Housing Supply in the 2010s

Salim Furth

MERCATUS WORKING PAPER

MERCATUS CENTER
George Mason University
3434 Washington Blvd., 4th Floor, Arlington, Virginia 22201
www.mercatus.org

Salim Furth. "Housing Supply in the 2010s." Mercatus Working Paper, Mercatus Center at George Mason University, Arlington, VA, February 2019.

Abstract

Geographic patterns of housing growth in the United States result from the interaction of demand, density, and regulatory institutions. This paper presents a model of urban redevelopment in which opportunity costs and regulatory hurdles hinder the occurrence, timing, and density of redevelopment. Regulations may inhibit a land parcel's redevelopment through costs, delays, constraints, or rent control, each of which is shown to affect potential redevelopment differently. Unlike in models of greenfield development, the local distribution of existing uses will be a key determinant in the response to a shock or regulation, and housing supply curves will have an inverted S shape. Nonparametric estimation of housing supply growth from 2012 to 2018 as a function of demand growth and density yields a residual growth surplus for each census tract. The tract-level surpluses can be aggregated by metro area, town, or density cohort to indicate the relative growth-friendliness of local institutions and unobserved conditions. Growth surpluses offer a summary of implicit local barriers to growth that is more detailed and robust than measured supply elasticities. At the jurisdiction level, growth surpluses are negatively and significantly correlated with some measures of regulatory intensity, notably the de jure density restrictions in the Housing Regulation Database of Massachusetts Municipalities.

JEL codes: R12, R31

Keywords: housing supply, residential construction, urban growth, zoning, land use regulation, rent, elasticity of housing supply, nonparametric estimation, institutions

Author Affiliation and Contact Information

Salim Furth
Senior Research Fellow
Mercatus Center at George Mason University
sfurth@mercatus.gmu.edu

Housing Supply in the 2010s

Salim Furth

Introduction

The price of housing varies far more across American metropolitan areas than do wages, construction costs, or commute times. In 2012, the median home in the Los Angeles metro area would have rented for $2,175 a month, compared with only $1,165 a month in metro Atlanta.[1] Atlanta added twice as much housing to its supply over the subsequent six years, while the rent gap grew. Institutional and physical barriers to housing growth play important roles in explaining why high and rising prices did not lead to higher housing growth in Los Angeles.

The construction of new housing takes place in a complex institutional environment: retail transaction costs are high, tax incentives advantage owner-occupants, regulators set limits to both the stock and the flow of housing, and geographic variation in regulations, terrain, and demand is constant.

This paper measures how much of each location's residential growth rate may be owing to institutional barriers. For each census tract, the paper estimates an expected growth rate using the increase in demand and the existing density. The residual, or "surplus," growth may be owing to the relative effectiveness of local institutions.

In addition to confirming some well-known patterns, the geographic fineness of the analysis enables characterizing results for individual jurisdictions and distinguishing between densification and expansion as sources of growth. Texas, it turns out, densifies more readily than California and is friendlier to outward expansion.

---

[1] Zillow Rent Index, January 2012, https://www.zillow.com/research/data/. Rents are listed in current (2012) dollars.

Figure 1 offers some motivation for the approach previously described. Recent housing demand and supply growth have differed substantially across census tracts of different densities. While rent has risen most in the densest places, growth in housing supply has had a nonmonotonic relationship with density, suggesting that there are institutional as well as physical determinants of how and where growth occurs.

Figure 1. Rent and Housing Supply Growth across the Density Distribution

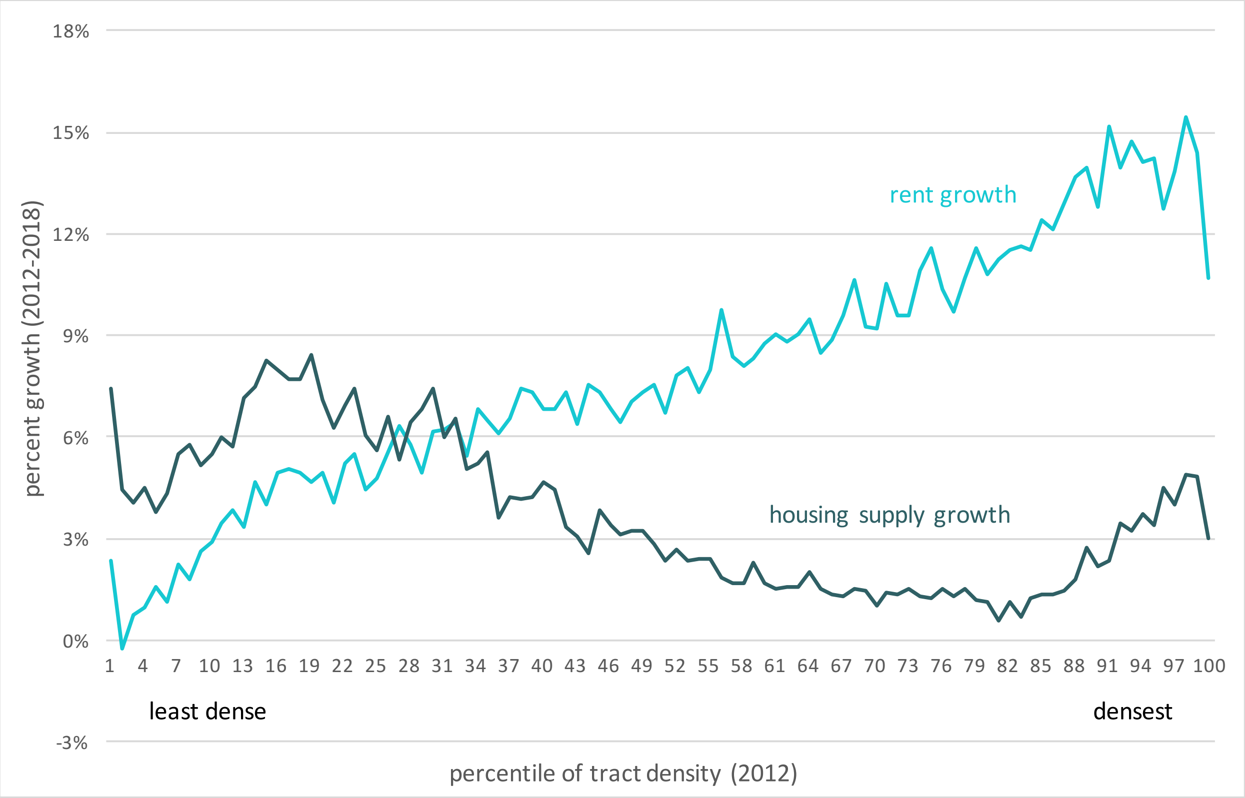

Note: Tracts are weighted so that every percentile has the same number of residential addresses.

Sources: Author's calculations based on Zillow, ZRI Time Series: Multifamily, SFR, Condo/Co-op (dataset), accessed October 2018, https://www.zillow.com/research/data; and US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

This paper contributes to an active literature on the importance of the institutions that regulate land use by estimating predicted and actual growth at the census tract level and by aggregating the tract-level surpluses or deficits to municipalities, counties, and metropolitan areas. Rather than directly measuring regulatory institutions such as zoning laws, political fragmentation, or regulatory delays and then inferring their impact on housing outcomes, this approach creates a near-universal measure of outcomes that can, among other uses, evaluate the predictive power of regulatory metrics.

The paper presents a model of parcel-based urban development, in which each parcel has an existing use and, conditional on redevelopment, an optimal density.[2] Five varieties of regulations can cause deviations from the laissez-faire equilibrium. Although the model yields implications for each type of regulation, the implications may not be testable, given that regulations are almost always observed in complex bundles.

Growth surplus outcomes can help explain the degree to which low elasticities of supply in coastal metropolitan areas are a product of their density and lack of hinterlands, whether Texas cities are pro-growth in dense areas or only in greenfields, and whether measures of regulatory input are helpful in predicting those output patterns across jurisdictions.

The next section reviews some of the key papers that informed this effort. The following section briefly describes the institutional structure of land use decisions in the United States. The study then presents a model of land use redevelopment decisions and analyzes the distortions caused by five forms of regulation. The following three sections describe the empirical strategy, data sources and choices, and estimation results of this paper.

---

[2] Here and throughout the paper, optimality is private, not public or social.

Literature

This paper builds in part on insights from the spatial urban model developed by Alonso (1964), Mills (1967), and Muth (1969) and restated along with many extensions in Duranton and Puga (2015). The spatial model emphasizes explanations for widely observed urban patterns, such as the decline of density and price with distance from a central business district. The spatial model is grounded in a market-based view of urban development, in which prices determine the size, density, and use patterns of a city.

The Alonso-Mills-Muth model has been the dominant theoretical framework behind empirical work since Mayer and Somerville (2000b). Among the empirical studies in this tradition are Mayer and Somerville (2000a); Green, Malpezzi, and Mayo (2005); Quigley and Raphael (2005); Vermeulen and Rouwendal (2007); and Saiz (2010). As Wheaton, Chervachidze, and Nechayev (2014) note, the spatial model is closely related to the stock-flow model of housing that developed earlier (see, e.g., DiPasquale and Wheaton 1994). Liu, Rosenthal, and Strange (2018) point out that most applications of the spatial model are focused on horizontal growth at the suburban fringe. An exception is Meen and Nygaard (2011), who show that existing patterns of land use are determinants of subsequent development.

Empirical studies of metropolitan housing markets, including many of the papers just cited, yield a wide range of quantitative outcomes. Many report elasticities of supply, but the definitions of elasticity vary from paper to paper,[3] as do the time periods over which the elasticities are defined. In some models, such as that of Saiz (2010), elasticity of supply is

---

[3] There are three approaches: the elasticity of housing stock with respect to price, the elasticity of building permits with respect to change in price, and (in papers from a different theoretical tradition) the elasticity of building permits with respect to price.

constant across price changes; in others, including the present paper, the model does not imply a constant elasticity supply curve.

In this paper, I extend a model in the spirit of Alonso-Mills-Muth to explicitly consider how five different types of regulation would affect development and, to the extent possible with available data, to test those implications against observed outcomes. Some previous studies have distinguished among types of regulation. Mayer and Somerville (2000a) distinguish cost-additive regulations from regulatory delays and show that in a multistage development process, regulatory delays could have counterintuitive results. Paciorek (2013) incorporates time-to-build lags in a structural model of the effects of regulation and land scarcity. Murphy (2017) incorporates fixed and variable cost parameters, which include regulatory and other costs, in a parcel-level estimation of flow-difference supply elasticities in the San Francisco Bay Area.

Many empirical papers, including Green, Malpezzi, and Mayo (2005); Quigley and Raphael (2005); Ihlanfeldt and Mayock (2014); Wheaton, Chervachidze, and Nechayev (2014); and Albouy and Ehrlich (2018), test the explanatory power of measured regulations on estimated outcomes such as supply elasticities. This paper includes such an exercise at the jurisdictional level, using three previously published measures of local land use regulation and focusing on jurisdiction-level implications.

