What's Lost in the Aggregate: Lessons from a Local Index of Housing Supply Elasticities

Anthony W. Orlando & Christian L. Redfearn

December 2018

Abstract

A growing divergence between supply and demand has led to the current "housing affordability crisis," but little is systematically understood about the local factors that created this divergence. We estimate the first county-leveel and municipality-level index of housing supply elasticities in major urban counties across the United States, using a structural vector autoregression model with sign restrictions to identify the effect of a positive demand shock. We find that these supply elasticities are lower than previous estimates, illustrating the degree to which supply is restricted in local housing markets. We document significant heterogeneity across the country and within the metropolitan area. The supply curve is less elastic in lower-income, higher-density central counties and cities, potentially pushing construction out to the periphery and contributing to urban sprawl.

Orlando can be reached at aworlando@cpp.edu and Redfearn can be reached at redfearn@usc.edu. We thank Antonio Bento, Raphael Bostic, Morris Davis, and seminar participants at the Northeast Political Science Association, Urban Affairs Association, and Urban Economics Association for helpful comments. All errors are the authors' alone.

1 Introduction

Across the United States, urban land prices have been growing significantly faster than incomes in recent decades, leading many economists and planners to declare a "housing affordability crisis".[1] At the root of this crisis is a growing shortage of housing units (Morrow 2013, Taylor 2015). Economic theory predicts that prices will rise as long as supply fails to keep up with demand, and the empirical literature has overwhelmingly found evidence of this effect. To date, however, this literature has focused on averages across metropolitan areas, whereas supply restrictions are actually enacted at the county and municipality levels. This paper is the first to identify where supply is being restricted the most *within* the metro—and therefore, which counties and cities are driving this affordability crisis.

Consider the growth of housing units from 2000 to 2012 in Los Angeles County, shown in Figure 1. It is clearly not a homogeneous story. At the municipality level, housing unit growth ranges from 0% to 27.7%. From this graph, however, it is impossible to identify how much of this variation is due to differences in the supply curve—and not simply from differences in housing demand. It may be the case, for example, that Calabasas and Pasadena, with their high growth rates, are simply more desirable locations than low-growth cities like Downey and Inglewood. The empirical economist faces the challenge of determining how much supply *responds* to these different demand curves in each city—in other words, how *elastic* is the supply curve.

In this paper, we construct housing price and quantity indices to calculate these elasticities. For prices, we use county-level and municipality-level data from Zillow from 1996 to 2017. For quantities, we use annual building permit data from the U.S. Department of Housing and Urban Development to supplement total housing unit counts from the decadal Census. With these two indices, we calculate housing supply elasticities (a) for the counties in the five largest metropolitan statistical areas (MSAs) in the United States and (b) for the municipalities in Los Angeles County, using a sign-restricted vector autoregression (VAR) model that identifies positive demand shocks as simultaneous increases in demand and supply. These structural shocks allow us to draw causal inference from the impulse responses of both quantities and prices, which we then use to calculate elasticities. This methodology improves on previous estimates that suffered from endogeneity bias, and it reveals that those estimates were likely overestimating the price elasticity of housing supply by orders of magnitude.

We find that "short-run" supply elasticities range from 0.01 to 0.42 at the county level and from 0.01 to 0.20 at the municipality level.[2] These estimates are consistent with recent econometric advances in similar markets where investment is costly to reverse and production is therefore hesitant to respond to year-to-year signals (Kilian 2009, Kilian and Murphy 2012). While the majority of these estimates are clustered in the 0 to 0.05 range, there is a long right tail in the

[1] See, for example, Sinai (2014), Albouy, Ehrlich, and Liu (2016), Watson, Steffen, Martin, and Vandenbroucke (2017), Wetzstein (2017).

[2] We show "long-run" supply elasticities in Appendix A, but it is not clear yet how reliable this methodology is for such estimates that extend over several years.

distribution, suggesting that many counties' and cities' supply curves are significantly more elastic than the average.

We exploit this heterogeneity to understand the driving forces behind the housing shortage. We find evidence that the old adage "location, location, location" still matters: Figure 7 shows that elasticities are lowest in the core of the metropolitan area and increase as one travels northeast and northwest, away from the city center and the valuable coastal land. This map provides striking context for Figure 1. It suggests that the central cities are restricting housing supply the most and pushing it out to the periphery, contributing to urban sprawl, longer commute times, less agglomeration, and suboptimal allocation of housing.

In cross-sectional analysis, we find that many of the typical explanations for housing supply elasticity (or lack thereof) are not statistically significant: housing price, population, and even regulations as measured by the commonly used Wharton Residential Land Use Regulatory Index (WRLURI). Only two factors are significant: The housing supply curve is more elastic in cities with less density and higher income, consistent with the urban sprawl story illustrated in Figure 7.

These findings contribute, first and foremost, to the literature on supply elasticities in housing and land markets. Few papers have attempted to directly estimate these elasticities, likely due to data constraints as well as endogeneity challenges. These early estimates used a simple ordinary least-squares approach, often with lagged independent variables (Smith 1976, Pryce 1999, Mayer and Somerville 2000, Malpezzi and Maclennan 2001, Green, Malpezzi, and Mayo 2005). Harter-Dreiman (2004) and Wheaton, Chervachidze, and Nechayev (2014) took a step closer to our approach with a vector error correction series of equations, but even that model cannot identify structural shocks.

The second approach in the literature has been to measure housing supply restrictions with survey-based measures of land-use regulation and satellite measures of topographic constraints. The most widely-used indices on regulation and geography are Gyourko, Saiz, and Summers (2008) and Saiz (2010), respectively. These types of measures have overwhelmingly shown that inelastic housing supply leads to higher prices (Quigley and Raphael 2005, Glaeser, Gyourko, and Saks 2005, Schuetz 2009, Saiz 2010, Gyourko, Mayer, and Sinai 2013, Ganong and Shoag 2015). More recent structural approaches have concluded that these restrictions have significant negative effects on economic growth and productivity, particularly for less-skilled workers (Nieuwerburgh and Weill 2010, Moretti 2013, Hsieh and Moretti 2017, Parkhomenko 2017). These efforts all employ indirect proxies for elasticity, however, and they are only interested in averages for each metropolitan statistical area (MSA) as a whole.

This paper, in contrast, directly measures elasticities at the county and municipality levels within the MSA. As such, it also speaks to the broad literature on neighborhood sorting, giving new evidence of how supply restrictions factor into a household's decision about where to live in a given metro. This question draws its original inspiration from the famous Tiebout (1956) model showing how households sort into the municipality with their preferred balance of taxes and public goods and services. Equally important is the literature developed around the Roback (1982) model showing how households trade off amenities, house prices, and wages. Recent work has extended these neighborhood dynamics to include crime rates (Bayer, McMillan, Murphy, and Timmins 2016), path dependence (Malone and Redfearn 2016), pollution (Heblich, Trew, and Zylberberg 2016), public transit (Baum-Snow 2007, Waxman 2017), racial preferences (Bayer, Fang, and McMillan 2014), and school quality (Bayer, Ferreira, and McMillan 2007). Housing supply restrictions are a key tool that neighborhoods use to achieve sorting along all these dimensions, making this paper an important step toward a more comprehensive theory of these inter-city (intra-metro) equilibria.

The remainder of the paper is organized as follows. Section 2 derives theoretical predictions from the canonical urban economics model to motivate our empirical work. Section 3 presents our data and methodology for constructing local indices of housing prices and quantities. In Section 4, we use these indices to calculate supply elasticities. Section 5 analyzes the spatial variation in these elasticities, including their correlation with the traditional explanations in the literature. Section 6 concludes.

2 Motivating Theory

We begin by modeling the supply side of the housing market according to standard urban theory. Let $S = N/l$ represent the capital-land ratio, or the structural density, at a given location. Following Brueckner's (1987) classic unified model, developers maximize the profit function

$$l\big(ph(S) - iS - r\big) \ , \tag{1}$$

where p is the rent per square foot paid by consumers, $h(S)$ is the floor space, i is the rental price of capital, and r is the land rent. In equilibrium, each developer earns zero total profit:

$$ph(S) - iS = r \ , \tag{2}$$

and developers maximize their profits by choosing S:

$$ph'(S) = i \ . \tag{3}$$

We can conduct comparative statics by differentiating these equations with respect to variables we care about, such as commuting cost, distance from the central business district, consumer income, and utility, all represented by ϕ:

$$\frac{\partial r}{\partial \phi} = h\frac{\partial p}{\partial \phi}, \tag{4}$$

$$\frac{\partial S}{\partial \phi} = -\frac{h'}{ph''}\frac{\partial p}{\partial \phi}. \tag{5}$$

If $h'(S) > 0$ and $h''(S) < 0$ according to the law of diminishing returns, then $\frac{\partial p}{\partial \phi}$, $\frac{\partial r}{\partial \phi}$, and $\frac{\partial S}{\partial \phi}$ all have the same sign. Anything that increases prices, therefore, should increase structural density.

The canonical model of urban economics thus predicts that higher prices per square foot should be associated with more structural density, both across space and over time. This is consistent with our conversations with developers, who seek higher density on more expensive land, where less density would not earn a high enough return per square foot to cover the cost basis.

Importantly, there are no frictions in this model. It hinges on the assumption that developers *can* choose S with certainty. In the real world, however, a number of obstacles prevent this freedom of choice: land-use regulations, topographical constraints, lack of capital, construction delays, and economic uncertainty. When the supply process is stifled by these factors, price growth often will not be met with a sufficient increase in density to meet the demand. Formally,

$$\%\Delta P_{i,t} > \%\Delta Q^S_{i,t} \ , \tag{6}$$

where $Q^S_{i,t}$ is the quantity supplied in locality i in year t and $P_{i,t}$ is the price at which it is supplied. The more imbalanced this equation becomes, the lower will be the price elasticity of supply,

$$\varepsilon_S = \frac{\%\Delta Q^S_{i,t}}{\%\Delta P_{i,t}} = \frac{\Delta Q^S_{i,t}/Q^S_{i,t-1}}{\Delta P_{i,t}/P_{i,t-1}} = \frac{P_{i,t-1}}{Q^S_{i,t-1}} \times \frac{\Delta Q^S_{i,t}}{\Delta P_{i,t}} \ . \tag{7}$$

By measuring this elasticity in each city, we can therefore impute the degree to which the supply process is being restricted because density cannot keep up with prices.

3 Local Housing Prices and Quantities: Data and Trends

3.1 Local Housing Price Index

Before calculating housing supply elasticities, the researcher must make two decisions: how to calculate price and quantity indices over time and how to calculate elasticities from those indices. Until recently, these decisions have seemed nearly insurmountable in the housing market.

The first calculation requires a reliable measure of all housing values in a given geography. The classic Case and Shiller (1989) approach forms a repeated-sale index based on multiple transactions of the same property. While this approach makes the return a true measure of price growth in a given asset over time, it did not control for changes in the asset itself, such as modifications to or depreciation of the house. Even more concerning, it limits the sample size to the unrepresentative subset of properties that transact multiple times. Other researchers create hedonic indices that expand the sample to all properties that transact at least once, controlling for observable building characteristics. The sample still is not representative of all the properties that an investor could purchase; it only gives the value of those that did transact, a decision that clearly involved selection bias (Englund, Quigley, and Redfearn 1999). Moreover, controlling for the observable characteristics alone is probably inferior to actually comparing the same house over time.

