

Is the Rent Too High? Aggregate Implications of Local Land-Use Regulation

Devin Bunten*

UCLA

January 11, 2016

Abstract

Highly productive U.S. cities are characterized by high housing prices, low housing stock growth, and restrictive land-use regulations (e.g., San Francisco). While new residents would benefit from housing stock growth due to higher incomes or shorter commutes, existing residents justify strict local land-use regulations on the grounds of congestion and other costs of further development. This paper assesses the welfare implications of these local regulations for income, congestion, and urban sprawl within a general equilibrium model with endogenous regulation. In the model, households choose from locations that vary exogenously by productivity and endogenously according to local externalities of congestion and sharing. Existing residents address these externalities by voting for regulations that limit local housing density. In equilibrium, these regulations bind and house prices compensate for differences across locations. Relative to the planner's optimum, the decentralized model generates spatial misallocation whereby high-productivity locations are settled at too-low densities. The model admits a straightforward calibration based on observed population density, expenditure shares on consumption and local services, and local incomes. Welfare and GDP would be 1.4% and 2.1% higher, respectively, under the planner's allocation. Abolishing zoning regulations entirely would increase GDP by 6%, but lower welfare by 5.9% due to greater congestion.

*Email: devinbunten@ucla.edu. This research was supported by Award Number T32AG033533 from the National Institute on Aging. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institute on Aging or the National Institutes of Health. The title of this paper was inspired by Yglesias (2012). For most recent version, see www.devinbunten.com/jmp.

1 Introduction

Neighborhoods within productive, high-rent regions like the San Francisco Bay Area have very strict controls on housing development and very limited new housing construction. The Bay Area—home to Silicon Valley—is incredibly productive, with a GDP per capita of over \$92,000, higher than all countries except tiny Luxembourg and oil-rich Norway and Qatar. Even for workers in non-tech sectors, wages in the region are higher than elsewhere in the US: median earnings for those with just a high school diploma are 12% higher in the Bay Area. Despite these high wages the region of 6.1 million people permitted just 20,046 new housing units in 2014, consistent with a growth rate of 0.8% annually. Of these, many were located in areas with few existing neighbors, either on former industrial land or previously undeveloped locations at the periphery.¹ The evidence suggests that locally-determined regulatory constraints are a substantial impediment to development (Glaeser and Gyourko, 2002, 2003; Glaeser et al., 2006). Existing residents justify these constraints by appealing to the costs of new development, including increased vehicle traffic and other types of congestion, and claim that they see few, if any, of the benefits from new development.² Instead, these benefits accrue to new potential residents, or to society at large.

While residents may pass local regulations to deal with local externalities of congestion, in aggregate the implications stretch beyond local borders: regions with highly-regulated municipalities experience less-elastic housing supply (Glaeser et al., 2006; Saiz, 2010). Land-use regulations in high-productivity cities constrain the expected supply response to sustained high rents, and these regulations are the result of local endogenous political processes.³ Understanding the welfare effects of zoning regulations involves quantifying both the aggregate housing supply costs as well as the benefits that lead households to pass zoning regulations.

This paper builds a spatial equilibrium model of housing and endogenous regulation to assess the welfare implications of locally-determined zoning. To motivate households' preference for zoning, I introduce local externalities of agglomeration (e.g., residents can share fixed-cost infrastructure or a diversity of restaurants) and congestion (e.g., traffic or limited on-street parking) into a standard Rosen/Roback-style model of heterogeneous locations. In line with the externalities under study, I identify these locations with *neighborhoods*.⁴ Existing residents feel the effects of these externalities from new development, but do not see the benefits

¹In San Francisco, Mission Bay and the Dogpatch are exemplars of redevelopment on former industrial land (http://sf.curbed.com/archives/2014/05/29/mapping_25_projects_transforming_mission_bay_and_dogpatch.php).

²For example, the group Livable Boulder supports a ballot initiative that they argue will “ensure that City levels of service are not diminished by new development”, due to concerns about “huge buildings, blocked views of the mountains, more congestion, proposals to change the unique character of many Boulder neighborhoods” (<http://livableboulder.org/>).

³That regulations are determined locally—at the level of a neighborhood, city district, or municipality—is argued by Hills and Schleicher (2011, 2014); Schleicher (2013); and Monkkonen and Quigley (2008).

⁴Neighborhood productivity differences within and across cities arise from differential access to employment opportunities, transportation networks, or productive amenities.

of higher local output. They thus use local political processes to address these externalities by establishing *zoning laws* that limit development. The trade-off between agglomeration (i.e., sharing of fixed costs and services) and congestion ensures that households prefer a positive but limited number of co-residents. I use the model to identify the wedges between the zoning equilibrium and the planner’s optimal allocation of households to locations. Compared to the utilitarian planner’s allocation, high-productivity locations are developed less intensively, low-productivity locations are developed more intensively, and too many (low-productivity) locations are opened in a decentralized equilibrium.

The key model components—fixed costs, congestion, and endogenous regulation at the local level—are grounded in empirical evidence. The fixed costs of infrastructure and local services can explain the sharp edge of development visible in metropolitan areas, in contrast with the predictions of the standard monocentric city model of a smooth gradient approaching zero density at the fringe.⁵ Local congestion resulting from fixed quantities of street parking, limited space for children’s play, and concerns about sunshine, shadows, and wind permeate anti-development discourse.⁶ Finally, zoning laws in most cities tend to be the result of neighborhood- or municipality-level political processes (Hills and Schleicher, 2014; Monkkonen and Quigley, 2008), rather than broad metropolitan outcomes.⁷

I model a large set of locations analogous to the neighborhoods or municipalities of the national economy. Opening a location to settlement necessitates payment of a fixed cost (e.g., of extending a sewer line) that can be shared more broadly with more residents. Balancing this incentive to share fixed costs, households in a location receive disutility from the addition of further co-residents. These forces of agglomeration and congestion, respectively, constitute the *endogenous amenities* provided by a location. In light of these forces, local residents choose *zoning laws* to restrict the maximum number of housing units that can be built by a competitive construction sector. Distributions of house prices, zoning laws, and location choices are jointly determined in equilibrium. Households are *ex ante* homogeneous and fully mobile, and so house prices adjust to ensure that households are indifferent between all occupied locations. In particular, productive locations have binding zoning restrictions that cause house prices to become significantly higher than the marginal cost of construction, an empirical finding documented in Glaeser and Gyourko (2002, 2003); Glaeser et al. (2005b,a); Cheshire and Hilber (2008) and Koster et al. (2012).

I next analyze the welfare implications of local zoning laws by comparing the equilibrium allocation with the optimal allocation chosen by a constrained planner.⁸ When choosing

⁵For a graphic example, see Figure 4 in the appendix.

⁶A recent example from a zoning debate in Santa Monica is <http://smdp.com/santa-monica-beach-town-dingbat-city/147819>.

⁷Relatedly, Aura and Davidoff (2008) argue that housing and land demand elasticities imply that any particular municipality would be unable to increase supply sufficiently to move down housing demand curves to lower prices. The converse of this finding is that municipal (or neighborhood) governments have limited ability to act as monopolists and raise prices by simply restricting supply.

⁸The planner is *constrained* in that they maximize welfare subject to a spatial equilibrium constraint: identical households must receive identical utility, regardless of location.

the intensity of development within a location, the planner considers the local forces of agglomeration and congestion as well as the value of placing the marginal household in a location of higher or lower productivity. When passing zoning laws, households only consider the local amenity forces and thus restrict development too heavily in productive locations. This underdevelopment induces *sprawl*: the opening of too many unproductive locations, consistent with e.g. Fischel (1999). I further show that the planner’s allocation can be implemented through a zoning reform where residents choose all zoning laws at the national level.

The model enables a straightforward calibration. I calibrate the utility function to the consumption share on non-local goods, excluding housing and local services. For the local sharing parameter, I calibrate to the share of spending on local services. The congestion cost curvature is chosen to match the average population density of urban census tracts.⁹ For location productivity, I use Census tract income data, adjusted for average household observable characteristics (Hsieh and Moretti, 2015). I calibrate construction-sector productivity to match the ratio of price to construction costs in high-price locations like New York and San Francisco where the effects of zoning laws have been well-studied (cf. Glaeser et al., 2005b).

The model provides a unified framework for addressing the full array of costs and benefits to reforming local zoning policy by implementing the optimal allocation. Avent (2011) and Hsieh and Moretti (2015) argue that these local growth patterns have been quite costly for national growth, while incumbent residents argue that new development would itself be costly. My main exercise is to quantify these costs while accounting for the changes to endogenous amenities resulting from the planner’s allocation. The planner’s allocation raises welfare and GDP by 1.4% and 2.1%, respectively. The more intensive development of productive locations necessary to increase GDP would also increase the congestion disamenity received by the residents of these regions. The median resident experiences an increase in density of 3.6%; increases as great of 10–15% are experienced by the most productive locations. More intensive development reduces the total number of locations opened to development by 3%. The planner’s allocation thus features less *sprawl*, with less fixed-cost expenditure and commensurately higher consumption.

A second exercise offers a quantitative assessment of an alternative policy: zoning abolition, wherein developers are free to build housing without constraint.¹⁰ This policy results in over-building in high-productivity areas as construction firms do not internalize congestion effects. Model GDP increases by 6%, but increased congestion causes welfare to decline by 5.9%. The extra productivity comes arises because productive regions like San Francisco see

⁹While the parameter governs the shape of the congestion cost function, in equilibrium it has a first-order effect on the optimal zoning laws and thus on average density. The approach is similar to the tradition of using wage shares to calibrate labor and capital exponents in a production function: in equilibrium, these shape parameters have first-order effects on levels.

¹⁰This is the option implicitly studied by Hsieh and Moretti, who find that output would increase by 13.5%.

population inflows of 50% or more. The policy implication is that zoning should be relaxed in productive locations, but not abolished.

