Houston, You Have a Problem: How Large Cities Accommodate More Housing

Anthony W. Orlando & Christian L. Redfearn

September 2022

Abstract

We document the spatial distribution of new net housing units in large and growing metropolitan areas from 1970 to 2010. While fundamental to both urban and housing economics, little is known about how housing supply evolves to accommodate growth as metropolitan areas expand and mature. We find that metropolitan areas with ample proximal greenfields add new housing differently than those whose periphery is no longer a substitute for more desirable locations in the core. Interestingly, we see patterns in "pro-growth" MSAs that echo those of "highly-regulated" MSAs, but lagged several decades. The empirical results point to falling elasticity as a function of growth, suggesting that rising prices are an inherent feature of large and growing metropolitan areas. This secular trend poses potential challenges to many urban housing policies aimed at enhancing affordability.

Orlando can be reached at aworlando@cpp.edu and Redfearn can be reached at redfearn@usc.edu. We thank Lu Han (editor), two anonymous referees, Evan Mast, Jake Wegmann, and seminar participants at the Association for Public Policy Analysis & Management and the Symposium on Housing Policy and Displacement for helpful comments. We are also grateful to Heonuk Ha for valuable research assistance. All errors are the authors' alone.

1 Introduction

In recent years, housing supply has struggled to keep up with demand in large and growing cities. Nowhere has this trend been more striking than California. During the 1970s and 1980s, California provided an out-sized share of the country's new units. This production declined significantly after 1990. Now, it is known more for its growing prices than its growing population. In its place, Texas has taken the mantle as the largest source of population growth in the United States. From recent aggregate statistics, it may appear that these two states are opposites, but aggregate and short-term statistics can be deceiving. We look within metropolitan areas over 40 years to examine where housing stocks are actually changing. In doing so, we uncover several striking regularities. In particular, we see marked differences in how built-out metropolitan areas accommodate growth relative to those with ample inventories of developable land. Moreover, we see all of our sample metropolitan areas evolve toward a similar development pattern as these inventories decline. This finding calls into question a common model used in urban economics to explain the differences in price levels between more- and less-regulated metropolitan statistical areas (MSAs), namely different linear supply curves. These are often treated as exogenous and static. In the short-run, this assumption might be appropriate, but how do MSAs accommodate more housing over longer horizons? Are there really different processes of adding net new units—or rather, one process that evolves over time?

In this paper, we outline an alternative, unified conceptualization that encompasses both "pro-business" and "anti-development" regimes. Our analysis proposes hypotheses that show how these regimes can follow similar patterns despite different institutions. We then test these hypotheses by documenting common empirical regularities across both types of regimes. We do not propose a full theoretical model here. Rather, we describe a brief conceptual framework—based on recent findings that challenge previous assumptions in the literature—to motivate our empirical work, which is the focus on this study.

Informally, we begin with the stylized fact that unmet demand is most-inexpensively delivered on low-cost land at the periphery of the commuting shed, known as a "greenfield" site. This type of development uses low-cost, low-density construction methods. However, in productive and desirable urban areas, low-cost land, especially close to jobs and retail, is quickly consumed, pushing single-family home builders farther away from the amenities that make these urban area attractive. Eventually, this progression reaches a limit where commuting back to these amenities is too costly. At this point, we can say that the greenfield land is effectively "built out," and developers are forced to look inward to more expensive land closer to the core. When this "infill" development encompasses a larger share of new housing supply, the marginal cost of supply a new housing unit will increase, and the elasticity of supply will fall. Thus, even in the absence of different regulatory regimes, a more mature MSA—with more population and more density—will appear to have a steeper supply curve because large and growing urban markets naturally progress in this direction: The supply curve for all such MSAs is convex, bending upward as demand shifts to the right.

This potential dynamic motivates our empirical efforts with five hypotheses that we test using the Neighborhood Change Database (NCDB), tracking Census tracts from 1970 to 2018 in the four largest MSAs in California and the four largest MSAs in Texas. The first two tests focus on how different MSAs—of different sizes and densities—accommodate new housing units. The remaining three tests reveal the evolution of new housing supply over time as these MSAs fill up the available land.

First, we show that larger, denser MSAs build a higher proportion of net new units (NNUs) in denser neighborhoods than smaller, less-dense MSAs. This finding shows that larger, denser MSAs have less proximal developable land to build low-cost housing, forcing developers to face the challenges of infill development in denser neighborhoods.

Second, we find that smaller, less-dense MSAs build most NNUs on the periphery. In larger, denser MSAs, in contrast, the NNUs are scattered throughout the entire MSA, especially the core. This evidence further shows the preference for builders to choose proximal greenfield land when available, and the need to look to infill locations when it is not.

Third, we follow each Census tract over time, and we estimate similar frequencies of transitioning to a higher level of density across all MSAs, indicating that they are all on the same supply curve. The difference between the MSAs is not their ease of growing denser, but rather their initial starting point on the curve. This distinction is more consistent with the convex supply curve we propose than with the distinct linear curves that the literature has previously used to conceptualize the housing supply process in different MSAs.

Fourth, we formalize this change over time by comparing regression results that explain the characteristics of tracts that result in more NNUs. In particular, we contrast the impact of distance from the central business district (CBD), lagged housing unit density, and lagged development activity on NNUs in 1980 and again in 2010. Across all eight sample MSAs—and under two distinctly different regulatory regimes—the impact of distance has fallen dramatically. In 2010, distance from the CBD is associated with fewer new units than in 1980. Lagged density and lagged building activity have also become less of a deterrent to more development. The pattern is consistent with infill development becoming more attractive as developable greenfield land has dwindled.

Finally, we document the progression of smaller, less-dense MSAs as they use up their low-density (greenfield) neighborhoods and eventually turn inward. Over time, their NNU distribution begins to resemble the larger, denser MSAs, building less on the periphery and more in denser neighborhoods—and with fewer and fewer tracts participating in adding to the aggregate housing supply.

This evolution of MSAs over time traces out the housing supply curve as it becomes increasingly steep with growing demand—and foreshadows the challenges that may soon befall smaller, less-dense MSAs as they converge toward their larger, denser counterparts. These findings suggest important policy implications. If high costs are largely a function of the lack of available greenfield land and the resultant need to provide housing via infill, then there may be less scope for regulatory reforms in California to ease housing development at scale. Likewise, the current "pro-business" land-use regimes in Texas are not likely to prevent price rises like those in more restrictive cities as their rapidly growing MSAs have left little proximal greenfields. Indeed, we have already begun to see these price rises in mid-sized cities during the pandemic, when the increase in housing demand clearly encountered some supply constraints (Li and Zhang 2021). Our goal here, however, is not to provide a full theory or direct test of prices, for which we do not have data. Rather, it is a more modest descriptive exercise to understand the location of new housing units, which we believe merits its own separate analysis.

This analysis takes a large and mature literature in a new direction. Since DiPasquale (1999) asked the question "Why don't we know more about housing supply?", we have learned much about housing supply. First, urban economists have found that the cost of new housing supply varies across cities primarily due to land prices, not construction costs (Glaeser and Gyourko 2003, Glaeser, Gyourko, and Saks 2005, Gyourko and Saiz 2006, Furman 2015, Gyourko and Krimmel 2021). These high land values create a "wedge" between the price of new housing supply and the replacement cost of construction. As land becomes increasingly expensive, housing prices will appreciate faster than the rate of inflation necessary to pay for building materials and labor. Second, they have pinpointed two culprits for this wedge: land-use regulations and topography (Quigley and Raphael 2005, Gyourko, Saiz, and Summers 2008, Saiz 2010, Gyourko, Mayer, and Sinai 2013, Gyourko, Hartley, and Krimmel 2021). Regulations restrict the amount of new supply that can be built on a given plot of land and increase the time and effort necessary to undertake new construction. Similarly, topography makes it impossible to build profitably on some plots of land and increases the time and effort necessary to build on others. Together, these factors explain much—but not nearly all—of the variation in housing supply costs and production across cities. Our study does not challenge these findings but rather adds a new factor that, we believe, merits consideration alongside these two cost drivers.

Overwhelmingly, these studies focused on MSAs as indivisible atomic units. Few studies compared neighborhoods within MSAs. Explicitly or (more often) implicitly, neighborhoods were treated as (nearly) perfect substitutes: A housing unit not supplied in one neighborhood would simply emerge somewhere else in the MSA at a similar price point (Glaeser and Ward 2009). Our paper does not accept this assumption a priori. We take the location of new housing supply within the MSA as one of the most important decisions facing developers—and one of the most serious challenges for housing affordability. If there are systematic differences between neighborhoods—for example, between "greenfield" and "infill"—then they will not be perfect substitutes for housing supply, and much (if not most) of the important variation in production costs will be lost by aggregating up to the MSA level.

In questioning this assumption, we are contributing to a new strand of the literature that is beginning to focus on sub-MSA supply elasticities. Baum-Snow and Han (2021) drill down to the most local level with data on commuting patterns in the largest U.S. cities that allow them to instrument for housing demand across Census blocks. They find that housing supply is significantly less elastic in denser neighborhoods closer to the city center. We take this finding as a key starting point in our analysis, as it suggests that developers cannot substitute easily and costlessly between neighborhoods with different densities and distances from the urban core. Orlando and Redfearn (2022) replicate this finding at the county level using a structural vector autoregression approach—and they extend beyond the largest cities to most of the country, including rural areas. Their dynamic approach allows them to estimate changes in elasticities over time, revealing declining supply elasticity across most of the U.S. This pattern is not confined to expensive coastal states like California. Many (more affordable) cities in Texas and throughout the South are exhibiting declining elasticities as well. The fact that these cities have higher initial levels of elasticity may create the appearance of permanent institutional differences in cross-sectional studies, but this longitudinal methodology reveals that they are evolving more similarly to California than previously understood. We take this finding as another key starting point.

Our paper proceeds in several steps. In Section 2, we provide a larger context for these issues by discussing our sample: the four largest MSAs in the two largest states: California and Texas. Section 3 outlines a loose theoretical basis for seeing a common dynamic among the eight MSAs, rather than seeing them as different static supply curves based on more- and less-restrictive land use regimes. In Section 4, we document a set of a common trends by stages of maturation, with MSAs in Texas enjoying high elasticities largely because they have readily available greenfield land. Section 5 reports regression results that show marked shifts in the characteristics of tracts that receive new housing units as the local greenfield supply declines. Section 6 concludes.

2 Motivating Facts: A Tale of Two States

Over the last century, California and Texas have housed an astonishing amount of population growth—but along two very different trajectories. At the beginning of the 20th century, Texas had twice as many residents as California. But by 1934, they were roughly equals, and since then, California has remained the larger state. By 1966, at the peak of its relative growth, California had 80% more residents than Texas. Figure 1 shows the annual population race between these two states from 1900 to 2020. Nowhere is there evidence of a sustained slowdown. Both states have grown steadily—Texas at an annual rate of 1.9%, California at 2.8%—for the last 120 years. Today, California has approximately 10 million more residents than Texas, though the gap is narrowing.