This calculation also requires a long enough time period to draw statistically significant conclusions at a given geographic level. While the literature has been able to make statements about metropolitan areas, they rarely drill down to the municipality level where they have lacked sufficient observations. It is only now that we have at least a decade worth of high-quality transaction data to calculate indices for most counties and cities within a large MSA such as Los Angeles.

Zillow owns the most comprehensive dataset of housing values in the United States. Not only do they have all publicly recorded transactions over time, but they can combine these transactions with listing values and other non-transaction data that they collect on their website, which is now the predominant site for buyers and sellers in today's housing market.[3]

Zillow uses these data to calculate a median home value index (ZHVI) in five steps: First, they calculate raw median sale prices for all properties, whether they transacted or not, with $r_{i,j}(t)$ representing the raw median price for market segment i in geographic region j at time t.[4] Second, they adjust for any residual systematic error in region j at time t,

$$b_j(t) = \text{Median}\frac{z_j(t-1) - s_j(t)}{s_j(t)} \ , \tag{8}$$

where $s_j(t)$ is a vector of the actual sales prices transacted and $z_j(t-1)$ is Zillow's estimate of those properties' value in the period before they transacted. The adjusted median $u_{i,j}(t)$ will correct for this error in Zillow's estimates by incorporating the new sales data about those properties into the raw median price:

$$u_{i,j}(t) = \frac{r_{i,j}(t)}{1 + b_j(t)} \ . \tag{9}$$

Third, they apply a five-term Henderson (1916) moving average filter to reduce noise. Fourth, they adjust for seasonality with a decomposition proposed by Cleveland, Cleveland, McRae, and Terpenning (1990), where the time series is broken down into seasonal, trend, and remainder components,

$$U(t) = S(t) + T(t) + RE(t) \ , \tag{10}$$

and then the seasonal component $S(t)$ is subtracted. Finally, Zillow deletes all time series that have too few observations, too much volatility, or too many outliers, gaps, or jumps to meet their standard of quality control.[5]

The resulting data are available in a variety of forms, from time series of particular building types to different quantiles. We employ the median home value per square foot index, which captures all residential buildings and standardizes prices as a function of size for the best comparability. We download these indices at the county and municipality levels.[6] These data are available at a monthly frequency. We aggregate up to the annual frequency by averaging over the twelve monthly observations for each city in a given year.

[3]This predominance has become particularly strong since Zillow's merger with its largest competitor, Trulia, in 2015 (Kusisto and Light 2015).

[4]The market segment is the type of building—single family, condo, etc.

[5]For more details, see `https://www.zillow.com/research/zhvi-methodology-6032/`.

[6]To download these and other data from Zillow, go to `https://www.zillow.com/research/data/`.

Since we are interested in local policy and planning decisions, we narrow our focus to the counties available in the five largest MSAs in the country: New York, Los Angeles, Chicago, Dallas, and Houston. These five cases offer a wide variety of urban spatial structures and planning codes, as well as regional diversity, covering all four Census regions: the Midwest, the Northeast, the South, and the West.

Following this county-level analysis, we then take a deep dive into one major urban county, the lowest level of government within which municipalities operate. At this level, any variation in the cross-section of housing supply elasticities must represent differences in the municipalities themselves, not in counties or metropolitan areas or other larger geographies.[7] We focus on Los Angeles County for its size and its variety of different geographies, topographies, and neighborhood characters. It is one of the most ideal laboratories within which to study cross-sectional variation in municipalities, as it encompasses over 80 municipalities from 1996 to 2015.

3.2 Local Housing Quantity Index

To our knowledge, building permits are the only *annual* data on housing quantity available at a local level. Municipalities, counties, townships, and other towns issue and record building permits on an ongoing basis, and the Census Bureau surveys 9,000 of these "permit-issuing places" every month. They collect information on the amount of new privately-owned single-family homes, as well as the number of new units in two- to five-unit residential buildings and the overall dollar value of construction. They ask local officials to describe permits valued at least $1 million and to list the owner or builder. Once a year, the Census Bureau also surveys another 11,000 permit offices.[8] The annual estimate is therefore a larger and more accurate sample, which is one of the reasons why we use it in this paper.[9] Of course, not everyone responds to the sample, so the Census Bureau has to impute roughly 19% of the monthly units and 7% of the annual units by assuming that the missing localities issued permits at the same rate as the sampled localities.

Figure 2 compares single-family home permits, as well as all residential permits, to the respective housing prices.[10] In both cases, it appears that the real estate booms of the late 1980s and early 2000s were very different. The former was an instance of unusually high building activity with mild price appreciation, while the latter exhibited unusually high price appreciation with strong building activity that peaked lower and earlier than prices. This last point is crucial. One of the leading theories in urban economics alleges that prices rose so high because supply was restricted, but this does not appear to be the case in this graph. On the contrary, from 1992 to 2004, housing quantity

[7] Landvoigt, Piazzesi, and Schneider (2015) take a similar approach to the housing market(s) of San Diego.

[8] These numbers refer to recent years. When the survey was initiated in 1959, the number of offices was half of what it is today. The Census Bureau selects the 9,000 monthly survey respondents using a stratified sampling methodology where the most populous areas are chosen with certainty and the rest of the localities are chosen with a sampling factor of 1 in 10.

[9] We also use annual data because the monthly *price* data are more volatile and likely to be driven by outliers, especially at the municipality-level, where very few transactions may occur in a given month—and those transactions may not be representative of the average housing value across all properties.

[10] Prices were calculated by the authors using transaction-level data from DataQuick to estimate further back in time than Zillow will allow. These are not the prices used in the rest of the paper.

growth tracked housing price growth almost exactly. It had no trouble keeping up, in spite of these alleged restrictions. It was only in the last couple years of the "bubble" that prices outpaced quantity. This is unsurprising, given Orlando's (2018) findings that permit activity typically slows a year or two before a recession hits.

The starting point is 1990, when the Census of Population and Housing gives the total housing units.[11] Since 1990, the housing stock has grown less than 1% per year, for a total of only 10% in 22 years. This evidence is consistent with the literature, as well as anecdotal accounts from the real estate market, suggesting that it has been difficult to build in Los Angeles, with construction falling far behind the nation's (and the state's) population growth (Quigley and Raphael 2005, Morrow 2013, Dovey 2015). It does not appear, however, that construction has been on a downward trend over these recent decades, contradicting claims that supply has become *more* restricted over time, particularly during the recent "bubble" period.[12]

4 Calculating Housing Supply Elasticities

We begin by considering the classic supply-and-demand equations, which represent our ideal goal for estimation:

$$Q^S_{i,t} = \varepsilon^S_i P_{i,t} + u^S_t \ , \tag{11}$$

$$Q^D_{i,t} = \varepsilon^D_i P_{i,t} + u^D_t \ , \tag{12}$$

where quantity and price variables are expressed in natural logs to achieve the percent changes in Equation 7. Unfortunately, $Q^S_{i,t}$ and $Q^D_{i,t}$ are difficult to measure in the real world. If we make some simplifying assumptions, however, we can estimate a similar system of equations with standard housing indices. Specifically, we assume that positive demand shocks manifest themselves in increasing price and increasing quantity, following a standard Marshallian model (Marshall 1890). A structural vector autoregression (SVAR) model allows us to simulate such a shock using the past behavior of prices and quantities. The responses of quantity, $\%\Delta Q^S_{i,t}$, and price, $\%\Delta P_{i,t}$, to these impulse shocks will allow us to calculate the price elasticity of supply, $\varepsilon^S_{i,t}$.

4.1 Sign-Restricted Vector Autoregression Model

First, consider a two-variable reduced-form vector autoregression (VAR) model:

$$\begin{bmatrix} Q_t \\ P_t \end{bmatrix} = \begin{bmatrix} \alpha^Q_{Q,1} & \alpha^Q_{P,1} \\ \alpha^P_{Q,1} & \alpha^P_{P,1} \end{bmatrix} \begin{bmatrix} Q_{t-1} \\ P_{t-1} \end{bmatrix} + \begin{bmatrix} \alpha^Q_{Q,2} & \alpha^Q_{P,2} \\ \alpha^P_{Q,2} & \alpha^P_{P,2} \end{bmatrix} \begin{bmatrix} Q_{t-2} \\ P_{t-2} \end{bmatrix} + \begin{bmatrix} e_{Q,t} \\ e_{P,t} \end{bmatrix} , \tag{13}$$

$$\text{or} \quad \mathbf{Y}_t = \mathbf{A}_1 \mathbf{Y}_{t-1} + \mathbf{A}_2 \mathbf{Y}_{t-2} + \mathbf{E}_t \ , \tag{14}$$

[11] See https://www.census.gov/prod/cen1990/cph2/cph-2-6.pdf. The Census does not report the total number of single-family homes at the county or city level. We therefore do not estimate growth rates for single-family homes.

[12] These claims have been made most famously by Edward Glaeser and Joseph Gyourko. See, for example, Glaeser and Gyourko (2003), Glaeser (2004), and Glaeser, Gyourko, and Saks (2005).

where $\mathbf{Y_t} = \big(Q_t, P_t\big)$.[13] This model extracts the relationship between price and quantity over time, but in this form, the relationship is endogenous, making causal identification impossible. The trouble is that the innovations in the model, $\mathbf{E}_t$, are not "economically meaningful or fundamental," to quote Uhlig (2005). They are simply prediction errors. We need to transform them *into* fundamental innovations, or structural shocks, $\mathbf{U}_t$,

$$\mathbf{BE}_t = \mathbf{U}_t \ , \tag{15}$$

using a matrix, $\mathbf{B}$, that weights the VAR residuals to identify *only* the fundamental innovations. In our case, for example, we are looking for the response of the variables to demand shocks, which we identify as a simultaneous increase in both price and quantity. We therefore need a matrix $\mathbf{B}$ that excludes any impulse vector that does not result in positive responses in the first period.[14] This type of model is called a "sign-restricted VAR."

This approach solves the structural identification problem, but it creates a new problem: "model identification" (Preston 1978). There is no unique matrix $\mathbf{B}$. A *set* of possible solutions exists, and so we say that sign-restricted VARs are "set-identified." A growing literature has proposed various methods to select the most appropriate impulse response from the set of admissible models.[15] In this paper, we employ the simplest approach, Uhlig's (2005) rejection method, which begins by jointly drawing the VAR parameters from a Normal-Wishart posterior and an $n \times 1$ vector from a uniform distribution over the unit sphere. It then computes the impulse response. If the signs all satisfy the exclusion restrictions, it keeps the draw. If not, it rejects the draw. It repeats this process for as many draws as the researcher desires. In selecting our final point estimate, we follow the standard approach in the literature and use the median response of each variable from the distribution of remaining draws.

