myopic models comes in

²⁹It is not the case, for example, that all large, or elastic, or productive cities respond the same way. There are other central U.S. cities (including Texas) with nontrivial differences between myopic and dynamic models.

Table 5: Comparison of Models: Elasticity of Policy Function in Estimated Model

			Policy Function: Land					
		States \ Model	Myopic			Dynamic		
City	Params. (dynamic)	Natl Inc.: Local Inc.	Low	Med	High	Low	Med	High
San Francisco	$c_1 = 3,327.91$	Low	-20.69	-21.70	-22.56	-11.70	-13.02	-14.15
	$c_2 = 767.47$	Med	1.21	0.00	-1.03	1.39	0.00	-1.20
	$\phi = 0.1198$	High	23.30	21.89	20.70	13.32	11.85	10.58
Boston	$c_1 = 820.90$	Low	-27.18	-28.50	-29.62	-29.18	-31.29	-33.08
	$c_2 = 532.67$	Med	1.58	0.00	-1.35	2.33	0.00	-1.99
	$\phi = 0.0297$	High	30.59	28.75	27.18	30.83	28.28	26.11
New York	$c_1 = 391.86$	Low	-25.23	-26.47	-27.51	-22.18	-23.88	-25.31
	$c_2 = 272.71$	Med	1.49	0.00	-1.26	1.89	0.00	-1.58
	$\phi = 0.0510$	High	28.47	26.71	25.23	23.40	21.32	19.59
Austin	$c_1 = 2,752.60$	Low	-24.02	-25.18	-26.18	-24.18	-25.30	-26.25
	$c_2 \approx 0$	Med	1.38	0.00	-1.19	1.33	0.00	-1.13
	$\phi = 0.0984$	High	27.01	25.40	24.02	27.10	25.56	24.24

			Policy Function: Housing					
		States \ Model	Myopic			Dynamic		
City	Params. (dynamic)	Natl Inc.: Local Inc.	Low	Med	High	Low	Med	High
San Francisco	$c_1 = 3,327.91$	Low	-20.69	-21.70	-22.56	-21.30	-22.40	-23.34
	$c_2 = 767.47$	Med	1.21	0.00	-1.03	1.32	0.00	-1.12
	$\phi = 0.1198$	High	23.30	21.89	20.70	24.27	22.73	21.42
Boston	$c_1 = 820.90$	Low	-27.18	-28.50	-29.62	-34.54	-36.49	-38.16
	$c_2 = 532.67$	Med	1.58	0.00	-1.35	2.35	0.00	-2.00
	$\phi = 0.0297$	High	30.59	28.75	27.18	40.06	37.31	34.98
New York	$c_1 = 391.86$	Low	-25.23	-26.47	-27.51	-29.52	-31.11	-32.45
	$c_2 = 272.71$	Med	1.49	0.00	-1.26	1.95	0.00	-1.64
	$\phi = 0.0510$	High	28.47	26.71	25.23	34.17	31.84	29.90
Austin	$c_1 = 2,752.60$	Low	-24.02	-25.18	-26.18	-24.08	-25.24	-26.23
	$c_2 \approx 0$	Med	1.38	0.00	-1.19	1.38	0.00	-1.18
	$\phi = 0.0984$	High	27.01	25.40	24.02	27.06	25.46	24.08

NOTES: The table reports the percent change in the land (upper panel) and housing (lower panel) policy functions (see (6a), (6b)) with respect to changes in demand conditions for each time-horizon specification of the model. Unlike Table 1, these are derived from the equilibrium model (not single agent) using estimates from actual U.S. data. There are four example locations reported; more results available upon request. There are two demand states, local and national income, structured in a 3x3 matrix of low/medium/high states; the differences are taken with respect to the center cells (medium/medium).

counterfactual simulation, when the researcher must put more faith in the quality of the model and has less discipline from actual data. This section demonstrates the differences between the dynamic and myopic models in various policy scenarios. The point of these exercises is to illustrate how the models differ and when and how severely using the wrong model will affect the predicted answer. The simulations are examples of the types of policies that might be evaluated with a model like this, but do not attempt seriously to answer any specific policy questions.

Each simulation will compare policy functions from the counterfactual scenario to a baseline using the estimated parameters (and RER for the dynamic model) from the data. For the dynamic model, a new RER and associated value functions must be found for the new policy regime. The simulations begin with all locations at their mean states (A_j, H_j, y, H_N, Y) and then simulate forward for ten periods (years). Each simulation produces a lot of data, so for brevity I present the results for some archetypal cities. The results present, heuristically speaking, a difference-in-difference evaluation of the models: the difference *between* models in the difference *within* the models' baseline and counterfactual simulations. The question is, how do the models predict changes in the agents' behaviors as their environments change?

6.4.1 A Tax on Housing Output

The first simulation is a tax on new housing services constructed by the builder. That is, the builder's return function becomes $(1 - \tau)R_{tj}i$. This could represent an actual tax on revenues, or an implicit one, such as a regulation that builders set aside some portion of their new construction to be sold below market value, in that the supplier adds some quantity of housing services i but receives only a fraction $1 - \tau$ of the revenue.³⁰ The tax is assumed the same magnitude in all locations for simplicity of exposition, and is set to $\tau = 0.1$.

Figure 2 displays the results for a selection of locations: a large, high rent, and inelastically supplied market, New York, and a smaller one, San Francisco, and a large, low rent, elastically supplied market, Atlanta, and a smaller one, Nashville. The figure displays the relative policy function under the counterfactual scenario plotted over the ten years of simulation; for example, a position of -0.1 indicates that there is ten percent less construction under the counterfactual than in the baseline. The upper panel plots housing services added, and the lower plots land used in construction of those housing services. The dynamic model is plotted in a blue solid line and the myopic in a green dashed line.

The direct effect of the policy is obvious—all locations build less in all periods. The initial impact is slightly larger than ten percent, with larger effects in the more elastic markets where densities are lower and materials are a higher share of costs.³¹ As time progresses, each market

³⁰This simulation is intended to mimic in a generic way an “affordable housing” policy imposed on builders.

³¹This can be seen by comparing the analytical relative policy functions, $\frac{(1-\tau)R-cc-\frac{\phi}{\partial h}+\frac{\partial V}{\partial h}}{R-cc-\frac{\phi}{\partial h}+\frac{\partial V}{\partial h}}$. Of course, prices and continuation values change between the counterfactual and baseline.

“bounces back” to some degree, since relatively higher housing prices result from the reduced supply remaining from initial construction levels. (Both dynamic and myopic specifications are equilibrium models, after all.) The bounce-back is larger in the more elastic markets.

Qualitatively, these effects are by design common to the myopic and dynamic models, though quantitative differences emerge in heterogeneous ways across markets. First, there are large differences between the models in the amount of land employed, seen in a comparison of the upper and lower panels of the figure. In the myopic model, the reduction is proportional to the change in housing services. In the dynamic model, the reduction is much smaller because the tax reduces option values, and land becomes less dear, so more is used than would be in the no-tax baseline; so, buildings are less dense. The gap between models is greatest in the most land constrained markets, where land option values have larger import (greater V_A) for the policy functions. The second difference (a by-product of the first) is that the counterfactual affects housing services less severely in the dynamic model. The direct effect is to build less, but the indirect effect is that land is less valuable, so more is used, causing a rebound in housing services offered. This rebound effect is more pronounced in the more elastic markets, where the builders’ value functions are higher since the flow of future dividends is larger (through higher quantities and lower costs, not higher prices).³² Hence the gap between the myopic and dynamic models in housing services is larger for the more elastic markets. In San Francisco, the housing service policies are effectively overlapped, though with the intercept gap in the land policy functions. In Atlanta and Nashville, the difference in slope between models drives the predicted housing services policies apart.