Treating zoning as endogenous enables the inference of welfare losses from zoning in the absence of detailed data on the wedges between housing prices and marginal costs of construction at the location level. Were such data available, the wedges could be treated as exogenous and the welfare costs of zoning would be the costs of moving the wedges from their observed to their optimal level. In principle, observed density could be used to proxy for underlying regulations. The relationship between productivity and zoned density, however, is confounded in the data by factors like exogenous amenities—which make land more expensive and increase density—and long-lived investments, both in building stock and in infrastructure such as subway networks. Instead of attempting to account for these myriad forces, I model zoning as an equilibrium object, infer the unobserved wedge from the model, and calculate the loss in welfare accordingly. The key qualitative prediction of this model is that, in high-productivity areas, density does not rise fast enough relative to productivity.

Endogenous zoning has a further benefits: it offers predictions on the likely outcomes of different housing policy and zoning reform. First, policies which ignore the endogeneity of zoning may be unable to meet their stated goals. The federal government provides large housing subsidies throughout the income distribution, and increased may appear to offer an outlet for reducing the burden of high costs in expensive cities. Within the specified model, a housing subsidy would increase household willingness to pay in expensive areas. However, endogenous zoning driven by existing residents will not respond to these subsidies by creating more supply, and so the policy will not have the intended effect of making housing more affordable, nor will it stimulate additional supply in zoning-constrained locations.

Second, the model provides guidance for the outcomes of direct zoning reforms. Can a single location allow more development—increasing housing supply—and hope to bring down prices? Consider an infinitesimal neighborhood in the model that relaxes its zoning restriction and allows more development. The outside option—the value of locating elsewhere—will be unaffected, as a single location is insufficient to move the price gradient. Instead, house prices in the neighborhood will decline only to the extent that additional development makes congestion worse. This echoes the complaints of homeowners who fear new development will hurt their house values.¹¹

Finally, could a successful *upzoning*, where a set of productive locations increase the zoned capacity or even abolish zoning restrictions, harm landowners in less productive locales? Suppose that an entire productive region—perhaps a city, or a state—chooses to allow more intensive development. Many new households will flow in, abandoning the less productive locations they once occupied. For the least-productive locations, the initial population outflows will make the fixed costs too burdensome, and all households will choose to leave. If

¹¹For example, This claim is offered by an advocate for a moratorium on development in the Mission District of San Francisco here: <https://medium.com/@danancona/putting-market-fundamentalism-on-hold-432ecf1aab3c>.

profits from development of a location are tied to that location, perhaps via homeownership, then residents of less productive locations may lose out from this policy reform.

1.1 Related Literature

This paper draws on work that has established a strong relationship between zoning and the elasticity of housing supply (Fischel, 1999; Mayer and Somerville, 2000; Glaeser and Gyourko, 2002; Glaeser et al., 2005b,a; Saiz, 2010; Koster et al., 2012; Turner et al., 2014).¹² This literature has shown that zoning laws restrict the supply response to high house prices, and that cities that do zone restrictively have systematically higher housing prices. Following this literature, zoning laws in this paper shape local amenities (cf. Turner et al., 2014) and regional housing supply (cf. Fischel, 1999; Mayer and Somerville, 2000).

Several authors in diverse contexts have modeled zoning as an endogenous outcome of political processes (Hamilton, 1975; Fischel, 1987; Monkkonen and Quigley, 2008; Fischel, 2009; Hilber and Robert-Nicoud, 2013; Ortalo-Magné and Prat, 2014; Hills and Schleicher, 2011; Schleicher, 2013; Hills and Schleicher, 2014; Fischel, 2015). I implement the findings of (e.g.) Fischel (2009) and Hills and Schleicher (2011) that zoning laws are actively determined by highly-engaged utility-maximizing households at the local level—municipalities, neighborhoods, or even direct neighbors.¹³ In some contexts, zoning enables households to implement the planner’s allocation, usually by resolving problems of free-riding on local expenditure. I complement this literature by embedding a set of heterogeneous locations into a general equilibrium model, so that zoning is locally optimal but has aggregate external costs. Incorporating the endogeneity of local zoning enables a welfare calculation in the absence of detailed local data and also reflects the consensus of the literature (cf. Fischel, 2009, Hills and Schleicher, 2011, and Hills and Schleicher, 2014).

My work builds on Hsieh and Moretti (2015), who also link strict zoning and changes in aggregate output. They study the productivity effects of wage dispersion across metropolitan areas in a Rosen-Roback model, and attribute the increase in this dispersion (and resulting loss in output) to a decrease in the elasticity of housing supply. From this basis, I introduce local externalities and calculate welfare losses from zoning relative to the planner’s optimum, in addition to output losses from spatial misallocation. My paper differs in terms of the economic interpretation of zoning: here, it is an endogenous limit on neighborhood development instead of a change in the elasticity of metropolitan housing supply. As noted, this endogeneity serves two purposes: it enables me to infer the costs of welfare in the absence of location-specific data on marginal costs and housing prices and it allows me to study the likely local outcomes of different policy reforms.

This paper is related to others that study the conditions under which zoning laws and

¹²Gyourko and Molloy (2014) present an in-depth review of the economics of housing regulation.

¹³Lawsuits seeking to halt development in the name of various ills are increasingly common (Ganong and Shoag, 2013).

other local government interventions may be theoretically optimal (Stull, 1974; Hamilton, 1975; Henderson, 1991; Hochman et al., 1995; Rossi-Hansberg, 2004; Calabrese et al., 2007; Chen and Lai, 2008; Allen et al., 2015). These papers focus on a variety of externalities to motivate zoning, including nuisance, production, and free-riding interactions that I abstract from. Of course, the optimal zoning scheme depends upon the nature of the externalities under consideration. Similarly to some of these papers, I characterize the zoning regime consistent with maximizing welfare given the externalities studied. This paper also addresses the optimal distribution of population between communities who share local public goods (Flatters et al., 1974; Arnott and Stiglitz, 1979; Hochman et al., 1995). Like Hochman et al. (1995), I find that jurisdictions that are well-positioned to deal with one externality may not be optimal in the presence of other inter-location effects.

More broadly, this paper relates to others that study the spatial determinants of aggregate productivity and welfare (Albouy, 2012; Desmet and Rossi-Hansberg, 2013; Allen and Arkolakis, 2014; Behrens et al., 2014; Eeckhout and Guner, 2015; Fajgelbaum et al., 2015; Hsieh and Moretti, 2015). I also account for the joint spatial distributions of household location choices, incomes, and housing prices while incorporating the effects of zoning regulations. As such, it relates to Glaeser et al. (2005a), Van Nieuwerburgh and Weill (2010), Gyourko et al. (2013), and Diamond (2015), among others.

The rest of the paper is organized as follows. Section 2 presents the model and the analytic results. Section 3 calibrates a quantitative version of the model to match features of the data. Section 4 quantitatively analyzes the welfare costs of local zoning and the no-zoning counterfactual relative to the planner’s optimal allocation. Section 5 concludes.

2 Model

This section describes the model. I begin by outlining the environment. I then study the determination of equilibrium in three steps. First, I study the *local equilibrium* that results from the actions of households and firms, who take as given the local zoning rules as well as the endogenous outside option and level of profit. Second, I study the *local zoning choice* of initial residents. The zoning law restricts the level of housing production available to construction firms. Third, I study the economy-wide *general spatial equilibrium* in which the aggregate variables are determined: the endogenous outside-option level of utility and aggregate profits of construction firms.

Environment The spatial environment consists of a mass of locations of measure 1 with productivity indexed by $x \in [0, 1]$. These locations represent the set of potential neighborhoods in the entire economy. The neighborhood with $x = 1$ is the most productive in the country, e.g. downtown San Francisco. Locations with x less than but close to 1 include less-productive locations in the same area—for instance, Oakland, a train ride from downtown

San Francisco. A similar x could also correspond to the best locations in less-productive regions, like downtown Denver. Locations with x close to zero would be at the far fringes of metro areas: either a short commute to an unproductive job or a very long commute to a better job.

Production takes place at each location according to a linear production function. Productivity is given $y(x)$ with $y(0) = \underline{y}$, $y'(x) > 0$, and $y(1) = \bar{y}$. Each location begins empty and available for settlement by a measure 1 of households, whose preferences are described below. Opening a location to urban settlement requires a fixed cost F . The set of locations that are open in equilibrium is an endogenous outcome.¹⁴

I now turn to the description of agents and their maximization problems.

2.1 Local Equilibrium

2.1.1 Household and Construction Firm Problems

The economy consists of three types of agents. First, there is a measure 1 of households who choose location and consumption to maximize utility, taking as given location-specific rents, population, and the endogenous outside option (i.e., the maximum attainable utility from making the optimal choice of location).

Second, there is a representative construction firm that chooses the housing production level to maximize at each location profits, taking as given rents and the local zoning law. I represent a zoning law as a maximum allowable number of housing units.¹⁵

Third, some locations begin with a measure 0 of *initial residents*. The initial residents choose local zoning laws and consumption to maximize utility, while taking as given a pre-fixed initial rent.¹⁶ Initial residents have rational expectations about the equilibrium choices that households and firms will make given the zoning law. In locations with no initial residents, construction firms maximize profits without constraint.

¹⁴As a technical matter, it is possible that fixed and construction costs are sufficiently large that there is no feasible distribution of households under which total output is sufficient to cover these costs. In that case the economy can be thought of as *agricultural*: no fixed costs are paid, and each location is occupied by a solitary household who produces $y_A < \underline{y}$. In the urban equilibrium, fewer locations will be opened. A final case is that the endogenous outside option of the urban economy is equal to the utility earned from agricultural, in which case some locations would remain agricultural. Restricting y_A to be sufficiently small effectively rules out this equilibrium.

¹⁵As each household consumes a single housing unit, the choice of zoning law restricts both the number of housing units and the number of households.