Land Use Institutions

Urban land use in the United States is determined by a combination of technological, economic, and regulatory institutions operating within a context in which homes are simultaneously a key lifestyle good and the dominant asset in most middle-class portfolios. Land use decisions in the United States are affected by tax incentives and mortgage subsidies in favor of single-family, owner-occupied homes and, above all, by “Euclidean” zoning (Fischel 2015) that severely constrains the choice set of most landowners and developers. The financial incentives that confront local governments lead them to distort the market in favor of offices and retail buildings, which contribute to the local tax base without requiring extensive service expenditures (Estill and Means 2018). Low-income families with young children are the costliest potential residents for local governments, and officials zone accordingly (Hanushek and Yilmaz 2015). Perhaps as a consequence, office and retail vacancy rates are persistently double the vacancy rates in commercial apartment buildings (National Association of Realtors 2018, exhibit 3.2).

Even where external constraints are not imposed, the sequential interactions among the “stack” of institutions required to create new housing militate in favor of familiar forms. Housing construction at any significant scale involves institutional investors, land assembly, private and public planning, regulatory approval, general construction contractors, subcontractors, real estate agents, mortgage-issuing banks, reinsurers (private or public), and—finally—retail customers. Deviations from proven concepts require that each actor in the stack be persuaded to join the innovator in taking a risk. It is no wonder that even modest innovations in the structure of suburban neighborhoods are difficult to spread, as noted by Duany, Plater-Zyberk, and Speck (2010, 216). Understanding the effects of regulatory constraints requires placing such constraints in the context of this broader institutional structure.

Zoning and other land use rules are a hybrid stock-flow system. Zoning codes, which in most cities are overhauled infrequently,[4] set limits on the stock of housing: each parcel has a maximum density that it cannot exceed without receiving a variance. But the functional approval process is often decided on a case-by-case basis, requiring bureaucratic time and political will

---

[4] For example, the author’s home city recently reexamined its entire zoning code for the first time since 1958.

(Manville and Osman 2017). Some jurisdictions explicitly cap the flow as well as the stock of housing; Dain (2005) details the annual, and even monthly, permit restrictions of 28 Massachusetts municipalities.

A stock model of zoning implies that a high-demand metropolitan area will quickly reach its zoned maximum and then cease growing. That situation may apply to a few select neighborhoods, but no metro area—even those famously stingy with permits—has ceased net growth in a high-price environment. A granular examination of restrictive cities shows that most census tracts experienced some positive construction flows every year. Even in the restrictive San Francisco metro area, 73 percent of census tracts experienced positive growth over seven years, and 29 percent added at least 2 percent to their housing stock.

In cities that are clearly constrained yet growing, regulatory institutions serve as a bottleneck through which piecemeal permitting must pass. As Morrow (2013) shows (see figure 2), the estimated zoned population capacity of Los Angeles fell precipitously until it reached 114 percent of the population around 1990. Then it reversed its decline and began growing, keeping pace with actual population growth. Qualitative study does not indicate that LA became pro-growth after 1990; rather, the LA government uses a very tight zoning code to force most new developments through a variance process that officials can control (Brasuell 2013).

Figure 2. Zoned Capacity and Population in Los Angeles

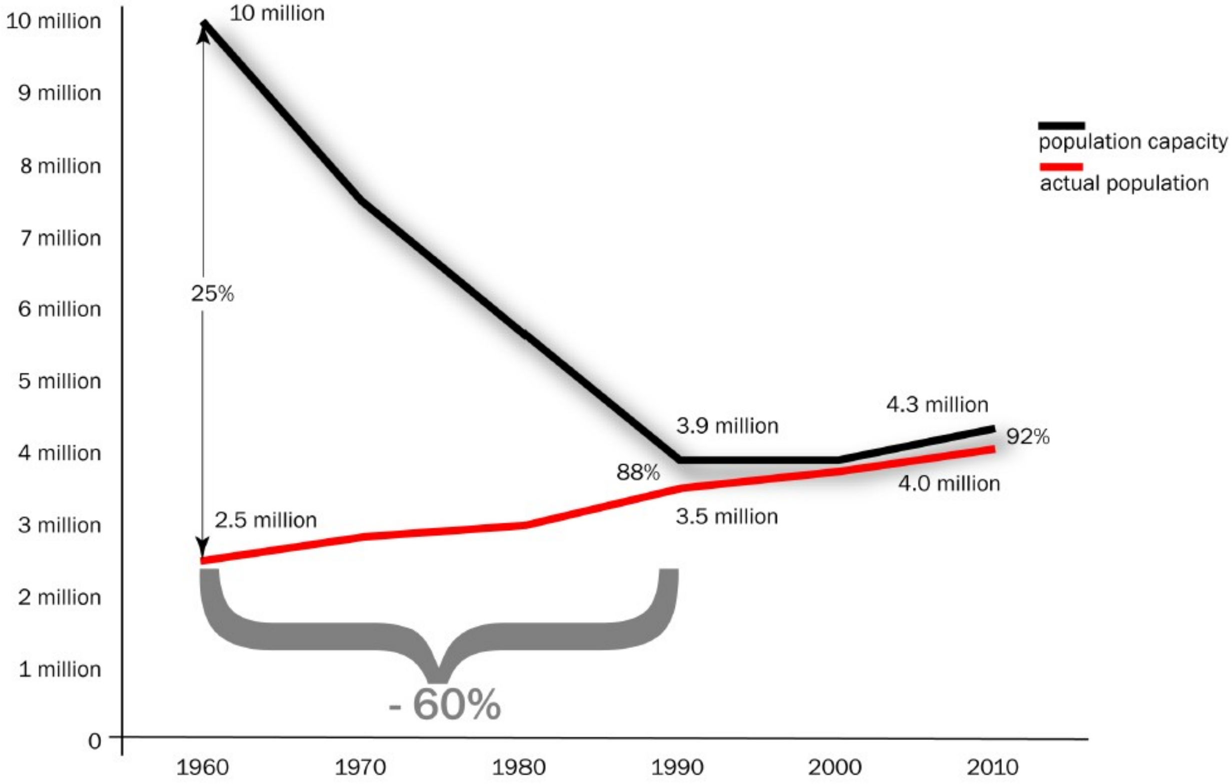

Note: "Zoned capacity" is what the estimated population of the city could have become if every parcel were built up to its maximum legal density. Zoned capacity fell as parcels were downzoned.

Source: Gregory D Morrow, "The Homeowner Revolution: Democracy, Land Use and the Los Angeles Slow-Growth Movement, 1965–1992" (PhD diss., University of California at Los Angeles, 2013), fig. 1-1: "Down-Zoning versus Population Growth." Data are from the US Census and all 104 Community Plans (cumulative population capacity).

Land use institutions can also be responsive to both levels and changes in size and price. Political pressures against growth mount as growth becomes more disruptive (Dyble 2007); political pressures in favor of growth mount as the profitability of potential projects rises (Molotch 1976) and (in a few recent cases) as renter groups coalesce to support growth for market reasons (McCormick 2017). The result is an institutional constraint, endogenous to prices and growth, that interacts with the markets for construction and land to exhibit an overall elasticity of housing supply. In this I follow Vermeulen and Rouwendal (2007), who make clear that they interpret the long-run supply elasticity as primarily reflecting the “price responsiveness of the body of institutions that supply residential land” (13).

In the following section, I present a simple model of residential land redevelopment that gives predictions for the effects of several different types of regulation.

Model

The supply of housing in a metropolitan area is the sum of the housing supplied in each parcel of real estate in that area. The owners of each parcel may reevaluate their housing supply by taking into account the rental value of existing and prospective land uses and the costs of transaction, demolition, and construction. Because of the sunk costs of durable construction, most parcels persist in the same land use intensity for decades. When parcels are built up, the owner maximizes expected future profits subject to regulatory constraints. The increasing marginal construction cost of building makes the construction of tall buildings suboptimal at most locations. However, explicit or implicit regulatory constraints may bind at much lower densities than the limits imposed by construction costs.

The following model is in the tradition of Alonso (1964), Mills (1967), Muth (1969), and Duranton and Puga (2015), but with substantial departures. A salient feature of urban growth models in this tradition is spatiality, which is dispensed with entirely in my treatment. Unlike adaptations of the urban growth model by Mayer and Somerville (2000b) and Saiz (2010), my approach allows redevelopment of urban parcels and several margins for regulatory impact on the market. The cost of slackening these assumptions is that my model offers no closed-form solutions for city size or rate of growth.

Construction

Consider an unregulated metropolitan area, K, composed of parcels of fixed (but not necessarily uniform) size indexed by k. Each parcel has an existing use (use 0) with an associated expected stream of future net rental income, $r_{k,\tau}^{0}$ for $\tau = t, t+1, \dots, \infty$, which is bounded below by the outside option of leaving the property vacant. The initial use of each parcel may be residential or otherwise. I assume that a vacant property has zero net rent.

The owner may also redevelop the parcel. Assume that housing is the only use under consideration in parcels where construction takes place. Furthermore, assume that housing units are identical in value and divisible. The expected net rent stream per housing unit, if the parcel is redeveloped, is $r_{k,\tau}$ for $\tau = t+1, \dots, \infty$. In the period of construction, however, a share of the net rent, $CT_k \leq 1$, is lost to construction time. Assume that CT_k is invariant to the number of housing units being built and does not vary over time. Construction time here includes all time during which the parcel can be used neither for its original nor its ultimate use; net rental income is assumed to be zero during the period of construction.

For the moment, let period length be set such that construction time $CT_k = 1$ for all k; this assumption will be relaxed later.

Denote construction costs of h units of housing on parcel k as a differentiable function $CC_k(h)$, denominated in net present value terms and construed to include costs of construction-related transactions, design, physical construction, and the time cost of money spent on these costs. For simplicity, assume that construction costs are time invariant. The financial sector is abstracted here: interest rates and default possibilities are all collapsed into the construction cost function.

Let construction costs rise with density; that is, $\frac{dCC_k(h)}{dh}$ is positive and increasing in h. Informally, one may add that the cost curve becomes steep once the density of housing units requires that building height exceed the limits of conventional softwood framing.

Owner's Problem

Owners are risk neutral, discount future periods at a rate β, and maximize expected profits subject to construction costs and regulatory constraints. Since each owner is too small to affect metropolitan housing supply, each takes prices as given. The owner of parcel k solves a dynamic optimization problem, balancing expected net rent from redevelopment against construction costs and the net rent from the existing use. Choosing not to reconstruct the parcel in the current period leaves open the option value of construction in the following period.[5]

In an unregulated market, the owner solves the following optimization problem:

$$V_{k,t} = \max\left\{V_{k,t}^{wait}, V_{k,t}^{build}\right\}$$

$$V_{k,t}^{wait} = r_{k,t}^{0} + \beta E\left[V_{k,t+1}\right]$$

$$V_{k,t}^{build} = \max_{h}\left\{(1 - CT_k)hr_{k,t} - CC_k(h) + E\left[\sum_{\tau=t+1}^{\infty} \beta^{\tau-t} hr_{k,\tau}\right]\right\} \tag{1}$$

Since $CT_k = 1$ by construction, and letting $h^* = \arg\max_h V_{k,t}^{build}$, this expression can be simplified to

$$V_{k,t} = \max\left\{r_{k,t}^{0} + \beta E\left[V_{k,t+1}\right] - CC_k(h^*) + E\left[\sum_{\tau=t+1}^{\infty} B^{\tau-t} h^* r_{k,\tau}\right]\right\}. \tag{2}$$

---

[5] I suppress the possibility of a second redevelopment of the property in the distant future. Such an option would add mathematical complexity without altering the qualitative results.