Since World War II, there appear to be two distinct periods of growth: an early period, dominated by California, and a later period of slower growth, in which Texas has begun catching up. Figure 2a shows the first period from 1950 to 1990, with California's population growing 2.6% per year and Texas's population growing 2.0% per year. Around 1990, we see a break in this relationship. From 1990 to 2020, Figure 2b shows Texas growing slightly slower, at 1.8% per year, and California growing significantly slower at 0.9% per year. This structural shift is consistent with historical evidence of the "slow growth" movement in California restricting housing supply through a combination of exclusionary zoning, environmental regulations, and opposition from increasingly empowered local residents (Morrow 2013). It is important to note, however, that population growth has slowed in Texas at the same time, suggesting broader forces at work.

Despite the current contrasts, these two states have long been been magnets, attracting higher population growth than the rest of the country. Since 1960, the overall U.S. population has only grown approximately 1%, half as fast as California and Texas. As a result, these residents now comprise a larger share of the total population than they did in the mid-twentieth century. Figure 3 shows these population shares from 1960 to 2020. Here, the structural break in California is particularly evident. Note, however, that California has not *lost* ground since 1990. On the contrary, it has merely slowed down to the average rate of the rest of the country. Thus, this graph provides no evidence that California has been less desirable or more restrictive than the rest of the country. It simply has stopped *gaining* ground, relative to the other states. Texas, in contrast, has continued gaining ground steadily from 5.3% of the U.S. population in 1960 to 8.8% today. Still, it has a long way to go to catch up to the 12% share currently held by California. At this rate, Texas will not catch up to California until at least 2055.

The narrowing gap between California and Texas is also evident in its major cities. From 2000 to 2019, Los Angeles, the largest city in California, added approximately 497,000 new residents, while Houston, the largest city in Texas, added 2.3 million. For smaller cities, however, the Texas catch-up is less clear. The next two largest cities, Dallas and San Antonio, have added 1.2 million new residents, while California's next two largest, San Francisco and San Diego, have added 1.1 million. As a result, Figure 4 shows that the rank ordering of these cities has remained stable throughout the early 21st century. Texas cities are growing slightly faster, but at this rate, it will take them a long time to change the rank ordering significantly.

In fact, the race between these cities appears to be slowing down. The city-level population growth from 2000 to 2010, shown in Figure 5a, was noticeably faster for most cities than it has been from 2010 to 2020, shown in Figure 5b. In both decades, the fastest-growing cities are Austin, Houston, and Sacramento, while the slowest-growing cities are Los Angeles, San Francisco, and San Diego – and all but the smallest among them (Austin) display slowing growth. Notice that the divide is not cleanly Texas versus California. The smaller California city of Sacramento is growing faster than most Texas cities, just as the smaller Texas city of Austin is growing fastest of all, and the slowest-growing cities are three of the five largest in this sample. In other words, city *size* appears to be at least as strong a factor as the *state* in which it is located.

Could it be that Texas is growing relatively faster not because it is institutionally or culturally different from California? Or rather is it because its cities are simply smaller and have been better able to accommodate new growth using low-cost technology via greenfield development? This possibility motivates much of the following analysis.

3 Housing Supply: A Tale of Two Curves?

How do these large and growing MSAs accommodate new demand? In this section, we consider the conventional wisdom, as embodied in the urban economics literature, and we propose a subtle but powerful reframing to interpret what is happening in California and Texas.

Most of the literature assumes that supply curves in growing cities are linear and differ over time and across cities (Glaeser, Gyourko, and Saks 2006).[1] As Figure 8 shows, holding demand constant, this model results in higher prices—and lower supply elasticity—for markets with steeper supply curves. Implicitly, this model is arguing that different housing markets have different *institutions*, i.e. that they operate according to different production functions resulting in different marginal costs. Researchers have demonstrated that different supply elasticities are, in fact, associated with different legal and institutional environments,

[1] In cities with declining populations, Glaeser and Gyourko (2005) argue for a kinked supply curve, which is vertical when demand shifts to the left. This consideration does not apply to our analysis, which is focused only on growing cities.

evidence that is consistent with this approach.

However, it is impossible to observe the actual supply curve in a market, as we cannot measure the supply response to demand in all states of the world necessary to trace out the curve itself. Thus, the previous literature can only identify the supply elasticity observed *at a particular point in time*. If supply responses vary across the curve as demand fluctuates over time, the standard empirical approach will mistakenly attribute a static observation to a constant relationship, when in fact supply is dynamic and possibly nonlinear.

To be clear, the mainstream view does not claim that supply curves are fixed over time. On the contrary, supply curves can and do change. However, it is precisely this change that leads to declining elasticities over time, as documented by Aastveit, Albuquerque, and Anundsen (2020) and Orlando and Redfearn (2022). Because supply curves are linear, this model can only explain changing elasticities in a given city by changing the slope of the curve over time. There is no way for the elasticity to change significantly *without* graphing a new curve.

Figure 9 proposes an alternative model that is equally consistent with the empirical evidence in the literature. In this model, many different markets can exist on one unified supply curve, which is convex and therefore nonlinear. If the economist measures supply elasticities for two markets with different levels of housing demand, they will measure two very different elasticities, which may be identical to the elasticities of the two linear curves in the old model. Even though the empirical estimates are the same, however, the underlying mechanisms are very different. In this second model, the differences in prices and quantities need not be explained solely by differences in institutional environments. The two markets can face the same production function. However, because they have such different levels of demand, the production function yields very different costs. The elasticity, in other words, changes in response to demand itself.

To be clear, this model does not imply that *only* demand matters. On the contrary, the elasticity of supply is critical. Differences in prices and quantities arise from different equilibria along the supply curve. Therefore, changes in prices and quantities can be attributed to a combination of demand factors *and* the elasticity of supply at a particular point on the curve. As a result, the supply elasticity will change along with the demand curve.

Note that we are referring here to the *levels* of housing demand. Thus, a large, dense city will have higher demand than a smaller, less dense city—and therefore, its equilibrium will exist farther to the right, on a steeper section of the supply curve. This does not mean that the large, dense city has faster *growth* in demand. On the contrary, smaller, less dense cities often have much faster population growth, but they still have a long way to go for their demand curve to catch up to the *level* of housing demand in the larger, denser cities. Thus, our empirical tests will treat size and density, rather than the growth in population, as the principal determinants of the level of demand.

It is also important to acknowledge that multiple supply curves can still exist. There is no reason to assume that *all* cities exist on one supply curve. For instance, different topographical constraints are likely to result in different *degrees* of curvature, as a more constrained city will face a curve that steepens faster as their demand curve shifts to the right.[2] This is another reason to believe that the supply curve exhibits convexity, as a linear curve cannot incorporate such constraints that come into play as the demand curve shifts.

To make this model more concrete, consider the following explanation: Excess housing demand is met either by (a) low-density development on "greenfield" lots or (b) high-density development on "infill" lots. High-density development can be more detached housing on smaller parcels, but typically it means more attached housing, spanning from duplexes to massive apartments and condominiums. This type of development is more expensive for two reasons.

First, new research shows that construction cost-per-square-foot tends to increase as buildings grow in height. Although the marginal cost curve is nonlinear with local minimums,

[2]In the empirical work below, we avoid this complication somewhat by only focusing on Census tracts where housing exists. Therefore, we exclude land where the topographical constraints prohibit development altogether.

the cost per square foot of a mid-size building (4 stories) is significantly higher than a shorter building, and the cost per square foot of a high-rise building (8+ stories) is significantly higher than anything shorter.[3] These increases primarily arise from changes in capital inputs. Taller buildings require more expensive materials for the "substructure" (foundations, slab on grade, basement), as well as more sophisticated technologies for the "services" (elevators, plumbing, energy supply, etc) (Eriksen and Orlando 2022b). Both of these upgrades also require more skilled labor.

In fact, the role of construction costs is so important that it typically accounts for the majority of development costs for multifamily housing structures. This fact was not well-understood until recently because researchers have tended to focus on single-family housing, where the data are more widely available and land value appears to be the primary driver of rising prices (Glaeser, Gyourko, and Saks 2005, Gyourko and Saiz 2006). In contrast, new simulations by Eriksen and Orlando (2022a) show that the high capital-to-land ratio of multifamily housing changes the calculation. Land values no longer dominate the cost of construction materials, even in the most expensive cities. Multifamily construction also is more likely to pay prevailing wages, which increase construction costs by over 8%. Together, these higher capital and labor costs represent a discontinuous change in the production function that must occur when development shifts from greenfield, single-family housing to infill, multifamily housing.

Second, previous research has shown that infill development occurs on more challenging land. In high-density neighborhoods, developers face up to 40% higher costs due to land assembly and holdouts (Cunningham 2013, Brooks and Lutz 2016). The legacy of older, sub-optimal buildings—that once were, but no longer are, the "highest and best use" of the land—create inefficient parcel shapes and high demolition costs (Hornbeck and Keniston 2017). The high density of the neighborhood also indicates highly desirable amenities, which motivate existing landowners to preserve their exclusive access by preventing an influx of newcom-

[3] The costs per *unit* may not always follow these patterns due to different size units, but this does not reflect differences in the input costs of construction itself.

ers. In other words, land-use regulation is endogenous, becoming more stringent as demand increases. We are not the first to note this endogeneity. Most researchers acknowledge it exists to some extent, but their goal is typically to treat this endogeneity as an identification problem to be overcome, rather than a central feature of the market to be explained. One early exception is Pogodzinski and Sass (1994), who argue that land-use regulations "follow the market" by increasing in response to rising prices, rather than vice versa. More recently, Davidoff (2016) shows that the most commonly used measures of supply constraints are so correlated with measures of demand that they cannot be considered independent causal factors. Hilber and Robert-Nicoud (2013) build a theoretical model consistent with these findings, where desirable amenities lead to more development, and owners of developed land have an incentive to restrict future development. Thus, desirable amenities, which generate high housing prices, also generate restrictive regulations. All of these challenges add up to higher costs.

Thus, even in the absence of different political regimes, greater density will necessitate higher-cost development. This conclusion, in turn, implies another: As cities grow more dense over time, their supply elasticity will decline, and the price needed to produce the marginal housing unit will increase. In other words, as the demand curve shifts to the right, it will trace out an increasingly steep—i.e. a convex—supply curve.

If this theory is accurate, we should see two stages in the evolution of a large and growing MSA: First, developers will meet demand with the low-cost option: single-family detached houses on greenfield lots. Suburbanization is the face of this stage. But at some point, greenfield lots will grow scarce. The periphery will be built out, at least as far as construction at scale is concerned. Then, in the second stage, developers will be forced to look throughout the larger commuting shed to meet demand. This second stage will require a switch to high-density infill development, with the higher costs described above.

This two-stage model suggests five natural hypotheses that we can test. The first two hypotheses describe how different large and growing MSAs accommodate net new units

(NNUs):

- **H1 (Density):** Larger, denser MSAs build a higher proportion of net new units (NNUs) in denser neighborhoods than smaller, less-dense MSAs.