¹⁶The initial rent can be thought of as the fixed mortgage payment of an existing homeowner, which does not fluctuate based on local rental conditions. It can also be thought of directly as rent control, where the price cannot be raised on existing tenants. Rent control is a common policy tool in expensive rental markets like San Francisco.

I proceed by defining the problems of the household and the construction firm and then the neighborhood equilibrium given zoning laws. Then I define the problem of the initial residents and their equilibrium choices of zoning law.

Household Problem A household in location x has preferences summarized by the utility function

$$u(c(x)) - v(n(x)).$$

The function $u(\cdot)$ represents preferences over consumption, $c(x)$. The function $v(\cdot)$ represents a congestion externality that depends on the neighborhood population, $n(x)$. Utility from consumption is twice continuously differentiable and strictly concave: $u'(c) > 0$ and $u''(c) < 0$. The congestion externality $v(\cdot)$ is bounded below, increasing, twice continuously differentiable, and convex: $v'(n) > 0$ and $v''(n) \geq 0$. Without loss of generality, I assume that $v(0) = 0$.

The household faces budget constraint

$$c(x) = y(x) + \Pi - \frac{F}{n(x)} - p(x),$$

where $y(x)$ and $p(x)$ represent the location-specific income and house price. Consumption is given $c(x)$. Households own a diversified portfolio of construction firms in all locations, and Π is the location-independent transfer of construction sector profits from the national economy. The fixed cost F is shared equally among all residents.

Households take as given local prices $p(x)$, total profits Π , and their endogenous outside option \bar{u} . Households are freely mobile and choose location x and consumption $c(x)$ to maximize utility subject to the budget constraint.

Construction Firm Problem Construction firms choose the development intensity $n^f(x)$ to maximize profits at each location subject to a local zoning constraint $\bar{n}(x)$. There is no pre-existing housing stock in any location. Taking the price $p(x)$ as given, the firm problem in location x is

$$\max_{n^f(x)} \Pi(x) = p(x)n^f(x) - \left(\frac{n^f(x)}{Z}\right)^2$$

subject to zoning constraint

$$n^f(x) \leq \bar{n}(x).$$

It is natural to think that construction costs within a location are convex due to the costs associated with building taller and denser structures. Construction firms earn positive profits from locations with positive construction. Profits are aggregated across all locations and redistributed equally to all households.¹⁷

¹⁷The cost function is equivalent to a housing production function that is Cobb-Douglas in building materials and land, with a coefficient of 1/2 on each term. Under this alternative production function,

2.1.2 Local Spatial Equilibrium

Consider a fixed $\bar{n}(x)$, a fixed \bar{u} , and a fixed Π . I define a *local spatial equilibrium* for location x to be a house price $p^*(x)$ and local population $n^*(x)$ such that households and firms solve their respective problems with households consuming their budget, the local housing market clears, and spatial equilibrium must hold. That is:

$$(1) \quad u \left(y(x) + \Pi - \frac{F}{n^*(x)} - p^*(x) \right) - v(n^*(x)) \leq \bar{u},$$

with equality if $n^*(x) > 0$. Housing market clearing implies $n^f(x) = n^*(x)$. Using this, the firm's problem yields a complementary slackness condition which states that either the zoning constraint binds or price is equal to the marginal cost of construction:

$$(2) \quad (p^*(x) - 2n^*(x)/Z^2) [\bar{n}(x) - n^*(x)] = 0.$$

The firm must also abide by the zoning constraint:

$$(3) \quad n^*(x) \leq \bar{n}(x).$$

A *local spatial equilibrium* in location x is a pair $\{n^*(x), p^*(x)\}$ satisfying Equations (1), (2), and (3).

Housing Equilibrium Selection The existence of a fixed cost creates complementarities in household location decisions: if some households go to a location, the fixed cost can be shared more broadly and the location becomes more attractive to other households. This complementarity implies that there can be multiple equilibria. Finally, the congestion externality implies that eventually an additional household will lower the utility of existing households as the congestion costs outweigh the sharing benefits. I show that there are up to three pairs of equilibrium population $n^*(x)$ and price $p^*(x)$ that may satisfy the definition of an equilibrium. Figure 1 characterizes graphically the set of potential equilibria. The *optimal location* curve, labeled $OL(n, x)$, shows the set of (n, p) points consistent with the household's maximizing choices, and is defined as follows:

$$(4) \quad OL(n, x) = \begin{cases} y(x) + \Pi - \frac{F}{n} - u^{-1}(v(n) + \bar{u}) & \text{if } n > 0 \\ [0, \infty] & \text{if } n = 0. \end{cases}$$

The curve labeled $MC(n)$ corresponds to the marginal cost of construction and is independent of location:

$$(5) \quad MC(n) = 2n/Z^2.$$

the profit redistribution outlined here corresponds to an assumption that each household owns a diversified portfolio of land. This share is within the range of estimates provided by Albouy and Ehrlich (2012).

Proposition 1. *The portion of the $OL(n, x)$ curve with $n(x) > 0$ is strictly concave, hump-shaped, and intersects the $MC(n)$ curve twice, not at all, or once at a point of tangency.*

Proof. See appendix. □




Figure 1: Marginal cost ($MC(n)$) and indifference ($OL(n, x)$) curves. Note that the $OL(n, x)$ curve includes the entire vertical axis above 0. The set of points potentially consistent with equilibrium, for which $OL(n, x) > MC(n)$, is the closed interval $(n_L(x), n_H(x))$ and the point 0.

If the zoning constraint lies to the right of $n_H(x)$ —the upper intersection of $MC(n)$ and $OL(n, x)$ —then both $n_H(x)$ and $n_L(x)$ are consistent with equilibrium. If the zoning constraint lies within the interval $(n_L(x), n_H(x))$ then both $n_L(x)$ —the lower intersection—and $\bar{n}(x)$ are consistent with equilibrium. If the zoning constraint lies to the left of $n_L(x)$, then the origin is the sole equilibrium. There is always an equilibrium at the origin: households and firms may expect a $p^*(x) = 0$ and $n^*(x) = 0$ and see these expectations fulfilled.

If there are multiple equilibria, I select the one with the largest population.¹⁸ This equilibrium selection has two virtues. First, it abstracts from coordination failures where both

¹⁸Note that this supremum is taken over a closed and bounded set. By continuity of the utility and construction cost functions, the supremum of the set of points consistent with equilibrium is, itself, a member of this set.

households and firms expect the location to be empty, and so it remains empty. Second, it is *stable* in the sense of economic geography models, i.e. a small deviation of population does not induce population movements that lead the location to a different equilibrium (cf. Krugman, 1991). This rules out $n_L(x)$ as an equilibrium population as it would be unstable.¹⁹ Concretely, both households and construction firms expect that all locations will be settled at the highest-population equilibrium.

The equilibrium population $n^*(x)$ and price $p^*(x)$ consistent with this selection criterion for a given zoning law $\bar{n}(x)$ are

$$(6) \quad \{n^*(x), p^*(x)\} = \begin{cases} \{0, 0\} & \text{if } MC(n) > OL(n, x) \forall n \in (0, \bar{n}(x)] \\ \{\bar{n}(x), OL(n, x)\} & \text{if } \bar{n}(x) \in [n_L(x), n_H(x)] \\ \{n_H(x), MC(n)\} & \text{if } \bar{n}(x) > n_H(x). \end{cases}$$

In the first case, there is no intensity of development $n(x)$ that abides by the zoning constraint and that delivers utility of at least \bar{u} . This could happen because $\bar{n}(x) < n_L(x)$ or because $OL(n, x) < MC(n)$ for every $n(x) > 0$. In the second case, the zoning constraint binds and the equilibrium price is consistent with household spatial equilibrium. In the third case, the zoning constraint does not bind and the equilibrium price is consistent with both household spatial equilibrium and (unconstrained) profit maximization.

2.1.3 Locally Optimal Zoning Choice

This section describes the determination of zoning laws within a location.

I assume that some locations have a measure 0 of *initial residents* who choose the zoning regulation, while some locations have no initial resident and are thus unregulated. The set inhabited locations will be endogenized later. Formally, they can choose the zoning constraint $\bar{n}(x)$; as Equation (6) makes clear, the choice of zoning law affects the local equilibrium.

For the subset of locations with no initial residents, the zoning law is effectively infinite: $\bar{n}(x) = \infty$. The equilibrium population in such locations is given by Equation (6): it will be $\{n_H(x), MC(n)\}$ if the $OL(n, x)$ and $MC(n)$ curves intersect and $\{0, 0\}$ if the $OL(n, x)$ curve is strictly lower than the $MC(n)$ curve for all positive populations.

Locations with initial residents The measure 0 of initial residents have identical preferences and productivity as the households described above, but they differ in two respects. First, the initial residents of location x face a fixed rent $p_0(x)$. I take this rent as given for now, but will endogenize it later. Second, they choose the zoning constraint $\bar{n}(x)$ that limits the maximum level of development within their location. The initial residents have rational

¹⁹See appendix for details on equilibrium stability.

expectations about the consequences of their zoning choice. Namely, they anticipate that a zoning constraint $\bar{n}(x)$ will induce the local spatial equilibrium of Equation (6). Like the households described above, the initial residents take as given the the endogenous outside option \bar{u} and the level aggregate profits Π .

Let $\theta(\bar{n}(x); \bar{u}, \Pi) = n^*(x)$ denote the population $n^*(x)$ from Equation (6) given a zoning constraint $\bar{n}(x)$ and aggregate variables \bar{u} and Π . The initial residents solve the following maximization problem:

$$\max_{\{c_0(x), \bar{n}(x)\}} u(c_0(x)) - v(\theta(\bar{n}(x); \bar{u}, \Pi))$$

subject to

$$c_0(x) = y(x) + \Pi - \frac{F}{\theta(\bar{n}(x); \bar{u}, \Pi)} - p_0(x)$$

We can rewrite the maximization problem in a more convenient form. Initial residents in location x act as if they were directly choosing the local population given \bar{u} and Π , subject to a household participation constraint.