The takeaway is that though the policy’s direction of impact is clear, empirical predictions can depend on which type of model is used. The dynamic and myopic models can differ in the effect on choice variables (factor intensities) and their indirect effects of construction levels, but in ways that are heterogeneous between markets and over time.³³

6.4.2 An Increase to Regulatory Burden


The next example simulation returns to a no-tax regime but considers an increase in the per period cost to housing construction, the assembly cost parameters c_1, c_2 . This reduces supply elasticity by making it more difficult to add housing in a given period, and is intended to represent, for instance, larger regulatory burdens on construction. Again for simplicity I use a common magnitude of ten percent increase across all locations.

Figure 3 uses the same organization and archetypal cities as the last simulation. In the

³²That value functions are lower in places with high rents may be counterintuitive. This is a consequence of modeling the agent as a developer-builder, and a single entity per metro area, not separate agents controlling individual parcels. It is better to be the developer of all of Atlanta than all of San Francisco, since the new market size is greater.

³³Of course, a heterogeneous policy like taxes that vary between markets would only complicate things further.

Figure 2: Changes to Policy Function in Builder Tax Simulation: Dynamic and Myopic Models (proportional differences over t periods)



NOTES: The figure plots policy function responses to a uniform housing tax on builders under both the dynamic and model specifications of the model. The vertical axis reports the proportional change (e.g. $-0.1 = 10\%$ lower) and the horizontal is the time elapsed since the change was instituted. Separate panels report the land and housing policy responses. See the main text for details of the simulation. Four example locations/agents are reported, and more results are available upon request.

myopic model, the cost increase has identical initial impact across cities and factor inputs. The degree of bounce back (from low initial supply leading to housing price increases in later periods) is similar across cities, but to a slightly steeper degree in larger and more elastic places. The dynamic model shows greater richness in the effects. As in the housing tax simulation, the impact on land is smaller, since the option value is lower and therefore the builders are less apt to retain it. The gap between models in land used is larger in the land constrained cities (New York, San Francisco). The greater willingness to use land, the indirect effect, means that the impact on housing is smaller in the dynamic model than in the static, though the gap between models can vary significantly based on density; the higher housing TFP places (e.g. San Francisco) have less gap than lower TFP (e.g. Nashville), since using relatively more land does not lead to that much more housing in San Francisco, while it does in Nashville. The dynamic model also exhibits heterogeneity between cities in the degree of bounce back over time, depending on the contribution of the land stock-dependent element c_2 versus c_1 parameter to per-period assembly costs (see values in Table 4). Those with more land stock dependence of costs (New York, Nashville) react more to the lower levels of stock in later periods.

In summary, like the tax simulation, the direction of the effect on housing services is obvious and qualitatively similar between the two models, but quantitative differences emerge in ways heterogeneous between locations.

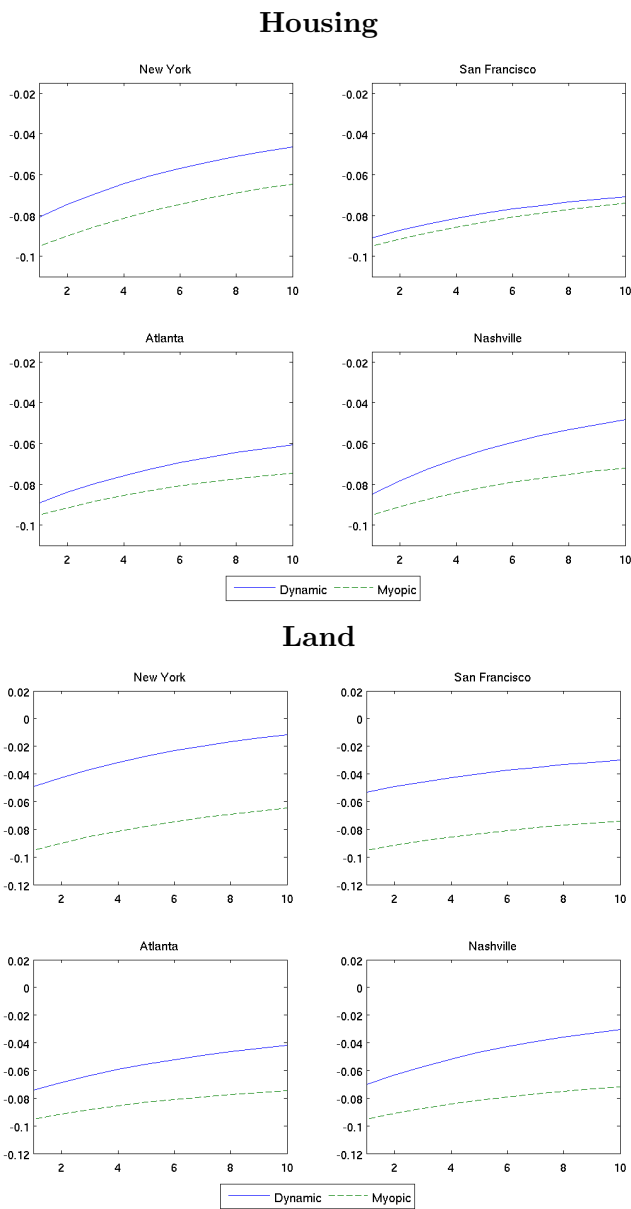
6.4.3 A Change in Location Amenities

The next simulation looks at altering location demand by changing the relative amenity parameters μ_j . Specifically, I reduce the amenity parameter of each city located on an ocean coast (Atlantic, Pacific, or Gulf of Mexico) by one standard deviation of 0.0276. This could represent a world in which coastal location is relatively less attractive because of, say, increased risk of storms or flooding brought on by climate change.

Like the previous simulations, this change to primitives requires a new solution of value functions and finding the associated RER. Unlike previous simulations, the change in parameters applies to only a subset of locations, though each will be affected in equilibrium. The simulation again starts from the mean values of location's states and walks forward ten periods. I swap out two archetypal cities in order to display a high rent, inelastically supplied coastal location (New York), and a low rent, elastically supplied one (Jacksonville, FL), and then a high rent, inelastically supplied inland city (Chicago) and low rent, elastically supplied one (Atlanta).

Figure 4 displays the results. Reduced demand for coastal locations leads to lower construction there. In equilibrium, inland locations see an increase in residual demand, so construction increases there. In the myopic model, the changes to housing and land employed are the same, whereas in the dynamic model, we again see the effect of option values—the changes to land employed are smaller in magnitude than housing. This can have varied indirect effects on hous-

Figure 3: Changes to Policy Function in Regulatory Burden Simulation: Dynamic and Myopic Models (proportional differences over t periods)



NOTES: The figure plots policy function responses to an increase to the assembly cost parameters (c_1, c_2) under both the dynamic and model specifications of the model. The vertical axis reports the proportional change (e.g. $-0.1 = 10\%$ lower) and the horizontal is the time elapsed since the change was instituted. Separate panels report the land and housing policy responses. See the main text for details of the simulation. Four example locations/agents are reported, and more results are available upon request.

ing construction, depending on the location's parameters. In the elastic locations (Jacksonville and Atlanta), the land effect mitigates the changes to housing. In the inelastic ones (New York and Chicago), the changes to building density actually amplify (to a small degree) the effect on housing.