The parcel's purchase value is thus equal to $V_{k,t}$, which is the present discounted value of the parcel given expected net rent and construction costs. Construction will only occur where long-term expectations of the return to construction exceed long-term expectations of the return to current use. Let the owner's margin be the difference between building and waiting:

$$M_{k,t} = V_{k,t}^{build} - V_{k,t}^{wait}$$

The distribution of $M_{k,t}$ across a metropolitan area will be an important determinant of where and how building activity takes place. For example, an area experiencing its first boom in decades will have a much larger stock of run-down, old buildings to redevelop than a metropolitan area that has been growing and redeveloping steadily.

Density

Conditional on redevelopment, a first-order condition imposes an optimal density, rising in expected rent level:

$$\frac{dCC_k(h^*)}{dh} = E\left[\sum_{\tau=t+1}^{\infty} \beta^{\tau-t} r_{k,\tau}\right] \tag{3}$$

Regulation may distort the cost or price schedule, or even make the optimal density illegal, in which case redevelopment may take place at a second-best density. There are, however, minimum and maximum densities at which redevelopment can economically occur. Consider a parcel k where the existing land use is h^0 units of housing, so that $r_{k,\tau}^0 = h^0 r_{k,\tau}$ for all τ, and $r_{k,\tau} = r_{k,t}$ for all $\tau > t$. It is an immediate result that any redevelopment must be at a higher density than the initial use and that there is no gain in delaying construction. To justify redevelopment, the net present value of the new density h^1 must exceed the net present value of the existing use:

$$\sum_{\tau=t+1}^{\infty} \beta^{\tau-t} h^1 r_{k,\tau} - CC_k(h^1) \geq \sum_{\tau=t}^{\infty} \beta^{\tau-t} h^0 r_{k,\tau} \tag{4}$$

or, expressed another way,

$$\sum_{\tau=t+1}^{\infty} \beta^{\tau-t}(h^1 - h^0) r_{k,\tau} \geq CC_k(h^1) + r_{k,\tau} h^0$$

If equation 4 cannot be satisfied, then construction does not take place. Since the first-order condition maximizes the left-hand side of equation 4, a nonempty range of profitable densities will always include h^*.

It is further apparent from equation 4 that the inequality is easier to satisfy for lower values of h^0. Thus, among parcels facing the same per-unit rent streams, those with lower original density are more likely to be redeveloped. Equation 4 less obviously imposes a maximum economical density as well. At densities above h^*, the left-hand expression shrinks because of the convexity of $CC_k(h)$.

As in the spatial model (see Duranton and Puga 2015), a uniform increase in log rent within a metro area will usually result in greater sprawl.[6] That is, because of the curvature of the

---

[6] "Usually" because one can imagine cases in which the marginal parcels are predominantly high density for idiosyncratic reasons.

construction cost function, higher prices justify a larger increase in density for low-density redeveloped parcels than for high-density ones.

Construction Timing

The timing of construction activity will be dictated not by changes in parcel value, but by changes in expected near-term rent. Consider a parcel in which $V_{k,t} > E\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} r_{k,\tau}^{0}\right]$. The high value must be justified by construction at some point, present or future. The owner can compare the returns to building h units in the present period or in the subsequent period:[7]

$$-CC_k(h) + E\left[\sum_{\tau=t+1}^{\infty} \beta^{\tau-t} h r_{k,\tau}\right] \lesseqgtr r_{k,t}^{0} - \beta CC_k(h) + \beta E\left[\sum_{\tau=t+2}^{\infty} \beta^{\tau-t-1} h r_{k,\tau}\right]$$

$$\beta h E\left[r_{k,t+1}\right] - (1-\beta) CC_k(h) \lesseqgtr r_{k,t}^{0} \qquad (5)$$

To engage in construction, the owner must choose to forgo rent either in the current or in the subsequent period. Delayed construction can be justified by a rise in redeveloped rent, a fall in existing-use rent, or a proportional rise in both that is sufficient to overcome the fixed wedge of construction costs.

Equation 5 highlights the importance of present-period rent in determining construction timing. If a vacancy temporarily lowers current rent, the time is ripe for redevelopment.

---

[7] The owner's decision will, in fact, be more complex: the optimal density from redevelopment in period $t + 1$ may be different from the optimum in t.

Regulation

Individual landowner decisions, and thus metro-wide outcomes, are distorted or constrained by diverse local and state regulations. The many forms of land-use regulation do not all have the same implications. In this framework, most regulations will be expressed as increases in construction costs, prohibitions on some densities or uses, delays, or rent control. Rewrite the value function of an owner choosing to build, equation (1), to include five potential regulations:

θ^T: regulatory delays

θ^{FC}: fixed financial costs

θ^{VC}: variable financial costs

θ^R: rent control

θ^D: density limits

The first four types of regulation are pecuniary and enter the owner's problem directly, while the last is a constraint. In this case, set the period length long enough that $CT_k + \theta^T < 1$ for all k.

$$\begin{aligned} V_{k,t}^{build} = \max_h \big\{ &(1 - CT_k - \theta^T)(1 - \theta^R) h r_{k,t} - CC_k(h) - \theta^{VC} h - \theta^{FC} + \\ &E\big[\textstyle\sum_{\tau=t+1}^{\infty} \beta^{\tau-1}(1-\theta^R) h r_{k,\tau}\big]\big\} \end{aligned} \tag{6}$$

s.t. $h \leq \theta^D$

It is immediate that every type of regulation lowers the value of $V_{k,t}^{build}$ and thus makes redevelopment less likely, shrinking some margins and pushing others below zero. Density limits lower margins proportionally more in high-rent places and by a fixed amount within areas facing a common rent (and thus a common h^*). Fixed costs shift the distribution of margins down by a constant. Thus, fixed costs can distort (upward) the observed density of new construction, even though they do not lead to intensification of any project, by pushing the margin of small projects below zero. Variable costs, delays, and rent control will shift the distribution of parcel margins $M_{k,t}$ downward proportionately and affect the optimal housing density choice, as one sees from the first-order condition of equation 6 with respect to h:

$$\frac{dCC_k(h)}{dh} = (1 - CT_k - \theta^T)(1 - \theta^R)r_{k,t} - \theta^{VC} + E\left[\sum_{\tau=t+1}^{\infty} \beta^{\tau-t}(1 - \theta^R)r_{k,\tau}\right] \tag{7}$$

One can imagine a regulation that is intended to increase density but that requires higher-than-economical density. Regulators hoping for high-density development could set a density floor $\underline{h} > h^*$, but at the cost of shrinking the number of parcels that can be economically redeveloped. The impact of such a regulation, broadly applied, on overall urban density would depend on the distribution of parcel parameters.

In addition, regulatory delays, by their nature, make housing supply less immediately responsive to demand shocks. This is especially true of transitory shocks, when the delay has a higher opportunity cost relative to the long-term value of the shock.

Regulations of the five varieties interact by influencing different margins of the investment decision. Consider a town where a developer can expect to go through a series of time-consuming zoning board hearings, pay a “planned unit development” side payment, pay “prevailing wages” to construction workers, face an affordable-housing mandate, and adhere to a binding minimum lot size. All five forms of regulation on their own reduce the number of parcels that will be redeveloped. But those owners who are most likely to be dissuaded by an added cost are also those who are most likely to be dissuaded by a delay or a constraint on density; the regulations work on different margins but put pressure on the same investment decision. At the same time, the effects of the five regulations together could stop redevelopments that would have survived any one regulation. Applied across the metropolitan area, each regulation also constrains supply and increases expected future rent, mitigating its own impact. Whether the combined effects are less than or greater than the sum of the separate effects depends on the distribution of parcel parameters.

Supply Curves and Elasticities

The solution to the owner's problem describes a parcel-specific supply curve that can be expressed in units either of housing supplied as a function of the expected rent level or of the increase in housing units supplied as a function of the increase in the expected rent. As shown, a typical parcel supply curve will be discontinuous: at some rent level, redevelopment becomes economical. Figure 3 illustrates supply curves for three parcels facing identical parameters but with differing existing uses. Conditional on redevelopment, the supply curve is described by the convex construction-cost curve. But redevelopment has lower opportunity costs for less intensely used parcels.

Parcel supply curves can be summed to create aggregate supply curves. In the case of the three supply curves shown in figure 3, the aggregate supply curve (on a per-parcel basis) will have an inverted S shape: steeper than the construction-cost curve at low rent growth, flat at two points in the middle of the diagram, and equal to the construction-cost curve when rent growth justifies development of all three parcels. In practice, massive redevelopment is rare, so observed supply curves of aggregated areas will only include the bottom and middle of the inverted S curve.

Figure 3. Parcel Supply Curves

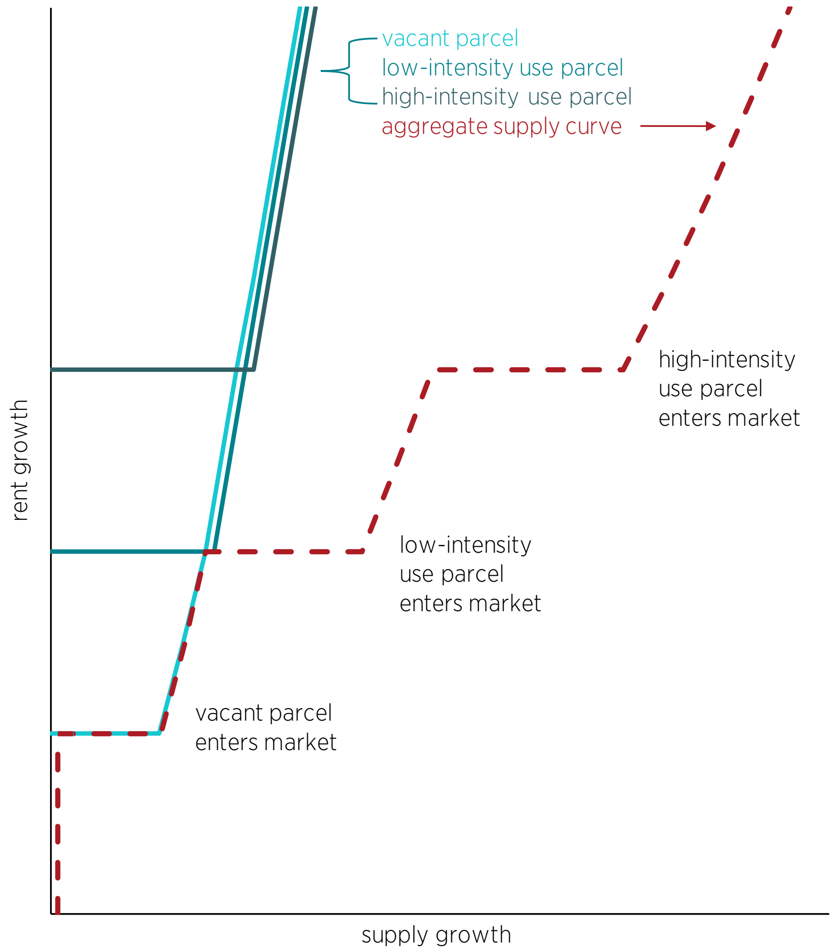

Note: This figure presents supply curves for three illustrative parcels and the aggregate supply curve, which is the horizontal sum of the three. When rent growth is low enough, none of the parcels are redeveloped. At higher rates of rent growth, some or all of the parcels may be redeveloped, starting with the least intensely used. The intensity of redevelopment in each parcel rises with rent. The parcel lines are slightly offset for illustrative purposes.