4.2 County-Level Results

Our first goal is to calculate supply elasticities for each county in the five largest MSAs in the United States. These elasticity estimates can best be understood as a combination of impulse responses. Figure 3, for example, shows the impulse response graphs for Los Angeles County, the core county of the Los Angeles MSA. It indicates that a one-standard-deviation demand shock leads to a 0.03% increase in housing units and a 2.83% increase in prices after on year. Plugging these percent changes into the standard elasticity equation yields a one-year ("short-run") elasticity of 0.010, as shown in Table 2. This elasticity is the lowest county-level estimate in the MSA, with 0.032 for Orange County, 0.034 for San Bernardino County, 0.048 for Ventura County, and 0.057 for Riverside County. It is not surprising to see the least elastic market in the core, where regulations

[13] The two-lag structure has been chosen because it minimizes the information criteria more often than alternative specifications while remaining small enough to allow sufficient degrees of freedom in a short time window.

[14] After the first period, there is no reason to maintain the exclusion restriction. We do not want to impose a strong prior that demand shocks persist past the first response period. Rather, we remain agnostic and let the model play out.

[15] See Fry and Pagan (2011), Kilian and Murphy (2012), and Baumeister and Hamilton (2015) for critical reviews.

tend to be the most stringent. It also fits with our expectations to see the two least elastic markets as the areas where housing density was already the highest, according to the 2000 Census. It is typically easier to develop in areas that are less "built out." This ranking of elasticities therefore accords with both economic theory and firsthand conventional wisdom in real estate. What is more surprising is the low magnitude of the estimates compared to previous literature, as discussed in the municipality-level section below.

These results are not universal across all markets, however. Table 1, for example, shows that the core Kings County actually has the highest elasticity in the New York MSA, 3 shows that the core Cook County has the lowest elasticity in the Chicago MSA, 4 shows that the core Dallas County has a low-to-mid elasticity within the Dallas MSA, and 5 shows that the core Harris County has the median elasticity within the Houston MSA. This final estimate is particularly interesting, as Houston is known for its lack of a traditional zoning code, which does not appear to make its housing supply more or less elastic than the surrounding neighborhoods, though it does have more elastic supply than the core counties in the other four MSAs.

4.3 Municipality-Level Results

Our second goal is to calculate supply elasticities for each municipality in Los Angeles County using the sign-restricted VAR(2) model. As in the county-level case, these municipality-level elasticities are a combination of the respective impulse responses for each time series. Figure 4, for example, shows the impulse response graphs for the core city of Los Angeles. It indicates that a one-standard-deviation demand shock leads to a 0.07% increase in housing units and a 3.29% increase in prices after one year. We define this one-year response as the *short run*. After four years, the quantity response has increased to 0.20%, while the price response has stabilized at 4.20%. Consistent with economic theory, it appears that there is a longer lag in production than prices, leading the elasticity to rise in the *long run*.[16]

As a comparison, consider the impulse response graphs for the suburb city of Calabasas in Figure 5. Calabasas sits on the periphery of the metropolitan area, where land is cheaper, terrain is flatter, and density is lower. All else being equal, economic theory predicts that elasticities should be higher as a result. Consistent with this expectation, the sign-restricted VAR reveals a higher increase in units (0.31%) and a smaller increase in prices (2.50%) in Calabasas than in Los Angeles in the short run. In the long run, Calabasas price growth catches up to Los Angeles, but its quantity growth is nearly three times as high.

Equation 7 allows us to combine these impulse responses into supply elasticities for each city. Table 6 lists the one-year elasticities ranked from highest to lowest. At the top are the cities where the supply curve is the flattest—that is, where the market responds to a demand shock more by building housing units and less by increasing prices. These municipalities tend to be some of the furthest suburbs from the center of the metropolitan area. At the bottom are the cities where

[16] Orlando (2018) demonstrates this lag with a time-series analysis of the production process, leading from permits to starts to completions.

construction responds very little, if at all, to demand shocks. These municipalities appear mostly to hail from low-income neighborhoods in South Central, a historically segregated and high-crime region. The most general takeaway from this table as a whole is that the estimates are very low. Even in the most elastic city, a one-standard-deviation demand shock never leads to more than a 0.55% increase in housing units. For most cities, the response is less than 0.10%. This finding is consistent with a long literature suggesting that inelastic supply is to blame for high housing costs in California (Quigley and Raphael 2005, Morrow 2013).

These estimates are so low, in fact, that they suggest that previous literature has significantly overestimated the price elasticity of supply at a local level. Though no one has attempted to measure this elasticity at a municipal level, the MSA-level estimates are instructive. Green, Malpezzi, and Mayo (2005), one of the early canonical works, regresses the annual change in the housing stock on the lagged first differences of housing prices. As we discussed earlier, this methodology suffers from endogeneity and therefore lack of interpretable identification. It suggests that Los Angeles had a supply elasticity of 3.73 from 1979 to 1996. Our estimates begin in 1996, but even so, it is unlikely that they decreased so drastically between these time periods as to be in the 0 to 0.2 range. The 3.73 estimate is particularly suspect because it places Los Angeles in the middle third of the national range of elasticities, which does not square with the MSA's high price growth, stringent regulations, and challenging topography relative to the national average. Saiz (2010), the most widely cited recent work, estimates a much lower elasticity for Los Angeles, 0.63, making it the second-most inelastic MSA in the country. These elasticities are indirect measures, however, combining regulatory and geographic constraints rather than directly estimating the slope of the supply curve in response to exogenous demand shocks, as this paper attempts to do.

These low estimates are more consistent with recent econometric work in oil markets, where sign-restricted VARs have been used to identify the response of oil prices and production to supply and demand shocks. This structural approach has led to a consensus estimate for the short-run price elasticity of oil supply between 0 and 0.2, precisely the range identified in this paper for the short-run price elasticity of housing supply (Kilian and Murphy 2012). It is unlikely that this similarity is mere coincidence, as it is consistent with economic theory in both markets. Producers need time to respond to demand shocks, and once they make a decision, it is costly to reverse. As a result, they tend to respond only to *persistent* shocks that change expectations about the future (Kilian 2009). It is unlikely, therefore, that the true price elasticity of housing supply is as high as previous estimates have suggested.

5 Cross-Sectional Analysis

The heterogeneity within the metropolitan area is significant. Figure 6 gives a sense of how much variation is lost in the aggregate. It shows normal and kernel density estimates of these distributions. While there is a high clustering around 0.025 for municipality-level elasticities and 0.500 for county-level elasticities, there are also long tails to the right, which are not captured by the average or the median. These tail suggests that some cities are increasing supply much faster in response to price signals. In these counties and cities, we may find lessons to overcome the barriers that stand in the way of housing production and affordability in many metropolitan areas.

Why does this heterogeneity exist? To return to our original questions, which cities are the most restrictive and why? Table 7 reveals some potential answers. It shows bivariate regressions in the year 2000 with the natural logs of price (median price per square foot from Zillow in 2000), density (population density from the 2000 Census), population (from the 2000 Census), median household income (from the 2000 Census), and the Wharton Residential Land Use Regulatory Index (Gyourko, Saiz, and Summers 2008). Surprisingly, only two of these factors—which are typically considered to be important determinants of a city's housing supply function—are statistically significant. More elastic housing supply exists in cities with less density and more income.

It is also important to note that the R^2 never exceeds 12% in any of these five regressions. On their own, none of these factors explains the vast majority of the variation in housing supply elasticity within the metropolitan area. Together, they may account for more, but we must conclude that we still do not understand most of this heterogeneity. We have much to learn about how supply responds differently to demand in different cities.

One important factor appears to be location. Figure 7 shows a heat map of the short-run elasticities, with darker colors indicating more elastic cities.[17] From this picture, it is clear that housing supply is the most restricted in the city center—and least restricted in the northern periphery. (The southern periphery is valuable coastal land.) It is simply hardest to build in the most central locations, where it is already most dense, consistent with the regressions in Table 7. It is no wonder, therefore, that Los Angeles is such a sprawled metropolitan area. The core is restricting supply and pushing it out to the periphery. The "location, location, location" adage, it turns out, has some merit.

Another way to test the robustness of this conclusion is by comparing the county-level elasticities to their preexisting housing density, which serves as a useful proxy for the city center in the traditional urban economic model of spatial structure. Figure 8 shows that this finding is not constant across all MSAs. While Los Angeles, Chicago, and Dallas exhibit a strong negative correlation between housing density and supply elasticity, Houston does not show any linear pattern, and New York actually has a positive correlation. Within metropolitan areas, it seems that the core areas are more likely to be less elastic, but it also depends on other factors that future versions of this paper will explore.

Across metropolitan areas, the picture is more straightforward. Figure 9 shows all counties in these five largest MSAs on the same scatterplot, and higher housing density is clearly associated with less supply elasticity. The relationship appears to be nonlinear, however, with the slope declining at higher housing densities. Beyond a certain level of housing density—somewhere around 1,000 units per square foot, it seems—supply elasticity cannot go much lower. This is a new and important finding that deserves further research.

[17] The lightest color is land for which we do not have any elasticities.

6 Conclusion

In this paper, we document local housing supply elasticity. Our intuition rests on the idea that when land prices rise, developers want to respond by building more intensively. Their ability to do so varies significantly across municipalities within the metropolitan area. We show significant heterogeneity in housing unit growth across county borders in the five largest U.S. MSAs and across city borders in Los Angeles County, and we find that much of this heterogeneity can be explained by variation in the price elasticity of supply.

We compute the first county-level and municipality-level indices of housing supply elasticities that have been constructed to our knowledge. We apply a structural vector autoregression model with sign restrictions to identify a positive demand shock. Though this model has been successfully applied in other contexts, particularly oil markets, it is an important advance for urban economics, which has traditionally suffered from endogeneity bias in its elasticity estimates. Using this sign-restricted VAR, we find that previous work has overestimated housing supply elasticities. On average, our elasticities are very low—in the 0 to 0.4 range, similar to findings in oil markets—but there is a long right tail in the distribution, suggesting that several counties and cities are responding to demand with much more elastic supply than others.

We conduct a cross-sectional analysis that reveals these elasticities to be mostly unrelated to price, population, or regulations as measured by the Wharton index. Cities and counties with lower density tend to have more elastic supply, as do cities with higher incomes. The most striking finding, however, is the clustering of elastic cities in locations on the outskirts of the city, far away from the city center and the valuable coastal land. The supply curve is least elastic in the city center, it appears, pushing demand out to the periphery.

We demonstrate the importance of structurally identifying shocks to accurately estimate the slope of the supply curve in response to demand shocks—and by extension, the degree to which construction activity can be attributed to differential supply elasticities that can be conflated with differences in housing demand based on cities' desirability. We also demonstrate that the aggregate findings of the past have masked important heterogeneity at the local level. Future research should strive to disentangle this variation across counties and municipalities and to help identify policies that can increase supply elasticity and reduce the burdens associated with the growing housing shortage. It should also extend this approach to other counties, cities, and time periods to document how elasticities vary depending on terrain, regulatory environment, and other institutional factors. Hopefully, this paper is only beginning as we unravel the more nuanced story that lies at this deeper level.

Table 1: County-Level Supply Elasticities in New York MSA

County	One-Year Elasticity	Housing Density
Passaic	0.024	918
Nassau	0.029	1,598
Sussex	0.032	108
Westchester	0.033	807
Essex	0.040	13
Dutchess	0.041	132
Bronx	0.046	11,675
Suffolk	0.052	573
Middlesex	0.052	884
Monmouth	0.054	510
Union	0.057	1,868
Orange	0.059	150
Richmond	0.062	2,804
Rockland	0.065	545
Queens	0.065	7,482
Putnam	0.065	152
Ocean	0.069	391
Bergen	0.071	1,451
Hudson	0.077	5,154
Hunterdon	0.082	105
Morris	0.096	372
Somerset	0.099	368
Kings	0.116	13,184

Notes: One-year impulse responses to sign-restricted VAR of one-standard-deviation positive demand shocks, divided to calculate price elasticity of supply, for all available counties in the New York metropolitan statistical area. Housing density is defined as units per square feet of land area.