Like the supply side interventions above, the qualitative effects are mechanically similar between the two models, but the quantitative effects can in some instances be quite different. The dynamic model exhibits a richness the myopic one lacks when the general equilibrium imposes substitution effects from dynamic externalities.

6.4.4 A Change in Location Housing Stocks

The final simulation is the most simple. The thought experiment is a sudden loss of housing stock in eastern coastal locations,³⁴ resulting from, say, a catastrophic event like a hurricane. For simplicity, I use the same ten percent loss in all East Coast locations.³⁵ To make ceteris paribus comparisons, I assume the loss in stock does not affect the primitive cost of land assembly nor demand for the locations through, for instance, the amenity parameters or total urban population. Thus, instead of changing parameters or other primitives, this simulation merely changes state variables, so the value functions and RER need not be updated. Therefore, the housing and land effects are identical within the dynamic model as well as the myopic, so I only report the results for housing services. The initial states are at the means, except for the loss of housing stock in affected cities.

Figure 5 reports the results for two directly affected cities, New York and Boston, and two unaffected, San Francisco and Atlanta. There is nationally a rebuilding effort since demand for housing remains, so all places see an uptick in construction that wanes over time as stock returns. The rebuilding is greatest in the coastal locations which have newly available land. However, this newly available land comes with newly available option value which mutes the rebuilding response. Thus, the rebuilding under the dynamic model is less than the myopic model, where it is effectively proportional to the loss. This insight highlights the potential importance of path persistence in models with durable capital that myopic specifications do not capture.

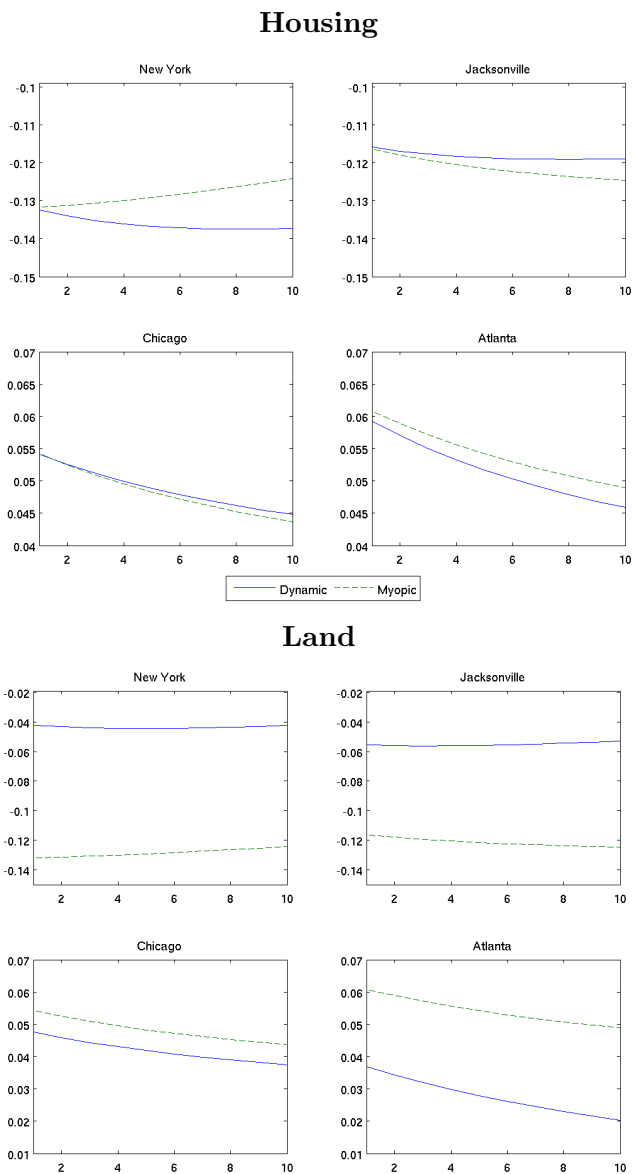
7 Conclusion

This paper suggests a path forward for incorporating dynamics into empirical spatial equilibrium models. The main insight is that models with spatial indifference already embed a natural starting point for incorporating dynamics: the level of reservation utility can be treated as a

³⁴Results for other coasts and catastrophes are available upon request for those feeling dismal.

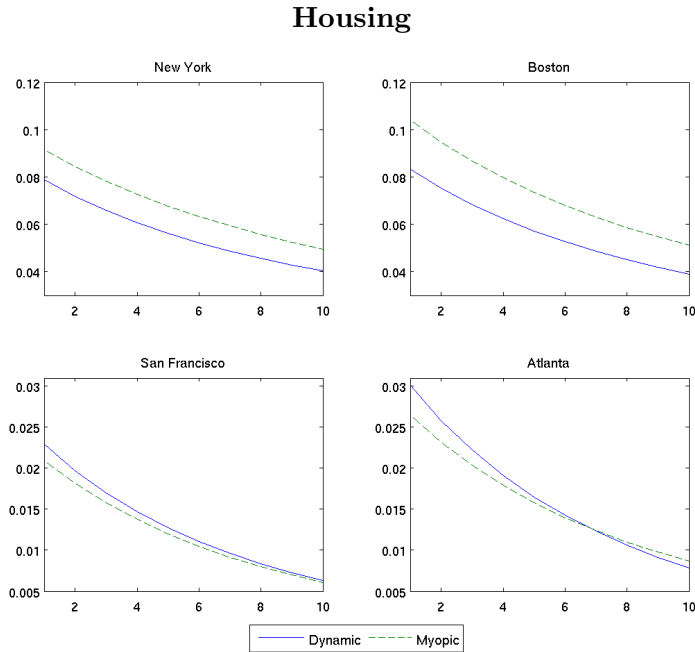
³⁵These are, by size: New York, Philadelphia, Washington DC, Boston, Baltimore, Providence, Norfolk, and Jacksonville.

Figure 4: Changes to Policy Function in Coastal Amenity Simulation: Dynamic and Myopic Models (proportional differences over t periods)



NOTES: The figure plots policy function responses after a change to the amenity value of coastal located cities under both the dynamic and model specifications of the model. The vertical axis reports the proportional change (e.g. $-0.1 = 10\%$ lower) and the horizontal is the time elapsed since the change was instituted. Separate panels report the land and housing policy responses. See the main text for details of the simulation. Four example locations/agents are reported, and more results are available upon request.

Figure 5: Changes to Policy Function in Coastal Stock Loss Simulation: Dynamic and Myopic Models (proportional differences over t periods)



NOTES: The figure plots policy function responses to a capital loss in east coast cities under both the dynamic and model specifications of the model. The vertical axis reports the proportional change (e.g. $-0.1 = 10\%$ lower) and the horizontal is the time elapsed since the change was instituted. See the main text for details of the simulation. Four example locations/agents are reported, and more results are available upon request.

state variable. This evades a potentially very complicated curse of dimensionality in spatial models with heterogeneity. The second insight is that existing techniques for solving rational expectations models can be readily applied to empirically implement dynamic spatial equilibrium models. The paper demonstrates the techniques on a relevant example of an economy of interrelated heterogeneous housing suppliers, showing that approximation errors are small and illustrating how dynamic and myopic specifications of the same model can depart from one another.