The shape of the supply curve clearly rules out a constant elasticity of supply. Observed supply curve slopes will depend on demand growth and the existing intensity of use. In the

Results section later in this paper, observed supply elasticity is reported at the metro level and denotes the logarithmic increase in the number of housing units divided by the increase in median gross rent over the same time period:

$$\epsilon^{S}_{K,t_0,t_1} = \frac{\ln(H_{K,t_1}) - \ln(H_{K,t_0})}{\ln\left(r^{gross}_{K,t_1}\right) - \ln\,(r^{gross}_{K,t_0})} \tag{8}$$

The expected relationship between observed supply elasticity and initial price level is ambiguous. Note that ϵ^{S}_{K,t_0,t_1} is defined in terms of current gross rent, while owners in the model respond to expected net rent. Since some costs of providing housing are fixed, an equal increase in log gross rent will result in a larger increase in log net rent in ex ante cheaper places. At the same time, as Glaeser and Gyourko (2005) and others have noted, cities with home prices below construction costs display low elasticities of supply, because housing cannot be readily removed from the market.[8]

Thus, a parcel-level model of redevelopment implies that aggregate supply growth will increase with demand growth and decrease in baseline density, but in nonlinear ways. In the Results section, I estimate a nonparametric relationship between supply growth, demand growth, and density.

---

[8] Closed-form solutions for supply elasticities in some previous papers (such as Saiz 2010) have relied on the assumption that metropolitan areas are growing and have their growth limited by construction costs on vacant lots in every period. But when such an area is recovering from a period of decline, its prices can rise substantially without inducing much construction. The author is fond of joking that in grad school he bought a $200,000 structure for $28,000, on land that was therefore worth −$172,000.

Demand, Spot Market, and Equilibrium

Households and landlords contract for period rental services in a spot market. Demand is an exogenous schedule of households willing to rent housing units in the metropolitan area, with demand decreasing as rent increases. Housing units in the same metropolitan area are substitutes, with local amenities (including spatial factors such as commuting time) constant across time and log-additive in price. In implementation, I will assume that demand has a constant elasticity, ϵ^D, that is common across metropolitan areas and over time.

Landlords can offer housing at the spot price, withhold it vacant if the rental rate is below the variable costs of providing rental services, or take it off the market to redevelop the parcel.

In order to close the model and define an equilibrium, one must specify landlord expectations of future rent levels. An equilibrium is characterized by a period market rental rate that induces a vacancy rate and a distribution of redevelopment across the universe of parcels.

Aggregate Effects of Regulation and Model Predictions

In addition to their localized effects, regulations can have equilibrium effects. Since the first-order effect of all five noted forms of regulation is to decrease supply, one should expect vacancies to fall and prices to rise as a result of higher regulation. If parts of a metropolitan area are especially constrained, owners of other, less constrained parcels will respond by increasing their own densities to take advantage of the higher rent, thus partially counteracting the effects of the regulation.

In a theoretical housing market with marketwide regulations but no density controls, market rent would (given sufficient demand) eventually rise high enough to overcome those barriers. Variable costs, delays, and rent-control regulations dampen responsiveness proportionately, so a rent increase must be proportionately larger to induce the same response. Thus, for high enough rent, these variable pecuniary regulations leave unchanged the responsiveness of the housing market *in proportional terms*. And thus, as in traditional spatial models, variable pecuniary regulations should primarily affect the size and price of a metropolitan area, not its rate of growth. For “expensive enough” metropolitan areas, the existence of variable pecuniary regulations should not affect supply responsiveness. The strong consensus that supply responsiveness is much lower in (and generative of) the highest-cost metropolitan areas strongly suggests that density constraints play a large role.

Speculatively, fixed costs may help explain the phenomenon of “missing middle” housing (see Parolek 2015). Fixed costs will have relatively little effect on the supply responsiveness of large apartment complexes and suburban subdivisions but can, at the same time, be onerous for smaller projects, including “missing middle” housing and single-home redevelopments. Without any metric that can convincingly be construed as a strictly fixed cost, this speculation is not testable.

At the jurisdiction level, however, the model’s implications are completely different. Jurisdictions, except where they constitute a large share of the metropolitan area, are price takers, so one should not expect rent to rise to overcome local variable pecuniary regulations. In otherwise-similar jurisdictions, those with lower costs will experience higher growth. In fact, a jurisdiction with a particular regulatory portfolio should expect to experience more growth if it is among highly regulated neighbors.

The model offers several predictions that can be tested where data are available and housing prices are above replacement level, including the following:

- For a given demand growth, the rate of housing supply growth will decline with density unless regulatory barriers differ systematically by density.
- For a given density, housing supply growth will rise with demand growth.
- For a given demand growth and density, jurisdictions and metropolitan areas with fewer constraints on density will grow faster.
- For a given demand growth and density, jurisdictions with fewer variable pecuniary regulations will grow faster, but the growth rates of metropolitan areas with high-enough prices will not be systematically affected.

Empirical Approach

To take the model to aggregated data, I make four additional assumptions:

1) Landlords believe all rent changes are permanent.
2) Tracts of a given density have a common underlying distribution of parcel densities.
3) Cross-sectional log differences in rent within each zip code are constant.
4) All metropolitan areas face a demand curve with an elasticity of −2/3.[9]

Under those assumptions, the researcher can compute demand growth and identify which tracts had growth greater or lower than expected conditional on demand growth and tract density. I label the resulting residuals "growth surpluses" and interpret them as deviations from the outcome in a tract with national-average institutions among such tracts. The final two

---

[9] The academic literature on the price elasticity of demand exhibits a fairly strong consensus that housing demand is moderately inelastic to unit elastic. In a recent paper, Albouy et al. (2016) argue for an elasticity of −2/3, which I adopt. Papers covering different eras and samples show comparable results, as Albouy et al. show in a review of the literature reaching back to 1857. The argument for unit elasticity was made recently by Davis and Ortalo-Magné (2011), based on the stability of housing as a share of expenditure over time.

assumptions imply that a local demand shock, in quantity terms, equals the metro-area change in occupied residences plus 2/3 times the local price change.

Given the inverted S shape of growth rates across the density distribution in figure 1 and the (perhaps unrelated) inverted-S-shaped expected supply curve, linear regression analysis is likely a poor fit. I instead opt to use nonparametric estimation in two dimensions: density and demand growth. To avoid adding another dimension, I account for changes in vacancy rate mechanically, implicitly assuming that vacancy reduction occurs before demanders ask for any new housing. Parcels with below-replacement costs for housing will respond less to demand shocks, since prices can rise without new construction being worthwhile. Thus I partition the sample into above- and below-replacement cost tracts. To identify areas that are priced below the cost of construction, I adjust Zillow price-per-square-foot data by a cross-sectional measure of construction cost published by the RSMeans Company. To account for housing quality, I also adjust for the number of bathrooms per square foot, a bespoke metric for which I heartily thank the economists at Zillow.

The size of the dataset makes computation costly, so I assign each tract to a histogram bin, using 200 weighted-quantile bins for density and 40 for demand. The above-replacement and below-replacement partitions are estimated independently. Each bin's growth rate is estimated from a two-dimensional additive Epanechnikov kernel. Tracts in one's own bin thus have the highest weight, and nearby bins further contribute to the estimates. For each tract, other tracts in the same Combined Statistical Area (CSA) are excluded from the estimation.[10] At these parameters, the estimation takes about four hours to execute on a personal computer. At that computational cost, using an iterative approach to optimize the estimation parameters was

---

[10] The exclusion of own-CSA tracts may sound computationally costly, but it can be implemented with a one-time subtraction.

prohibitive, so I used rules of thumb and a few trials at different parameters to settle on the preferred parameters.

Having predicted an expected growth rate from similarly dense tracts facing similar demand shocks in other metropolitan areas around the country, one can compute the number of new residences built in excess of expectation. Surplus growth is actual minus predicted growth in each census tract. Local surpluses can be summed across geographic areas, such as jurisdictions, metropolitan areas, or deciles of density within each metropolitan region.

In policy analysis, surplus growth can help evaluate regional approaches to growth. For example, the metropolitan areas of the Pacific Northwest have attempted to discourage fringe expansion and encourage urban density—have they succeeded? Are Texas metropolitan areas, in fact, paragons of land use permissiveness, or do they merely have lots of low-density land to build up?

Data

My units of analysis are primarily census tracts, areas of varying physical size that include 1,700 residential addresses on average. I take core-based statistical areas (CBSAs) to represent housing markets. Table 1 gives summary statistics for key variables.

Table 1. Census Tract Summary Statistics

	10th percentile	Mean	90th percentile	Std dev	N
Residential addresses, 2012:1	773	1,702	2,737	835	72,595
Residential address growth, *ln*	−2.8%	4.0%	14.2%	11.3%	72,218
Vacancy rate, 2012:1	0%	3.7%	8.9%	4.5%	72,272
Vacancy rate, 2017:4	0%	3.0%	7.9%	4.4%	72,218
Rent, 2012:1	\$931	\$1,522	\$2,311	\$646	57,974
Rent, 2018:1	\$926	\$1,645	\$2,574	\$775	61,916

Notes: All rows except the first are weighted by 2012:1 residences. Rent is Median Zillow Rental Index for all homes, including multifamily homes, and is deflated to 2018:1 dollars using the deflator for personal consumption expenses. There are 73,057 census tracts in the United States. Std dev = standard deviation; N = number of observations; ln = natural log.

Sources: Author's calculations based on Zillow, ZRI Time Series: Multifamily, SFR, Condo/Co-op (dataset), accessed October 2018, https://www.zillow.com/research/data; US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

Quantities

For quantities demanded and supplied, I use the Department of Housing and Urban Development (HUD) Aggregated USPS Administrative Data on Address Vacancies. Since 2005, the US Postal Service has shared its administrative tabulations of postal addresses with the US Department of Housing and Urban Development. The data reached their current form in 2010:2. They are reported quarterly, and addresses are distinguished as residential, business, or other, and as occupied, vacant, or "no-stat." The no-stat category is not clearly defined, but it appears to include post office boxes, buildings under construction, and a large share of addresses in vacation towns. I find that subtracting no-stats from the totals leads to a closer match with the census housing totals in 2010. Investigating one city (Seattle) in detail, I found large, sudden changes in the number of no-stat addresses that occurred mainly in census tracts containing a post office. Large drops in one post office tract's no-stats were matched by large increases in another, nearby post office tract. These patterns suggest that excluding no-stats leads to a closer correspondence between postal addresses and residential units.

For the final quarter of 2017, the data contain 130 million residential addresses, 9.7 million business addresses, and 6.5 million other addresses. "Other" addresses pose another interpretive challenge: the largest concentrations of other addresses occur in some (but far from all) centers of government employment and colleges. A cross-sectional regression of the number of other addresses on the number of residential and business addresses has an R^2 of 0.22 and suggests that a census tract has an additional "other" address for every 3 business and every 42 residential addresses.

To compute tract density, I inflate business addresses by a factor of 5, reflecting the generally larger size of businesses relative to homes. The results are robust to different choices of inflation factor, and I would be glad to find a source to discipline this parameter. The density of residential addresses is positively related to the density of nonresidential addresses (the correlation coefficient is 0.68), and unweighted density is correlated at 0.95 with density, as I compute it.

The data also include some administrative discontinuities. For example, every census tract in Montour County, Pennsylvania, had a large decrease in residential addresses and an even larger increase in no-stat addresses in the second and third quarters of 2017 that coincide with a readdressing initiative in the county.[11] Although there are few such instances large enough to draw a researcher's attention, it is an open question how much of the tract-level variance in address quantities is because of administrative changes that do not correspond to an economic change.