- **H2 (Distance to CBD):** Smaller, less-dense MSAs build most NNUs on the periphery. In larger, denser MSAs, the NNUs are scattered throughout the entire MSA, especially the core.

The remaining three hypotheses describe how the housing supply evolves over time in these large and growing MSAs:

- **H3 (Neighborhood Densification):** Because all MSAs are on the same supply curve, their neighborhoods exhibit similar probabilities of transitioning to a higher level of density. The difference between the MSAs is not their ease of growing denser—as separate linear curves would suggest—but rather their initial starting point, or equilibrium, on the convex curve.

- **H4 (NNU Location):** As smaller, less-dense MSAs use up their low-density (greenfield) neighborhoods, they are forced to turn inward, and their NNU distribution begins to resemble the larger, denser MSAs. In other words, they build less on the periphery and more in denser neighborhoods.

- **H5 (Greenfield Disapperance):** As smaller, less-dense MSAs grow larger, the supply of low-density tracts dwindles. Their density distribution begins to resemble the larger, denser MSAs.

In the next two sections, we examine these hypotheses by uncovering where new housing units are located in both California and Texas—and how these patterns are changing over time. A corollary to the hypotheses above, in the context of these two states, is that these patterns should be driven less by the states themselves—representing the distinct political and regulatory institutions—and more by the size and density of the individual cities.

4 Accommodating New Units

To test our five hypotheses, we use the Neighborhood Change Database (NCDB), a longitudinal dataset created by GeoLytics to follow U.S. Census tracts from 1970 to 2018. It reconciles changing tract boundaries to allow researchers to follow each neighborhood as it changes over the entire time period. Specifically, this dataset allows us to focus on the net supply of housing, not amount of activity in housing construction. By using the latest tract maps as our universe, we know that every tract in our data is developable because we know that by the end of the sample period all the tracts are populated. We can assume in 1970 that an empty tract is empty because supply and demand fundamentals have not yet results in construction, not that these tracts have local conditions that prohibit development. Including non-developable land could complicate interpreting the statistical analysis. In this section, we use these data to test the first two hypotheses, documenting how different large and growing MSAs accommodate NNUs.

For this analysis, we focus on four largest metropolitan areas in each of Texas and California. As described earlier, these two states are infamous for their divergent approaches to land-use regulation, but their most important commonality is their growth. All eight urban areas are large and growing, and provide excellent case studies to learn about how they have added housing units over time.

4.1 Hypothesis 1: Density

Where do large and growing MSAs build NNUs? How does the location of NNUs differ between larger, denser MSAs and smaller, less-dense MSAs? Table 1 reports the aggregate housing supply and the average location of NNUs across the four largest MSAs in California and the four largest MSAs in Texas. In the first three columns, we can see that all eight have experienced remarkable growth between 1970 and 2015. Growth of the MSAs in Texas outstrips those in California by a large margin, but the smaller MSAs in California are not far behind.

The second set of columns reports the average density of Census tracts weighted by new units. This calculation reveals what is typical for a tract from the perspective of the new units. Tracts that add no new units get zero weight, while tracts that received more than their share get larger weights. In Los Angeles, the largest, densest MSA in our sample, the new units were located within tracts with an average density of 5,848 units/mile^2–a level of typical density far in excess of the tracts where new units were added among the four Texas MSAs today. The three large, coastal, and constrained MSAs experienced pronounced increases in the tract unit-density of new units. This implies that the typical new unit is now appearing in much denser tracts than earlier. Three of the Texas MSAs and land-locked Sacramento look like they have continued to add new units in low-density tracts, largely at the periphery.

These results are consistent with Hypothesis 1: The average NNU in the MSAs with less developable land is built amid significantly denser neighborhoods. We will show further evidence in Section 5 when we track the smaller MSAs as they begin to converge to their denser counterparts.

4.2 Hypothesis 2: Distance to CBD

Where are these NNUs located spatially within the commuting shed? Again, are there noticeable differences between smaller, less-dense MSAs and larger, denser MSAs? We answer these questions by mapping the NNUs built in Houston and Los Angeles as respective examples of these two categories of MSAs.[4]

Figure 10 maps the tracts of Houston during four windows of development. Each dot represents the centroid of a distinct Census tract, and the color indicates the tract's share of NNUs that are added in each decade. From 1970 to 1980, the map shows a clear spatial pattern of actual unit loss in the core, while the majority of the new units in Houston during

[4]The maps of the other case cites are too many to include, but we will be making them available online.

this period were clustered in the near west suburbs. This suburbanization continues clearly and consistently for three decades. Each decade witnesses further sprawl and a greater share of the new supply further away.

(In 2000-2010, the pattern is less clear. More of the new supply has returned to the inner suburbs, while new net unit growth in the far exurbs slows dramatically. We will return to this change in Section 5 when we test Hypothesis 4.)

Contrast this pattern, in a less-dense MSA, with the maps for Los Angeles in Figure 11. By 1970, much of Los Angeles is already built out. Figure 11a reveals that the majority of the new supply filled in Orange County and the Inland Empire. By 2010-18, the map shows a remarkable diversity of locations for new housing supply—much of it clustered, but not in any regular ring at the periphery. There are large pockets of densification to the west of Downtown Los Angeles, more to the northeast, and even some significant new supply at the very southern end of the region—100 miles from Downtown Los Angeles. Consistent with Hypothesis 2, the NNUs are scattered throughout the commuting shed, rather than on the periphery, as they were for Houston.

5 The Evolution of New Housing Supply

Despite the differences between MSAs documented thus far, they share some important similarities. First, new development appears clustered. This is true of suburbanizing Houston, and it is true of gentrifying Los Angeles. Second, the process moves outward to the edge then turns inward. The edge is, of course, a function of the location of the amenities and the time to commute to them, but in both cases, there is an apparent limit to the distances households are willing to drive. In this section, we document this progression explicitly, using the NCDB to examine Hypotheses 3 through 5 regarding the move toward infill development as MSAs become increasingly dense.

5.1 Hypothesis 3: Neighborhood Densification

What is the probability of a given neighborhood transitioning to a higher level of density? Do these probabilities differ across MSAs? In Tables 2 and 3, we investigate these "transition frequencies" in California and Texas, respectively. The tables contain six sections, corresponding to six time periods: what percent of Census tracts transition from one category to another from 1970 to 2010, 1970 to 1980, 1980 to 1990, 1990 to 2000, 2000 to 2010, and 2010 to 2015. The first column indicates the starting point, and the top row indicates the ending point. So, for example, the cell corresponding with "Exurb" in the first column and "Urban" in the top row indicates the percent of tracts that transitioned from exurban to urban during this time period.

Focusing on the full 1970-2010 time span in both tables, the upper diagonal, representing non-decreasing density, is fully populated with significant transitions. For example, Table 2 shows that 29.7% of rural tracts in California in 1970 remain rural 45 years later, but more than a third of these rural tracts in 1970 become moderate- to high-density tracts, as measured by the sum of the last three columns in the first row. Few tracts transition backwards to lower density. Below these first sets of results, the five one-step transition frequencies by decade show how little rapid change there is: The large majority of tracts remain as is or migrate up one category of density.

Comparing California to Texas, the transition frequencies are remarkably similar, as predicted by Hypothesis 3. This suggests a natural progression as MSAs grow. What *is* markedly different between the two tables is the very clear abundance of low density land in the Texan MSAs. Note, for example, that 23.5% of California tracts are classified as urban as early as 1970 (in the Share column of Table 2), while only 2.1% of Texas tracts begin in the urban category (in the Share column of Table 3). Even as late as 2010, only 6.7% of Texas tracts are urban, compared to 33.8% in California. While the transition probabilities are similar, the starting point for the two populations is significantly different.

The other striking pattern in the tables is the declining transition frequencies, with over

90% of tracts in both California and Texas unchanging in the most recent subperiod, as indicated by the diagonal in the final section of both tables, where the cells indicate Rural-Rural, Exurb-Exurb, and so on. This echoes the finding of long-run rank persistence in housing submarkets (Malone and Redfearn 2018). It is in this way that the notion of a convex supply curve becomes more directly clear. Elasticities are declining because the stock of proximal developable land is dwindling.

Table 4 shows how these two trends—rising density, declining proximal developable land—are linked. The first two columns measure the increase in density as the percent of tracts with more than 2,000 housing units per square mile. In 1970, the majority of all tracts in both Los Angeles and San Francisco have unit densities above this threshold, while none of the four Texas MSAs are close to this level. San Diego started 1970 with a density closer to that of 2015 San Antonio, and by the end of the data, it too has a majority at higher density. Our thesis, however, turns on the availability of tracts at low-density that would afford development of new housing units at scale. In Texas, the four MSAs all have significant tracts that are sparsely developed—half of all tracts in San Antonio and Austin and a third in Dallas and Houston. Our thesis assumes that this inventory of greenfield land is the low-cost option for new housing supply, and to accommodate the significant growth between 1970 and 2015, Texan MSAs consumed a lot of it—with the share of low-density tracts dropping to 5-10% in the bottom four cells of the last column in the table. This kind of tract was already scarce in three of the four California MSAs, and it has become even rarer today.

5.2 Hypothesis 4: New Unit Location

These densification trends are consistent with a pattern of new housing development in which the periphery is used first and then infill when the periphery is built out.[5] Do we see this pattern across all large and growing MSAs? If so, we should expect to see that distance

[5] This does not mean that there isn't developable land at their exurb fringes. Rather, it is associated with a distance at which households will no longer pay enough to justify new development.

to the CBD is positively associated with the location of NNUs when MSAs are smaller and less dense, as they add new units on the fringes, and then the relationship should disappear or flip when the fringes grow too far to expand farther. We should also expect to see that density is a negative correlate in the first stage and be less so in the second stage.

We test these predictions in Tables 5 and 6 with regressions for each of the eight MSAs, dividing their net new construction into the periods of 1980-1990 and 2010-2015. The dependent variable is the number of NNUs. The covariates include the distance to the CBD, the lagged unit density in the same tract a decade earlier, the average income, and the lagged flow of new units from the previous decade.

For all eight MSAs, the regressions strongly support Hypothesis 4. Over time, the coefficients on distance almost all move from positive to negative or insignificant. Consider, for example, the change in the first two columns of Table 5. From 1980 to 1990, the distance coefficient is 96.96; from 2010 to 2015, it is -16.89. This same transition occurs across the rest of that row, switching from positive in 1980-1990 to 2010-2015. Texas is not unique in this respect. The exact same pattern occurs in Table 6 for California and most of its cities. Tracts from the CBD are more likely to receive NNUs when there is ample developable land. But when the inventory of developable land dwindles, NNUs arrive closer to the CBD and/or throughout the commuting shed.

The density coefficients tell a similar story. Although the coefficients are generally negative across the density row in both tables, the size of the coefficients decrease substantially from 1980-1990 to 2010-2015 in most cities. When vacant land was still close, developers sought out tracts with much lower density, but later the correlation became less pronounced. Fewer units are being constructed in greenfield areas where density is very low.