Proposition 2. *The locally optimal zoning constraint solves:*

$$(7) \quad \max_{\{c_0(x), \bar{n}(x)\}} u(c_0(x)) - v(\bar{n}(x))$$

subject to the budget constraint

$$(8) \quad p_0(x) + c_0(x) + \frac{F}{\bar{n}(x)} = y(x) + \Pi$$

and the participation constraint

$$(9) \quad u\left(y(x) + \Pi - \frac{F}{\bar{n}(x)} - 2\bar{n}(x)/Z^2\right) - v(\bar{n}(x)) \geq \bar{u}.$$

Proof. See appendix. □

Note that the participation constraint is a function of the marginal cost of housing, rather than the equilibrium price. The participation constraint is satisfied if and only if $\bar{n}(x) \in [n_L(x), n_H(x)]$ as defined above. If the participation constraint is non-binding, Equation (6) states that equilibrium population $n^*(x)$ will equal the initial resident choice $\bar{n}(x)$. The logic of the constraint is as follows: for $\bar{n}(x) < n_L(x)$, the equilibrium population will be 0. However, the initial residents prefer to share the fixed cost F with a positive population, and so they will not choose $\bar{n}(x) < n_L(x)$. If they choose $\bar{n}(x) > n_H(x)$, the equilibrium population will be $n_H(x)$, and so limiting the initial resident choice to being below $n_H(x)$ does not restrict their potential payoffs.

In general, the participation constraint may or may not bind. If the participation constraint does not bind, then the zoning constraint will bind in the local equilibrium: the equilibrium house price is above the firm's marginal cost. If the participation constraint does bind, then the zoning constraint does not bind: the house price is equal to marginal cost.

When the participation constraint does not bind,²⁰ the initial resident choice of zoning choice is the $\bar{n}(x)$ that solves

$$(10) \quad u' \left(y(x) + \Pi - \frac{F}{\bar{n}(x)} - p_0(x) \right) \frac{F}{\bar{n}(x)^2} - v'(\bar{n}(x)) = 0.$$

The first term is the marginal benefit of sharing the fixed cost F more broadly. The second term is the marginal cost of congestion. As u is concave and v convex, the expression is strictly decreasing and the $\bar{n}(x)$ that solves the equation is unique.²¹

When the participation constraint binds, the initial resident's optimal zoning choice is by definition not in the interior of the set of points consistent with the participation constraint. By concavity of the initial resident problem, the zoning choice in these cases is given by the endpoint of the set that offers greater utility to the initial residents:

$$(11) \quad \arg \max_{\bar{n}(x) \in \{n_L(x), n_H(x)\}} u \left(y(x) + \Pi - \frac{F}{\bar{n}(x)} - p_0(x) \right) - v(\bar{n}(x)).$$

2.2 General Spatial Equilibrium

Define a *stable general spatial equilibrium* in this economy to be an endogenous outside option \bar{u} , a level of profit Π , a set of open locations \mathcal{X} , local populations $n^*(x)$, local prices $p^*(x)$, and zoning laws $\bar{n}(x)$ such that the zoning choices $\bar{n}(x)$ and local equilibrium outcomes $\{p^*(x), n^*(x)\}$ solve household, firm, and initial resident problems and are consistent with the population constraint

$$(12) \quad \int_0^1 n^*(x) dx = 1,$$

with the definition of profits

$$(13) \quad \Pi = \int_0^1 \left[p^*(x)n^*(x) - \left(\frac{n^*(x)}{Z} \right)^2 \right] dx,$$

²⁰The constraint states that the utility delivered by the location *with house price equal to marginal cost* is greater or equal to \bar{u} . If the constraint is non-binding then Equation (6) states that the equilibrium house price will be greater than marginal cost.

²¹As there is a measure 0 of initial residents, an inflow of $\bar{n}(x)$ households ensures that the new households will be the majority of the community. This raises the question of whether they would seek *ex post* to hold a new vote and modify the zoning law. However, the new residents would not have a strong preference to modify the law. Given an outside option \bar{u} , households expect that rents will adjust to make them indifferent across locations regardless of their vote. As such, households do not strictly prefer any alternative zoning law for their location. The pre-fixed rent of the initial residents eliminates this feedback mechanism.

with the definition of \mathcal{X}

$$(14) \quad x \in \mathcal{X} \iff n^*(x) > 0,$$

and with the stability condition

$$(15) \quad p^*(x) = p_0(x) \quad \forall x \in [0, 1].$$

The stability condition imposes restrictions on the prices faced by initial residents. This natural restriction is informed by a notion of dynamic stability and enables the static model to approximate the steady-state of a corresponding dynamic overlapping generations model, wherein a subset of agents inherit a fixed rent from the previous period's equilibrium.

2.2.1 Characterization of General Equilibrium

I proceed by characterizing the triplet $\{\bar{u}, \Pi, \mathcal{X}\}$ as well as optimal zoning $\bar{n}(x)$, local population $n^*(x)$, and price gradient $p^*(x)$ consistent with general equilibrium. Focusing on the *stable equilibrium* described above, I consider the natural restriction that $p_0(x) = p(x)$. Two preliminary results will be useful in characterizing the general equilibrium.

Proposition 3. *The set of occupied locations \mathcal{X} is $[\underline{x}, 1]$ for some threshold location \underline{x} . For the threshold location \underline{x} , the $OL(n, \underline{x})$ and $MC(n)$ curves are tangent at a unique level of population that depends on \bar{u} .*

Proof. See appendix. □

This tangency condition proves useful for pinning down the price and population gradients. The intuition for tangency is as follows: from Proposition (1), the two curves cross twice, once with tangency, or not at all. For a given location x , if they do not cross the location must be unoccupied: no level of population can deliver utility \bar{u} . If they do cross or meet with tangency, then $n^*(x)$ must be strictly positive. Otherwise, there would be no initial residents, no zoning constraint, and developers would find it optimal to build (and households to settle) to the local equilibrium population of $n_H(x)$. Thus x must be in \mathcal{X} for all locations where the $OL(n, x)$ curve meets the $MC(n)$.

The tangency of $OL(n, \underline{x})$ and $MC(n)$ at the threshold location \underline{x} also implies equality between $OL(n, \underline{x})$ and $MC(n)$. Thus Proposition (1) offers two additional conditions that will be used to pin down the price and population gradients. First, tangency implies

$$(16) \quad \frac{F}{n^*(\underline{x})^2} - [u^{-1}]'(v(n^*(\underline{x})) + \bar{u}) \times v'(n^*(\underline{x})) = 2/Z^2.$$

Second, equality implies

$$(17) \quad p(\underline{x}) = 2n(\underline{x})/Z^2.$$

The first condition pins down the threshold population consistent as a function of \bar{u} . The second condition pins down the price at the threshold as a function of this population.

Lemma 1. Define \bar{n} to be the population that maximizes the $OL(n, x)$ curve. Given outside option \bar{u} , the level of population \bar{n} is unique and independent of location.

Proof. See Appendix. □

Lemma 2. Define location $\hat{x}(\bar{u}, \Pi)$ as follows:

$$\hat{x}(\bar{u}, \Pi) = \begin{cases} x \text{ s.t. } OL(\mathcal{N}(\bar{u}), \hat{x}) = MC(\mathcal{N}(\bar{u})) & \text{if } OL(\mathcal{N}(\bar{u}), 1) \geq MC(\mathcal{N}(\bar{u})) \\ 1 & \text{otherwise.} \end{cases}$$

Then $\hat{x}(\bar{u}, \Pi)$ is unique and in the range $[\underline{x}, 1]$.

Proof. See appendix. □

Proposition 4. In a stable general equilibrium, the following characterizes the local equilibrium populations and optimal zoning. For all locations $x > \hat{x}$, $n^*(x) = \bar{n}(x) = \mathcal{N}(\bar{u})$. For all locations $x \in [\underline{x}, \hat{x}]$, the population $n^*(x)$ and optimal zoning $\bar{n}(x)$ are given by the upper intersection of the $OL(n, x)$ and $MC(n)$ curves. For all $x \geq \underline{x}$, $p^*(x)$ is given by the value of the $OL(n, x)$ curve evaluated at $n^*(x)$.

Proof. See appendix. □

The $OL(n, x)$ curve reflects the willingness to pay for a location as a function of its population: if a mix of congestion and sharing offers a higher level of utility, households are willing to pay more. As $p_0(x) = p(x)$ in the stable equilibrium, initial resident utility is maximized where the willingness to pay is maximized: at the peak of the $OL(n, x)$ curve. As the productivity $y(x)$ only shifts the $OL(n, x)$ curve vertically, the preferred zoning choice is independent of location. Hence for locations with $x > \hat{x}$, zoning laws bind in equilibrium at the population $\mathcal{N}(\bar{u})$.

For locations $x \leq \hat{x}$, the level of population $\mathcal{N}(\bar{u})$ preferred by the initial residents is inconsistent with local equilibrium: the price that would induce households to settle at population $\mathcal{N}(\bar{u})$ is too low to induce construction firms to build $\mathcal{N}(\bar{u})$ units of housing. Vice versa, any price high enough to induce firms to build $\mathcal{N}(\bar{u})$ will be higher than households are willing to pay, given local wages $y(x)$. In this case, the equilibrium population adjusts downwards: the house price falls, marginal costs fall, and so long as $x \geq \underline{x}$, the location settles at a smaller equilibrium population. Because the initial resident problem is concave, the participation constraint binds and the local equilibrium population (and optimal zoning law) is the upper intersection of the $OL(n, x)$ and $MC(n)$ curves. This case is shown in Figure (2).