It would be nice if researchers had some rule of thumb that a dynamic model will always have, say, effects ten percent larger or smaller, etc., but it is clear that no such rule exists in general. The gap between model specifications can vary by context, by the attributes of the agent in question, and over time as the states evolve. Thus, researchers are left to do the hard work of solving dynamic equilibrium models in order to have confidence in the empirical predictions.

This paper aims to offer good news in this regard. While full solutions for economies of heterogeneous markets may quickly become infeasible, approximation techniques are readily available and likely introduce minimal error. This paper illustrates a numerical method for solving models of this class. It cannot rule out in general that there exists some model for which this approximation performs poorly, but it at least offers opportunity for researchers to test and experiment. By suggesting a path forward, these methods could be the next step in empirical spatial equilibrium models, making dynamics more attainable than is apparent at first glance.

References

- AHFELDT, G. M. AND D. P. McMILLEN (2014): “New Estimates of the Elasticity of Substitution Between Land and Capital,” Mimeo.
- ALBOUY, D. (2009): “The Unequal Burden of Federal Taxation,” *Journal of Political Economy*, 117, 635–667.
- (2012): “Are big cities bad places to live? Estimating quality of life across metropolitan areas,” Mimeo.
- (2016): “What are cities worth? Land rents, local productivity, and the total value of amenities,” *Review of Economics and Statistics*, 98, 477–487.
- ALBOUY, D. AND G. EHRLICH (2016): “Housing Productivity and the Social Cost of Land-Use Restrictions,” University of Michigan.
- ALBOUY, D. AND B. STUART (2014): “Urban population and amenities,” Mimeo.
- ALONSO, W. (1964): *Location and land use. Toward a general theory of land rent.*, Cambridge, Mass.: Harvard Univ. Pr.
- ARNOTT, R. J. AND F. D. LEWIS (1979): “The transition of land to urban use,” *The Journal of Political Economy*, 161–169.
- BISHOP, K. C. (2008): “A Dynamic Model of Location Choice and Hedonic Valuation,” Olin School of Business, Washington University in Saint Louis.
- BLACK, D. AND V. HENDERSON (1999): “A Theory of Urban Growth,” *Journal of Political Economy*, 107, 252–284.
- BLACK, D. A., N. KOLESNIKOVA, AND L. J. TAYLOR (2014): “Local Labor Markets and the Evolution of Inequality,” *Annu. Rev. Econ.*, 6, 605–628.
- BROXTERMAN, D. A. AND A. M. YEZER (2015): “Why does skill intensity vary across cities? The role of housing cost,” *Regional Science and Urban Economics*, 55, 14–27.
- BRUECKNER, J. K. (1987): “The structure of urban equilibria: A unified treatment of the Muth-Mills model,” *Handbook of regional and urban economics*, 2, 821–845.
- BULAN, L., C. MAYER, AND C. T. SOMERVILLE (2009): “Irreversible investment, real options, and competition: Evidence from real estate development,” *Journal of Urban Economics*, 65, 237–251.

- CAPOZZA, D. AND Y. LI (1994): “The Intensity and Timing of Investment: The Case of Land,” *The American Economic Review*, 84, 889–904.
- CAPOZZA, D. R. AND R. W. HELSLEY (1990): “The stochastic city,” *Journal of Urban Economics*, 28, 187–203.
- CAPOZZA, D. R. AND Y. LI (2002): “Optimal land development decisions,” *Journal of Urban Economics*, 51, 123–142.
- COMBES, P.-P., G. DURANTON, AND L. GOBILLON (2012): “The production function of housing: Evidence from France,” *Processed, University of Toronto*.
- CUNNINGHAM, C. R. (2006): “House price uncertainty, timing of development, and vacant land prices: Evidence for real options in Seattle,” *Journal of Urban Economics*, 59, 1 – 31.
- (2007): “Growth controls, real options, and land development,” *The Review of Economics and Statistics*, 89, 343–358.
- DAVIS, M. A., J. D. FISHER, AND M. VERACIERTO (2013): “Gross Migration, Housing and Urban Population Dynamics,” Federal Reserve Bank of Chicago.
- DAVIS, M. A. AND F. ORTALO-MAGNE (2011): “Household Expenditures, Wages, and Rents,” *Review of Economic Dynamics*, 14, 248–261.
- DESMET, K., D. K. NAGY, AND E. ROSSI-HANSBERG (2015): “The Geography of Development: Evaluating Migration Restrictions and Coastal Flooding,” Working Paper 21087, National Bureau of Economic Research.
- DESMET, K. AND E. ROSSI-HANSBERG (2010): “On Spatial Dynamics,” *Journal of Regional Science*, 50, 43–63.
- (2014): “Spatial Development,” *American Economic Review*, 104, 1211–43.
- DIAMOND, R. (2016): “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000,” *American Economic Review*.
- EPPLE, D., B. GORDON, AND H. SIEG (2010): “A New Approach to Estimating the Production Function for Housing,” *American Economic Review*, 100, 905–924.
- GLAESER, E. L., J. GYOURKO, E. MORALES, AND C. G. NATHANSON (2014): “Housing dynamics: An urban approach,” *Journal of Urban Economics*, 81, 45 – 56.
- GYOURKO, J. AND A. SAIZ (2006): “Construction Costs and the Supply of Housing Structure,” *Journal of Regional Science*, 46, 661–680.

- HENDERSON, J. V. (1974): “The sizes and types of cities,” *The American Economic Review*, 64, 640–656.
- KRUSELL, P. AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.
- LEE, D. (2005): “An Estimable Dynamic General Equilibrium Model of Work, Schooling, and Occupational Choice,” *International Economic Review*, 46, 1–34.
- LEE, D. AND K. I. WOLPIN (2006): “Intersectoral Labor Mobility and the Growth of the Service Sector,” *Econometrica*, 74, 1–46.
- LJUNGQVIST, L. AND T. SARGENT (2001): *Recursive Macroeconomic Theory*, Cambridge, MA: MIT Press.
- MA, L. (2016): “Learning in a Hedonic Framework: Valuing Brownfield Remediation,” Mimeo., University of Kentucky.
- MANGUM, K. (2015): “Cities and Labor Market Dynamics,” W. J. Usery Workplace Research Group Paper Series, 2015-2-3.
- MILLS, E. S. (1967): “An aggregative model of resource allocation in a metropolitan area,” *The American Economic Review*, 57, 197–210.
- MORTEN, M. AND G. BRYAN (2015): “Development and the Spatial Allocation of Labor: Evidence from Indonesia,” Mimeo., Stanford Univ.
- MORTEN, M. AND J. OLIVEIRA (2016): “Paving the Way to Development: Costly Migration and Labor Market Integration,” Working Paper 22158, National Bureau of Economic Research.
- MURPHY, A. (2013): “A Dynamic Model of Housing Supply,” Mimeo., Arizona State University.
- MURPHY, A. AND K. C. BISHOP (2016): “Valuing Time-Varying Attributes Using the Hedonic Method: When is a Dynamic Approach Necessary?” Mimeo., Arizona State University.
- MUTH, R. F. (1969): *Cities and Housing: The Spatial Pattern of Land Use*, University of Chicago Press.
- NENOV, P. T. (2015): “Regional reallocation and housing markets in a model of frictional migration,” *Review of Economic Dynamics*, 18, 863 – 880.
- NIEUWERBURGH, S. V. AND P.-O. WEILL (2010): “Why Has House Price Dispersion Gone Up?” *Review of Economic Studies*, 77, 1567–1606.