Finally, the path dependence of NNUs declines. During the early periods of suburbanization, NNUs were clustered near past development, measured by lagged NNUs from the previous period. This correlation loses much of its explanatory power over time. Looking at the "new units" row in both tables, the coefficients are more likely to be higher in the

1980-1990 period than the 2010-2015 period. Over time, NNUs are moving to new neighborhoods.

5.3 Hypothesis 5: Greenfield Disappearance

While this evolution is compelling, it focuses only on the changing location of new development, not the land they left behind. This final piece of the puzzle is crucial: It demonstrates whether greenfield land is filling up, driving developers to look throughout the commuting shed. Do they turn to infill development at the same time that low-density tracts are disappearing? This is the key connection underlying the convex supply curve.

In Figure 12, we show this disappearance happening in all eight sample MSAs. Panel (a) reports the distributions of the tract densities for Los Angeles, Sacramento, San Diego, and San Francisco. Clearly visible is the declining shares of the lowest-density tracts from 1970 to 2010. Once developed, the tracts are reallocated from low-density to higher-density tracts. This is visible in the consistent upward movement of the densities above 150 units/square mile. Compared to Texas, however, these California cities began with a much smaller share of low-density tracts. This is evident in Panel (b), where the share of low-density tracts begins much higher and decline more sharply in the MSAs of Austin, Dallas, Houston, and San Antonio. Although the California MSAs are larger and more dense, the progression of the distributions is very similar in both states.

This convergence is clearer in more local data. Figure 13 repeats the same exercise but uses the tracts of Houston and Los Angeles as comparisons. The scant share of low-density tracts in the Los Angeles MSA is obvious, but still there have been clear declines decade after decade. Houston, by comparison, begins 1970 with a large stock of low-density tracts that are consumed via development decade after decade until the share of low-density tracts is almost equal to the share of tracts at 200 units/square mile.[6] Even Austin, which is smaller

[6] This evidence is consistent with Orlando and Redfearn (2022), who find that the supply elasticity declines in Harris County, the central county of the Houston MSA, from 0.32 in 1980-1994, to 0.25 in 1990-2004, to 0.15 in 2000-2016.

and growing more rapidly, exhibits this pattern. In Panel (b) of Figure 14, Austin looks like Houston did only a decade or two earlier. The progression over time is clear in the history of these three cities.

Finally, as a test of size versus regulatory regime, consider Sacramento. As a California city, it is institutionally very different than Texas, yet it has similar size and density to many of the Texas cities. In Panel (a) of Figure 14, it behaves more like the MSAs of Texas than the other three in California. Its density distribution appears to be governed more by size and density—and especially the availability of proximal developable land—than any "anti-development" attitude that may be unique to California.

6 Conclusion

The importance of understanding housing supply has not waned in the two decades since DiPasquale (1999) first prompted the field to investigate it. The inability to produce enough housing has central implications for household quality of life, for firms' location choice and profitability, and indeed for the country's economy, as density has been key to modern growth (Hsieh and Moretti 2019). But much of what we know about housing supply is aggregate and static in nature. This paper documents where these units were being located and how the housing supply process has been changing over time.

We have focused on eight large, growing metropolitan areas in two states that are famous for their differences in culture and regulation. Despite the differences, we see a common dynamic in which growth alone can cause declining housing supply elasticity. We document a slow-moving shift in how the housing stock adjusts to new demand. It appears to be a natural byproduct of today's metropolitan maturation. Regulation alone cannot explain these trends, nor can a reputation of easy development guarantee enough housing supply when proximal greenfield land supply dwindles. This conclusion has important implications for the future of housing affordability as MSAs grow larger and denser. Our findings suggest that Texas is moving more toward California-style outcomes with regard to housing supply—and therefore faces the prospect of rising house prices in the years ahead. Recall that while Houston has been eager to build more housing over the last 20 years, so too was Los Angeles 40 years ago.

Indeed, we may already be seeing evidence of this transformation, as the COVID-19 pandemic has accelerated the rightward shift of the demand curve in small-to-mid-sized cities. These cities have not accommodated this new demand by simply building more housing without marginal cost increases. On the contrary, home price appreciation has been *faster* in these smaller cities than in cities like Los Angeles and San Francisco, suggesting that they are encountering the convexity in their supply curves (Li and Zhang 2021).

The results presented in the paper are largely descriptive but instructive. They suggest that new theoretical approaches are needed to describe the evolving supply process for new housing in large, growing cities, as well as new empirical methodologies that are not reliant on outdated assumptions about exogenous linear supply curves. It will not be appropriate, for instance, to assume fixed effects with regard to a metropolitan area's supply process without some care. Future research should unpack the internal dynamics of metropolitan housing markets, which are neither homogeneous nor static. As our urban populations continue to grow in the decades to come, these dynamics will determine where they reside and, as a direct result, how they—and we all—thrive.

Table 1: Changing Housing Stock and the Mean Density of New Units

State	CBSA Name	Housing Units			Avg. Density for NNUs		
		1970	2015	Growth	1980	2010	Growth
California	San Francisco	1,096,511	1,763,363	61%	1,564	6,138	292%
	Sacramento	292,793	881,401	201%	1,501	1,423	−5%
	Los Angeles	2,927,713	4,539,566	55%	3,210	5,848	82%
	San Diego	425,244	1,180,806	178%	1,788	3,044	70%
Texas	San Antonio	260,601	857,732	229%	746	1,188	59%
	Houston	734,428	2,415,643	229%	1,510	1,551	3%
	Austin	100,705	747,159	642%	1,228	1,008	−18%
	Dallas	804,825	2,586,323	221%	1,587	1,567	−1%

Table 2: California: Transition Frequencies by Decade

Type	Share	Rural	Exurb	Low-Sub	Hi-Sub	Urban
1970-2010						
Rural	20.3	29.7	35.0	22.1	10.6	2.6
Exurb	19.5	1.5	28.0	37.4	23.8	9.4
Low Suburb	16.9	0.1	1.0	35.3	51.0	12.6
High Suburb	19.7	0.0	0.1	2.5	61.0	36.4
Urban	23.5	0.0	0.0	0.0	2.5	97.5
1970-80						
Rural	20.3	65.1	29.5	4.4	0.8	0.2
Exurb	19.5	1.0	56.2	33.4	8.0	1.4
Low Suburb	16.9	0.1	0.7	58.5	37.9	2.8
High Suburb	19.7	0.0	0.0	3.1	80.9	16.1
Urban	23.5	0.0	0.0	0.0	4.3	95.7
1980-90						
Rural	13.8	75.2	22.5	1.6	0.7	0.0
Exurb	17.1	0.5	68.1	25.7	4.9	0.8
Low Suburb	17.8	0.0	0.7	73.9	24.2	1.3
High Suburb	25.0	0.0	0.1	1.3	85.0	13.7
Urban	26.3	0.0	0.0	0.1	1.4	98.5
1990-2000						
Rural	10.6	81.7	16.4	1.5	0.4	0.0
Exurb	14.8	1.4	84.3	11.9	2.2	0.3
Low Suburb	18.1	0.0	1.7	85.5	12.8	0.0
High Suburb	26.8	0.1	0.0	3.5	90.0	6.5
Urban	29.7	0.0	0.0	0.1	1.8	98.2
2000-10						
Rural	9.0	70.6	20.4	5.0	2.9	1.1
Exurb	14.7	2.9	72.5	16.1	4.5	4.0
Low Suburb	18.4	0.0	2.0	80.7	14.3	3.1
High Suburb	27.3	0.1	0.2	2.9	86.8	10.0
Urban	30.7	0.1	0.0	0.0	2.7	97.2
2010-15						
Rural	6.8	97.2	2.8	0.0	0.0	0.0
Exurb	12.9	0.3	97.7	2.0	0.0	0.0
Low Suburb	18.4	0.0	1.2	94.8	4.0	0.0
High Suburb	28.1	0.0	0.0	1.7	95.4	2.9
Urban	33.8	0.0	0.0	0.0	1.2	98.8

Table 3: Texas: Transition Frequencies by Decade

Type	Share	Rural	Exurb	Low-Sub	Hi-Sub	Urban
1970-2010						
Rural	56.0	25.1	36.2	25.8	9.4	3.4
Exurb	21.2	0.5	36.5	39.7	15.1	8.3
Low Suburb	14.4	0.0	6.2	57.7	27.8	8.4
High Suburb	6.2	0.0	2.8	22.7	51.4	23.2
Urban	2.1	0.0	0.0	8.2	26.2	65.6
1970-80						
Rural	56.0	68.6	27.4	3.1	0.9	0.0
Exurb	21.2	1.1	59.5	30.9	5.7	2.9
Low Suburb	14.4	0.0	3.6	69.8	20.7	6.0
High Suburb	6.2	0.0	0.6	9.9	70.7	18.8
Urban	2.1	0.0	0.0	0.0	18.0	82.0
1980-90						
Rural	41.6	72.8	25.1	1.6	0.5	0.0
Exurb	27.4	0.6	67.4	26.0	5.0	1.1
Low Suburb	17.9	0.0	4.7	75.5	16.8	3.1
High Suburb	8.9	0.0	0.7	12.7	76.5	10.1
Urban	4.1	0.0	0.0	0.0	12.6	87.4
1990-2000						
Rural	32.1	76.6	23.2	0.2	0.0	0.0
Exurb	29.3	0.4	79.0	19.7	0.5	0.3
Low Suburb	21.9	0.0	3.6	89.4	6.6	0.4
High Suburb	11.6	0.0	0.3	10.9	81.3	7.6
Urban	5.2	0.0	0.0	0.6	10.3	89.1
2000-10						
Rural	24.6	71.5	23.9	4.5	0.1	0.0
Exurb	31.5	2.6	72.4	21.4	3.4	0.2
Low Suburb	26.7	0.0	3.4	81.1	14.3	1.2
High Suburb	11.5	0.0	0.8	8.5	79.3	11.4
Urban	5.7	0.0	0.6	1.1	9.9	88.4
2010-15						
Rural	18.4	95.5	4.5	0.0	0.0	0.0
Exurb	29.7	0.2	94.4	5.4	0.0	0.0
Low Suburb	30.6	0.0	1.0	94.9	4.1	0.0
High Suburb	14.6	0.0	0.0	2.8	93.8	3.4
Urban	6.7	0.0	0.0	0.0	1.9	98.1

Table 4: Shares of Exurban/Urban Tracts: Percent of Tracts Above/Below Unit Densities

State	CBSA Name	% > 2,000/mi^2		% < 50/mi^2	
		1970	2015	1970	2015
California	Los Angeles	57.5	75.2	5.3	2.4
	Sacramento	9.7	43.2	28.1	6.6
	San Diego	22.8	54.1	16.1	4.0
	San Francisco	50.9	67.5	7.5	1.7
Texas	Austin	4.6	26.6	55.7	10.6
	Dallas	9.0	28.4	35.7	5.0
	Houston	12.7	33.1	32.6	5.5
	San Antonio	11.8	24.9	48.8	10.9