Characterization of aggregate variables: \bar{u} and Π . As noted, the equilibrium levels of population $n^*(x)$ and prices $p^*(x)$ described in Proposition (4) are consistent with the local equilibrium, optimal zoning, and stable equilibrium conditions for a given pair $\{\bar{u}, \Pi\}$. The general equilibrium \bar{u} and Π are the pair for which the local equilibria $n^*(x)$ and $p^*(x)$ and




Figure 2: Optimal location $OL(n, x)$ curves for threshold occupied location \underline{x} , \hat{x} , and other locations with the marginal cost $MC(n)$ curve. The lowest curve, $OL(n, \underline{x})$, is the threshold location and its equilibrium population will be given by the point of tangency with the $MC(n)$ curve. For the next curve, the equilibrium population is given by the intersection on the right. For the top three curves, the initial resident zoning choice and equilibrium population is given by $\mathcal{N}(\bar{u})$. Note that all of the $OL(n, x)$ curves reach their peak at $\mathcal{N}(\bar{u})$; the equilibrium population at each location is as close to this as possible subject to ensuring that $OL(n, x) \geq MC(n)$.

locations \underline{x} and \hat{x} described above are consistent with the population constraint and profit definition.

The threshold conditions from Proposition (3) implicitly define the threshold location \underline{x} as a function of \bar{u} and Π : $\underline{x}(\bar{u}, \Pi)$. The profit definitions states

$$\Pi = \int_{\mathcal{X}} \left(p^*(x)n^*(x) - \left(\frac{n^*(x)}{Z} \right)^2 \right) dx.$$

Substituting the equilibrium conditions, this becomes

$$(18) \quad \Pi = \int_{\underline{x}(\bar{u}, \Pi)}^{\hat{x}(\bar{u}, \Pi)} \left(\frac{n_H(x)}{Z} \right)^2 dx + \left(\int_{\hat{x}(\bar{u}, \Pi)}^1 y(x) dx + \Pi - \frac{F}{\mathcal{N}(\bar{u})} - u^{-1}(v(\mathcal{N}(\bar{u})) + \bar{u}) \right) \mathcal{N}(\bar{u}) - \left(\frac{\mathcal{N}(\bar{u})}{Z} \right)^2.$$

From above, $n_H(x)$ is the upper intersection of the $OL(n, x)$ and $MC(n)$ curves and is pinned

down as a function of Π and \bar{u} . From the population constraint,

$$(19) \quad 1 = \int_{\underline{x}(\bar{u}, \Pi)}^{\hat{x}(\bar{u}, \Pi)} n_H(x) dx + \int_{\hat{x}(\bar{u}, \Pi)}^1 \mathcal{N}(\bar{u}) dx.$$

The characterization of a stable general equilibrium is completed by an outside option \bar{u} and a profit Π that solve Equations (18) and (19). I do not yet have a formal existence proof. However, given functional forms for $u(c)$ and $v(n)$, it is straightforward to solve the equations and characterize the equilibrium numerically.

2.2.2 Qualitative Predictions

The model can be used to generate a set of empirical predictions regarding density, house prices, and marginal costs. Regarding density, the model predicts that the equilibrium density will be uncorrelated with local productivity for locations where zoning is binding. For all tracts, it should be positive. For the set of urban census tracts, the correlation between income and density is -0.02 . Adjusting for skill differences, in a process to be detailed below, moves the correlation to 0.02 . After adjusting for county-level averages—which may result from common investments in infrastructure, or exogenous amenities—the correlation moves to 0.07 . These findings are consistent with the model.

The model predicts that house prices will perfectly offset productivity differences in locations with binding zoning. Of course, housing prices respond to many factors not included in the model, so the data is unlikely to show a perfect fit. To test this prediction, I compute a residual house price by regressing tract-level housing costs on the same observable characteristics that I use to adjust income. Then, I correlate this measure with tract-level income, adjusted for observable characteristics. The correlation is 0.42 for owner-occupied housing costs, and 0.38 for rental costs. These correlations are not inconsistent with the model. In addition to the aforementioned exogenous amenities, unmodeled forces that may affect this relationship include the quality of the housing stock, length of ownership for homeowners, and housing subsidies or rent-control laws.

Finally, the model predicts that marginal costs will be constant across locations with binding zoning. Together, these last two predictions imply that house prices in the most expensive locations will be well above marginal costs of construction. This prediction is well-supported by the (cf. Glaeser et al., 2005a).

2.3 Optimality: The Constrained Planner’s Problem

To provide a welfare benchmark with which to contrast the zoning equilibrium this section describes the problem of a constrained planner. The planner chooses the set \mathcal{X} of locations to open, allocates population $n(x)$ among these locations, and allocates consumption $c(x)$

to households in order to maximize welfare. The planner is free to transfer output across locations. The planner faces standard population and aggregate resource constraints as well as a spatial equilibrium constraint. This final constraint restricts the planner to choosing allocations that deliver identical utility to each household regardless of location; in this sense, the planner is *constrained*. The spatial equilibrium condition implies that the welfare criterion to be maximized is simply \bar{u} . The set of locations to open can be simplified to the choice of a threshold location \underline{x} .²²

The planner's problem is thus given by the following:

$$\max_{\{\bar{u}, c(x), n(x), \underline{x}\}} \bar{u}$$

subject to the spatial equilibrium constraint

$$u(c(x)) - v(n(x)) = \bar{u} \quad \forall x \in [\underline{x}, 1],$$

the population constraint

$$\int_{x=\underline{x}}^1 n(x) dx = 1,$$

and the aggregate resource constraint

$$\int_{x=\underline{x}}^1 \left[(y(x) - c(x)) n(x) - F - \left(\frac{n(x)}{Z} \right)^2 \right] dx = 0.$$

The population $n(x)$ can be thought of as the *intensive* margin of development while \underline{x} is the *extensive* margin of development. Note that each opened location is subject to a spatial equilibrium condition. Given the choice of intensive margin $n(x)$, the location-specific spatial equilibrium constraints ensure the planner engages in transfers of output such that household consumption offsets the level of congestion and each household receives utility \bar{u} .

The Lagrangian associated with the problem is:

$$(20) \quad \mathcal{L} = \bar{u} + \int_{x=\underline{x}}^1 \lambda(x) [u(c(x)) - v(n(x)) - \bar{u}] dx + \mu \left[1 - \int_{x=\underline{x}}^1 n(x) dx \right] \\ + \Lambda \int_{x=\underline{x}}^1 \left[(y(x) - c(x)) n(x) - F - \left(\frac{n(x)}{Z} \right)^2 \right] dx.$$

Here, $\lambda(x)$ is the Lagrange multiplier on the spatial equilibrium constraint at location x , μ is the multiplier on the population constraint and Λ is the multiplier on the resource constraint. Taking first order conditions for $n(x)$ and $c(x)$ and rearranging, the planner

²²The planner will never choose to open location x_1 if location $x_2 > x_1$ has not already been opened. Hence choosing the set \mathcal{X} from within the set of locations $[0, 1]$ amounts choosing the lowest-productivity location to open: \underline{x} .

weighs the following objects against one another when choosing the intensive margin of development $n(x)$:

$$(21) \quad \Lambda [y(x) - c(x) - 2n(x)/Z^2] - \Lambda n(x) \frac{v'(n(x))}{u'(c(x))} - \mu = 0$$

The first term is the shadow resource value of the marginal household, the second term is the shadow cost of congestion, in terms of resources, and the third term is the shadow cost of having one fewer household to allocate elsewhere. The marginal household in a location adds output—net of consumption and the construction cost—but also increases congestion for the $n(x)$ households already there.

To recall, the first order condition for the choice of local zoning was

$$\frac{F}{n(x)} - n(x) \frac{v'(n(x))}{u'(c(x))} = 0.$$

Both the planner and the initial residents consider the role of the congestion externality. While the initial residents weigh this externality against the value of sharing the local fixed cost more broadly, the planner weighs it against the resource value of allocating a marginal household to location x and the shadow value of the binding population constraint. The first term in Equation (21) represents the net output of allocating a marginal household to location x . Under the flat gradient that arises from the local zoning equilibrium, the marginal value would be higher in productive locations. The planner will thus allocate more households to such locations. This is the key intensive-margin wedge between the optimal solution and the allocation with local zoning.

While initial residents ignore aggregate effects, they do consider the marginal value of sharing the fixed cost more broadly. The planner ignores this margin: the fixed cost is paid when the location was opened and should not affect the intensive margin. This is the second wedge between the optimal solution and the allocation with local zoning. In short, the initial residents ignore the aggregate effects of their choice to restrict housing supply.

The planner instead considers the fixed cost F at the extensive margin. As shown in Equation (22), the planner weighs the net output of opening the marginal location \underline{x} against the shadow value of assigning $n(\underline{x})$ households to this location.

$$(22) \quad \Lambda \left[n(\underline{x}) (y(\underline{x}) - c(\underline{x})) - \left(\frac{n(\underline{x})}{Z} \right)^2 - F \right] - n(\underline{x}) \mu = 0.$$

Recall that in the local zoning allocation outlined previously, the location in which the $OL(n, x)$ and $MC(n)$ are tangent becomes the threshold. Restating, this condition was met where

$$\frac{F}{n^*(\underline{x})^2} - [u^{-1}]' (v(n^*(\underline{x})) + \bar{u}) \times v'(n^*(\underline{x})) = 2/Z^2.$$

This location arises through the general spatial equilibrium, rather than through being chosen directly by any agent. In choosing the optimal threshold, the planner weighs the value and costs of opening a marginal location. The differential outcomes for the threshold \underline{x} highlight an additional externality generated by the initial resident choice of local zoning: restrictive zoning ensures that too many locations are opened in equilibrium, necessitating the payment of fixed costs to open locations that would not be paid under the optimal allocation.