- NOVY-MARX, R. (2007): “An Equilibrium Model of Investment Under Uncertainty,” *Review of Financial Studies*, 20, 1461–1502.
- PACIOREK, A. (2013): “Supply Constraints and Housing Market Dynamics,” *Journal of Urban Economics*, 77, 11–26.
- POTERBA, J. (1992): “Taxation and Housing: Old Questions, New Answers,” *AEA Papers and Proceedings*, 82, 237–242.
- QUAH, D. (2002): “Spatial agglomeration dynamics,” *American Economic Review Papers and Proceedings*, 92, 247.
- RAPPAPORT, J. (2004): “Why are population flows so persistent?” *Journal of Urban Economics*, 56, 554–580.
- (2005): “How does labor mobility affect income convergence?” *Journal of Economic Dynamics and Control*, 29, 567 – 581.
- ROBACK, J. (1982): “Wages, Rents, and Quality of Life,” *Journal of Political Economy*, 90, 1257–1278.
- ROSSI-HANSBERG, E. AND M. L. J. WRIGHT (2007): “Urban Structure and Growth,” *The Review of Economic Studies*, 74, 597.
- RUGGLES, S., K. GENADEK, R. GOEKEN, J. GROVER, AND M. SOBEK. (2015): “Integrated Public Use Microdata Series: Version 6.0 [Machine-readable database]. Minneapolis: University of Minnesota,” .
- SAIZ, A. (2010): “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*, 125, 1253–1296.
- STOKEY, N. L., R. E. LUCAS JR, AND E. C. PRESCOTT (1989): *Recursive methods in economic dynamics*, Harvard University Press.
- TITMAN, S. (1985): “Urban Land Prices Under Uncertainty,” *The American Economic Review*, 75, 505–514.
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. VANROY (2008): “Markov Perfect Industry Dynamics With Many Firms,” *Econometrica*, 76, 1375–1411.
- WHEATON, W. C. AND W. E. SIMONTON (2007): “The Secular and Cyclic Behavior of ‘True’ Construction Costs,” *Journal of Real Estate Research*, 29, 1–25.

WOOLDRIDGE, L. W. (2002): *Econometric Analysis of Cross Section and Panel Data*, Cambridge, MA: MIT Press.

YOON, C. (2014): “Estimating a Dynamic Spatial Equilibrium Model to Evaluate the Welfare Implications of Regional Adjustment Processes: The Decline of the Rust Belt,” *International Economic Review*.

A Data and Estimation Details

A.1 Housing Stocks and Flows

I use annual county population estimates from the Census, 1980-2011, aggregated to CBSAs to measure city size and national (urban) population. To measure the housing stock for each metro area over time, I use single family housing units by county from the decennial Censuses of 1980, 1990, 2000, and 2010, aggregated to the CBSA level. For intercensal years, I allocate the decadal change in the housing stock by the level of building permit activity in the CBSA, as collected by the Census and provided by Housing and Urban Development’s State of the Cities Database (SOCD). As in Glaeser et al. (2014), I use permitting activity as an index because these are a noisy estimate of actual building activity, and do not necessarily sum to the change in housing stock. The index allocates permits by their expected arrival to the inventory of housing, which may vary spatially and temporally, using annual regional summaries of permit-to-completion rates and times from the Census. The allocated units comprise “construction” by city-year.

To measure input factor intensities, I use detailed data on home size and lot size by location and vintage year of construction. This information comes from microdata of U.S. county tax assessor records compiled by real estate data firm Dataquick/Corelogic. The tax assessor files are a single cross section of urban counties from 2011 to 2012, but the property records contain year of construction. Thus, I can measure the housing and land intensity per unit by the place and time of entry into the housing stock. The measure of housing-to-land density is the actual flooring area ratio (FAR), living area to lot size, for a particular structure. Unfortunately, these data are not available for multi-unit buildings, so throughout I focus on single family homes.³⁶ Populations are scaled by the fraction of the metro area living in single-family homes taken from a five percent subsample of the decennial census (Ruggles et al. (2015)).

Then, to get the total stock of living area H and land employed in housing, A , I multiply the number of units added each city and year by the living area and lot sizes measured from the tax assessor records for that city and year of construction:

$$H_{jt} = (1 - \delta)H_{j,t-1} + units_{j,t} \cdot liv_area_{j,t}, \quad A_{jt} = (1 - \delta)A_{j,t-1} + units_{j,t} \cdot lotsize_{j,t}$$

where j indexes locations and t time. δ is a depreciation rate to be calibrated below. I treat stock in 1980 as the initial condition because it is the first year of permitting data available in the SOCD. The estimation will target the policy functions for each location, $i_{jt} = H_{jt} - (1 - \delta)H_{j,t-1}$, $a_{jt} = A_{jt} - (1 - \delta)A_{j,t-1}$.

The capital component of costs, cc , is obtained using physical construction cost data from

³⁶Construction activity data also tend to be more precise for single family than multifamily, as there is less measurement error in the counting of units.

the RS Means Company. The data report material with installation (labor and equipment) cost, on the basis of square foot of living area, annually by city for 1988-2013. Below I described how I map these costs to capital costs of the model.

A.2 Calibrations

Some preliminaries are calibrated outside the model. The data come at annual frequency, and transition processes are set accordingly. The builders' common discount rate is set to $\beta = 0.95$. Local per capita income is taken from the regional economic accounts data by the Bureau of Economic Analysis. The cities' annual income processes are found to have unit roots, and I take the empirical distribution of annual differences to be the shock distribution.

To connect a flow cost of residential housing to sales values available to a builder, I use the commonly known user cost method (see e.g. Poterba (1992)). I find the implied rental rate, $r = uc \cdot v$, where uc is the user cost rate and v is the house value. Poterba (1992) suggests the user cost formula $uc = [(1 - \kappa_t)(m + \kappa_p) + \psi + \nu_r] - \nu_g$, where κ_t is the income tax rate, m is the nominal mortgage rate, κ_p is the property tax rate, ψ is maintenance cost and depreciation, ν_r is the risk premium associated with housing, and ν_g is expected inflation. Calibration of these follows Poterba (1992), Albouy (2009), and my own estimate of average appreciation by market.

Housing values are set by the sales of newly constructed homes using the Dataquick/Corelogic microdata on deed transactions. For most counties of the metro areas in the data, transaction registers from deeds records can be matched to the assessor records using a unique property identifier. Then, I obtain home values per square foot by averaging transaction prices from 2004 and 2005 sales of newly constructed homes as identified by the year built field. In counties for which no transactions data were available, I used median value for homes built 2005 or later in the 2008-2010 American Community Survey (ACS). The values were converted to 2000 dollars using the Consumer Price Index (less housing) and averaged to the CBSA using housing units by county as weights. Values over time are pegged to the Federal Housing Finance Administration's (FHFA) all-transactions housing price index for the CBSA. Finally, the user cost method is applied to yield $r_{jt} = uc_j v_{jt}$ on a per square foot basis.