Table 5: Texas Regime Shift Regressions, by Census Tract

Model	I	II	III	IV	V	VI	VII	VIII	IX	X
Time Period	80-90	10-15	80-90	10-15	80-90	10-15	80-90	10-15	80-90	10-15
Subsample	TX	TX	Aus	Aus	Dal	Dal	Hou	Hou	SA	SA
Intercept	103.34	142.66	56.36	122.00	150.19	182.17	-394.18	196.36	-125.78	77.82
	(1.95)	(6.02)	(0.31)	(1.62)	(1.70)	(5.70)	(3.79)	(3.53)	(0.71)	(1.60)
ln(Dist to CBD)	96.96	-16.89	125.62	-7.45	76.91	-21.93	174.97	-21.69	204.87	-15.21
	(7.74)	(3.63)	(2.79)	(0.49)	(3.79)	(3.69)	(7.53)	(2.02)	(4.48)	(1.34)
ln(Density(t-1))	-54.43	-18.80	9.96	-14.74	-56.70	-19.66	-35.23	-28.35	-12.24	-11.85
	(9.10)	(6.82)	(0.41)	(1.66)	(6.12)	(5.55)	(3.25)	(4.57)	(0.60)	(2.07)
Avg Income	0.00377	0.00021	-0.00016	0.00011	0.00579	0.00013	0.00800	0.00023	-0.00681	0.00047
	(5.35)	(2.61)	(0.07)	(0.39)	(5.24)	(1.46)	(6.95)	(1.47)	(3.14)	(2.19)
New Units(t-1)	0.181	0.131	0.039	0.113	0.255	0.089	0.123	0.178	0.387	0.066
	(9.0)	(32.9)	(0.5)	(8.4)	(7.2)	(17.7)	(4.4)	(24.5)	(5.3)	(5.6)
R^2	0.129	0.274	0.048	0.191	0.115	0.213	0.220	0.384	0.270	0.102
df	3,189	3,189	345	345	1,309	1,309	1,068	1,068	452	452

Notes: Tract-level data obtained from U.S. Census Bureau for population density, average income, housing units, and distance to central business district.

Table 6: California Regime Shift Regressions, by Census Tract

Model	I	II	III	IV	V	VI	VII	VIII	IX	X
Time Period	80-90	10-15	80-90	10-15	80-90	10-15	80-90	10-15	80-90	10-15
Subsample	TX	TX	Aus	Aus	Dal	Dal	Hou	Hou	SA	SA
Intercept	148.21	48.79	-12.83	57.90	-284.44	63.01	-61.88	25.63	179.03	48.77
	(5.17)	(5.48)	(0.30)	(4.70)	(3.20)	(1.93)	(0.65)	(1.05)	(2.67)	(2.09)
ln(Dist to CBD)	61.27	-2.98	92.71	-4.04	121.79	-2.80	123.90	0.60	36.17	-3.30
	(10.07)	(1.94)	(11.64)	(1.94)	(5.62)	(0.49)	(5.86)	(0.14)	(2.78)	(0.93)
ln(Density(t-1))	-35.59	-4.41	-24.60	-4.53	24.65	-10.62	-14.63	-1.21	-33.57	-5.50
	(11.98)	(4.37)	(5.58)	(3.21)	(2.22)	(2.77)	(1.31)	(0.41)	(5.17)	(2.31)
Avg Income	2.1e-04	-4.2e-05	1.4e-04	-1.1e-04	8.0e-03	1.0e-04	1.5e-03	-1.2e-04	9.9e-04	7.3e-05
	(0.50)	(1.55)	(0.30)	(3.17)	(3.67)	(0.62)	(0.88)	(1.16)	(1.05)	(1.36)
New Units(t-1)	0.166	0.038	0.057	0.040	0.019	0.020	0.232	0.053	0.124	0.032
	(13.00)	(17.29)	(3.46)	(12.38)	(0.36)	(2.82)	(5.05)	(10.54)	(4.50)	(6.77)
R^2	0.139	0.058	0.136	0.050	0.115	0.043	0.175	0.154	0.131	0.055
df	5,006	5,005	2,920	2,919	479	479	622	622	970	970

Notes: Tract-level data obtained from U.S. Census Bureau for population density, average income, housing units, and distance to central business district.

Figure 1: California vs. Texas Population, 1900-2020

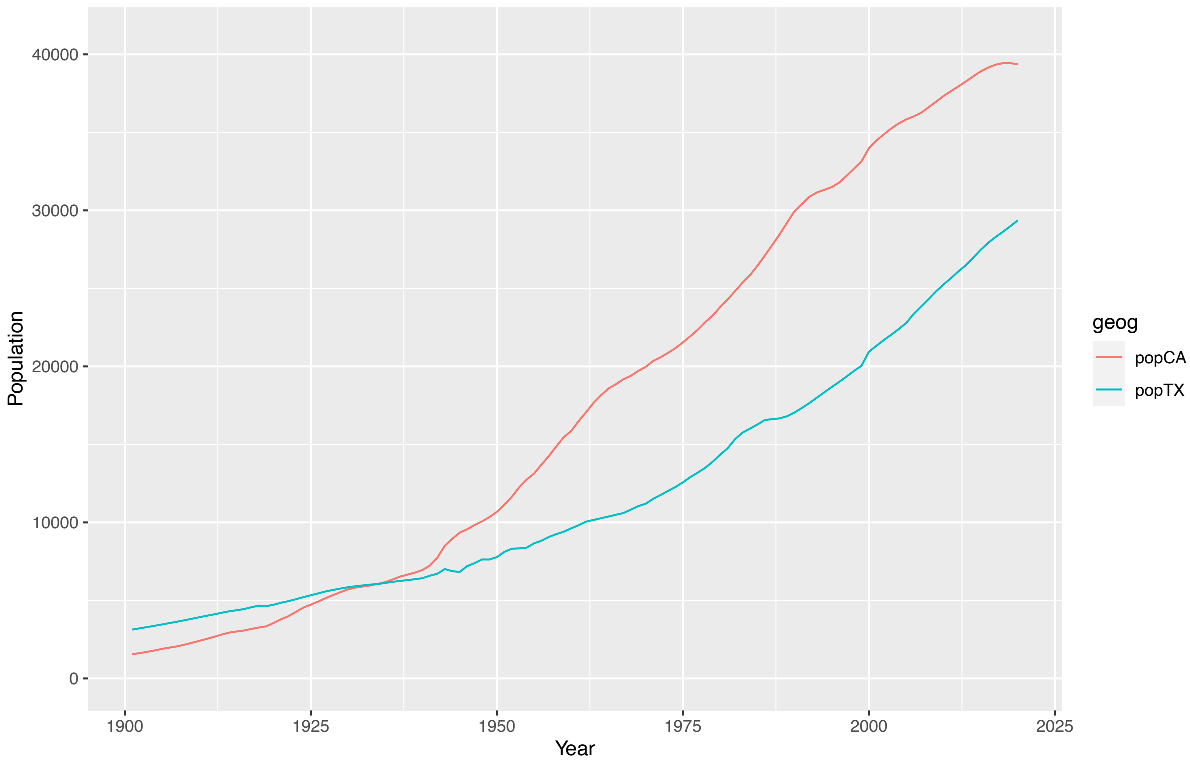

Notes: Population data obtained from U.S. Census Bureau.

Figure 2: California vs. Texas Population Trends, 1950-2020

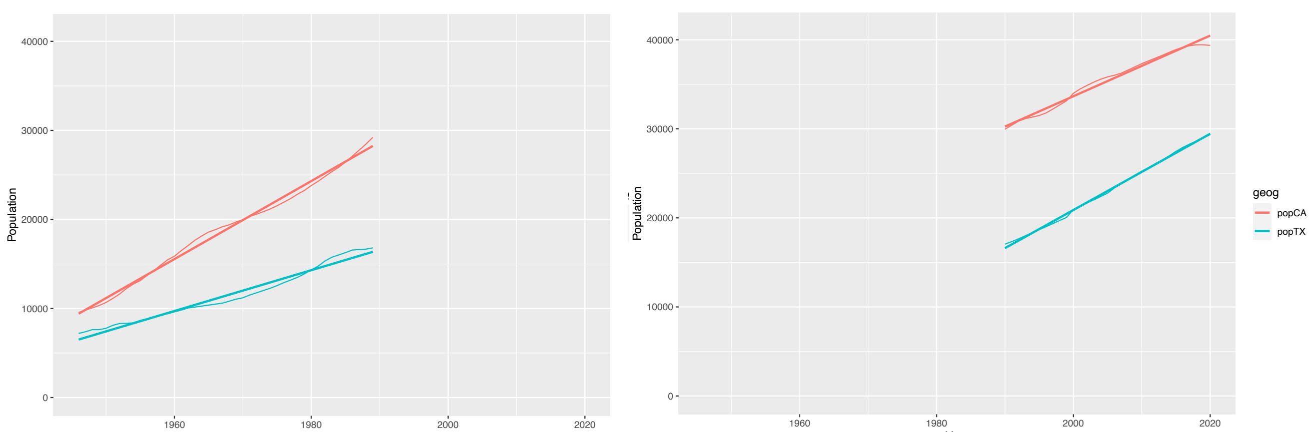

(a) 1950-1990 (b) 1990-2020

Notes: Population data obtained from U.S. Census Bureau.

Figure 3: California vs. Texas Share of U.S. Population, 1960-2020

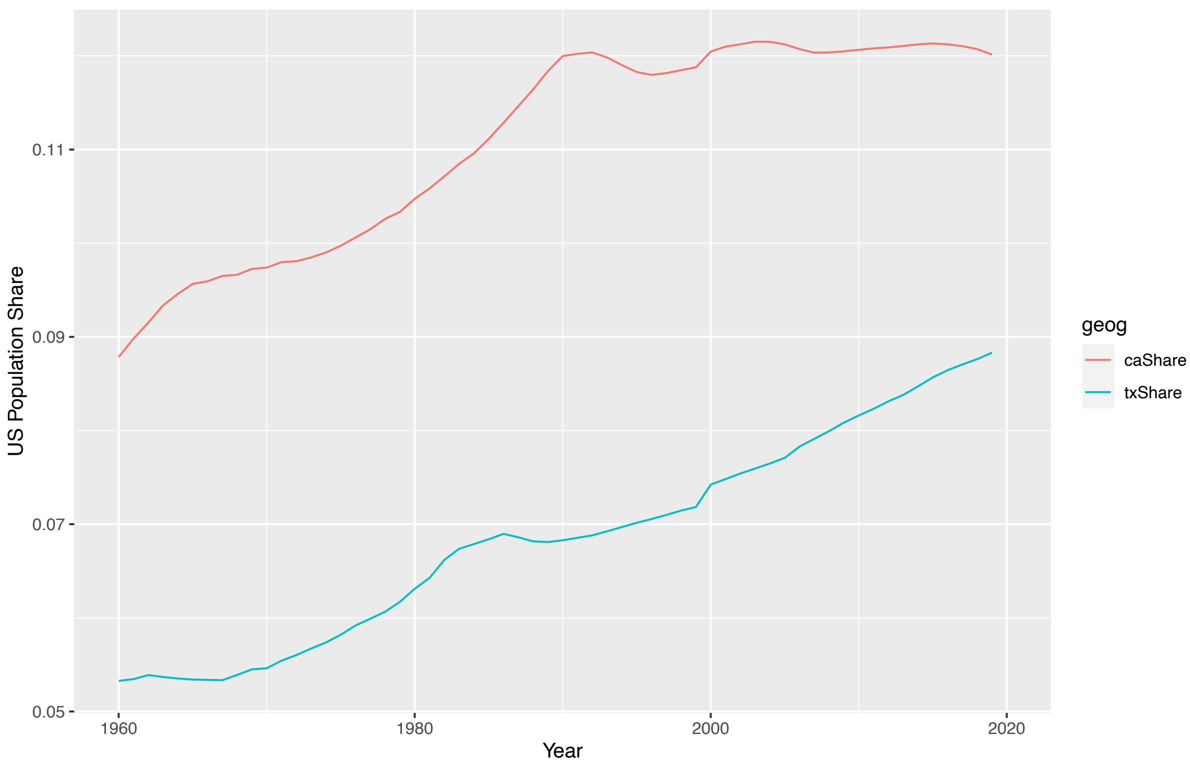

Notes: Population data obtained from U.S. Census Bureau.