2.3.1 Decentralization: Socially Optimal Zoning

The planner's allocation can be decentralized through the adoption of an alternative zoning regime. In this decentralization, the full measure 1 of households votes on the set \mathcal{X} of locations to open and the gradient of zoning laws $\bar{n}(x)$ for each location x in \mathcal{X} . In so doing, households have rational expectations over the spatial equilibrium induced by the choices they make. Households therefore choose zoning laws $\bar{n}(x)$ and the set of opened locations \mathcal{X} to maximize the endogenous outside option \bar{u} . Households choose these laws subject to the population constraint and to the local and general spatial equilibria conditions.

Proposition 5. *The utility-maximizing choices for the set of opened locations \mathcal{X} and the set of zoning laws $\bar{n}(x)$ are identical to those chosen by the planner. The equilibrium price gradient and aggregate profits implement the planner's choice of consumption. This set of instruments allows households to fully implement the planner's allocation.*

Proof. See appendix. □

Intuitively, households have the option of choosing the same zoning as the planner and will do so. Prices will adjust such that households are indifferent across all locations. With identical population across locations as the planner's allocation, total output and aggregate construction and fixed costs are also identical. By resource balance, aggregate consumption must also be identical. By spatial equilibrium, consumption must make households indifferent between locations, and therefore equilibrium consumption is identical to the planner's allocation. The zoning laws and set of opened locations chosen by households in this problem can therefore be described as *socially optimal zoning*, as opposed to the *locally optimal zoning* described previously.

The profits paid to households are a crucial mechanism for ensuring that this spatial equilibrium will also be consistent with construction firm behavior. Note that household spatial equilibrium pins down relative prices, but not levels. Because profits are distributed to households regardless of location, an increase of ε in the absolute level of prices at each location leads to an increase in household profits by ε . This increase leaves household consumption identical. This mechanism ensures that there exists a price gradient consistent with local equilibrium from the firm perspective.

3 Calibration

3.1 Data

Four key model components must be calibrated to data: the magnitude of the local fixed cost, the shape of utility from consumption, the disutility of congestion, and the distribution of location-specific productivity. This section describes the moments in the data that the calibration seeks to match. In matching the moments, I take the set of urban U.S. Census tracts to be the empirical counterpart of model locations.

For the utility of consumption, I consider $u(c) = c$. Linear utility follows related papers (Van Nieuwerburgh and Weill, 2010; Davis and Dingel, 2014).²³

For the disutility of congestion, I consider $v(n) = \gamma n^\eta$. To calibrate γ , I match the fraction of income spent on consumption goods as measured in the Consumer Expenditure Survey of the Bureau of Labor Statistics. As γ controls the relative utility weights of consumption and congestion, a higher γ will induce households to open more locations and spend more on housing and fixed costs rather than consumption. To ensure the data is well-matched to the model concept, calibrate the consumption share of income to spending on tradeable goods. These spending shares range from 52–61% for different income deciles, and average 58% overall.

As noted previously, an alternate interpretation of the local fixed cost has residents gaining utility from a greater diversity of monopolistically competitive local services, each of which is subject to a fixed cost.²⁴ Using this interpretation, I calibrate the fixed cost F to match the share of consumer spending on local services. In particular, I use the Consumer Expenditure Survey categories for food away from home, personal services, medical services, and fees and admissions. These spending shares range from 7–10% for different income deciles, and average 8.5% overall.

I calibrate the congestion shape parameter η to match the average population density of urban census tracts, as identified by the Census Bureau for 2010. Given the functional form assumptions, the local equilibrium population for locations with binding zoning is given by

$$n = \left(\frac{F}{\gamma \eta} \right)^{1/(1+\eta)}.$$

As the zoning law will bind for a substantial fraction of locations, the shape parameter η has a first-order effect on average population density.²⁵ Intuitively, too low a congestion cost

²³See appendix for robustness with log utility.

²⁴In a local services model, expenditure will be precisely $F/n(x)$ if the elasticity of substitution between local service firms is 1 or if aggregated local services enter the utility function linearly, such that $u(c, S) = c + S$, where S is a CES aggregator of local services. More generally, local expenditure will be $F/n(x)^\sigma$, where σ is the elasticity of substitution between local services.

²⁵The intuition is similar in spirit to using wage shares to calibrate the labor and capital exponents in a production function: in equilibrium, the parameters can have first-order effects on levels.

Param	Value	Description	Model Target	Data	Model
γ	1/215	Congestion weight	Consumption sh. of expend	52%–61%	55%
F	0.1265	Fixed cost	Local services sh. of expend	7%–10%	7.8%
η	7	Congestion curvature	Ave. urban density	5,100/sq mi	5,900/sq mi
Z	2.7	Construction prod.	p/MC in best locations	5.2–5.8	5.6

Table 1: The classification of tracts as *urban* follows the Census Bureau’s 2010 classification. The consumption share of expenditure is taken from the Consumer Expenditure Survey for 2012, and is equal to consumption net of spending on housing and local services. The price-to-marginal cost ratio follows the approach of Glaeser et al. (2005a).

would cause households to live at too-high densities in productive locations, compared to the data.²⁶

I calibrate construction firm productivity to match the ratio of price per square foot to cost per square foot in the most productive places: Manhattan, San Francisco, or Silicon Valley. For this exercise, average price per square foot data is taken from Zillow, and cost data is taken from RS Means. Both report cost data at the level of counties and, in some instances, for sub-county regions.

The distribution of income is taken from census tract median incomes, adjusted for observables in a process that follows Hsieh and Moretti (2015). I take individual-level data from IPUMS on income, education, race, and gender. I then run a regression using the following equation:

$$y_i = X_i\beta + \varepsilon_i,$$

where y_i is the income of individual i and X_i is the vector of observable characteristics. I then take the estimate of β and calculate the residual income \tilde{y} of each tract ℓ :

$$\tilde{y}_\ell = y_\ell - X_\ell\beta,$$

where y_ℓ is the measured average income per worker and X_ℓ is the fraction of the tract with the given observable characteristics.

Table (1) summarizes the calibration targets and the model fit. For three of the four parameters, the calibrated parameter is within the range from the data. The search of the parameter space failed to find a set of values for which the density was closer to the data while maintaining consistency with the other parameters. At the same time, it is plausible that households have reasons beyond productivity for choosing locations.

²⁶Note that it is not necessary to take a stand on the precise productivity of *non-urban* locations, whose measured incomes may not be a good guide to those earned by new households. It is sufficient to assume that non-urban households have lower productivity than those occupied at urban population levels.

4 Quantitative Results

The key quantitative question this paper addresses is: what are the welfare costs of locally-determined zoning laws? In the baseline calibration, implementing the planner’s allocation would increase welfare by 1.4%. Consumption would increase by 2.4% and GDP by 2.1%, but increased congestion would mitigate these gains.²⁷ The planner would close the 3% lowest-productivity locations, and population density at the 95th percentile location would increase by 18%. The density at the median location would increase by 2%. These results are summarized in Table (2), along with results from two additional counterfactuals.

The first counterfactual, labeled *Market* in the table, corresponds to an allocation with no zoning, and construction firms free to build without constraint. In the language of the model, the equilibrium for each location is the upper intersection of the $OL(n, x)$ and $MC(x)$ curves. This counterfactual corresponds to that studied by Hsieh and Moretti (2015), and I find an increase in GDP of 6%, half of the 13.5% that they report. Part of the different may be due to the presence in this model of congestion externalities, which limit the willingness of households to crowd into productive locations. Moreover, the estimates are of the same order of magnitude despite using different methodologies.

While the GDP increase from zoning abolition is substantial, the increased in congestion is as well. The median location sees an increase in population density of 15%, and the increase is almost 50% at the 95th percentile. Accordingly, welfare declines by almost 6%, despite the large increase in consumption. This result suggests that the productivity gains to zoning abolition put forth by Hsieh and Moretti (2015) will not, in fact, increase welfare.

The second counterfactual, labeled *Upzoning* in the table, corresponds to an allocation with no zoning in locations with the productivity greater than the 95th percentile. As census tracts average about four thousand inhabitants, these locations are home to approximately ten million residents. For scale, the San Francisco Bay Area is home to five million urban residents. Under this allocation, GDP increases by 2% while welfare declines by 2%. This contrasts with Hsieh and Moretti (2015), who find large productivity gains of 9.5% from increasing the population of the best cities.

The median rent, also in Table (2), provides insight into the changing tradeoffs with each reform. Three key forces determine the median rent: the productivity of the median resident, the congestion faced by the median resident, and the outside option \bar{u} . Under the decentralized version of the planner’s allocation, house prices *increase* by 3% for the median household. Each location is a little more crowded—and so rents for each location fall, between 1–5%—but the median household is now in a more productive location and thus their rent increases. The market equilibrium also sees the median household in a more productive location, but this location is now much more crowded. Correspondingly, rents fall by 16%. Finally, in the upzoning case, the median household is in a more productive location with

²⁷Recall that utility is quasilinear, so the welfare and consumption figures can be compared meaningfully.

Relative Outcome	Planner	Market	Upzoning
Welfare	1.4%	-5.9%	-2.0%
GDP	2.1%	6.0%	2.0%
Consumption	2.4%	5.6%	2.1%
Median Rent	2.9%	-16%	-1.8%

Table 2: All results are relative to the main zoning results. The planner results give the gains from moving to the optimal allocation. The market results give the counterfactual with no zoning laws. The upzoning results give the counterfactual with no zoning laws in the 5% highest-productivity locations.

the same level of congestion. However, they are slightly worse off, and so they are unwilling to bid as much for housing, and rents fall.

Given the calibrated parameters, the main policy implication of the model is to allow zoning laws to be chosen at a higher geographic level—preferably, the national level. This would ensure that the productivity effects of zoning laws are internalized and would enable the implementation of the planner’s optimal allocation.

The fundamental role of congestion in shaping preferences and outcomes points towards a second policy implication beyond implementing the planner’s regime: introducing reforms that could lower the cost of congestion to the neighborhood. Understanding the specific component of congestion that drives externalities will identify the appropriate form of mitigation. If neighborhood congestion is driven by vehicle traffic, then perhaps rapid transit would enable more-dense development by reducing the costs of congestion imposed by new development. Congestion-mitigation efforts like transit can be quite costly, and within a context of spatial equilibrium, the benefits would be felt widely. As with local zoning, the current transit-planning regime may not take into consideration the external effects of transit.