---

[11] See the Montour County website, "Montour County 911 Re-Addressing Information," http://montourco.org/Pages/911-Re-Addressing-Information.aspx, accessed April 6, 2018.

As its title and previous use in research suggest,[12] the Administrative Data on Address Vacancies is helpful for measuring local vacancy rates. The concept of vacancy here has a clear real-world meaning: no one is picking up the mail. In the context of this paper, one can think of vacancies as the gap between quantities supplied and demanded. The vacancy rates in this dataset are of the same order of magnitude as vacancy rates produced by the Census Bureau in surveys. In the median CBSA, the 2012 vacancy rate was 3.8 percent and the 2018 vacancy rate was 3.4 percent. Vacancy rates tend to be lower and to have fallen more in large metropolitan areas.

Prices

For price data, I use Zillow's rental index, which is published monthly as a three-month moving average at the zip code level. The Zillow index is the median of frequently updated predictions of every housing unit's market rental value, designed to minimize prediction error. Thus, even homes that are never rented are included in the index. Zillow's approach is more sensitive to changing composition in an area than a repeat rent index would be, but less sensitive to the changing composition of which types of units are on the market at any given time. In any case, repeat rent indices are not available for narrow geographies (Ambrose, Coulson, and Yoshida 2015). Zillow also provides home-price and home-price-per-square foot indices, which I use for robustness checks.

Because prices are available by zip code but not by census tract, I use the geographical correspondence files provided by HUD to map prices from zip codes to census tracts. In urban areas, a zip code contains roughly five census tracts, and their boundaries are not coterminous. Even worse, zip codes are not completely stable in their boundaries. The correspondence files

---

[12] The data have primarily been used for vacancy-related research, such as by Silverman, Yin, and Patterson (2013).

are based on the share of each tract's addresses that are in each zip code. Thus, the researcher cannot easily distinguish between a case in which new construction is concentrated in one zip code of a split census tract and a case in which construction is split between the two but part of the tract is reassigned from one zip code to another. A large share of the quarterly shifts in how many addresses of a tract are in one zip code versus another zip code appears to be due to footloose no-stat addresses. HUD and the United States Postal Service could improve this data source by creating correspondence files based only on occupied and vacant addresses (i.e., not "no-stat" ones).

I take three steps to minimize noise in the mapping. First, I use data only from 2012:1 and on, avoiding a major realignment of zip codes that took place between 2011:4 and 2012:1. Second, I average across all available quarters to create a single, invariant crosswalk. Third, I use an algorithm to identify tracts with large zip code shifts that cannot be accounted for by no-stats and assign them no price data, which shrinks my sample by about 5 percent. Growth patterns in excluded tracts do not appear to differ substantially from nearby included tracts (see table 8).

Finding a virtue in necessity, I note that having zip code–level rent data mitigates the bias owing to composition changes in individual tracts. Furthermore, zip code–level rent changes are highly correlated with metro-area changes, consistent with the view that metropolitan areas are coherent housing markets.

Results

In this section, I present some of the results of this outcome-based approach to the data on housing demand and supply. The first subsection reports some descriptive statistics and the metropolitan supply elasticities implied by the assumption that supply curves did not shift in the

2012–2018 period. The next subsection describes growth surplus data at several levels of aggregation. The final subsection compares growth surpluses at the jurisdiction level with several regulatory metrics, reviewing the theoretical predictions made earlier.

Descriptive Statistics and Supply Elasticities

Table 2 shows Zillow data coverage, log housing supply growth,[13] log real rent growth,[14] and supply elasticity for 27 large metropolitan areas. Identification of supply elasticity as the ratio of supply growth to rent growth requires the assumption that supply curves—in the institutional sense intended here—were constant during the relevant period. For the reported metropolitan areas, data coverage (that is, Zillow Rent Index availability) is relatively strong. The patterns that emerge are well known: southern CBSAs grew rapidly; West Coast CBSAs had steep rent increases. It will perhaps surprise readers in the Washington, DC, area that real rent fell there, and housing demand growth was merely modest. For Washington and Virginia Beach (where rent also fell), I do not report the supply elasticity, which is negative under my identifying assumption. Real rent rose nationally during the sample period; this has not always been the case. It is easy to imagine an era when positive growth is accompanied by falling real rent, rendering supply elasticities meaningless. Surplus growth, which this paper introduces, is robust to times and places where prices and quantities have differing signs.

---

[13] In this section's tables, housing growth is measured across all census tracts, not just those for which Zillow Rental Index data are available.

[14] To be internally consistent, rent growth in this case is averaged up from the census tracts included in my sample rather than measured directly.

Table 2. Observed Supply Elasticities

	ZRI coverage (%)	Supply growth (%)	Rent growth (%)	Supply elasticity
San Francisco	94.2	2.8	29.5	0.10
Los Angeles	95.5	2.5	18.8	0.13
New York	79.5	2.6	11.3	0.23
Boston	95.8	3.8	15.7	0.24
Portland	92.9	6.2	23.7	0.26
Cincinnati	90.8	2.3	7.6	0.30
Seattle	97.5	7.7	24.3	0.32
Denver	94.6	8.6	25.8	0.33
Riverside, CA	93.1	3.4	10.1	0.34
Miami	91.6	4.2	10.7	0.40
Atlanta	87.4	5.7	11.4	0.50
Salt Lake City	92.9	8.6	15.8	0.54
Minneapolis	97.6	4.9	8.8	0.56
Phoenix	89.9	6.9	11.9	0.58
Dallas	91.4	9.1	15.3	0.60
Nashville	95.2	10.9	17.2	0.63
Orlando	91.3	10.0	14.3	0.70
Kansas City	90.4	4.7	6.6	0.71
San Antonio	92.8	10.6	13.0	0.81
Chicago	91.2	2.5	3.1	0.83
Charlotte	84.4	10.2	11.6	0.87
Houston	92.6	9.9	9.2	1.07
Austin	94.2	15.9	13.8	1.15
Philadelphia	91.3	2.5	1.9	1.34
Las Vegas	89.3	7.8	5.4	1.46
Washington, DC	93.0	6.6	−0.9	
Virginia Beach	94.0	4.9	−5.3	

Note: ZRI = Zillow Rental Index.

Sources: Author's calculations based on Zillow, ZRI Time Series: Multifamily, SFR, Condo/Co-op (dataset), accessed October 2018, https://www.zillow.com/research/data; US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

As noted in the section Supply Curves and Elasticities, supply elasticities in metropolitan areas where housing prices have fallen below the cost of construction are expected to be low, even when those areas experience a demand increase. Table 3 shows that supply elasticities are low in many rebounding CBSAs in which the average home price per square foot was well below estimated construction cost in 2012. Most of these are small metropolitan areas with ample land availability that show large price increases unaccompanied by growth. Thus, supply elasticity is of limited value as a broad measure of regulatory barriers to growth.[15]

Table 3. Supply Elasticities in Rebounding Metropolitan Areas

	ZRI coverage (%)	Supply growth (%)	Rent growth (%)	Supply elasticity
Salem, OH	100.0	−0.3	4.2	
Muskegon, MI	96.4	0.9	9.3	0.10
Merced, CA	97.7	2.9	20.8	0.14
Sebring, FL	87.9	2.4	15.7	0.15
Madison, IN	100.0	1.1	6.5	0.16
Kokomo, IN	94.6	0.8	5.0	0.16
Del Rio, TX	99.9	1.7	9.7	0.18
Canton, OH	89.9	1.2	4.9	0.25
Clewiston, FL	100.0	4.5	16.2	0.28
Arcadia, FL	100.0	4.6	16.0	0.29
Kendallville, IN	100.0	4.1	14.1	0.29
Bellefontaine, OH	100.0	1.4	4.4	0.32
Ardmore, OK	97.2	1.9	4.4	0.43
Lakeland, FL	95.9	7.6	13.3	0.58
Warrensburg, MO	96.3	1.8	3.0	0.58
Huntington, IN	100.0	2.6	4.1	0.63
Kingsville, TX	100.0	5.7	6.1	0.94
Las Vegas, NV	89.3	7.8	5.4	1.46

Note: ZRI = Zillow Rental Index.

Sources: Author's calculations based on Zillow, ZRI Time Series: Multifamily, SFR, Condo/Co-op (dataset), accessed October 2018, https://www.zillow.com/research/data; US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

Table 4 confirms that downward supply elasticities were usually very small in metropolitan areas where both housing stock and rent declined.

Dense census tracts are unevenly distributed across US metropolitan areas. Among the top percentile of tracts, 78 percent are in the New York metro area. The borough of Queens has more top-percentile census tracts (52) than San Francisco (44), Cook County, Illinois, (38), or indeed any metro area other than New York. Manhattan has 422 top-percentile tracts.

---

[15] In tables 3 and 4, the metropolitan areas are selected for their high ZRI coverage rates.

Table 4. Supply Elasticities in Declining Metropolitan Areas

	ZRI coverage (%)	Supply growth (%)	Rent growth (%)	Supply elasticity
Rolla, MO	100.0	−0.4	−22.8	0.02
Erie, PA	86.8	−0.5	−16.2	0.03
Tallahassee, FL	91.2	−0.1	−2.4	0.03
Williamsport, PA	98.3	−0.6	−17.4	0.04
Olean, NY	99.9	−0.8	−21.5	0.04
Somerset, PA	97.7	−0.9	−14.3	0.06
Great Bend, KS	95.3	−0.4	−5.8	0.08
Brownwood, TX	90.0	−0.7	−7.9	0.09
Duncan, OK	100.0	−0.6	−5.7	0.10
Marion, OH	100.0	−0.5	−4.2	0.11
Vineland, NJ	98.9	−1.1	−8.8	0.13
Lebanon, MO	100.0	−1.5	−4.4	0.34
Bedford, IN	100.0	−1.6	−4.1	0.40
Ashtabula, OH	100.0	−3.4	−3.3	1.03
Jesup, GA	100.0	−6.5	−2.4	2.70

Note: ZRI = Zillow Rental Index.

Sources: Author's calculations based on Zillow, ZRI Time Series: Multifamily, SFR, Condo/Co-op (dataset), accessed October 2018, https://www.zillow.com/research/data; US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

Table 5 shows New York's dominance of the top decile of census tracts, of which it contains over a third. Fortunately, the growth experience in New York's high-density tracts is similar to the non–New York average. As the rest of the table shows, single-metro area dominance is not a concern in the ninth decile or below.

Reversing the lens and examining metro composition, one sees that metro areas have their own dialects of development that result in differing density concentrations. In figure 4, the Boston and New Haven metropolitan areas each exhibit a "hump" at the third decile of density that is common across New England CBSAs. In metropolitan areas like Phoenix, where even working-class neighborhoods were built to accommodate automobile ownership, moderate-density census tracts make up the bulk of the housing stock. Smaller metropolitan areas, such as those of figure 4b, have a higher share of their census tracts in low-density hinterlands.