Figure 4: California vs. Texas City Populations, 2000-2020

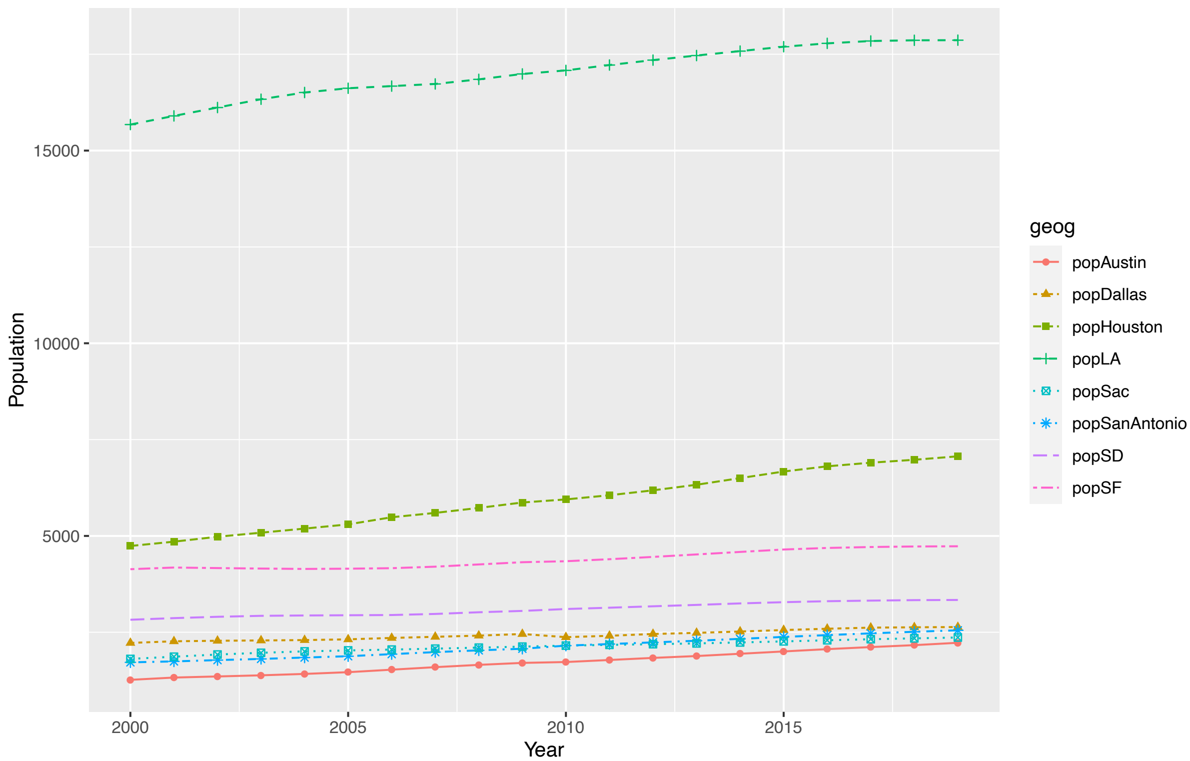

Notes: Population data obtained from U.S. Census Bureau.

Figure 5: California vs. Texas City Population Growth, 2000-2020

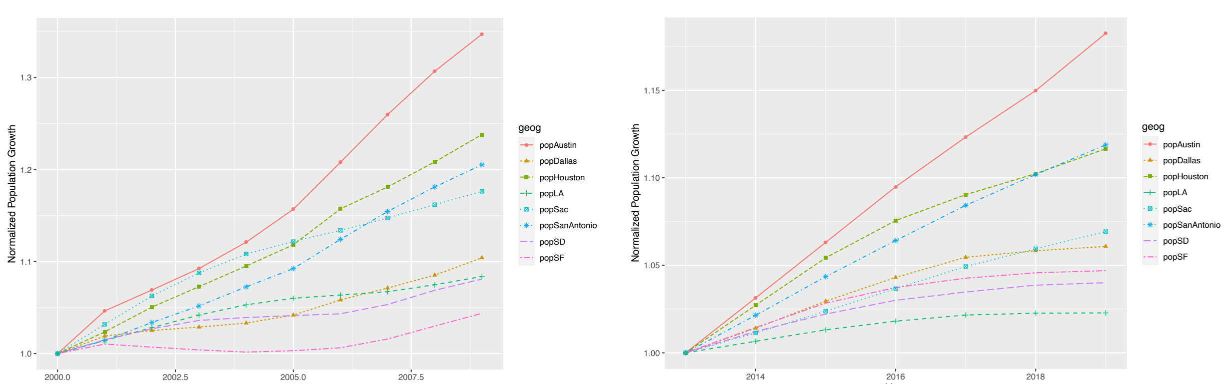

(a) 2000-2010 (b) 2010-2020

Notes: Population data obtained from U.S. Census Bureau.

Figure 6: California vs. Texas City Price Growth, 1996-2006

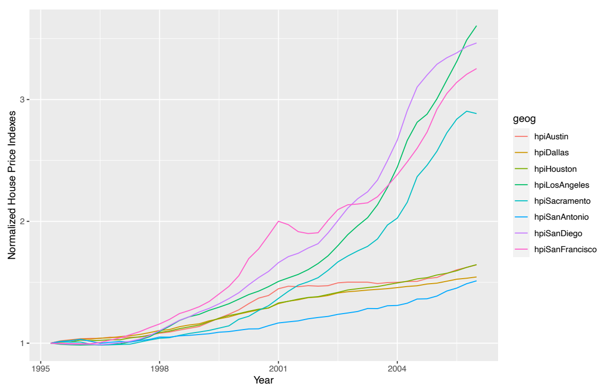

(a) By City

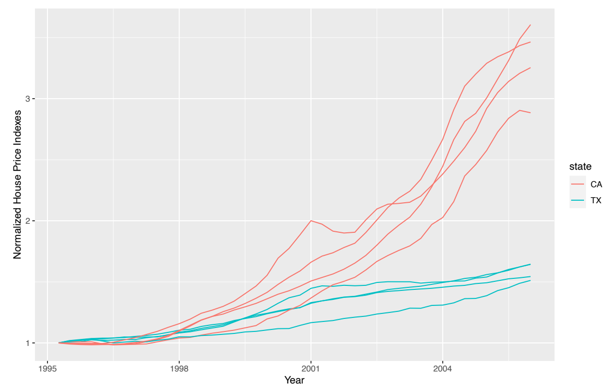

(b) Cities by State

Notes: Price indexes obtained from U.S. Federal Housing Finance Agency (FHFA).

Figure 7: California vs. Texas City Price Growth, 2014-2021

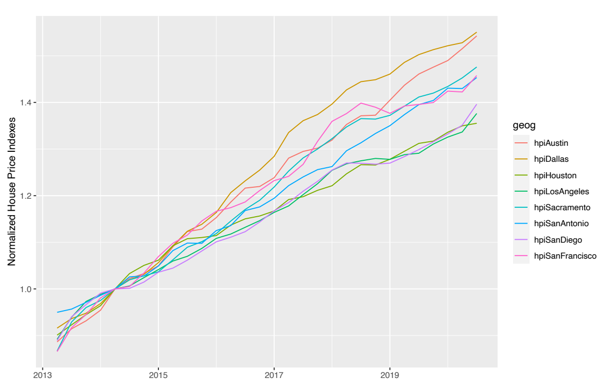

(a) By City

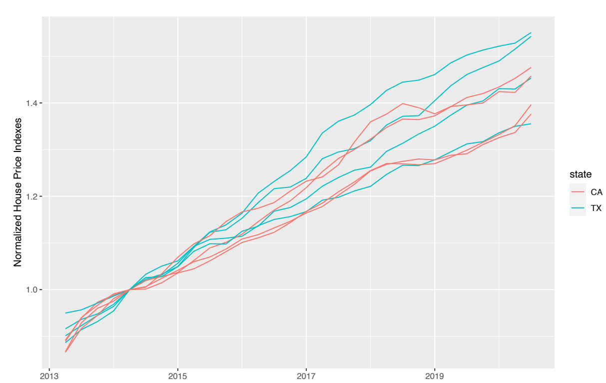

(b) Cities by State

Notes: Price indexes obtained from U.S. Federal Housing Finance Agency (FHFA).