Table 2: County-Level Supply Elasticities in Los Angeles MSA

County	One-Year Elasticity	Housing Density
Los Angeles	0.010	806
Orange	0.032	1,228
San Bernardino	0.034	30
Ventura	0.048	136
Riverside	0.057	81

Notes: One-year impulse responses to sign-restricted VAR of one-standard-deviation positive demand shocks, divided to calculate price elasticity of supply, for all available counties in the Los Angeles metropolitan statistical area. Housing density is defined as units per square feet of land area.

Table 3: County-Level Supply Elasticities in New York MSA

County	One-Year Elasticity	Housing Density
Cook	0.025	2,217
Lake (IN)	0.028	392
DuPage	0.055	1,006
Jasper	0.059	20
Porter	0.061	138
Lake (IL)	0.064	505
Kenosha	0.073	220
Kane	0.092	267
Will	0.119	210
Grundy	0.206	36
McHenry	0.362	154
Kendall	0.401	61

Notes: One-year impulse responses to sign-restricted VAR of one-standard-deviation positive demand shocks, divided to calculate price elasticity of supply, for all available counties in the Chicago metropolitan statistical area. Housing density is defined as units per square feet of land area.

Table 4: County-Level Supply Elasticities in New York MSA

County	One-Year Elasticity	Housing Density
Hood	0.047	45
Hunt	0.057	39
Parker	0.069	38
Dallas	0.085	971
Tarrant	0.119	655
Johnson	0.164	63
Ellis	0.170	42
Somervell	0.195	15
Kaufman	0.204	33
Denton	0.219	189
Collin	0.303	230
Rockwall	0.425	119

Notes: One-year impulse responses to sign-restricted VAR of one-standard-deviation positive demand shocks, divided to calculate price elasticity of supply, for all available counties in the Dallas metropolitan statistical area. Housing density is defined as units per square feet of land area.

Table 5: County-Level Supply Elasticities in New York MSA

County	One-Year Elasticity	Housing Density
Austin	0.062	16
Galveston	0.116	280
Harris	0.177	751
Brazoria	0.234	65
Montgomery	0.264	108

Notes: One-year impulse responses to sign-restricted VAR of one-standard-deviation positive demand shocks, divided to calculate price elasticity of supply, for all available counties in the Houston metropolitan statistical area. Housing density is defined as units per square feet of land area.

Table 6: Municipality-Level Supply Elasticities in Los Angeles County

City	Quantity Response	Price Response	Elasticity
Agoura Hills	0.0055	0.0272	0.202
Hidden Hills	0.0022	0.0158	0.142
Calabasas	0.0031	0.0250	0.124
Claremont	0.0020	0.0185	0.109
West Hollywood	0.0018	0.0187	0.097
Pasadena	0.0017	0.0176	0.096
La Verne	0.0021	0.0248	0.084
Temple City	0.0022	0.0269	0.082
Glendora	0.0016	0.0192	0.082
Santa Monica	0.0015	0.0189	0.081
Azusa	0.0026	0.0323	0.080
Santa Clarita	0.0023	0.0292	0.077
Santa Fe Springs	0.0028	0.0365	0.076
Glendale	0.0017	0.0237	0.070
San Dimas	0.0018	0.0271	0.067
Walnut	0.0015	0.0238	0.064
Monterey Park	0.0011	0.0172	0.063
South El Monte	0.0015	0.0267	0.055
Monrovia	0.0011	0.0201	0.052
Beverly Hills	0.0010	0.0200	0.048
Malibu	0.0005	0.0116	0.047
San Fernando	0.0020	0.0454	0.045
Hermosa Beach	0.0006	0.0146	0.042
El Monte	0.0012	0.0286	0.041
Diamond Bar	0.0008	0.0189	0.040
Redondo Beach	0.0006	0.0158	0.040
San Gabriel	0.0007	0.0173	0.039
La Mirada	0.0008	0.0220	0.038
Burbank	0.0008	0.0216	0.038
Inglewood	0.0011	0.0299	0.036
El Segundo	0.0010	0.0273	0.035
Alhambra	0.0008	0.0228	0.033
South Pasadena	0.0007	0.0224	0.032
Duarte	0.0008	0.0260	0.031
Montebello	0.0008	0.0260	0.030
Gardena	0.0009	0.0288	0.030
Commerce	0.0012	0.0419	0.029
Rancho Palso Verdes	0.0005	0.0191	0.028
Sierra Madre	0.0005	0.0200	0.027
Rolling Hills Estates	0.0005	0.0168	0.027
Torrance	0.0006	0.0239	0.027
Manhattan Beach	0.0005	0.0180	0.027
Palos Verdes Estates	0.0005	0.0171	0.026
Baldwin Park	0.0008	0.0321	0.026
Pomona	0.0008	0.0305	0.026
Rosemead	0.0006	0.0237	0.025
Bell Gardens	0.0008	0.0374	0.022
Los Angeles	0.0007	0.0329	0.021
Covina	0.0007	0.0331	0.020
South Gate	0.0008	0.0425	0.020
Long Beach	0.0005	0.0271	0.019
Carson	0.0005	0.0267	0.018
San Marino	0.0003	0.0205	0.016
Artesia	0.0005	0.0363	0.015
Hawaiian Gardens	0.0007	0.0448	0.015
La Canada Flintridge	0.0003	0.0226	0.014
Lynwood	0.0006	0.0432	0.013
Lawndale	0.0004	0.0284	0.013
Norwalk	0.0004	0.0338	0.012
Culver City	0.0002	0.0174	0.012
Avalon	0.0004	0.0441	0.009
Bellflower	0.0003	0.0332	0.009
Whittier	0.0002	0.0237	0.008
Paramount	0.0002	0.0247	0.007
Downey	0.0002	0.0371	0.006

Notes: One-year impulse responses to sign-restricted VAR of one-standard-deviation positive demand shocks, divided to calculate price elasticity of supply, for all available municipalities in Los Angeles County.

Table 7: Cross-Sectional Analysis of Short-Run Supply Elasticities in LA County

	(1)	(2)	(3)	(4)	(5)
Price	0.0000				
	(0.80)				
Density		-0.0027^{**}			
		(-2.80)			
Population			-0.0078		
			(-0.78)		
Income				0.0004^{**}	
				(2.84)	
WRLURI					0.0154
					(1.47)
Constant	0.0355^{**}	0.0640^{***}	0.0452^{***}	0.0203^{*}	0.0356^{**}
	(2.97)	(7.81)	(9.77)	(2.14)	(3.45)
Observations	65	65	65	65	25
R^2	0.0101	0.1107	0.0096	0.1132	0.0864
p-value	0.4248	0.0068	0.4370	0.0061	0.1538

OLS regressions of housing supply elasticities on city-level variables. t statistics in parentheses. Income, density, population, and income collected by Census Bureau and geocoded by GeoLytics.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 1: Municipal Housing Unit Growth in Los Angeles County, 2000-2012

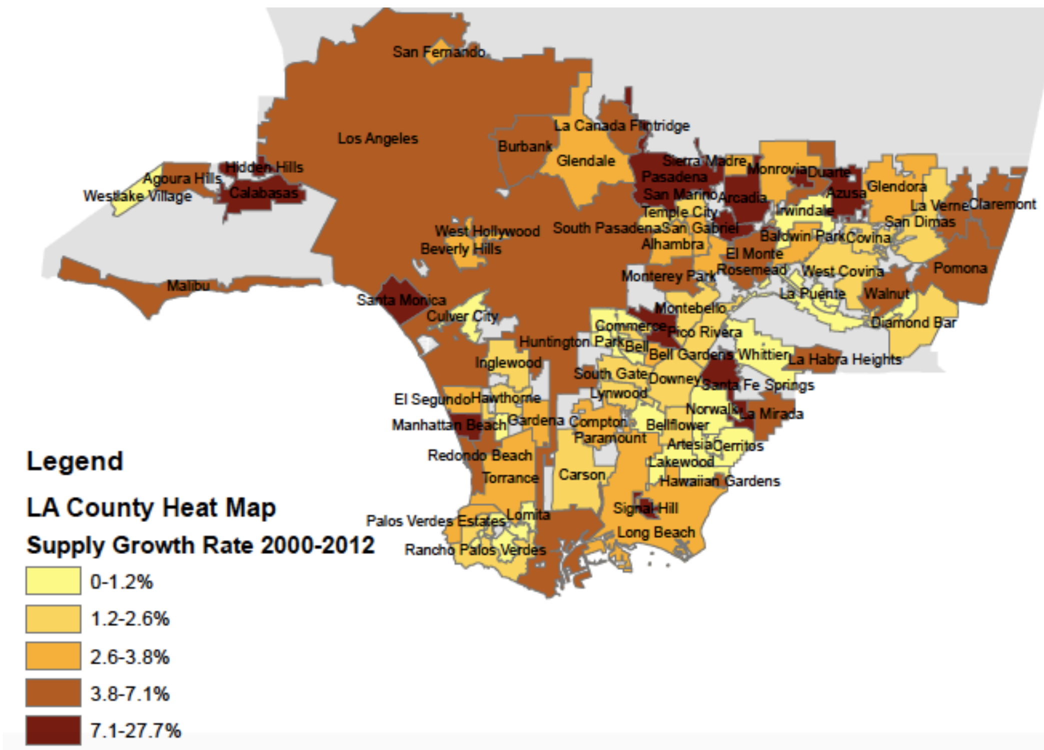

Notes: Percentage growth in number of housing units by municipality in Los Angeles County. Original housing unit data obtained from 2000 Census. Growth rates calculated using annual building permit data from U.S. Census Bureau.

Figure 2: Los Angeles County Permits vs. Prices, 1988-2012

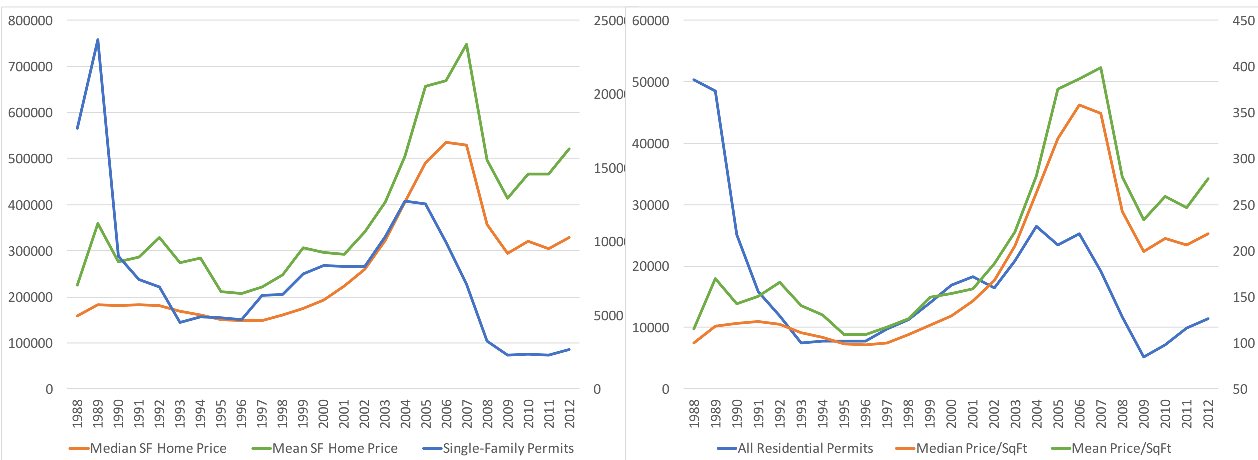

(a) Single-Family Homes (b) All Residential

Permits on left axis; prices in US dollars ($) on right axis, calculated using transaction-level data from DataQuick to estimate further back in time than Zillow will allow. These are not the prices used in the rest of the paper.