Table 5. Metro Shares of Each Density Decile

10th decile	%	*5th decile*	%
New York	34.68	New York	4.59
Los Angeles	12.83	Chicago	3.34
Chicago	7.21	Washington, DC	2.71
San Francisco	4.90	Atlanta	2.54
9th decile		*4th decile*	
Los Angeles	12.31	New York	4.33
New York	8.45	Chicago	2.97
Chicago	5.79	Atlanta	2.67
Miami	4.65	Philadelphia	2.60
8th decile		*3rd decile*	
Los Angeles	7.45	New York	3.46
New York	4.32	Atlanta	2.59
Detroit	4.06	Boston	2.24
Chicago	3.57	Philadelphia	2.13
7th decile		*2nd decile*	
Los Angeles	4.23	New York	2.93
New York	4.03	Philadelphia	1.75
Chicago	3.86	Atlanta	1.66
Dallas	3.36	Washington DC	1.58
6th decile		*1st decile*	
New York	4.16	New York	1.08
Chicago	3.79	Washington, DC	1.08
Dallas	3.56	Columbus, OH	0.90
Los Angeles	2.54	Pittsburgh	0.83

Notes: Deciles of Census tracts are weighted so that each has the same number of residential addresses nationwide. Shares shown here are among the included metro areas.

Source: Author’s calculations based on US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

Figure 4. Density Profiles of Select Metropolitan Areas

Panel A. Large Metropolitan Areas ***Panel B. Smaller Metropolitan Areas***

Note: Census tracts are weighted so that each decile has the same number of residential addresses nationwide.

Source: Author's calculations based on US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

Surplus Growth

The concept of "surplus growth" in this paper is explicitly comparative. It does not rely on a production possibilities frontier but rather is intended to compare among existing institutional environments within the United States. To the extent that certain policies are shared almost everywhere in the United States—the protection of established single-family neighborhoods comes to mind—the method will not identify the costs of those policies. However, it may be of more use to policymakers, because every deficit is, by construction, balanced against a surplus somewhere else; some other city has a model that works a bit better. See the section Empirical Approach for details of the surplus growth methodology.

Table 6 breaks out the growth surplus for some large metropolitan areas by density, grouping the bottom four deciles, deciles 5–8, and the top two deciles to form three density categories.

Table 6. Growth Surplus by Density

Deciles:	1–4 (%)	5–8 (%)	9–10 (%)	All (%)
Austin	12.1	4.0	7.4	8.2
Las Vegas	18.8	5.9	0.4	5.9
Houston	12.3	1.6	1.9	5.0
San Antonio	13.4	1.1	0.2	4.9
Orlando	8.4	0.5	4.5	4.0
Dallas	8.5	0.9	2.4	3.5
Phoenix	9.6	0.7	0.5	2.8
Washington, DC	2.8	0.4	4.3	2.2
Nashville	0.9	2.4	18.7	1.9
Salt Lake City	7.0	0.0	1.3	1.8
Denver	6.6	-0.6	2.7	1.6
Virginia Beach	2.4	1.2	0.8	1.5
Miami	1.5	1.4	0.7	1.1
Charlotte	−1.6	2.2	16.0	0.3
Seattle	−6.3	0.2	6.6	0.3
Portland	−4.2	0.3	4.3	0.1
Atlanta	−0.9	0.2	2.5	−0.3
Kansas City	−0.2	−0.8	2.9	−0.4
Chicago	−0.4	−0.7	−0.7	−0.6
Minneapolis	−2.2	−0.9	4.1	−0.7
Philadelphia	−2.1	−0.6	−0.2	−0.9
Riverside, CA	−2.4	−0.3	−1.8	−1.2
Los Angeles	−3.7	−0.4	−1.6	−1.4
Cincinnati	−2.3	−1.2	−1.9	−1.7
New York	−4.6	−1.3	−1.2	−1.8
San Francisco	−7.2	−1.9	−2.3	−2.9
Boston	−7.1	−1.3	0.1	−2.9

Note: Census tracts are weighted so that each density decile has the same number of residential addresses nationwide.

Source: Author’s calculations.

Now we can draw a verdict on the big Texas cities: their growth is higher than expected in every density tranche. Austin in particular has a 4 percentage point growth surplus in the moderate-density tranche, which has very low growth nationwide.

At the opposite end of the scale, the Boston and San Francisco metropolitan areas have growth deficits of almost 3 percent. In Boston's case, the underperformance comes despite a slight growth surplus in its abundant high-density tracts. As noted previously, New England CBSAs have a high share of tracts in density decile 3; the region's resistance to adding housing in those areas helps explain why the Worcester, Providence, Springfield, and New Haven metro areas all join Boston with growth deficits larger than 2 percent. In fact, not a single New England metropolitan area with data has a positive growth surplus.[16]

Boston and San Francisco provide an interesting case comparison across various metrics of supply responsiveness. The two metro areas grew their housing stocks at about the same rate. But demand in San Francisco grew much more, leading to a smaller supply elasticity in the latter (0.10 versus 0.24 in Boston). Meanwhile, as figure 4a shows, Boston has far more tracts with low densities. Table 6 shows that Boston had smaller deficits in all three density tranches of tracts but because of its composition had an equally large overall growth deficit. When evaluating the possibilities for future growth, one should take into account San Francisco's lack of buildable land; when evaluating regulatory institutions, one should not count nonexistent hinterlands against a metro area.

Most major California metro areas performed below average in all groups, with notably poor performance in high-density tracts. Some smaller California metro and micropolitan areas

---

[16] To be fair, the fastest-growing metropolitan areas in the region are Portland, Maine, and Burlington, Vermont, neither of which has rent data.

did have overall growth surpluses, though: Hanford, Bakersfield, Visalia, Fresno, Clearlake, Madera, and Susanville.[17]

Seattle, Portland, and Minneapolis match the "smart growth" prototype, with above-predicted growth in the densest tracts and far-below-expected growth in low-density areas. Many college-town metropolitan areas are similar: Boulder and Fort Collins, Colorado; Bloomington, Indiana; Durham, North Carolina; and Knoxville, Tennessee, all follow the smart-growth pattern.

In many midsize cities, there are very few census tracts in the densest deciles. So the double-digit surplus growth in the top two deciles in Charlotte and Nashville is based on 5 to 10 census tracts each, only one of which is in the top decile. A strategy of urban density in those places requires many more neighborhoods to become dense, not just eye-popping rates of housing growth in a small downtown.

Growth in neighborhoods with established-suburb density was low, and the surpluses and deficits across metro areas were also small in the moderate-density tranche.

Not surprisingly, given the downward stickiness of housing supply, the distribution of growth and growth surpluses has a compressed left tail. Across the 100 largest metro areas with high data coverage, the median census tract had positive housing growth in 98 areas.[18] The highest median tract growth rate (14.3 percent) was in Myrtle Beach, South Carolina; Austin's 8.9 percent was the highest among major cities. The distribution of growth rates also reveals that California has exceptionally expansive growth-resistant areas: the median census tract growth rate in Los Angeles, San Diego, and San Francisco was about the same as in struggling Rochester and Buffalo, New York.

---

[17] Data were too sparse for Crescent City, El Centro, Eureka, Red Bluff, Sonora, and Ukiah.

[18] Ocala, Florida, had a negative median growth, and Pittsburgh's median growth rate was 0.

Construction is concentrated. Even in many fast-growing cities, most tracts experience modest amounts of growth. Median surplus growth was negative in 96 of the 100 large metropolitan areas; Myrtle Beach, South Carolina; El Paso and Austin, Texas; and Boise, Idaho, were the exceptions.

Table 7. Bay Area: Growth Surplus by County

	Housing supply	Growth surplus	
	2018:1	level	% of 2012 supply
Santa Clara County	608,348	147	0
Alameda County	534,242	−17,795	−3.4
Contra Costa County	402,288	−11,655	−3.0
San Francisco County	360,237	−2,174	−0.6
San Mateo County	261,804	−9,072	−3.5
Marin County	95,829	−5,822	−6.1
San Benito County	17,227	−402	−2.5

Note: Housing supply above includes only tracts with complete data.

Source: Author's calculations based on US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

Growth surpluses can be broken out by geographic subentities. Table 7 shows the growth surpluses for the seven counties in the two metro areas that make up the Bay Area. The data show that relative to average US institutions, the Bay Area's suburban counties have been uniquely anti-growth. Although San Francisco is a symbol of anti-housing politics, Marin County's deficit is nearly three times as large as San Francisco's.

Figure 5. Washington, DC, Census Tract Growth Surpluses

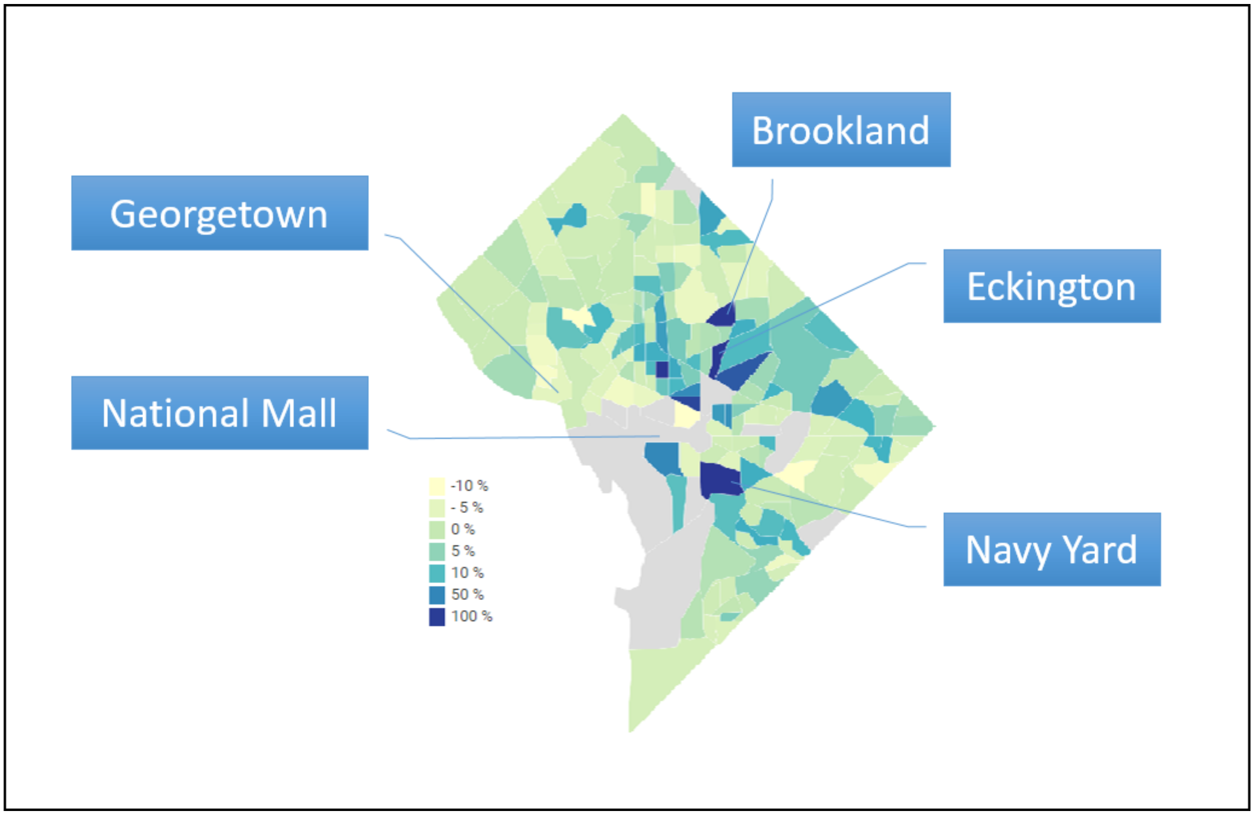

Source: Author’s calculations; map created with Datawrapper.