Figure 8: The Old Model: Separate Linear Supply Curves

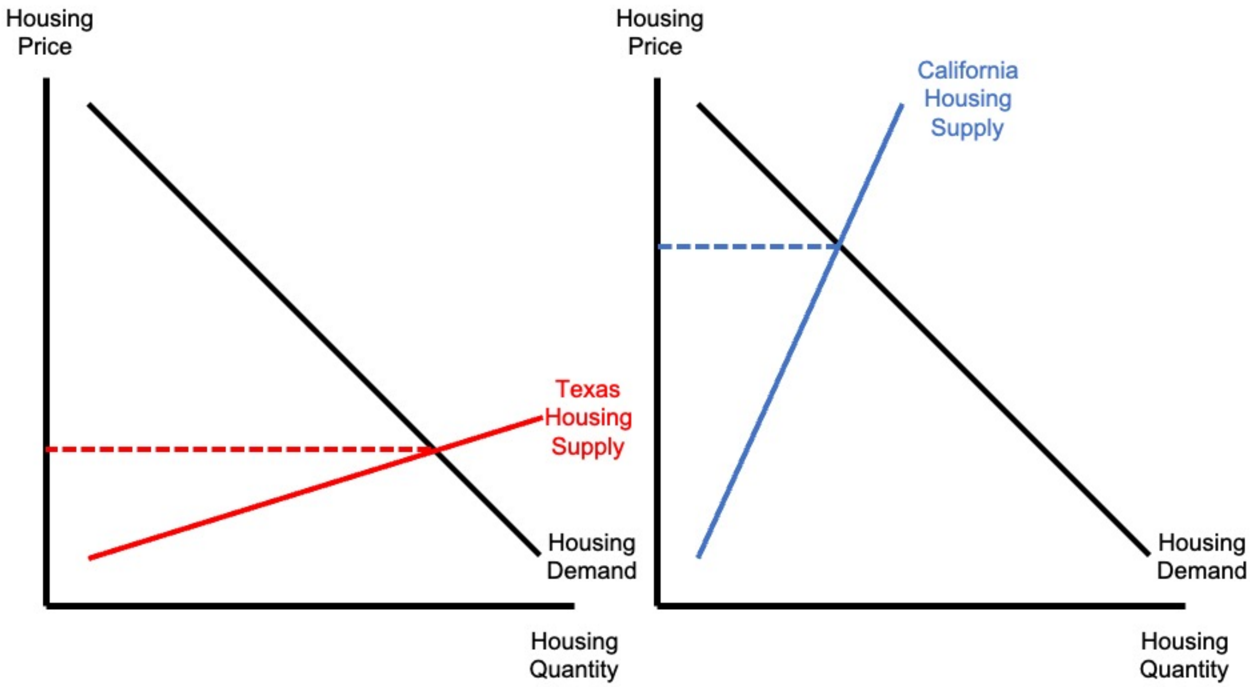

Figure 9: The New Model: Unified Convex Supply Curves

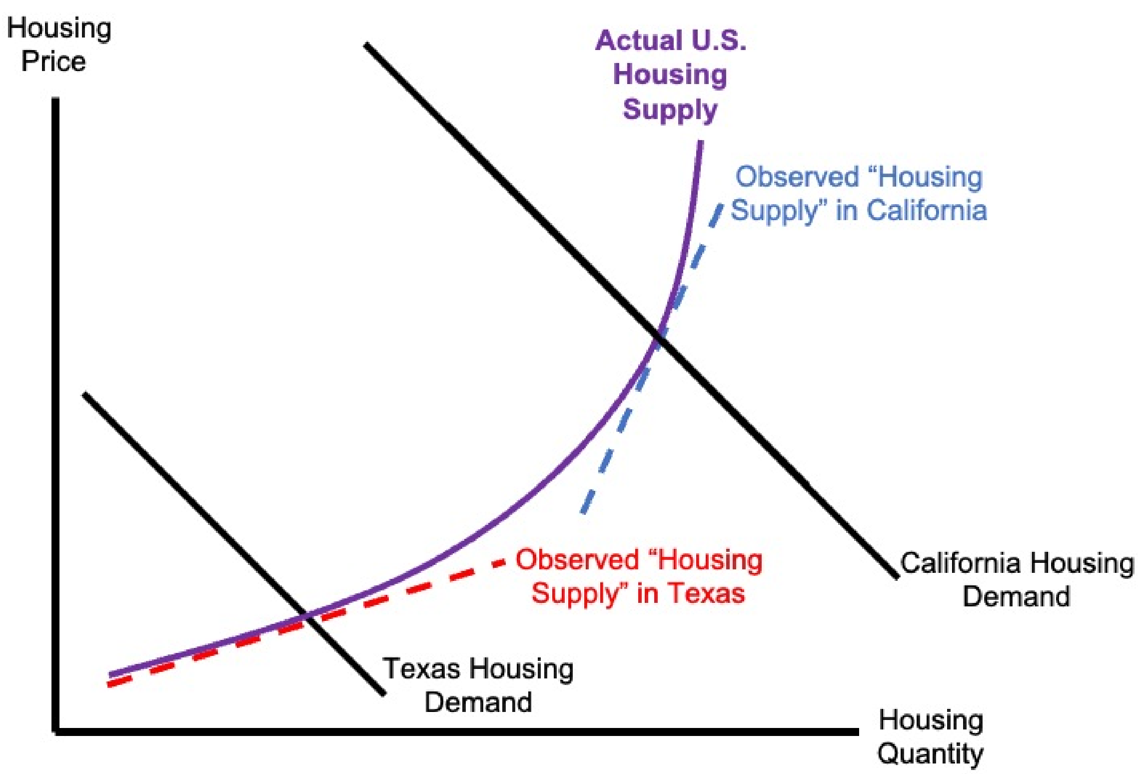

Figure 10: Share of Net New Units in Houston, by Census Tract

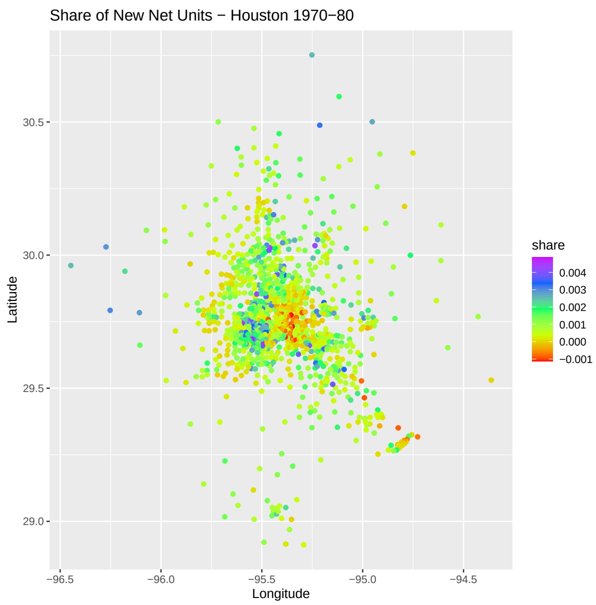

(a) 1970-1980

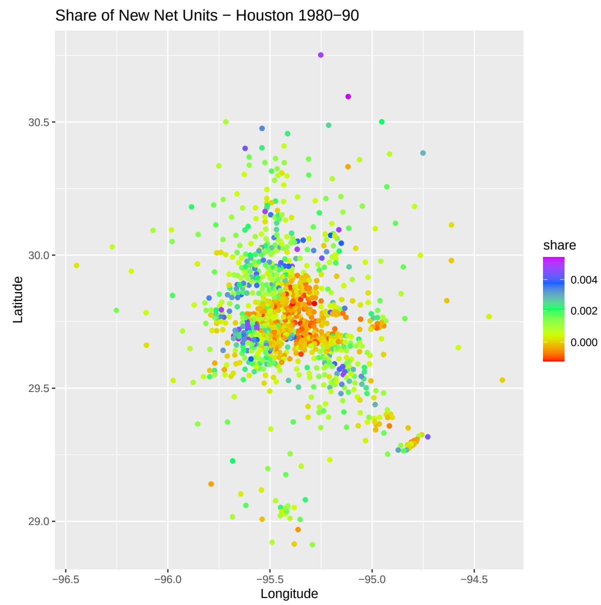

(b) 1980-1990

(c) 1990-2000

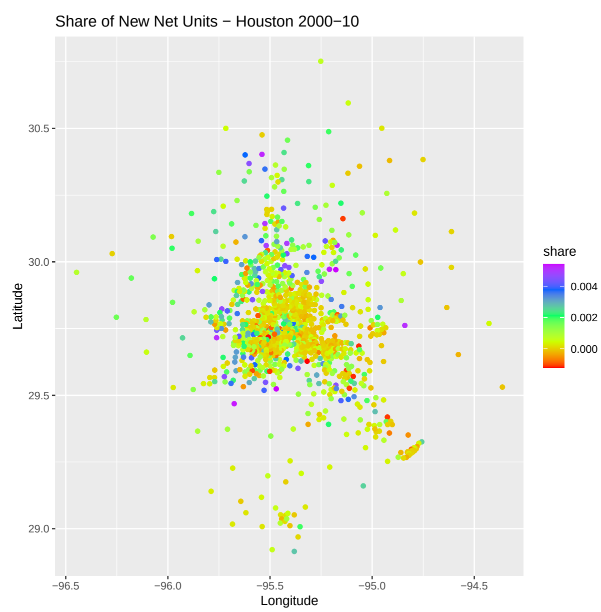

(d) 2000-2010

Notes: Housing unit data obtained from U.S. Census Bureau.

Figure 11: Share of New Net Units in Los Angeles, by Census Tract

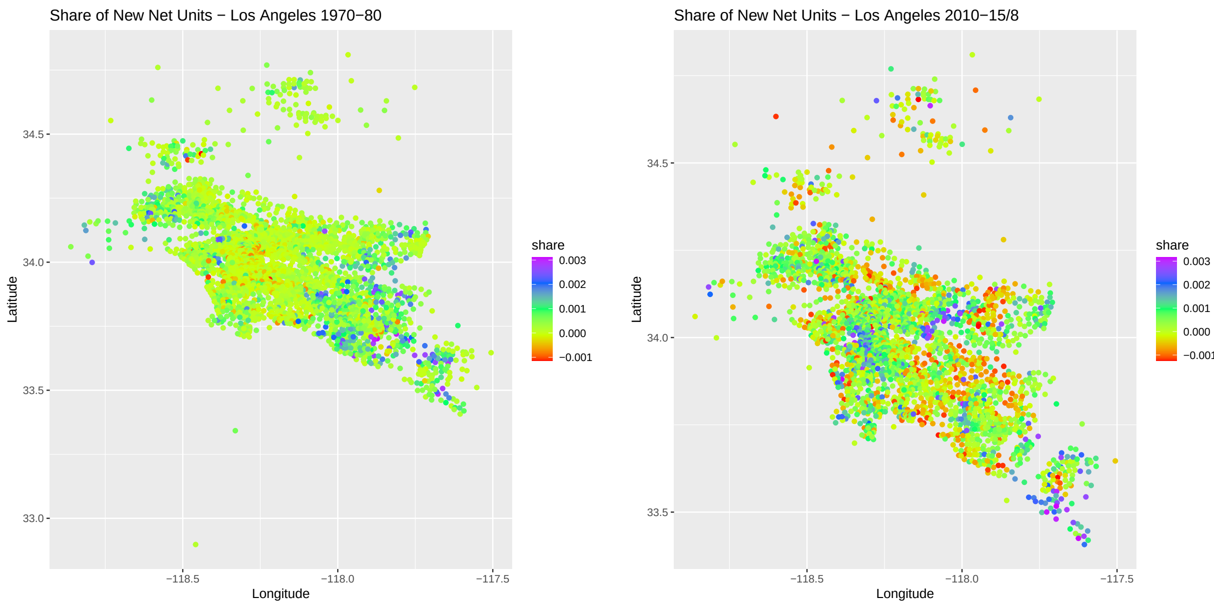

(a) 1970-1980

(b) 2010-2018

Notes: Housing unit data obtained from U.S. Census Bureau.