The model implies in (7) that housing expenditure constitutes a constant fraction of income, so that the utility parameter γ can be calibrated from expenditures on housing. Davis and Ortalo-Magne (2011) have shown using microdata on incomes and rents that average housing expenditures are remarkably consistent across metro areas, lending direct support for this type of utility function. However, the utility function has been challenged for its failure to hold across the income distribution (Black et al. (2014), Broxterman and Yezer (2015)). Considering the current model focuses on across city differences and uses homogenous agents, the functional specification seems appropriate for the context. Moreover, the analytical expressions for population and rent that Cobb-Douglas utility delivers are very convenient, not a trivial consideration for equations

that might otherwise need to be solved numerically millions of times during estimation. Using the mean expenditure share of 0.19 (s.d. of 0.03), γ is set to 0.23.

With rents, incomes, population and housing stocks, I turn to the utility conditions of the model to derive the city-specific amenity value. To recover these, I run the regression

$$\frac{1}{\gamma} \ln \frac{y_j}{y_Y} + \ln \frac{h_j}{h_N} = D'_j \mu + \sigma \ln \frac{y_j}{y_N}$$

where h_N is the national average housing services per person and D_j is a matrix of indicator variables for each city. The extra parameter σ comes from specifying the amenities in units of income and relaxes the assumption that population elasticity with respect to income be $\frac{1}{\gamma}$, which is empirically far too high.

The depreciation parameter, δ , is calibrated using a regression of the national time series of housing stock and permits $H_{N,t} - i_{n,t-1} = (1 - \delta)H_{n,t-1}$; I use the national series (not metros) because it is longer (1947 using the Statistical Abstract of the United States) and the permits data are less subject to measurement error. Note that this measures stock depreciation, not value depreciation. I find $1 - \delta = 0.989$, indicating that about one percent of the single family housing stock is destroyed each year.

A.3 Estimation Routine

Having obtained parameters general to all locations and the amenity terms, I turn to estimation of the cost parameters for the builders in each metro area. While the housing construction data are detailed, they are actually insufficient to separately identify all the parameters specified in the general formulation of the model. Estimation will therefore normalize some parameters in such a way as to maintain the major dimensions of heterogeneity between cities.

First consider the relative factor intensities. Intuitively, these are identified by the building densities observed in the data. Equation (4) is the model's density condition. The first challenge with this condition is that true capital is not actually observed (or easily defined, for that matter). The RS Means data provide the materials and installation cost for a square unit of finished housing; this is the bundled k component of i , but not units of k . The observable density condition is (5), housing services per unit of land. Hence, I set $\kappa = 1$ and subtract the capital cost component from the output price, so that (9) become $\frac{1}{1-\delta\beta}r - cc$, where cc is the construction cost from RS Means.³⁷

³⁷There is significant spatial heterogeneity in the costs, but virtually no evidence that costs fluctuated with the level of building activity. The finding that construction labor and materials are elastically supplied is common (see Gyourko and Saiz (2006), Wheaton and Simonton (2007)), but costs vary spatially and temporally with construction sector wages (the "installation" component). Discussions with data engineers at RS Means indicated the labor component determines fluctuations. To measure the correlation of construction costs with income, I run a regression of cc_{jt} on y_{jt} , pooling the data across cities and including city fixed effects: $cc_{jt} = cc_{0,j} + cc_y y_{jt}$, for which I find $cc_y = 0.6142$ (s.e. 0.0388).

The second challenge is that the parameters determining density—the TFP ϕ , the elasticity of substitution parameter α , and land cost ρ —are in practice difficult to identify.³⁸ The important feature to maintain is locational heterogeneity in the capital/land ratio, and in the dynamic model, its evolution with other state variables like output prices and land in use. Thus, I elect to set $\rho = 1, \alpha = 0.5$ for all locations, and estimate ϕ . The restrictive implication of this assumption is that densities rise faster with V_A in ceteris paribus dense places, but this feature is satisfied in the data.

Next consider the assembly costs. Intuitively, these are identified by level and elasticity of building activity (conditional on the density and construction material costs). The model's condition is (6a). I first need to specify the functional form for $c(A)$, which I set to be $\frac{1}{1+\nu}(c_1 + c_2 \frac{A_{jt}}{A_j})i^{1+\nu}$ subject to the constraints $c_1 \geq 0, c_2 \geq 0$. This form allows but does not impose that the assembly costs increase in the amount of land already in use (\hat{A}_j is the average land stock in the data, a normalization to make the parameters comparable across cities of different physical size). In practice, it was difficult to separately identify the scale and convexity parameters, so I set $\nu = 1$ (making costs quadratic), which allows heterogeneity in elasticity to come through the c_1, c_2 parameters.

After reducing the estimable parameters, the policy function has become

$$a = \left(\frac{(R - cc)\Phi - 1 + \beta(V_H\Phi + V_A)}{c_1 + c_2 \frac{A_{jt}}{A_j} \Phi^2} \right) \quad (15a)$$

$$\frac{i}{a} = \Phi = \phi(1 - \beta V_A)^{0.5} \quad (15b)$$

These form the joint moment conditions in the estimation of ϕ, c_1, c_2 for each city. The objective function is a vector of squared residuals for each city's T observations,

$$M = \sum_{t=1}^T \left(\begin{array}{c} ((\hat{i}_{jt}(c_1, c_2, \phi, \nu) - i_{jt})/i_{jt})^2 \\ ((\hat{a}_{jt}(c_1, c_2, \phi, \nu) - a_{jt})/a_{jt})^2 \end{array} \right) \quad (16)$$

where the errors are in percentages. The density moments can be noisy in some locations, so the residuals are weighted by a Gaussian kernel on their distance away from a quadratic trend. The estimates are then $(\hat{c}_1, \hat{c}_2, \hat{\rho}) = \text{argmin} M$. In practice, I find the parameters by first conducting a coarse grid search over a wide guess of values, and then a standard simplex-based minimization routine using the grid search outcome as starting values. This reduces concerns over finding local minima. Following Wooldridge (2002) on M-estimators, standard errors are calculated using numerical first and second derivatives to find the score and Hessian matrix for

³⁸The functional form causes the parameters have slightly different implications for density to exploit, but in city level data, there are simply not enough data points to distinguish these.

the objective function (16).

B Solution of the Value Function

B.1 Method

Any standard method for solving the agents' value functions can apply. For this model, I use projection methods to approximate the the value function in (14). This choice was the result of two considerations. First, projections tend to be computationally faster since they can be updated by a simple matrix operation, and there is a preference for speed since the value functions must be solved many times in estimation. Second, this particular model has closed form solutions for policy functions if one can evaluate the derivatives of the function functions (V_A, V_H in (6a), (6b)). With a linear-in-parameters value function, these are known in closed form, so the policy step is trivial.

The value function approximation is specified as

$$V(X) \approx \sum_k \lambda_k F_k(x) \tag{17}$$

where $F_k(x)$ are polynomial functions of the state space and λ_k are the parameters on these polynomial terms. After much specification testing, I found the polynomial with linear terms, second order interactions, and the square of the local income state to be a very good approximation. Adding more terms did not improve the fit of the function.

The projections are evaluated on a randomly drawn sample of state space basis points. For each location, I use 1000 basis points in total, 500 each from draws of a uniform distribution with bounds 25 percent outside those of the data, and the empirical distribution of the data. This evaluates the value function at points outside the areas observed in the data (which may be relevant for counterfactuals), but makes the basis thick in the neighborhood of the data.