At the tract level, census tracts show a great deal of localized variation in growth rates, and thus growth surplus. Figure 5 displays growth surpluses for Washington, DC, expressed as a percentage of each tract’s 2012 housing stock. The map also shows that a handful of tracts lack data availability. A few of these are almost exclusively parkland and public buildings; others are places where the zip code boundaries have changed too rapidly in the intervening years to allow confidence in the rent change. The latter were removed during the data-cleaning process described in the previous section. Since zip code changes may be more likely in places undergoing rapid change, there is a reasonable concern that my approach understates growth. In table 8, I compare growth rates of included and excluded tracts in metro areas with good coverage. Excluded tracts show slightly more dispersion in growth rates for most density deciles, with the largest gaps in the top and bottom deciles.

Table 8. Robustness: Growth Rates in Excluded Tracts

	Growth rate			
	Excluded tracts		Included tracts	
Density decile	25th	75th	25th	75th
1	−0.3%	16.7%	1.0%	10.5%
2	−0.6%	13.1%	1.5%	12.9%
3	−0.3%	10.9%	0.7%	12.3%
4	−1.0%	10.9%	0.2%	8.3%
5	−0.7%	7.1%	0%	5.1%
6	−0.7%	3.4%	−0.2%	3.4%
7	−0.8%	2.8%	−0.4%	2.3%
8	−1.5%	2.3%	−0.3%	2.2%
9	−1.6%	2.3%	−0.4%	2.5%
10	−0.4%	7.9%	0%	4.7%

Note: See text for inclusion criteria. Statistics are from metro areas in which total data coverage is at least 80 percent. Growth rates are reported at the 25th and 75th percentiles of the unweighted distribution of census tract growth rates.

Source: Author’s calculations.

An advantage of tract-level analysis is that it largely solves the problem of differing land availability. Saiz (2010) tackled this problem head-on, compiling a measure of buildable land within a 50 kilometer radius of the centers of 95 US metro areas; the tract-based approach implicitly incorporates the same physical barriers that Saiz measured. Census tracts give land area (as distinct from water). And tracts in very lightly inhabited areas—parks or wilderness—tend to be extremely large and to contribute very little to an area’s overall surplus or deficit growth even if they have rent data, which is the case for only 18 percent of tracts in the bottom percentile of density. For example, Census Tract 5135 in Santa Clara County, California, covers almost half the county’s land and includes just 328 residences. Estimation implies that the tract was expected to grow by 12.4 percent, or 40 homes. Instead, it added just 2 homes, thus contributing −38 to the county's surplus. By contrast, the two census tracts composing San Martin and its semi-rural surroundings—an area frozen in a state of half-suburbanization since the 1980s—ran a combined deficit of 363 residences on 35 square miles. In the aggregate, Santa Clara County achieved a surplus (see table 7) despite having a 10 percent deficit among tracts in the bottom four deciles of the national density distribution. Its existing built geography includes few of the exurban tracts where rapid growth has tended to take place, and it found ways to accommodate enough infill growth to offset its relative hostility to exurban expansion.

Reviewing Predictions

I previously gave four predictions of the model, most of which can at least be reviewed with these data, although I do not attempt to formally test them.

To measure regulatory intensity, I use data published by Gyourko, Saiz, and Summers (2008); Quigley, Raphael, and Rosenthal (2008); and Dain (2005). The first source is the Wharton Residential Land Use Regulatory Index (WRLURI). The Wharton survey includes detailed responses to a 2005 survey of regulatory institutions from employees of over 2,000 jurisdictions. I matched all of those jurisdictions to census tracts and have a minimum of 80 percent data coverage in over 1,700 jurisdictions. The second source includes the Berkeley Land Use Regulatory Index (BLURI), which has questions and methods that closely mirror WRLURI. The BLURI response rate was much higher within its narrow Bay Area research region. BLURI was published along with data from other surveys, including project-based surveys of developers and California Environmental Quality Act (CEQA) consultants in the Bay Area. The sample sizes on the latter surveys are very small. I averaged up among projects within each survey to create jurisdiction-specific metrics. Finally, the Pioneer/Rappaport Massachusetts Housing Regulation Database, prepared by Amy Dain and Jenny Schuetz, represents a close study of the published land use regulations of 187 cities and towns in Greater Boston.

Each of the sources includes many other indices and items. To discipline my findings, I selected those that seemed to best align with restrictions as they appear in the model and stuck with the originally selected metrics regardless of the outcome. A data-driven selection process would, no doubt, have yielded more results of statistical significance.

Prediction 1. For a given demand growth, the rate of housing supply growth will decline with density unless regulatory barriers differ systematically by density. Figure 6 shows that this prediction is not, in its unconditional form, borne out by the data. Although growth is fastest in low-density areas, as expected, high-density tracts have higher growth rates than tracts with densities from about the 40th to 90th percentiles experiencing the same rate of demand growth. This pattern is similar across the demand distribution, although it is somewhat noisy. The U-shaped pattern of growth is consistent with the view that established suburban neighborhoods have stronger institutional barriers to growth than do dense tracts.

Prediction 2: For a given density, housing supply growth will rise with demand growth. This pattern holds loosely across all density tranches, although there is some noise. Figure 6 exhibits the pattern: the higher-demand growth scatterplots are largely, but certainly not exclusively, above the lower-demand growth scatterplots. Figure 6 is computed using only tracts in which median price is above estimated construction cost, but the pictures look similar when all tracts are included.

Figure 6. Growth Patterns by Density for Selected Demand Deciles

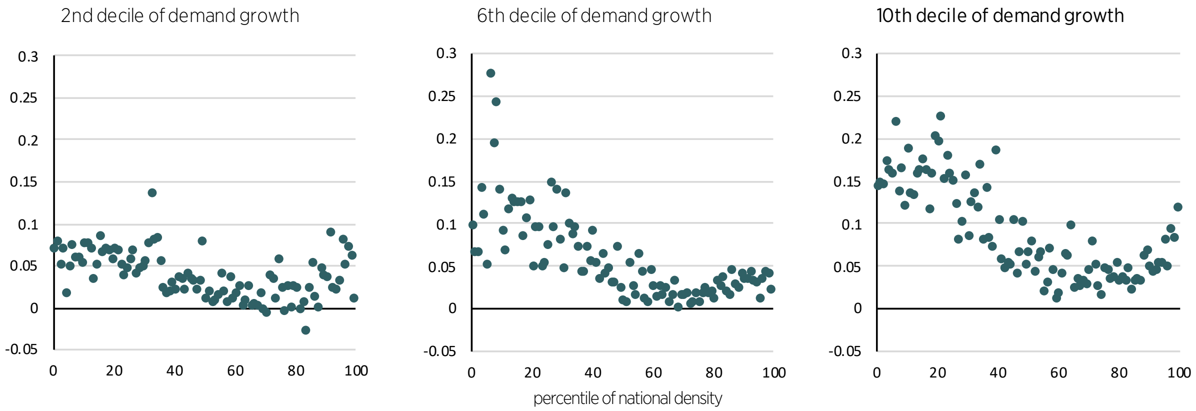

Note: Each panel shows the average growth rate of Census tracts in each percentile of the density distribution that experienced a similar rate of demand growth. The number of tracts underlying each data point differs, since density and demand growth are correlated.

Source: Author's calculations based on US Department of Housing and Urban Development, HUD Aggregated USPS Administrative Data on Address Vacancies (dataset), accessed October 2018, https://www.huduser.gov/portal/datasets/usps.html.

Table 9 uses the jurisdiction-level measures of regulation to evaluate predictions 3 and 4. In each case, I matched jurisdictions to census tracts and computed housing growth and surplus housing growth by jurisdiction, including only tracts with available data and in which the median home price was above estimated construction cost in 2012. Table 9 reports the correlations between each regulatory metric and growth outcomes, as well as the sample sizes, which vary from 13 to 1,319 jurisdictions. Statistical significance levels are reported for the correlations; these levels are not, however, adjusted for multiple hypothesis testing, nor do they take into account data quality concerns.

Table 9. Regulation Metrics and Outcomes: Correlations

		Supply growth	Growth surplus	N
Density constraint				
Wharton	Supply Restrictions	−0.02	−0.0622**	1311
Wharton	Density Restrictions	−0.02	−0.0991***	1319
Wharton	Open Space Index	0.1328***	0.1242***	1271
Berkeley	Constraints index	0.10	0.04	77
Pioneer/Rappaport	Septic rule	0.2983***	−0.1586**	185
Pioneer/Rappaport	Wetlands bylaw	−0.03	0.07	185
Pioneer/Rappaport	Wide front zone share	0.2781***	−0.1782**	185
Delay				
Wharton	Avg Delay Index	−0.0595**	−0.0831***	1290
Berkeley	Avg Delay Index	0.18	0.15	70
Berkeley/CEQA	Time required	−0.28	−0.44	13
Berkeley/Dev.	Time required	0.11	0.15	30
Cost				
Wharton	Exactions Index	0.099***	0.0778***	1270
Berkeley	Infrastructure Index	0.07	−0.05	77
Berkeley/Dev.	Regulatory process cost per unit	−0.09	0.02	30
Rent control				
Berkeley	Inclusionary index	0.14	0.2654**	77
Pioneer/Rappaport	Inclusionary mandate	0.09	0.11	185
Other				
Wharton	WRLURI	−0.0531*	−0.1321***	1252
Pioneer/Rappaport	Growth rate restriction	0.12	−0.04	185
Berkeley/CEQA	Attitude toward development	0.6434**	0.5152*	14

Notes: Correlations at the jurisdictional level, including only tracts with complete data and where 2012 prices were above the estimated cost of construction. * = statistically significant at the 10 percent level. ** = statistically significant at the 5 percent level. *** = statistically significant at the 1 percent level.

Sources: Joseph Gyourko, Albert Saiz, and Anita Summers, "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies* 45, no. 3 (2008): 693–729; Amy Dain, *Residential Land-Use Regulation in Eastern Massachusetts: A Study of 187 Communities* (Boston: Pioneer Institute for Public Policy Research and Rappaport Institute for Greater Boston, 2005); John M. Quigley, Steven Raphael, and Larry A. Rosenthal, "Measuring Land-Use Regulations and Their Effects in the Housing Market" (BPHUP Working Paper No. W08-004, Berkeley Program on Housing and Urban Policy, University of California, Berkeley, May 2008); author's calculations.

Prediction 3: For a given demand growth and density, places with fewer constraints on density will grow faster. The seven "density constraint" regulations and regulatory indices are expected to have negative correlations with surplus growth. Four of them do, all significantly, although none of them with a strong negative association. One measure, the Wharton Open

Space Index, has a positive, significant correlation with surplus growth. Two regulations that show the expected results are Pioneer/Rappaport's indicator variable for town septic rules that are stricter than the Massachusetts statewide rule and their data on residential zones that require at least 150 feet of width. Both rules tend to be adopted in towns with undeveloped land, so they have a positive correlation with unadjusted growth. But adjusting for density and demand, the correlations reverse. However, few indicators exhibit such a neat story, and the aggregate result should cast doubt on any associations within the data.

Prediction 4: For a given demand growth and density, jurisdictions with fewer pecuniary regulations will grow faster, but the growth rate of metro areas with high-enough prices will not be systematically affected. The "delay," "cost," and "rent control" indices and regulations in table 9 are all pecuniary; within a metro area, one expects that growth will be higher in places where the returns to investment face the fewest costs. Of the nine jurisdictional correlations with growth surplus, however, only three have the expected negative sign, and only one of those is significant.