Distributional Outcomes Abolishing zoning—whether nationally or just in the highest-productivity locations—cannot improve welfare under this calibration. This fact is driven in part by the curvature to the cost of congestion: increasing the intensity of development in productive places is simply too costly in terms of welfare. A second factor that plays an important role is the assumption that profits are shared equally among all households. When restrictive zoning drives up house prices in productive locations, the profits earned are redistributed across all locations. To understand the role played by this assumption, I now introduce a second calibration wherein construction-sector profits are not returned to households. This calibration mirrors the traditional urban economics assumption of *absentee landlords*, who collect rent but do not otherwise interact with households.²⁸

²⁸Eeckhout and Guner (2015) performs a similar exercise and identifies absentee landlords with the concentrated landholdings of the 1% wealthiest households. They suppose that the planner values a fraction of housing sector profits that is distributed to households, but not the fraction distributed to absentee landlords.

Relative Outcome	Planner	Market	Upzoning
Welfare	1.5%	-10%	-3.2%
Welfare ex profits	7.3%	23%	3.7%
GDP	2.6%	6.3%	1.9%
Consumption	9.3%	50%	12%
Median Rent	-5.3%	-42%	-2.1%

Table 3: All results are relative to the main zoning results, and all results ignore the welfare derived from profits—this allows the welfare gains *discounting profits* from the market allocation to exceed those from the planner’s. The planner results give the gains from moving to the optimal allocation. The market results give the counterfactual with no zoning laws. The upzoning results give the counterfactual with no zoning laws in the 5% highest-productivity locations.

Table (3) shows the analogous results under this model.²⁹ Here, the gains to GDP are similar to those before, and the losses to welfare of zoning abolition are even greater. However, the change to welfare discounting profits—that is, the welfare of the typical households—are positive. This exercise highlights the key role played by construction profits. Residents of productive places don’t actually enjoy the high productivity, they simply pay higher rents. When these profits are shared broadly, these rents are redistributed so that all households, regardless of location, share the output of high-productivity locations. When these profits are not shared broadly, households in less productive locations have more to gain from zoning abolition because moving into a productive location now earns them significantly more consumption.

4.1 Empirical Implications

The model focuses on just two of potentially many local externalities. Numerous facets of the model lend themselves to further empirical validation. Similarly, the simple model can be extended with greater heterogeneity, including of household productivity, of locational or household preferences over congestion costs and fixed costs, and of locational amenities.

Within the model, existing residents choose zoning regulations to maximize their amenity mix, trading sharing externalities against congestion. In doing so, they also maximize the rents that new households are willing to pay: maximizing amenities and maximizing land values are identical. However, local ownership of land (e.g., in the form of homeownership) may break this felicitous duality as landowners may seek to restrict the supply of new housing, a close substitute to their own asset. Empirically, it may be possible to distinguish between these intents using evidence from surveys or by identifying preferences over regulations that would increase supply but have a positive effect on local amenities (or vice versa).

²⁹For these results, I have recalibrated the model so that the new model-defined moments are consistent with the data. The consumption share is 54%, the local service share 7.7%, the construction sector markup is 7.4%, and 88% of locations in the data are opened.

Local variation in model parameters could be tested to examine the strength of their relationship to equilibrium outcomes. For instance, the fixed costs of development may be higher in regions that are more arid, or prone to extreme weather, due to the costs of necessary infrastructure. New development in the arid fringe of Los Angeles is quite dense by the standards of new development in rainier eastern cities like Atlanta.³⁰ All else equal, do locations with higher fixed costs see higher-density development, as the model would predict?

Similarly, some locations have seen dramatic changes to their productivity, but homeowners—like the initial residents of the model—don’t see any immediate change in their costs. The model predicts that locations with large positive shocks should see little new development but great increases in house price. Instead, the model predicts that new development would be concentrated in previously-rural locations at the fringes of newly productive regions. Anecdotally, Palo Alto—a key center of innovation in the San Francisco area—has grown from a population of 55,225 in 1980 to 64,403 in 2010. The city of Antioch, of similar geographic extent at the fringe of the region, grew from 43,559 residents to 102,372 over the same period. Part of this differential can be explained by the effects of durable capital (cf. Siodla, 2015); a careful test of this prediction could be challenging but insightful.

In Boulder, CO, a measure on the ballot in the fall of 2015 would have empowered dozens of local neighborhood groups with an easy path to halting unwanted development, within a city of just 100,000 residents. The outcome and voting patterns in the election could help to test model predictions. In particular, the model assumes locations are sufficiently small that residents can ignore the aggregate effects of their zoning decisions. Regulatory changes at higher levels of government are more likely to affect aggregate variables, such as the threshold occupied location. Residents who stand to benefit more from changes to aggregate variables would be more likely to oppose the measure. Renters, the young, and new arrivals may be more likely to benefit from aggregate changes within Boulder, relative to long-time homeowners with substantial equity dependent on the status quo.

Similarly, the model could be extended to include differential preferences over congestion: perhaps some households are more or less bothered by adding more neighbors. Empirically, households with families might be more sensitive to congestion, even conditioning on dwelling size. This extension could refine the estimates of welfare gains from adding density to productive locations.

5 Conclusion

This paper provides a unified framework to address the local and aggregate welfare effects of local land-use regulation. It provides empirically-grounded externalities that incent

³⁰For example, a recently-developed subdivision in Lancaster, CA was developed to approximately 5,000 people per square mile, while a similarly-new development in Powder Springs, GA is approximately half the density.

households to pass restrictive zoning laws that prevent new housing development at the neighborhood level. In their endogenous choice of zoning, households rationally ignore the aggregate implications that arise from location heterogeneity. Households thus bid up the price of artificially scarce housing in productive locations so that prices exceed the marginal costs of construction.

The endogeneity of zoning plays a critical role in overcoming the lack of sufficient data to identify the wedge between house price and marginal cost at the neighborhood level. A calibrated version of the model is used to perform a welfare calculation: how large are the aggregate losses to welfare from local zoning relative to the planner's optimum? The model provides an interpretation of these welfare costs without measuring these wedges directly. In the preferred calibration, welfare could be raised 1.4% lower by implementing the planner's allocation. Aggregate output would be raised 2.1%: one-third of the gains to output are negated by the increased congestion felt by residents of productive locations.

References

- Albouy, D. (2012). Evaluating the efficiency and equity of federal fiscal equalization. *Journal of Public Economics*, 96(9):824–839.
- Albouy, D. and Ehrlich, G. (2012). Metropolitan land values and housing productivity. Unpublished.
- Allen, T. and Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 1085:1139.
- Allen, T., Arkolakis, C., and Li, X. (2015). Optimal City Structure. Unpublished.
- Arnott, R. J. and Stiglitz, J. E. (1979). Aggregate land rents, expenditure on public goods, and optimal city size. *The Quarterly Journal of Economics*, pages 471–500.
- Aura, S. and Davidoff, T. (2008). Supply Constraints and Housing Prices. *Economics Letters*, 99(2):275–277.
- Avent, R. (2011). *The Gated City*. Amazon Digital Services.
- Behrens, K., Mion, G., Murata, Y., and Südekum, J. (2014). Spatial frictions. Technical report, Düsseldorf Institute for Competition Economics.
- Calabrese, S., Epple, D., and Romano, R. (2007). On the Political Economy of Zoning. *Journal of Public Economics*, 91(1):25–49.
- Chen, C.-S. and Lai, F.-C. (2008). Location choice and optimal zoning under cournot competition. *Regional Science and Urban Economics*, 38(2):119–126.
- Cheshire, P. C. and Hilber, C. A. (2008). Office Space Supply Restrictions in Britain: The Political Economy of Market Revenge. *The Economic Journal*, 118(529):F185–F221.
- Davis, D. R. and Dingel, J. I. (2014). The Comparative Advantage of Cities. Technical report, National Bureau of Economic Research.
- Desmet, K. and Rossi-Hansberg, E. (2013). Urban accounting and welfare. *American Economic Review*, 103(6):2296–2327.
- Diamond, R. (2015). The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000. *Manuscript*.
- Eeckhout, J. and Guner, N. (2015). Optimal Spatial Taxation: Are Big Cities Too Small? Technical report, Institute for the Study of Labor.
- Fajgelbaum, P. D., Morales, E., Serrato, J. C. S., and Zidar, O. M. (2015). State Taxes and Spatial Misallocation. Technical report, National Bureau of Economic Research.

- Fischel, W. A. (1987). *The Economics of Zoning Laws: A Property Rights Approach to American Land Use Controls*. Johns Hopkins University Press.
- Fischel, W. A. (1999). Does the American Way of Zoning Cause the Suburbs of Metropolitan Areas to be Too Spread Out? *Governance and Opportunity in Metropolitan America*, 15:1–19.
- Fischel, W. A. (2009). *The Homevoter Hypothesis*. Harvard University Press.
- Fischel, W. A. (2015). *Zoning Rules!* Lincoln Institute of Land Policy, Cambridge, Massachusetts.
- Flatters, F., Henderson, V., and Mieszkowski, P. (1974). Public goods, efficiency, and regional fiscal equalization. *Journal of Public Economics*, 3(2):99–112.
- Ganong, P. and Shoag, D. (2013). Why has Regional Income Convergence in the US Declined. *Manuscript, Harvard*.
- Glaeser, E. L. and Gyourko, J. (2002). Zoning’s Steep Price. *Regulation*, 25:24.
- Glaeser, E. L. and Gyourko, J. (2003). The Impact of Building Restrictions on Housing Affordability. *Federal Reserve Bank of New York. Economic Policy Review*, 9(2):21.
- Glaeser, E. L., Gyourko, J., and Saks, R. E. (2005a). Why Have Housing Prices Gone Up? *The American Economic Review*, 95(2):pp. 329–333.
- Glaeser, E. L., Gyourko, J., and Saks, R. E. (2005b). Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices. *Journal of Law and Economics*, 48(2):331–369.
- Glaeser, E. L., Gyourko, J., and Saks, R. E. (2006). Urban Growth and Housing Supply. *Journal of Economic Geography*, 6(1):71–89.
- Gyourko, J., Mayer, C., and Sinai, T. (2013). Superstar Cities. *American Economic Journal: Economic Policy*, 5(4):167–199.
- Gyourko, J. and Molloy, R. (2014). Regulation and Housing Supply. Technical report, National Bureau of Economic Research.
- Hamilton, B. W. (1975). Zoning and Property Taxation in a System of Local Governments. *Urban Studies*, 12(2):205–211.
- Henderson, J. V. (1991). Optimal regulation of land development through price and fiscal controls. *Journal of Urban Economics*, 30(1):64–82.
- Hilber, C. A. and Robert-Nicoud, F. (2013). On the origins of land use regulations: Theory and evidence from us metro areas. *Journal of Urban Economics*, 75:29–43.