Updating the value function is a computationally inexpensive step of finding

$$\lambda_k = [F(X) - \beta E(F(X))]^{-1} \pi(X) \tag{18}$$

which can be evaluated by a regression. This does however require an evaluation of the expected polynomial terms, although when using linear-in-parameters projections, these can be pre-computed.

The derivative of the value function is readily obtained by

$$V_A \approx \sum_k \lambda_k \frac{\partial F_k(x)}{\partial A}; \quad V_H \approx \sum_k \lambda_k \frac{\partial F_k(x)}{\partial H} \quad (19)$$

B.2 Computation Time

Computational times will of course vary with machines, software, and code structure, but I report my experience with computation for reference. I used a dual eight core workstation with 8mb of memory per core.

Computer time is not unduly burdensome. Projection methods have the disadvantage of not guaranteeing convergence—and hence, the researcher might spend time experimenting to find reasonable functions—but with linearity, they can be very fast. As noted, the value function evaluation is a simple matrix problem and policy step is a closed form under the approximation.

Conditional on a RER, the individual agent’s problem is effectively a single agent problem. A single solution of the value function problem took one to two seconds on the 1,000 basis point state space. Evaluation of the integral over future states, $E(F(X))$, by Monte Carlo integration can be somewhat expensive, lasting up to three minutes with 5,000 draws over 50 locations. However, Monte Carlo integration can be easily parallelized, cutting times to about one minute with pre-packaged Matlab parallelization (over 12 cores). Moreover, the matrix can be precomputed so it is not embedded in the value function iteration. (It does, however, have to be reevaluated for each iteration of the RER.)

Nested fixed point estimation of one location takes a few thousand iterations from grid search to minimization routine. With integration, value function iteration, and evaluation of the objective function, with steps parallelized when possible, estimation of the 50 locations takes about three hours. The RER had to be updated about a dozen times, so complete estimation took about two days run time.

For the counterfactual simulations, the RER has to be updated a from few up to a few dozen times, depending on how different is the counterfactual from the baseline. The housing tax and assembly cost simulations take about ten minutes. The coastal amenity simulation takes about eight minutes. The housing stock destruction simulation takes about four minutes.

C Additional Exhibits

Table C1: Approximation Error in The Estimated Model: Mean Policy Function Percentage Differences by Location (10^{-6})

Perturbation To: City / Perturb. Size:	Income			Housing Stock		
	0.05	0.10	0.25	0.05	0.10	0.25
New York, NY	0.2251	0.2280	0.2379	0.1184	0.1184	0.1187
Los Angeles, CA	0.0772	0.0781	0.0813	0.0409	0.0409	0.0409
Chicago, IL	0.1293	0.1307	0.1357	0.0684	0.0684	0.0685
Philadelphia, PA	0.1008	0.1018	0.1057	0.0533	0.0532	0.0529
Dallas, TX	0.0478	0.0484	0.0506	0.0253	0.0254	0.0254
Miami, FL	0.0517	0.0524	0.0545	0.0274	0.0274	0.0274
Washington, DC	0.0006	0.0006	0.0007	0.0004	0.0004	0.0004
Houston, TX	0.0329	0.0333	0.0349	0.0174	0.0174	0.0174
Detroit, MI	0.1034	0.1047	0.1095	0.0545	0.0546	0.0546
Boston, MA	0.0964	0.0978	0.1023	0.0509	0.0509	0.0510
Atlanta, GA	0.2012	0.2040	0.2133	0.1062	0.1062	0.1064
San Francisco, CA	0.0891	0.0902	0.0940	0.0471	0.0471	0.0471
Riverside, CA	0.0545	0.0552	0.0577	0.0286	0.0286	0.0287
Phoenix, AZ	0.0399	0.0404	0.0423	0.0211	0.0211	0.0211
Seattle, WA	0.0873	0.0885	0.0924	0.0461	0.0462	0.0463
Minneapolis, MN	0.0724	0.0735	0.0771	0.0382	0.0382	0.0383
San Diego, CA	0.0113	0.0115	0.0120	0.0060	0.0060	0.0060
St Louis, MO	0.0929	0.0942	0.0986	0.0490	0.0490	0.0491
Baltimore, MD	0.0365	0.0371	0.0389	0.0193	0.0193	0.0194
Pittsburgh, PA	0.0960	0.0974	0.1019	0.0506	0.0507	0.0507
Tampa, FL	0.0241	0.0244	0.0255	0.0127	0.0128	0.0128
Denver, CO	0.0042	0.0042	0.0044	0.0022	0.0022	0.0022
Cleveland, OH	0.0808	0.0819	0.0858	0.0426	0.0426	0.0426
Cincinnati, OH	0.0220	0.0223	0.0234	0.0117	0.0117	0.0117
Portland, OR	0.0302	0.0306	0.0320	0.0160	0.0160	0.0160
Kansas City, MO	0.0383	0.0388	0.0406	0.0202	0.0203	0.0203
Sacramento, CA	0.0075	0.0076	0.0080	0.0040	0.0040	0.0040
San Jose, CA	0.0472	0.0479	0.0502	0.0249	0.0249	0.0249
San Antonio, TX	0.0271	0.0275	0.0287	0.0142	0.0142	0.0143
Orlando, FL	0.0132	0.0134	0.0140	0.0070	0.0070	0.0070
Columbus, OH	0.0249	0.0252	0.0264	0.0131	0.0132	0.0132
Providence, RI	0.0184	0.0186	0.0195	0.0098	0.0098	0.0098
Norfolk, VA	0.0321	0.0326	0.0341	0.0169	0.0169	0.0170
Indianapolis, IN	0.0502	0.0509	0.0532	0.0265	0.0265	0.0266
Milwaukee, WI	0.0141	0.0143	0.0150	0.0075	0.0075	0.0075
Las Vegas, NV	0.0788	0.0799	0.0836	0.0415	0.0415	0.0416
Charlotte, NC	0.0128	0.0130	0.0136	0.0068	0.0068	0.0068
Nashville, TN	0.1061	0.1075	0.1125	0.0560	0.0560	0.0561
Austin, TX	0.0004	0.0004	0.0004	0.0002	0.0002	0.0002
Memphis, TN	0.0049	0.0050	0.0052	0.0026	0.0026	0.0026
Buffalo, NY	0.0101	0.0102	0.0107	0.0054	0.0054	0.0054
Louisville, KY	0.1165	0.1181	0.1236	0.0614	0.0615	0.0615
Hartford, CT	0.0128	0.0130	0.0136	0.0067	0.0067	0.0067
Jacksonville, FL	0.0205	0.0208	0.0217	0.0108	0.0108	0.0108
Richmond, VA	0.0033	0.0033	0.0035	0.0018	0.0018	0.0018
Oklahoma City, OK	0.0120	0.0121	0.0127	0.0068	0.0068	0.0068