Figure 3: SVAR Impulse Responses: Los Angeles County

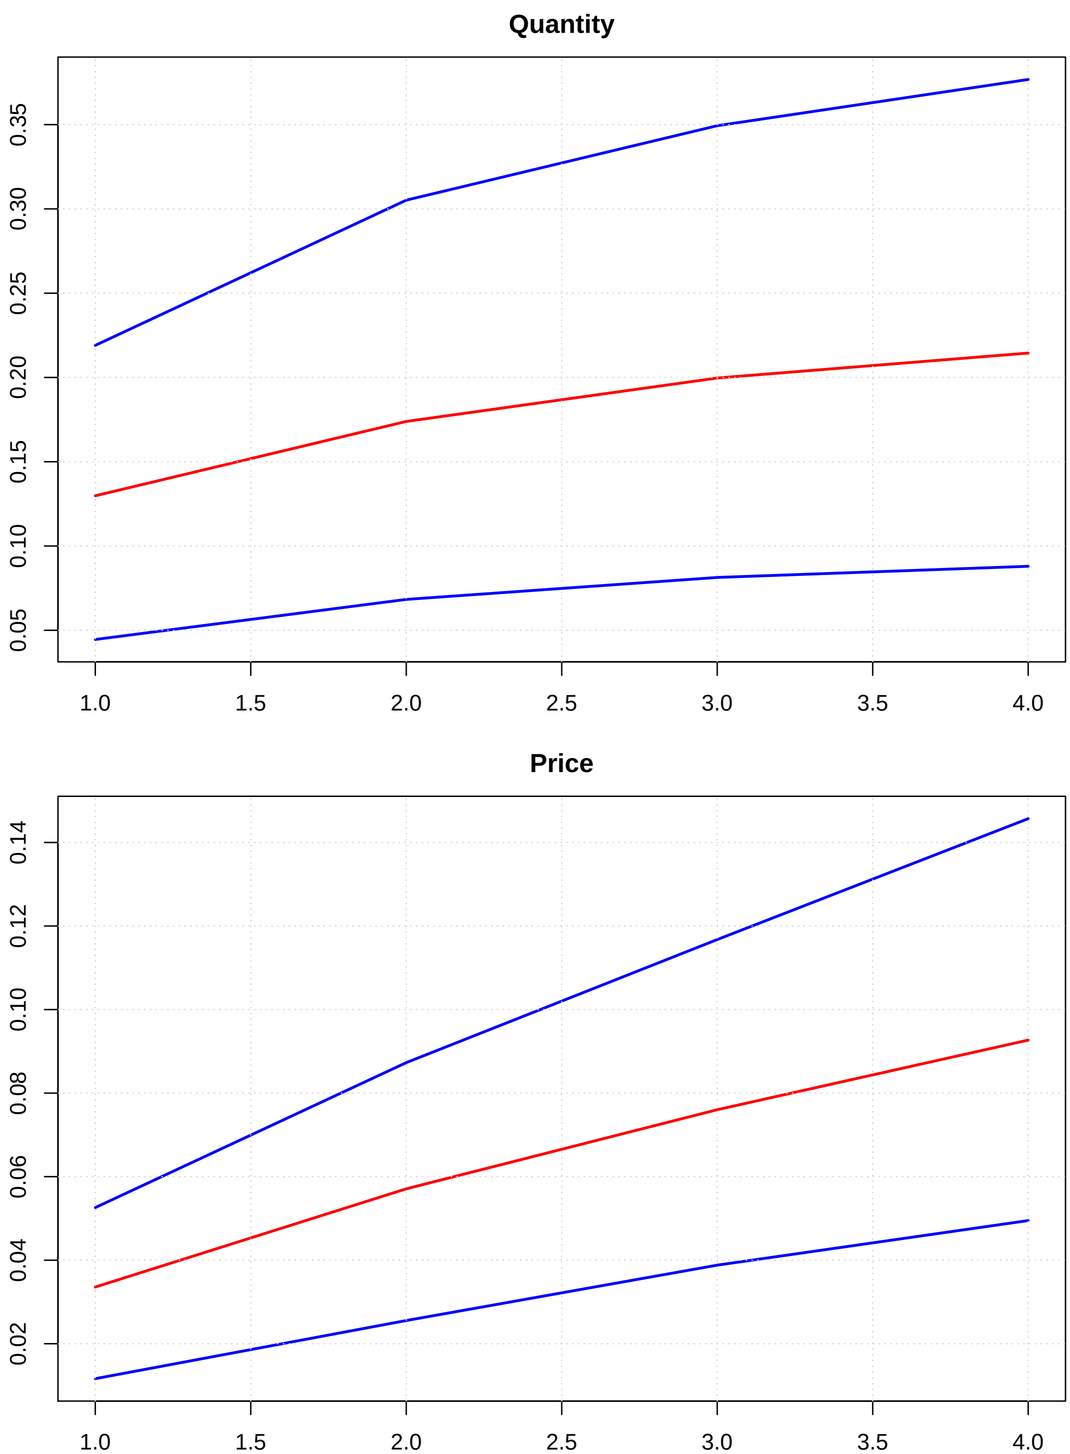

Impulse response to one-standard-deviation positive demand shock using sign-restricted VAR(2) model for Los Angeles County. Units are rough percentages because all variables are estimated using natural logs. Each period equals one year. Estimated using Uhlig's (2005) rejection method. Original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

Figure 4: SVAR Impulse Responses: City of Los Angeles

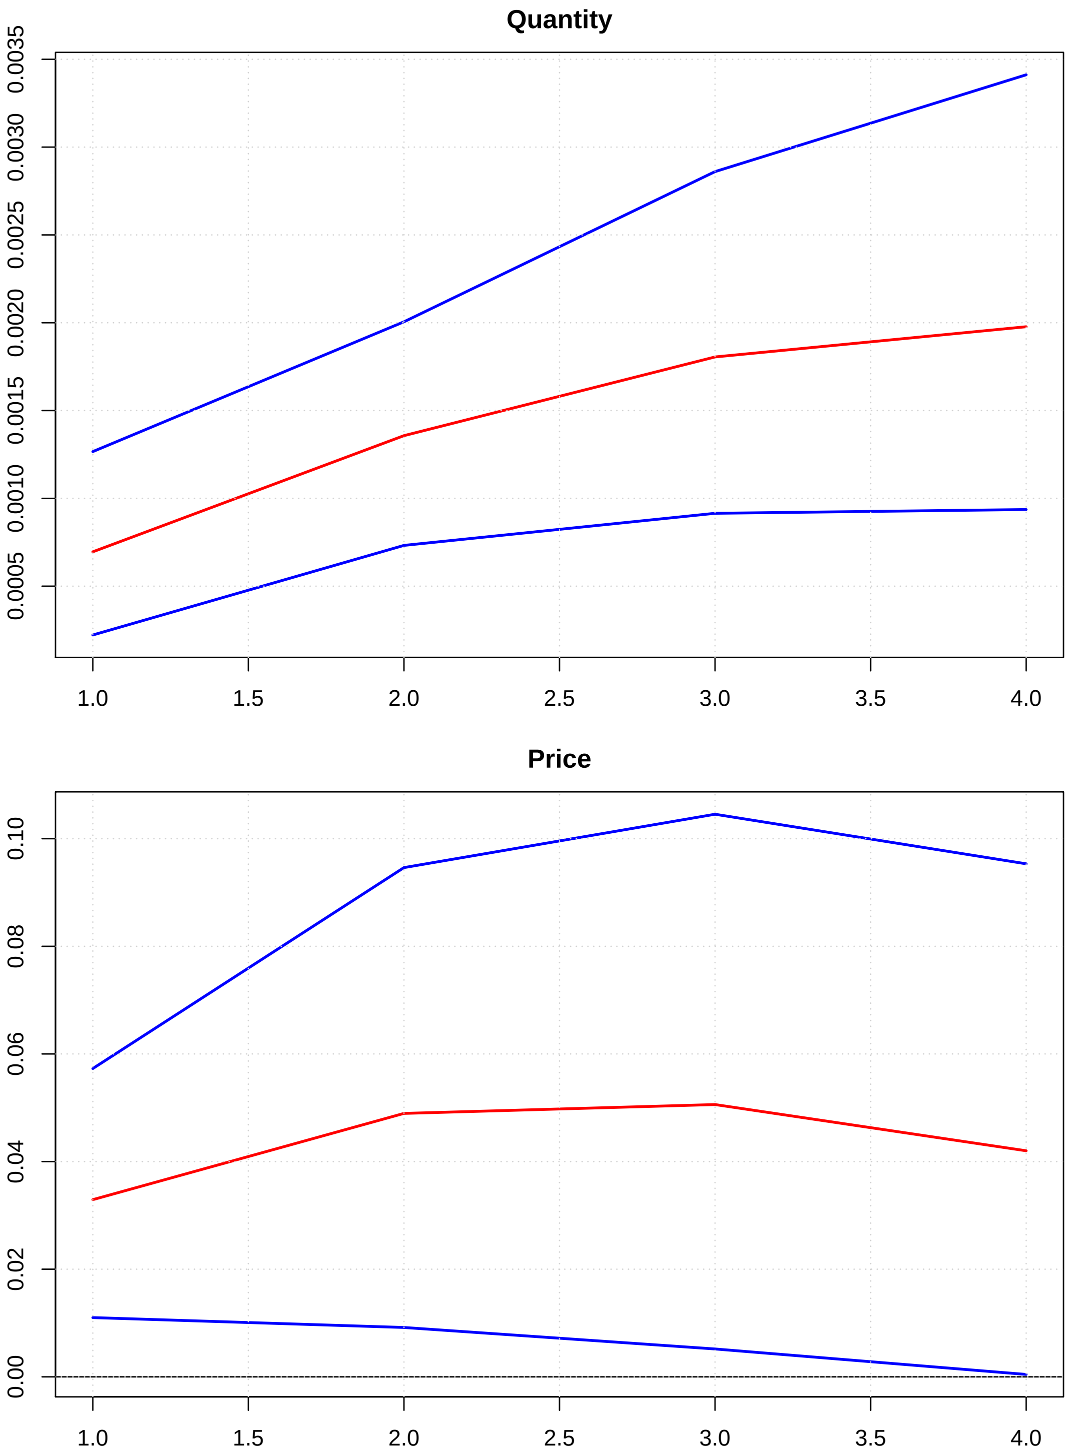

Impulse response to one-standard-deviation positive demand shock using sign-restricted VAR(2) model for city of Los Angeles. Units are rough percentages because all variables are estimated using natural logs. Each period equals one year. Estimated using Uhlig's (2005) rejection method. Original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