Figure 12: Convergence: Texas Joins California as Built-Out

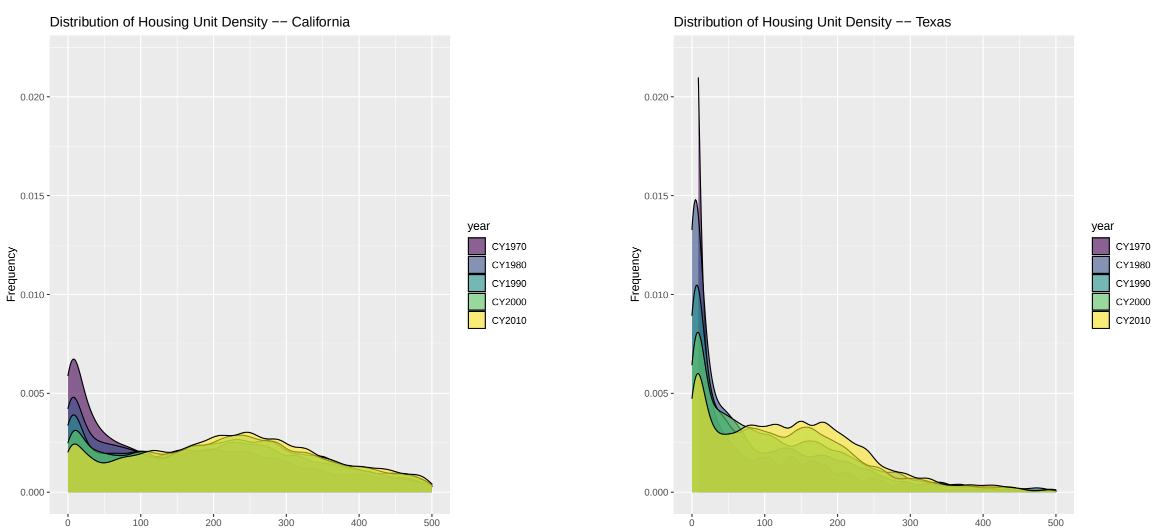

(a) Evolving California

(b) Evolving Texas

Notes: Tracts with density greater than 500 units/square mile are excluded.

Figure 13: Convergence: Houston Approaches Los Angeles as Built-Out

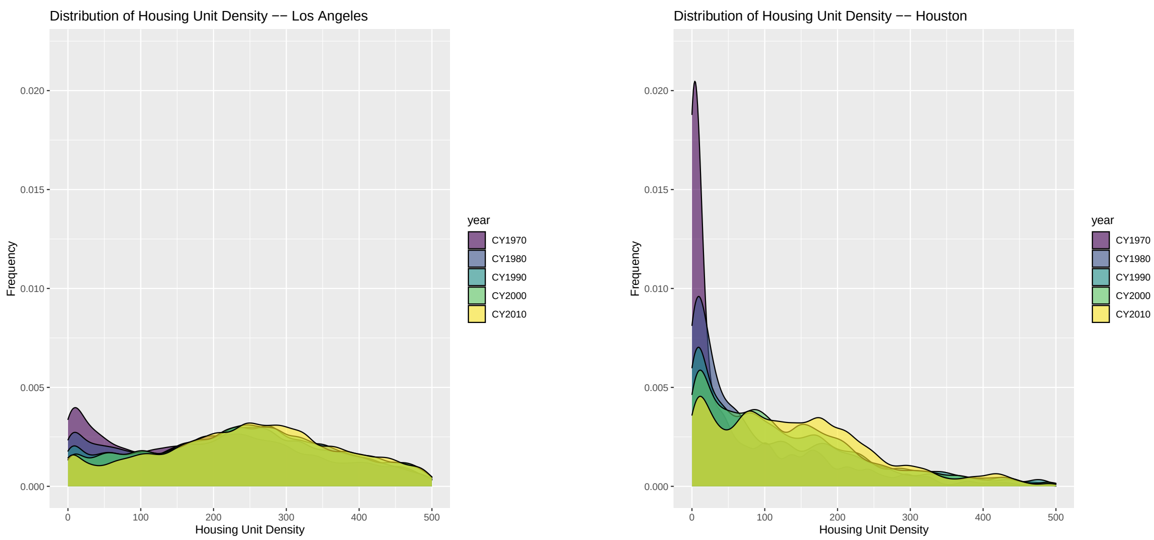

(a) Evolving Los Angeles (b) Evolving Houston

Notes: Tracts with density greater than 500 units/square mile are excluded.

Figure 14: Convergence: Sacramento Looks Like Houston; Austin Fills Rapidly

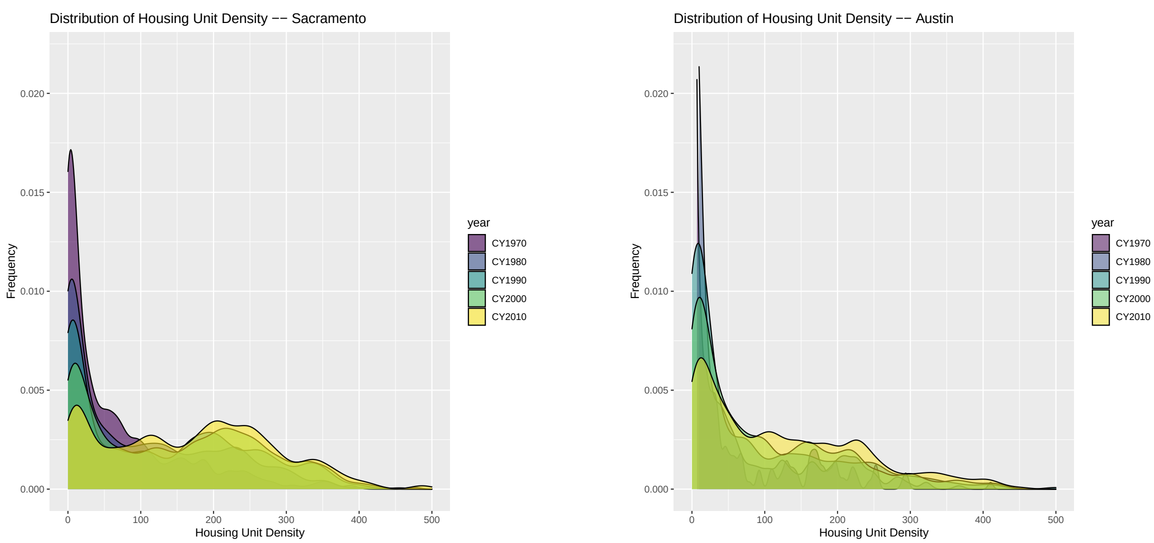

(a) Evolving Sacramento (b) Evolving Austin

Notes: Tracts with density greater than 500 units/square mile are excluded.

References

Aastveit, K. A., B. Albuquerque, and A. K. Anundsen (2020): "Changing Supply Elasticities and Regional Housing Booms," Bank of England Working Paper No. 844, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3520650.

Baum-Snow, N., and L. Han (2021): "The Microgeography of Housing Supply," Working Paper, https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbnxiYXVtc25vd3xneDoxNWI0MDAxNDA5YWEyNTJl.

Brooks, L., and B. Lutz (2016): "From Today's City to Tomorrow's City: An Empirical Investigation of Urban Land Assembly," *American Economic Journal: Economic Policy*, 8(3), 69–105.

Cunningham, C. (2013): "Estimating the Holdout Problem in Land Assembly," Federal Reserve Bank of Atlanta Working Paper 2013-19, https://www.atlantafed.org/-/media/documents/research/publications/wp/2013/wp1319.pdf.

Davidoff, T. (2016): "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors," *Critical Finance Review*, 5(2), 177–206.

DiPasquale, D. (1999): "Why Don't We Know More About Housing Supply?," *Journal of Real Estate Finance and Economics*, 18, 9–23.

Eriksen, M. D., and A. W. Orlando (2022a): "The Causes and Consequences of Development Costs: Evidence from Multifamily Housing," Working Paper, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4225004.

——— (2022b): "Returns to Scale in Residential Construction: The Marginal Impact of Building Height," *Real Estate Economics*, 50(2), 534–564.

Furman, J. (2015): "Barriers to Shared Growth: The Case of Land Use Regulation and Economic Rents," Prepared Remarks, The Urban Institute, https://obamawhitehouse.archives.gov/sites/default/files/page/files/20151120_barriers_shared_growth_land_use_regulation_and_economic_rents.pdf.

Glaeser, E. L., and J. Gyourko (2003): "The Impact of Building Restrictions on Housing Affordability," *Federal Reserve Bank of New York Economic Policy Review*, 9(2), 21–39.

——— (2005): "Urban Decline and Durable Housing," *Journal of Political Economy*, 113(2), 345–375.

Glaeser, E. L., J. Gyourko, and R. Saks (2005): "Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices," *Journal of Law and Economics*, 48(2), 331–369.

Glaeser, E. L., J. Gyourko, and R. E. Saks (2006): "Urban Growth and Housing Supply," *Journal of Economic Geography*, 6(1), 71–89.

Glaeser, E. L., and B. A. Ward (2009): "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston," *Journal of Urban Economics*, 65(3), 265–278.

Gyourko, J., J. S. Hartley, and J. Krimmel (2021): "The Local Residential Land Use Regulatory Environment Across U.S. Housing Markets: Evidence from a New Wharton Index," *Journal of Urban Economics*, 124, 103337.

Gyourko, J., and J. Krimmel (2021): "The Impact of Local Residential Land Use Restrictions on Land Values Across and Within Single Family Housing Markets," NBER Working Paper No. 28993.

Gyourko, J., C. Mayer, and T. Sinai (2013): "Superstar Cities," *American Economic Journal: Economic Policy*, 5(4), 167–199.

Gyourko, J., and A. Saiz (2006): "Construction Costs and the Supply of Housing Structure," *Journal of Regional Science*, 46(4), 661–680.

Gyourko, J., A. Saiz, and A. 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.

Hilber, C. A., and F. Robert-Nicoud (2013): "On the Origins of Land Use Regulations: Theory and Evidence from US Metro Areas," *Journal of Urban Economics*, 75, 29–43.

Hornbeck, R., and D. Keniston (2017): "Creative Destruction: Barriers to Urban Growth and the Great Boston Fire of 1872," *American Economic Review*, 107(6), 1365–1398.

Hsieh, C.-T., and E. Moretti (2019): "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, 11(2), 1–39.

Li, X., and C. Zhang (2021): "Did the COVID-19 Pandemic Crisis Affect Housing Prices Evenly in the U.S.?," *Sustainability*, 13(21), 12277.

Malone, T., and C. L. Redfearn (2018): "Shocks & ossification: The durable hierarchy of neighborhoods in US metropolitan areas from 1970 to 2010," *Regional Science and Urban Economics*, 69, 94–121.

Morrow, G. (2013): *The Homeowner Revolution: Democracy, Land Use and the Los Angeles Slow-Growth Movement, 1965–1992*. University of California, Los Angeles.

Orlando, A. W., and C. L. Redfearn (2022): "Housing Supply Elasticities: A Structural Vector Autoregression Approach," Working Paper.

Pogodzinski, J., and T. R. Sass (1994): "The Theory and Estimation of Endogenous Zoning," *Regional Science and Urban Economics*, 24(5), 601–630.

Quigley, J. M., and S. Raphael (2005): "Regulation and the High Cost of Housing in California," *American Economic Review*, 95(2), 323–328.

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