Table C2: Parameter Estimates With Standard Errors, Dynamic and Myopic Model

Location	Dynamic						Myopic					
	c1 est	c2 est	ϕ est	c1 se	c2 se	ϕ se	c1 est	c2 est	ϕ est	c1 se	c2 se	ϕ se
Residual	50.7	0.0	0.036	4.8	4.3	0.002	61.7	0.0	0.035	15.0	16.6	0.001
New York, NY	391.9	272.7	0.051	78.0	21.3	0.004	384.6	594.4	0.102	107.3	117.4	0.004
Los Angeles, CA	791.6	404.2	0.119	84.8	57.7	0.009	1,302.8	0.0	0.207	426.2	453.4	0.007
Chicago, IL	1,210.5	66.4	0.090	77.2	41.9	0.007	1,323.1	0.0	0.103	191.4	192.0	0.002
Philadelphia, PA	1,504.2	1,080.8	0.048	363.1	161.8	0.005	2,422.4	823.8	0.099	491.7	501.0	0.005
Dallas, TX	867.7	0.0	0.109	52.3	16.2	0.007	897.4	0.0	0.120	289.3	361.5	0.006
Miami, FL	1,938.8	379.6	0.131	200.7	129.4	0.016	2,464.9	0.0	0.205	485.8	519.7	0.006
Washington, DC	1,908.9	574.7	0.043	215.8	114.7	0.004	2,871.8	0.0	0.069	306.9	347.4	0.003
Houston, TX	1,030.2	0.0	0.129	57.6	20.4	0.009	1,068.1	0.0	0.148	548.5	629.2	0.008
Detroit, MI	1,679.7	195.2	0.055	176.9	89.1	0.004	1,946.1	0.0	0.060	225.7	211.6	0.002
Boston, MA	820.9	532.7	0.030	259.7	41.0	0.003	680.1	1,560.4	0.047	162.1	188.7	0.001
Atlanta, GA	569.8	118.7	0.045	96.3	61.2	0.004	790.1	0.0	0.054	175.0	226.2	0.004
San Francisco, CA	3,327.9	767.5	0.120	2,775.3	1,784.8	0.081	3,103.9	1,425.8	0.262	864.8	895.3	0.010
Riverside, CA	1,119.8	99.6	0.137	97.3	66.8	0.015	1,250.4	0.0	0.156	1,200.0	1,456.2	0.018
Phoenix, AZ	1,069.7	54.7	0.157	84.8	24.6	0.013	1,164.8	0.0	0.188	1,029.8	1,369.8	0.029
Seattle, WA	1,427.7	685.0	0.053	283.0	161.2	0.005	2,633.3	0.0	0.091	445.5	507.4	0.006
Minneapolis, MN	1,985.2	514.1	0.028	303.5	112.4	0.002	2,365.4	545.4	0.034	282.7	266.5	0.001
San Diego, CA	3,434.1	0.0	0.061	380.4	229.8	0.013	3,352.6	0.0	0.055	290.5	274.8	0.003
St Louis, MO	1,849.6	494.5	0.040	271.9	129.5	0.003	2,595.7	0.0	0.047	302.0	304.8	0.001
Baltimore, MD	5,439.3	1,540.7	0.046	1,126.2	560.5	0.007	6,163.5	1,959.0	0.083	1,218.2	1,230.4	0.003
Pittsburgh, PA	1,696.6	617.0	0.039	415.8	152.0	0.003	1,974.8	925.1	0.051	4,249.8	4,250.1	0.001
Tampa, FL	1,810.6	409.9	0.104	178.2	136.8	0.009	2,447.2	0.0	0.153	746.9	876.4	0.006
Denver, CO	3,063.8	252.0	0.071	264.4	131.7	0.006	3,586.7	0.0	0.094	476.7	507.5	0.003
Cleveland, OH	3,086.8	0.0	0.061	251.3	144.9	0.003	3,175.8	0.0	0.065	425.4	425.3	0.002
Cincinnati, OH	2,042.0	270.3	0.043	265.6	116.6	0.003	2,580.1	0.0	0.050	288.9	280.5	0.001
Portland, OR	1,634.2	388.3	0.074	325.2	152.2	0.007	2,499.1	0.0	0.111	475.3	479.7	0.006
Kansas City, MO	1,827.3	306.4	0.082	241.1	144.1	0.008	2,287.9	0.0	0.097	471.4	499.3	0.003
Sacramento, CA	1,574.2	808.8	0.079	306.9	227.9	0.010	2,675.8	0.0	0.116	628.2	779.9	0.006
San Jose, CA	3,619.9	6,497.7	0.073	3,397.0	1,777.8	0.015	0.0	13,457.6	0.240	4,167.4	4,404.8	0.010
San Antonio, TX	2,331.3	0.0	0.077	175.9	86.1	0.006	2,380.1	0.0	0.082	903.3	1,147.9	0.006
Orlando, FL	2,053.1	160.9	0.103	185.2	131.6	0.011	2,358.3	0.0	0.129	1,328.0	1,707.9	0.011
Columbus, OH	3,567.1	0.0	0.052	316.3	158.0	0.003	3,661.0	0.0	0.055	511.2	522.7	0.002
Providence, RI	1,505.2	518.7	0.034	445.0	35.3	0.003	756.2	2,255.7	0.042	170.1	178.5	0.001
Norfolk, VA	3,662.1	0.0	0.086	2,579.6	1,564.2	0.080	3,752.6	0.0	0.096	634.6	610.8	0.002
Indianapolis, IN	1,961.4	517.6	0.055	340.5	205.0	0.004	2,844.7	0.0	0.069	930.6	1,037.3	0.005
Milwaukee, WI	8,425.3	0.0	0.102	645.1	327.9	0.008	8,642.1	0.0	0.118	2,028.1	1,876.5	0.004
Las Vegas, NV	2,463.7	198.7	0.199	226.5	54.9	0.016	2,762.5	0.0	0.281	26,967.0	39,298.0	0.785
Charlotte, NC	1,546.4	826.1	0.054	392.1	298.4	0.006	2,784.2	0.0	0.074	3,789.4	4,361.8	0.011
Nashville, TN	1,504.6	508.6	0.028	513.1	183.7	0.003	2,658.7	0.0	0.034	500.6	550.2	0.002
Austin, TX	2,752.6	0.0	0.098	165.1	46.2	0.006	2,768.1	0.0	0.100	873.2	1,092.8	0.006
Memphis, TN	3,304.0	0.0	0.042	566.8	295.8	0.004	3,454.4	0.0	0.044	425.6	450.1	0.001
Buffalo, NY	4,933.5	0.0	0.035	1,045.4	560.1	0.003	4,995.0	0.0	0.035	89.0	89.8	0.001
Louisville, KY	2,277.0	0.0	0.017	3,152.6	1,656.0	0.004	2,183.3	0.0	0.017	472.2	436.7	0.000
Hartford, CT	3,657.7	716.7	0.035	669.8	280.5	0.003	4,992.5	0.0	0.040	625.3	607.4	0.000
Jacksonville, FL	3,683.1	831.2	0.087	353.0	218.0	0.008	4,962.2	0.0	0.125	1,154.7	1,377.6	0.005
Richmond, VA	3,456.8	0.0	0.034	394.3	201.2	0.002	3,386.3	0.0	0.033	304.4	308.6	0.001
Oklahoma City, OK	3,712.3	0.0	0.043	424.9	242.3	0.003	3,761.6	0.0	0.043	729.6	783.2	0.001
Birmingham, AL	1,367.0	518.6	0.020	540.1	116.4	0.002	3,498.9	0.0	0.023	916.1	800.6	0.001
Rochester, NY	7,866.7	0.0	0.027	1,018.0	633.8	0.001	7,924.4	0.0	0.027	1,132.8	1,153.4	0.000
Salt Lake City, UT	10,271.2	48.1	0.138	469.3	74.4	0.005	10,289.7	0.0	0.138	1,774.1	1,753.0	0.002

NOTES: The table reports the parameter estimates for each location under the dynamic and myopic specifications of the model. The full description of the data and estimation routine are in appendix A.