Figure 5: SVAR Impulse Responses: City of Calabasas

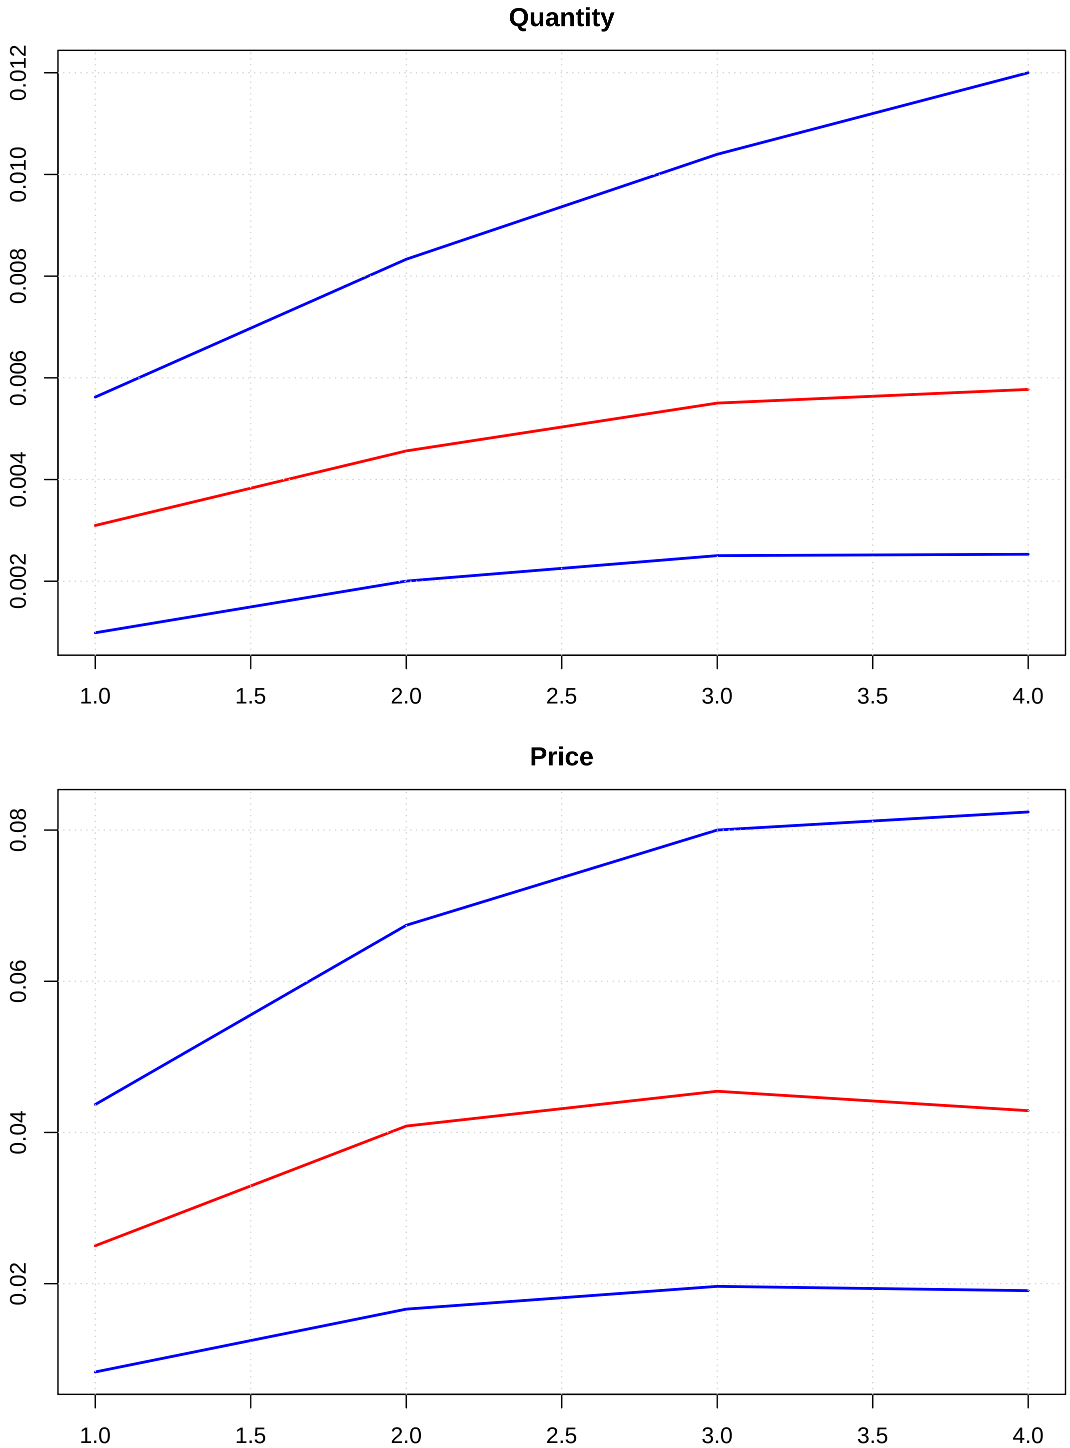

Impulse response to one-standard-deviation positive demand shock using sign-restricted VAR(2) model for city of Calabasas. Units are rough percentages because all variables are estimated using natural logs. Each period equals one year. Estimated using Uhlig's (2005) rejection method. Original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

Figure 6: Housing Supply Elasticity Distributions

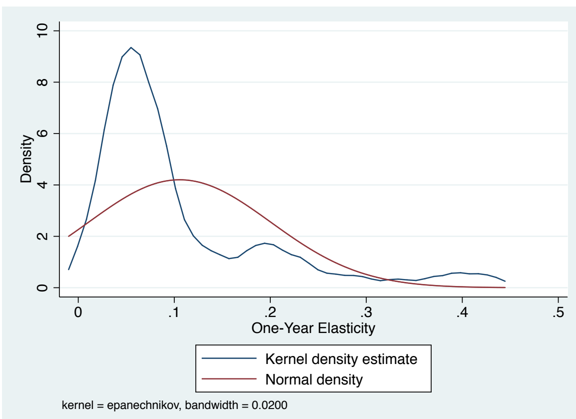

(a) Counties in Five Largest MSAs

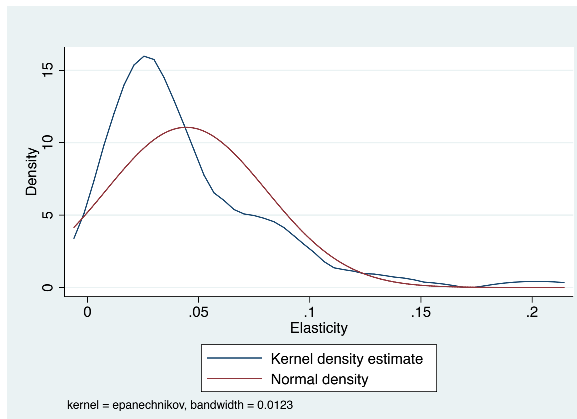

(b) Municipalities in Los Angeles County

Notes: One-year elasticities calculated using impulse responses from sign-restricted VAR(2) model for (a) each county in the five largest MSAs and (b) each city in Los Angeles County. Original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

Figure 7: Municipal Housing Supply Elasticities in Los Angeles County, 1996-2015

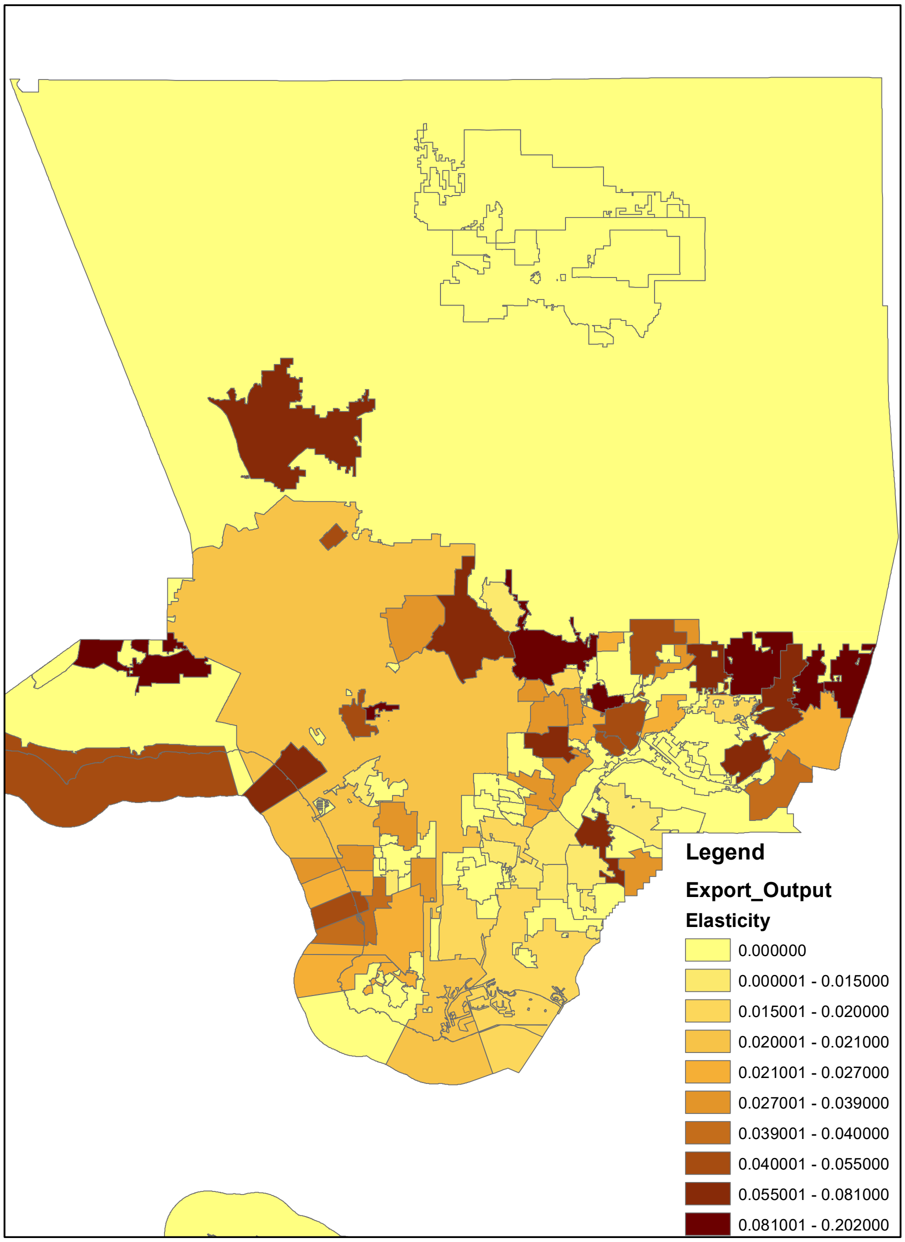

Notes: Short-run (one-year) elasticities calculated using impulse responses from sign-restricted VAR(2) model for each city in Los Angeles County. Original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow. Heat map colors divided by decile.