Since the WRLURI data cover many metropolitan areas, one can disaggregate the jurisdiction-level correlations between growth surplus and regulatory indices. For each metro area with at least five WRLURI respondents, I subtracted the metro WRLURI average from the jurisdictional WRLURI. The local residuals of all five WRLURI components are almost perfectly uncorrelated with jurisdictional growth surplus. Table 10 reports the decomposed correlations. In the case of pecuniary regulations, at least, this is the precise opposite of what theory predicts.

If we treat each metro area as a single observation, however, regulatory data are too sparse to meaningfully test the model's aggregate implications. The correlations between metro-level regulatory measures and metro-level growth outcomes are very sensitive to inclusion criteria.[19] Although they mostly align with the predictions, they are not statistically significant and are not reported here.

Table 10. Decomposed Wharton Index Correlations

	Supply growth	Growth surplus	N
Average Delay Index, deviation from metro average	0.08***	−0.02	1054
Exactions Index, deviation from metro average	0.06*	0.04	1038
WRLURI, deviation from metro average	0.17***	0.01	1021

Notes: Correlations at the jurisdictional level, including only tracts with complete data and where 2012 prices were above estimated replacement cost. Metro areas with fewer than five observations in the Wharton data are excluded. * = statistically significant at the 10 percent level. *** = statistically significant at the 1 percent level.

Sources: Joseph Gyourko, Albert Saiz, and Anita Summers, "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies* 45, no. 3 (2008): 693–729; author's calculations.

Given the relatively strong performance of the Pioneer/Rappaport data as a predictor of growth surpluses, researchers may be wise to invest more energy in de jure measurement of regulation. Although published restrictions can be waived and negotiated away, they can be observed by researchers with a high degree of accuracy. The Pioneer/Rappaport data have other

---

[19] Although several papers in this literature use the WRLURI as a measure of regulation at the metropolitan level, metro WRLURI is often based on small samples of municipalities covering small fractions of each metro area's land and homes. Since the within-metro variance of WRLURI is almost as high as the national variance, the use of metro WRLURI estimates that are based on sparse responses seems unwise. To identify metro areas with enough data points to consider them a remotely plausible measure of the prevailing rate of regulation within a CBSA, I considered those that had at least three responses from jurisdictions that together covered at least 30 percent of the housing units and 15 percent of the land area in the jurisdiction; there are only 41 such metro areas. Another eight large metro areas have at least 25 respondents. Taking into account that theory gives strong predictions only for areas where price is above the replacement cost of housing, and that I can only test those predictions where Zillow rent data are available, the testable sample is quickly pared down to 40 or fewer metro areas. Even among those, in some cases I would be testing the implications of regulations that cover a minority of the metro area on growth outcomes in another, partially overlapping minority of the area.

virtues: they cover all jurisdictions in the study area, Massachusetts cities and towns do not overlap, and those jurisdictions have static borders that are usually coterminous with census tracts.

Conclusion

Institutional and physical constraints to construction of new housing play important roles in the patterns of residential growth across the United States. As housing demand rises, rent in coastal metropolitan areas has become a major cause for concern among economists, politicians, and consumers. There and elsewhere, the slowest-growing metropolitan census tracts are those with the densities of established suburbs. Relatively rapid home construction took place in urban tracts between 2012 and 2018, and about half of residential growth in the United States could reasonably be considered infill.

This paper develops a model of urban growth that emphasizes the tradeoff between the rent from a parcel's existing use and the rent from the optimal density. Regulations may inhibit the parcel's redevelopment through costs, delays, constraints, or rent control, each of which are shown to affect the density, timing, and occurrence of redevelopment differently. Unlike in models of greenfield development, there may be few or many parcels near the margin of redevelopment, implying that the same demand shock or regulation may have different effects in differently built cities.

Surplus growth, which compares each tract's growth with that of like tracts nationwide, is intended to help uncover and measure the deep institutions governing the allocation of new residential growth in the United States. Unlike measured supply elasticities, growth surpluses naturally take into account the relative availability of buildable and lightly used land and the difficulty that even permissive metro areas face in accommodating rapid demand growth.

Growth-surplus patterns at the metro level comport with expectations, and jurisdiction-level growth surpluses have significant negative correlations with some, but not most, measures of regulation. This finding may merely be due to the datedness of the regulatory data, which were collected in the decade prior to the study period; newer and more complete data are needed to properly study the present epoch in US housing markets.

Patterns of growth surplus and deficit can nonetheless be used to point policymakers and analysts toward local institutional barriers to growth. Identifying and locating growth barriers is a first step toward addressing the crisis of housing affordability.

References

Albouy, David, and Gabriel Ehrlich. 2018. "Housing Productivity and the Social Cost of Land-Use Restrictions." *Journal of Urban Economics* 107: 101–20.

Albouy, David, Gabriel Ehrlich, and Yingyi Liu. 2016. "Housing Demand, Cost-of-Living Inequality, and the Affordability Crisis." NBER Working Paper 22816, National Bureau of Economic Research, Cambridge, MA.

Alonso, William. 1964. *Location and Land Use: Toward a General Theory of Land Rent*. Cambridge, MA: Harvard University Press.

Ambrose, Brent W., N. Edward Coulson, and Jiro Yoshida. 2015. "The Repeat Rent Index." *Review of Economics and Statistics* 97 (5): 939–50.

Brasuell, James. 2013. "Exceptions Rule: The Dirty Little Secret of L.A.'s Zoning Code." *Laws That Shaped L.A.* (blog), KCET, March 25.

Dain, Amy. 2005. *Residential Land-Use Regulation in Eastern Massachusetts: A Study of 187 Communities*. Boston: Pioneer Institute for Public Policy Research and Rappaport Institute for Greater Boston.

Davis, Morris A., and François Ortalo-Magné. 2011. "Household Expenditures, Wages, Rents." *Review of Economic Dynamics* 14 (2): 248–61.

DiPasquale, Denise, and William C. Wheaton. 1994. "Housing Market Dynamics and the Future of Housing Prices." *Journal of Urban Economics* 35 (1): 1–27.

Duany, Andres, Elizabeth Plater-Zyberk, and Jeff Speck. 2010. *Suburban Nation*, 2nd ed. New York: North Point Press.

Duranton, Gilles, and Diego Puga. 2015. "Urban Land Use." In *Handbook of Urban and Regional Economics*, vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, chap. 8, 467–560. Amsterdam: Elsevier.

Dyble, Louise Nelson. 2007. "Revolt against Sprawl: Transportation and the Origins of the Marin County Growth-Control Regime." *Journal of Urban History* 34 (1): 38–66.

Estill, John, and Tom Means. 2018. "Does Land Development Pay for Itself? A Critique of Cost-of-Community-Service Studies." Mercatus Working Paper, Mercatus Center at George Mason University, Arlington, VA, October.

Fischel, William A. 2015. *Zoning Rules!: The Economics of Land Use Regulation.* Cambridge, MA: Lincoln Institute of Land Policy.

Glaeser, Edward L., and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy* 113 (2): 345–75.

Green, Richard K., Stephen Malpezzi, and Stephen K. Mayo. 2005. "Metropolitan-Specific Estimates of the Price Elasticity of Supply of Housing, and Their Sources." *American Economic Review* 95 (2): 334–39.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008. "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies* 45 (3): 693–729.

Hanushek, Eric A., and Kuzey Yilmaz. 2015. "Land-Use Controls, Fiscal Zoning, and the Local Provision of Education." *Public Finance Review* 43 (5): 559–85.

Ihlanfeldt, Keith, and Tom Mayock. 2014. "Housing Bubbles and Busts: The Role of Supply Elasticity." *Land Economics* 90 (1): 79–99.

Liu, Crocker H., Stuart S. Rosenthal, and William C. Strange. 2018. "The Vertical City: Rent Gradients, Spatial Structure, and Agglomeration Economies." *Journal of Urban Economics* 106: 101–22.

Manville, Michael, and Taner Osman. 2017. "Motivations for Growth Revolts: Discretion and Pretext as Sources of Development Conflict." *City and Community* 16 (1): 66–85.

Mayer, Christopher J., and C. Tsuriel Somerville. 2000a. "Land Use Regulation and New Construction." *Regional Science and Urban Economics* 30 (6): 639–62.

———. 2000b. "Residential Construction: Using the Urban Growth Model to Estimate Housing Supply." *Journal of Urban Economics* 48 (1): 85–109.

McCormick, Erin. 2017. "Rise of the YIMBYs: The Angry Millennials with a Radical Housing Solution." *Guardian*, October 2.

Meen, Geoffrey, and Christian Nygaard. 2011. "Local Housing Supply and the Impact of History and Geography." *Urban Studies* 48 (14): 3107–24.

Mills, Edwin S. 1967. "An Aggregative Model of Resource Allocation in a Metropolitan Area." *American Economic Review* 57 (2): 197–210.

Molotch, Harvey. 1976. "The City as a Growth Machine: Toward a Political Economy of Place." *American Journal of Sociology* 82 (2): 309–32.

Morrow, Gregory D. 2013. "The Homeowner Revolution: Democracy, Land Use and the Los Angeles Slow-Growth Movement, 1965–1992." PhD diss., University of California at Los Angeles.

Murphy, Alvin. 2018. "A Dynamic Model of Housing Supply." *American Economic Journal: Economic Policy* 10 (4): 243–67.

Muth, Richard F. 1969. *Cities and Housing: The Spatial Pattern of Urban Residential Land Use*. Chicago: University of Chicago Press.

National Association of Realtors. 2018. “Commercial Real Estate Outlook: 2018.Q1.” Research Division report, National Association of Realtors, Washington, DC, February.

Paciorek, Andrew. 2013. “Supply Constraints and Housing Market Dynamics.” *Journal of Urban Economics* 77: 11–26.

Parolek, Daniel. 2015. “Missing Middle Housing: Responding to the Demand for Walkable Urban Living.” Working paper, Opticos Design, Berkeley, CA.

Quigley, John M., and Steven Raphael. 2005. “Regulation and the High Cost of Housing in California.” *American Economic Review* 95 (2): 323–28.

Quigley, John M., Steven Raphael, and Larry A. Rosenthal. 2008. “Measuring Land-Use Regulations and Their Effects in the Housing Market.” BPHUP Working Paper No. W08-004, Berkeley Program on Housing and Urban Policy, University of California, Berkeley, May.

Saiz, Albert. 2010. “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics* 125 (3): 1253–96.

Silverman, Robert Mark, Li Yin, and Kelly L. Patterson. 2013. “Dawn of the Dead City: An Exploratory Analysis of Vacant Addresses in Buffalo, NY 2008–2010.” *Journal of Urban Affairs* 35 (2): 131–152.

Vermeulen, Wouter, and Jan Rouwendal. 2007. “Housing Supply and Land Use Regulation in the Netherlands.” Tinbergen Institute Discussion Papers 07-058/3, Tinbergen Institute, Amsterdam, July.

Wheaton, William C., Serguei Chervachidze, and Gleb Nechayev. 2014. “Error Correction Models of MSA Housing ‘Supply’ Elasticities: Implications for Price Recovery.” MIT Department of Economics Working Paper No. 14-5, Massachusetts Institute of Technology, Cambridge, MA.