- Hills, R. and Schleicher, D. (2011). Balancing the “Zoning Budget”. *Case Western Reserve Law Review*, 42(81):112–118.
- Hills, R. M. and Schleicher, D. (2014). City Replanning. *George Mason Law & Economics Research Paper*, (14-32).
- Hochman, O., Pines, D., and Thisse, J.-F. (1995). On the optimal structure of local governments. *The American Economic Review*, pages 1224–1240.
- Hsieh, C.-T. and Moretti, E. (2015). Why Do Cities Matter? Local Growth and Aggregate Growth. Technical report, National Bureau of Economic Research.
- Koster, H. R., van Ommeren, J., and Rietveld, P. (2012). Bombs, Boundaries and Buildings: a Regression-Discontinuity Approach to Measure Costs of Housing Supply Restrictions. *Regional Science and Urban Economics*, 42(4):631–641.
- Krugman, P. (1991). Increasing Returns and Economic Geography. *The Journal of Political Economy*, 99(3):483–499.
- Mayer, C. J. and Somerville, C. T. (2000). Land Use Regulation and New Construction. *Regional Science and Urban Economics*, 30(6):639–662.
- Monkkonen, P. and Quigley, J. M. (2008). The Spatial Consequences of Autarky in Land-Use Regulation: Strategic Interaction or Parallelism? *Berkeley Program on Housing and Urban Policy*.
- Ortalo-Magné, F. and Prat, A. (2014). On the Political Economy of Urban Growth: Homeownership versus Affordability. *American Economic Journal: Microeconomics*, 6(1):154–81.
- Rossi-Hansberg, E. (2004). Optimal Urban Land Use and Zoning. *Review of Economic Dynamics*, 7(1):69–106.
- Saiz, A. (2010). The Geographic Determinants of Housing Supply. *The Quarterly Journal of Economics*, 125(3):1253–1296.
- Schleicher, D. (2013). City unplanning. *Yale Law Journal*, 122(7):1670–1737.
- Siodla, J. (2015). Razing San Francisco: The 1906 disaster as a natural experiment in urban redevelopment. *Journal of Urban Economics*, 89:48 – 61.
- Stull, W. J. (1974). Land Use and Zoning in an Urban Economy. *American Economic Review*, pages 337–347.
- Turner, M. A., Haughwout, A., and van der Klaauw, W. (2014). Land Use Regulation and Welfare. *Econometrica*, 82(4):1341–1403.

- Van Nieuwerburgh, S. and Weill, P.-O. (2010). Why Has House Price Dispersion Gone Up? *The Review of Economic Studies*, 77(4):1567–1606.
- Yglesias, M. (2012). *The Rent Is Too Damn High: What To Do About It, And Why It Matters More Than You Think*. Simon and Schuster.

A Additional Figures




Figure 3: Plot of housing prices and increases in the quantity of housing units in selected counties. In descending order, the top three counties are San Francisco, San Mateo, and Santa Clara: Silicon Valley.




Figure 4: A Google Earth satellite image of the built-up edge of Las Vegas. Note that the city edge is obvious: on the left, population densities are in the upper single-digit thousands while the right is undeveloped desert. While this desert is itself likely to fill in, the pattern of a sharp edge will likely remain. While desert development may entail unusually large fixed costs, similar patterns hold in regions with friendly climates.

B Proofs

Lemma 1. *Define \bar{n} to be the population that maximizes the $OL(n, x)$ curve. Given outside option \bar{u} , the level of population \bar{n} is unique and independent of location.*

Proof. From Proposition (1), the $OL(n, x)$ curve is strictly concave and twice continuously differentiable. The portion of the $OL(n, x)$ curve with $n(x) > 0$ is given:

$$OL(n, x) = y(x) + \Pi - \frac{F}{n} - u^{-1}(v(n) + \bar{u}).$$

The following first order condition is thus necessary and sufficient for the maximum:

$$(23) \quad \frac{F}{n^2} - [u^{-1}]'(v(n) + \bar{u})v'(n) = 0.$$

Call the population level that solves this \bar{n} . Given \bar{u} , the value of \bar{n} is constant and independent of location. Write as $\mathcal{N}(\bar{u})$ the strictly decreasing function that gives the level of population that solves Equation (23) as a function of \bar{u} . \square

Lemma 2. *Define location $\hat{x}(\bar{u}, \Pi)$ as follows:*

$$\hat{x}(\bar{u}, \Pi) = \begin{cases} x \text{ s.t. } OL(\mathcal{N}(\bar{u}), \hat{x}) = MC(\mathcal{N}(\bar{u})) & \text{if } OL(\mathcal{N}(\bar{u}), 1) \geq MC(\mathcal{N}(\bar{u})) \\ 1 & \text{otherwise.} \end{cases}$$

Then $\hat{x}(\bar{u}, \Pi)$ is unique and in the range $[\underline{x}, 1]$.

Proof. In the first case, the highest-productivity location has an $OL(n, 1)$ curve with maximum value that exceeds the marginal cost at the corresponding level of population \bar{n} . At the threshold occupied location \underline{x} , the $OL(n, \underline{x})$ and $MC(n)$ curves meet with tangency. The marginal cost curve has positive slope, and so this tangency must occur to the left of the maximum of the $OL(n, \underline{x})$ curve. Evaluated at $\mathcal{N}(\bar{u})$, $OL(\mathcal{N}(\bar{u}), x)$ is continuous and strictly increasing in x . Therefore there exists a unique $\hat{x}(\bar{u}, \Pi) \in [\underline{x}, 1]$ such that $OL(\mathcal{N}(\bar{u}), \hat{x}(\bar{u}, \Pi)) = MC(\mathcal{N}(\bar{u}))$.

In the second case, $OL(\mathcal{N}(\bar{u}), x) < MC(\mathcal{N}(\bar{u}))$ for all locations. Then $\hat{x}(\bar{u}, \Pi)$ is uniquely defined as 1 and is within the specified range. \square

Proposition 1. *In a stable general equilibrium, the following characterizes the local equilibrium populations and optimal zoning. For all locations $x > \hat{x}$, $n^*(x) = \bar{n}(x) = \mathcal{N}(\bar{u})$. For all locations $x \in [\underline{x}, \hat{x}]$, the population $n^*(x)$ and optimal zoning $\bar{n}(x)$ are given by the upper intersection of the $OL(n, x)$ and $MC(n)$ curves. For all $x \geq \underline{x}$, $p^*(x)$ is given by the value of the $OL(n, x)$ curve evaluated at $n^*(x)$.*

Proof. For locations $x > \hat{x}$, the $OL(n, x)$ curve is strictly greater than the $MC(n)$ curve, and therefore the equilibrium population can only be given by \bar{n} if the initial residents choose the

zoning law $\bar{n}(x) = \bar{n}$. I will show that this is the case. Note $OL(n, x) > MC(n)$ implies that the participation constraint will not bind, and so the solution to the initial resident problem is given by the first order condition

$$u' \left(y(x) + \Pi - \frac{F}{n} - p_0(x) \right) \frac{F}{n^2} = v'(n).$$

And in stable equilibrium, $p_0(x) = p(x)$ so substituting from the household spatial equilibrium condition gives

$$u' (u^{-1} (v(n) + \bar{u})) \frac{F}{n^2} = v'(n).$$

From above, the maximum of the $OL(n, x)$ curve is given by the n that solves

$$[u^{-1}]' (v(n) + \bar{u}) v'(n) = \frac{F}{n^2}.$$

By the definition of the derivative of the inverse, these are equivalent: for locations with non-binding participation constraints, the level of population $\mathcal{N}(\bar{u})$ that gives the maximum of the $OL(n, x)$ curve is identical to that which maximizes initial resident utility (given the stability condition that $p_0(x) = p(x)$). For all locations $x > \hat{x}(\bar{u}, \Pi)$, $\mathcal{N}(\bar{u})$ is consistent with the participation constraint and hence $n^*(x) = \bar{n}(x) = \mathcal{N}(\bar{u})$ for such locations. By Equation (6), the equilibrium price for such locations is given by $OL(\mathcal{N}(\bar{u}), x)$.

Consider location $x \in [\underline{x}, \hat{x}]$. From above, with $p_0(x) = p(x)$ the unconstrained choice of zoning that maximizes initial resident utility is the level of population that maximizes the $OL(n, x)$ curve. By definition, for locations with $x < \hat{x}(\bar{u}, \Pi)$ the maximum of the $OL(n, x)$ curve is below the $MC(n)$ curve and hence infeasible. By concavity of the initial resident problem, the participation constraint will bind and the local equilibrium population (and optimal zoning law) is given by the upper intersection of the $OL(n, x)$ and $MC(n)$ curves. \square