Figure 8: One-Year County-Level Housing Supply Elasticities by Metropolitan Area

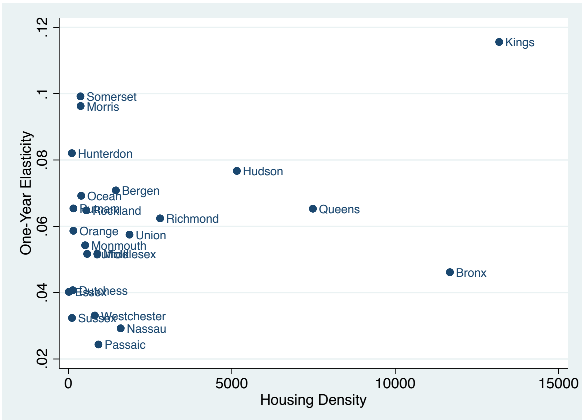

(a) New York MSA

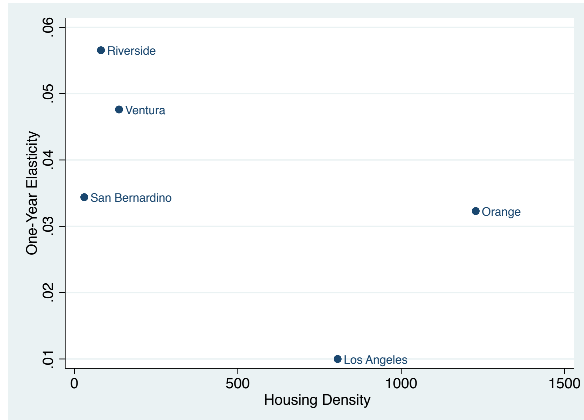

(b) Los Angeles MSA

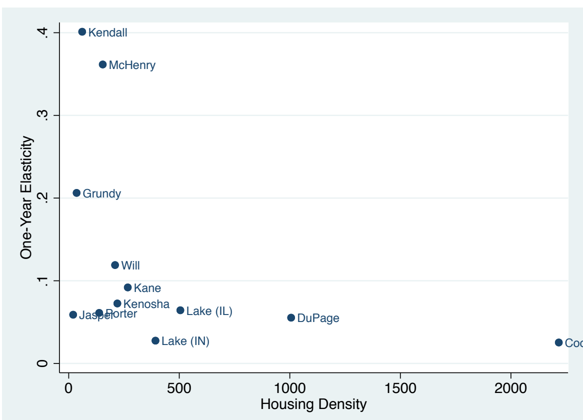

(c) Chicago MSA

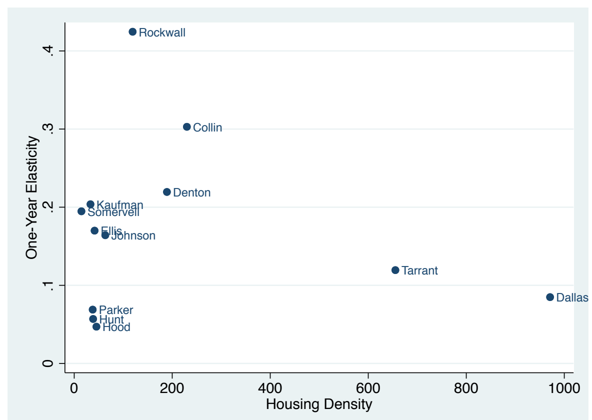

(d) Dallas MSA

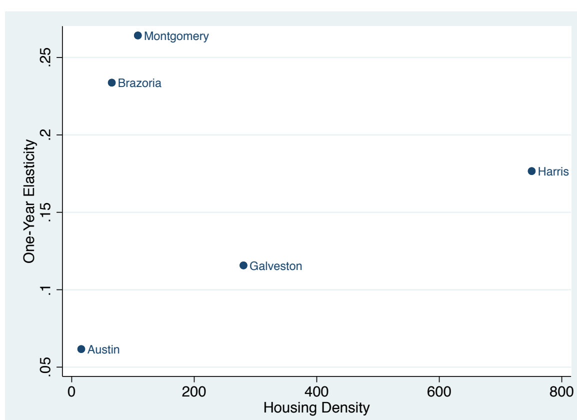

(e) Houston MSA

One-year elasticities calculated using impulse responses from sign-restricted VAR(2) model for each available county in the five largest U.S. metropolitan areas. Housing density (units per square foot of land area) and original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

Figure 9: One-Year County-Level Housing Supply Elasticities

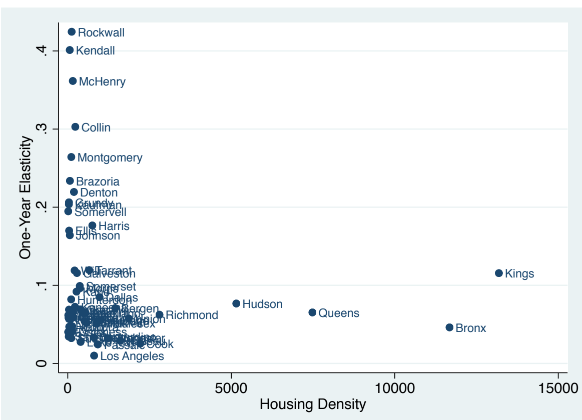

(a) Five Largest MSAs

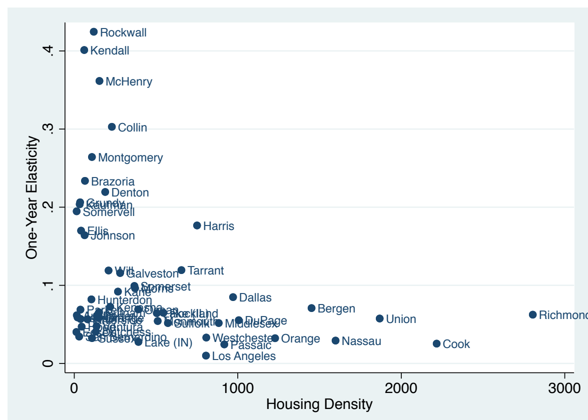

(b) Excluding Density $> 5,000$ Units/SqFt

One-year elasticities calculated using impulse responses from sign-restricted VAR(2) model for each available county in the five largest U.S. metropolitan areas. Housing density (units per square foot of land area) and original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

Figure 10: Long-Run Elasticity Distribution in Los Angeles County, 1996-2015

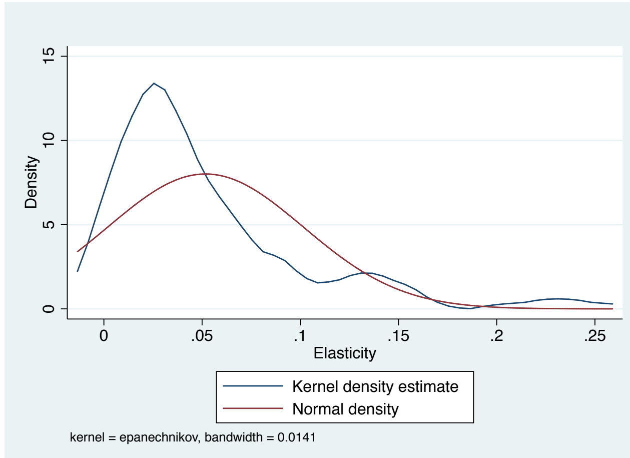

Notes: One-year elasticities calculated using impulse responses from sign-restricted VAR(2) model for each city in Los Angeles County. Original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

A "Long-Run" Supply Elasticities

Table 8 lists the four-year municipality-level elasticities, which are not much higher than the one-year elasticities shown above. In many cases, it appears that prices continue to rise, but quantity still has a very muted response. While there may be a lag in production, there simply appears to be very little construction relative to the pent-up demand. Figure 11 plots each city's long-run elasticity against its short-run elasticity. If the short-run elasticity captured the entire response, we would expect all the points to align on the 45-degree line. This is not exactly the case, but the majority of the points are near to it. More cities have higher elasticities in the long run, suggesting that quantity does catch up a bit to prices, but many cities experience the opposite, suggesting that the demand is simply never met and the housing shortage grows over time.

Table 8: Long-Run Supply Elasticities, 1996-2015

City	Quantity Response	Price Response	Elasticity
Hermosa Beach	0.0009	0.0038	0.245
Claremont	0.0024	0.0111	0.218
La Mirada	0.0020	0.0129	0.152
Diamond Bar	0.0010	0.0068	0.143
Temple City	0.0066	0.0467	0.140
Calabasas	0.0058	0.0429	0.135
Hidden Hills	0.0032	0.0238	0.133
Redondo Beach	0.0019	0.0160	0.117
Agoura Hills	0.0055	0.0487	0.113
La Verne	0.0029	0.0317	0.090
Santa Fe Springs	0.0056	0.0644	0.087
Pasadena	0.0036	0.0451	0.081
Pomona	0.0012	0.0155	0.080
South El Monte	0.0028	0.0363	0.078
West Hollywood	0.0078	0.1027	0.076
Azusa	0.0048	0.0676	0.071
Burbank	0.0026	0.0381	0.069
Torrance	0.0017	0.0247	0.067
Commerce	0.0024	0.0367	0.066
Glendale	0.0027	0.0525	0.051
San Gabriel	0.0011	0.0216	0.050
Rosemead	0.0013	0.0265	0.049
San Fernando	0.0041	0.0845	0.049
Los Angeles	0.0020	0.0420	0.047
Manhattan Beach	0.0014	0.0305	0.046
Santa Clarita	0.0019	0.0408	0.046
Sierra Madre	0.0012	0.0266	0.045
El Segundo	0.0017	0.0403	0.042
Glendora	0.0021	0.0498	0.042
Monterey Park	0.0011	0.0287	0.039
Walnut	0.0022	0.0591	0.038
Bell Gardens	0.0018	0.0506	0.036
Santa Monica	0.0020	0.0576	0.035
El Monte	0.0027	0.0762	0.035
South Pasadena	0.0015	0.0460	0.033
Beverly Hills	0.0009	0.0287	0.032
Malibu	0.0003	0.0098	0.030
Hawaiian Gardens	0.0019	0.0652	0.030
Rolling Hills Estates	0.0008	0.0279	0.028
Palos Verdes Estates	0.0007	0.0249	0.027
San Dimas	0.0014	0.0548	0.026
Inglewood	0.0013	0.0565	0.024
La Canada Flintridge	0.0009	0.0398	0.023
Artesia	0.0015	0.0692	0.022
Gardena	0.0009	0.0418	0.022
Lawndale	0.0005	0.0255	0.021
Long Beach	0.0010	0.0521	0.020
Rancho Palso Verdes	0.0003	0.0162	0.020
Covina	0.0011	0.0563	0.020
Lynwood	0.0008	0.0437	0.019
South Gate	0.0012	0.0657	0.019
Carson	0.0005	0.0335	0.016
Baldwin Park	0.0011	0.0681	0.016
Monrovia	0.0006	0.0445	0.014
Whittier	0.0010	0.0740	0.013
Bellflower	0.0007	0.0546	0.013
Montebello	0.0004	0.0352	0.013
San Marino	0.0008	0.0633	0.012
Downey	0.0006	0.0549	0.012
Norwalk	0.0003	0.0442	0.007
Avalon	0.0003	0.0461	0.007
Alhambra	0.0005	0.0708	0.006
Duarte	0.0002	0.0422	0.006
Culver City	0.0001	0.0319	0.003
Paramount	0.0000	0.0487	0.001

Notes: Four-year impulse responses to sign-restricted VAR of one-standard-deviation positive demand shocks, divided to calculate price elasticity of supply, for all available municipalities in Los Angeles County.

Figure 11: Long-Run vs. Short-Run Elasticity by City, 1996-2015

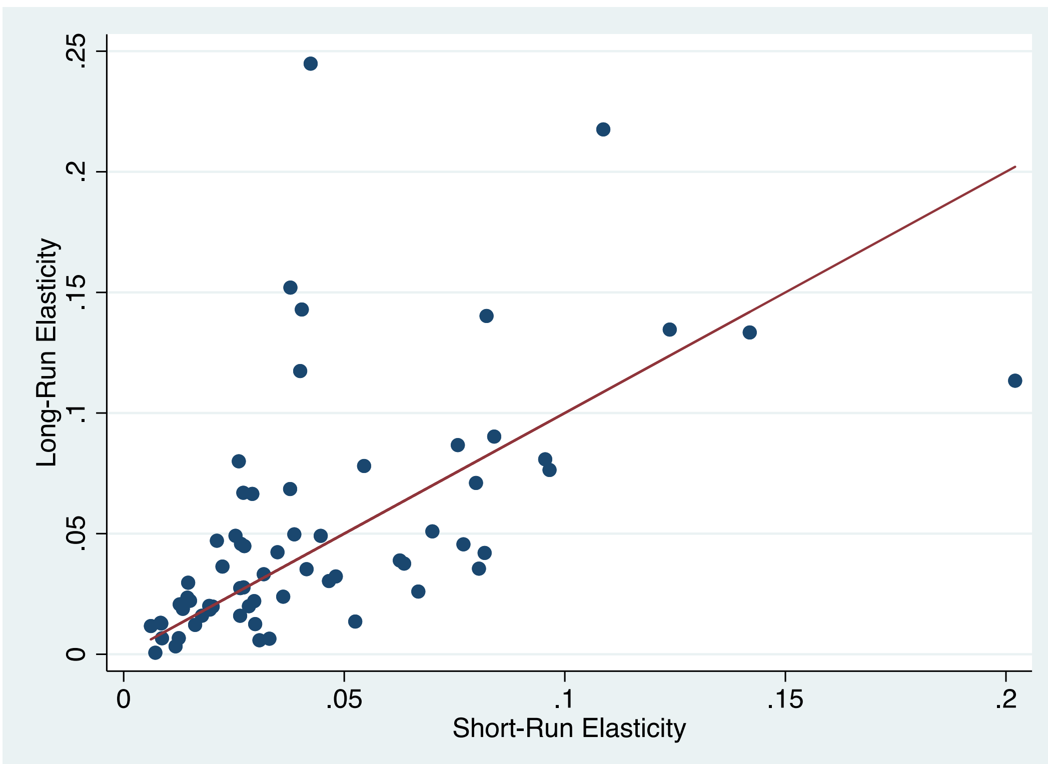

Elasticities calculated using impulse responses from sign-restricted VAR(2) model for each city in Los Angeles County. "Short-run" = 1 year. "Long-run" = 4 years. 45-degree line shows where cities would be if their elasticities did not change over time. Original housing unit data obtained from 2000 Census. Annual growth calculated using building permit data from U.S. Census Bureau. Price data obtained from Zillow.

References

Albouy, D., G. Ehrlich, and Y. Liu (2016): "Housing Demand, Cost-of-Living Inequality, and the Affordability Crisis," NBER Working Paper No. 22816.

Baum-Snow, N. (2007): "Did Highways Cause Suburbanization?," *Quarterly Journal of Economics*, 122(2), 775–805.

Baumeister, C., and J. D. Hamilton (2015): "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," *Econometrica*, 83(5), 1963–1999.

Bayer, P., H. Fang, and R. McMillan (2014): "Separate When Equal? Racial Inequality and Residential Segregation," *Journal of Urban Economics*, 82, 32–48.

Bayer, P., F. Ferreira, and R. McMillan (2007): "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 115(4), 588–638.

Bayer, P., R. McMillan, A. Murphy, and C. Timmins (2016): "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, 84(3), 893–942.

Brueckner, J. K. (1987): "The Structure of Urban Equilibria: A Unified Treatment of the Muth-Mills Model," in *Handbook of Regional and Urban Economics*, ed. by E. Mills, vol. 2. Elsevier Science Publishers B.V.

Case, K. E., and R. J. Shiller (1989): "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 79(1), 125–137.

Cleveland, R. B., W. S. Cleveland, J. E. McRae, and I. Terpenning (1990): "STL: A Seasonal-Trend Decomposition Procedure Based on Loess," *Journal of Official Statistics*, 6(1), 3–73.

Dovey, R. (2015): "Anti-Growth Lawsuits Are All the Rage," *Next City*, March 31, https://nextcity.org/daily/entry/environment–lawsuits–eir–bloomberg–bike–lanes.

Englund, P., J. M. Quigley, and C. L. Redfearn (1999): "The Choice of Methodology for Computing Housing Price Indexes: Comparisons of Temporal Aggregation and Sample Definition," *Journal of Real Estate Finance and Economics*, 19(2), 91–112.

Fry, R., and A. Pagan (2011): "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 49(4), 938–960.

Ganong, P., and D. Shoag (2015): "Why Has Regional Income Convergence in the U.S. Declined?," *Journal of Urban Economics*, 102, 76–90.

Glaeser, E. L. (2004): "Housing Supply," http://www.nber.org/reporter/spring04/glaeser.html.

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

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.

Green, R. K., S. Malpezzi, and S. K. Mayo (2005): "Metropolitan-Specific Estimates of the Price Elasticity of Supply of Housing, and Their Sources," *American Economic Review*, 95(2), 334–339.

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

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.

Harter-Dreiman, M. (2004): "Drawing Inferences About Housing Supply Elasticity from House Price Responses to Income Shocks," *Journal of Urban Economics*, 55(2), 316–337.

Heblich, S., A. Trew, and Y. Zylberberg (2016): "East Side Story: Historical Pollution and Persistent Neighborhood Sorting," Working Paper, `https://www.st-andrews.ac.uk/~wwwecon/repecfiles/4/1613.pdf`.

Henderson, R. (1916): "Note on Graduation by Adjusted Average," *Transactions of the American Society of Actuaries*, 17, 43–48.

Hsieh, C.-T., and E. Moretti (2017): "Housing Constraints and Spatial Misallocation," NBER Working Paper No. 21154.

Kilian, L. (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99(3), 1053–1069.

Kilian, L., and D. P. Murphy (2012): "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models," *Journal of the European Economic Association*, 10(5), 1166–1188.

Kusisto, L., and J. Light (2015): "Zillow, After Trulia Deal, Calls 2015 'a Transition Year'," *New York Times*, April 14.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015): "The Housing Market(s) of San Diego," *American Economic Review*, 105(4), 1371–1407.

Malone, T., and C. L. Redfearn (2016): "Stability & Change: The Durable Hierarchy of Neighborhoods in U.S. Metropolitan Areas from 1970 to 2010," Working Paper, `https://www.aeaweb.org/conference/2017/preliminary/paper/kYn7aNHb`.

Malpezzi, S., and D. Maclennan (2001): "The Long-Run Price Elasticity of Supply of New Residential Construction in the United States and the United Kingdom," *Journal of Housing Economics*, 10(3), 278–306.

Marshall, A. (1890): *Principles of Economics*. Macmillan and Co., Ltd.

Mayer, C. J., and C. T. Somerville (2000): "Land Use Regulation and New Construction," *Regional Science and Urban Economics*, 30(6), 639–662.

Moretti, E. (2013): "Real Wage Inequality," *American Economic Journal: Economic Policy*, 5(1), 65–103.

Morrow, G. D. (2013): "The Homeowner Revolution: Democracy, Land Use and the Los Angeles Slow-Growth Movement, 1965-1992," Unpublished Dissertation, University of California Los Angeles, http://escholarship.org/uc/item/6k64g20f.

Nieuwerburgh, S. V., and P.-O. Weill (2010): "Why Has House Price Dispersion Gone Up?," *Review of Economic Studies*, 77(4), 1567–1606.

Orlando, A. W. (2018): "An Econometric Analysis of Housing Supply Dynamics in the United States, 1959-2015," Working Paper.

Parkhomenko, A. (2017): "The Rise of Housing Supply Regulation in the U.S.: Local Causes and Aggregate Implications," Working Paper, `https://www.andrii-parkhomenko.net/files/Parkhomenko_JMP.pdf`.

Preston, A. J. (1978): "Concepts of Structure and Model Identifiability for Econometric Systems," in *Stability and Inflation: A Volume of Essays to Honour the Memory of A. W. H. Phillips*, ed. by A. R. Bergstrom, A. J. L. Catt, M. H. Peston, and B. D. J. Silverstone. Wiley.

Pryce, G. (1999): "Construction Elasticities and Land Availability: A Two-Stage Least-Squares Model of Housing Supply Using the Variable Elasticity Approach," *Urban Studies*, 36(13), 2283–2304.

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

Roback, J. (1982): "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 90(6), 1257–1278.

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

Schuetz, J. (2009): "No Renters in My Suburban Backyard: Land Use Regulation and Rental Housing," *Journal of Policy Analysis and Management*, 28(2), 296–320.

Sinai, T. (2014): "The Rental Affordability Crisis," *Wharton Public Policy Initiative Issue Brief*, 2(3).

Smith, B. A. (1976): "The Supply of Urban Housing," *Quarterly Journal of Economics*, 90(3), 389–405.

Taylor, M. (2015): "California's High Housing Costs: Causes and Consequences," Legislative Analyst's Office, State of California, http://www.lao.ca.gov/reports/2015/finance/housing-costs/housing-costs.pdf.

Tiebout, C. M. (1956): "A Pure Theory of Local Expenditures," *Journal of Political Economy*, 64(5), 416–424.

Uhlig, H. (2005): "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, 52(2), 381–419.

Watson, N. E., B. L. Steffen, M. Martin, and D. A. Vandenbroucke (2017): "Worst Case Housing Needs: 2017 Report to Congress," U.S. Department of Housing and Urban Development, https://www.huduser.gov/portal/sites/default/files/pdf/Worst-Case-Housing-Needs.pdf.

Waxman, A. R. (2017): "The Long Road to Work: Divergent Effects of Transportation Policies by Worker Skill in a Locational Sorting Model," Working Paper.

Wetzstein, S. (2017): "The Global Urban Housing Affordability Crisis," *Urban Studies*, 54(14), 3159–3177.

Wheaton, W. C., S. Chervachidze, and G. Nechayev (2014): "Error Correction Models of MSA Housing "Supply" Elasticities: Implications for Price Recovery," Working Paper, https://economics.mit.edu/files/9472.