

The Geography Channel of House Price Appreciation: Did the Manufacturing Decline Partially Cause the Housing Boom?

Greg Howard* Jack Liebersohn†

November 21, 2017

Abstract

Locational preferences contributed substantially to the United States' 2000-2006 housing boom. The increasing desirability of inelastic areas, much of which was due to manufacturing's decline, caused 28 percent of the rise in housing prices. We document population movements towards inelastic areas and create a new local rent index that shows rents and house prices co-moved. We show theoretically why an increase in desirability of inelastic areas would raise prices nationally. Our model also explains the geographic consequences of a national price-rent ratio change. Finally, we quantify the geography channel by creating new elasticity measures of that cover the entire country.

*University of Illinois, Urbana-Champaign, glhoward@illinois.edu.

†MIT Sloan, liebers@mit.edu. We would like to thank Dan Greenwald, Alp Simsek, Ivan Werning, Arnaud Costinot, David Autor, and Antionette Schoar for their comments. Jack would like to thank the Boston Federal Reserve Bank for their hospitality and the Becker Friedman Institute for financial support through the MFM Fellowship. Greg would like to thank the National Science Foundation for financial support through the Graduate Research Fellowship (Grant No. 2014136625). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. All errors are our own.

1 Introduction

House prices rose significantly in the early 2000s. The boom was spatially heterogeneous, with some areas seeing house prices more than double, and some areas seeing almost no increase. Economists have largely focused on two theories: an increase in credit supply¹ or a shift in expectations of future house prices.² Often, these two theories are judged based on the extent that they can match cross-sectional patterns, such as differences in price growth across regions.

In this paper, we argue that the cross-sectional patterns are not only a testing ground for different theories but are a crucial cause of the boom itself. For a variety of reasons, markets with an inelastic housing supply became desirable places to live. One such force is the decline in manufacturing, an industry highly concentrated in elastic areas (Liebersohn, 2017).³ In equilibrium, house prices rose nationally as people chose to move to inelastic areas.

Figure 1 illustrates the forces we highlight in this paper. Around 2000, there was a marked decline in manufacturing (Charles, Hurst and Notowidigdo, 2016; Autor, Dorn and Hanson, 2013). As would be expected in most urban models,⁴ this created a large gap in house prices between high-manufacturing areas, such as the Midwest, and the rest of the country. For market clearing to hold, this implied a rise in aggregate house prices because the Midwest has a relatively elastic housing supply.

¹For example, several papers have argued that the increase in subprime mortgages was a key cause of the boom, e.g. Mian and Sufi (2009). Other papers have argued that falling interest rates are particularly important such as Mayer and Sinai (2009), while others have emphasized changes to borrowing constraints, such as Greenwald (2016).

²For example, Glaeser and Nathanson (2017), DeFusco, Nathanson and Zwick (2017), and Foote, Gerardi and Willen (2012). Kaplan, Mitman and Violante (2017) also suggest changing expectations play a large role in the boom-bust cycle of house prices in a quantitative model. They also suggest a large role for credit supply in explaining other phenomenon during the time period.

³Davidoff (2016) further argues that housing supply elasticity is correlated with education, immigration, religion, and many other demand factors, concluding that using supply elasticity as an instrument for housing prices may not satisfy the exclusion restriction without controlling for demand shocks.

⁴Models in the tradition of Rosen (1979) and Roback (1982) imply that differences in housing prices reflect differences in wages across regions, and manufacturing decline is widely thought to be one of the main drivers of differential wage changes during the time period we study.




Figure 1: House Prices in the Midwest and the Rest of the U.S.

This paper makes several contributions. First, we show that changing demand for living in different regions explains many of the cross-sectional patterns of population and housing. We show a tight relationship between the growth in total housing units and population growth, and that this feeds through to house prices. We show that places that had lower initial manufacturing shares and higher initial costs of living experienced larger increases in housing quantities and prices. In contrast, there was no relationship between these variables and the amount of housing per capita, implying that population changes were the main force for cross-sectional housing demand. To provide further evidence that the cross-sectional differences are primarily due to changes in locational preferences, we show in a new dataset that the patterns of changes in rents are quite similar. Focusing on manufacturing, which experienced a widely-documented decline in labor demand during the early 2000s, we also show that the relative decline in house prices in manufacturing-intensive areas has lasted to at least 2015. The fact that the difference in house prices has persisted for many years following after the peak of the housing boom suggests that this difference is not due to factors specific to that period.

As a second contribution, we formalize a theory of why a shift in demand from high-elasticity to low-elasticity areas would increase national house prices. We show that a positive covariance between housing supply elasticity and the change in *local* housing prices leads to aggregate house price increases. The reason is that demand shocks cause a greater increase in prices in inelastic areas than elastic areas. We extend the basic model by describing another channel that operates through the supply of construction materials. Using formulas we derive from these models, we show how to quantify the effect of the decline in manufacturing on aggregate house prices through the geography channel.

In a similar vein, we show in a simple model how a change in the national price-rent ratio (for example, due to a change in interest rates) would change cross-sectional house prices and, in turn, affect national prices.⁵ The key force in this model is that if the price-rent ratio increases, then conditional on house prices, more expensive cities become more affordable. In equilibrium, the greater affordability causes higher demand and therefore a greater change in prices in the initially expensive cities. If, *ex ante*, rents and elasticity are negatively correlated, this will lead to an increase in aggregate prices.

Third, we construct a new housing supply elasticity measure for commuting zones (CZs) and show this is strongly negatively correlated to house price increases. Using the formula from the theory we developed, we estimate the geography channel increased house prices by at least 0.089 log-points, almost 30 percent of the total boom in real terms. Other measures of elasticity do not have the national coverage that this measure provides, which is essential to consider the aggregate effect that comes through changing relative demand across all regions (and in the model, this requires applying a national market clearing condition). Thus, we show that the correlation with manufacturing shares from 2000 can explain about a third of the geography channel.

Finally, we show that our theory predicts a housing boom during the time when the boom actually occurred.

⁵In our model, agents do not consume a larger quantity of housing, shutting down a channel which other papers have highlighted. The only choice agents have in the housing market is their choice of location.

This paper contributes to several strands of the literature on the role of housing in the macroeconomy. Most closely related is Charles et al. (2016), which establishes the parallel timing of the manufacturing decline and the housing boom, which had offsetting employment effects. To our knowledge, ours is the first paper to suggest this timing may not have been a coincidence, but that the manufacturing decline led to the housing boom. ⁶

Other papers have tried to quantify the macroeconomic effects of the elasticity of housing supply. Most similarly, Garriga, Hedlund, Tang and Wang (2017) argues that rural-to-urban migration and land use restrictions led to a housing boom in China. Hsieh and Moretti (2017) quantifies the effects that housing constraints have on aggregate productivity. Herkenhoff, Ohanian and Prescott (2017) suggests the recent decline in trend growth can be attributed to higher land-use restrictions. This is consistent with evidence from Ganong and Shoag (2017). Here, we suggest that these restrictions played a major role in the cyclical aspect of the macroeconomy as well.

Finally, there is a large overlap with papers that investigate the interaction between migration and housing. Van Nieuwerburgh and Weill (2010) suggests that changing wages can explain a significant share of house price dispersion from 1975-2005. Glaeser and Gyourko (2005) suggests housing markets may explain sluggish population responses. Davis, Fisher and Veracierto (2014) investigate this hypothesis and argue other migration frictions are more important. These models have their bases in Rosen (1979) and Roback (1982), as does ours. By closing the model, we add an aggregate market-clearing condition on housing from which we can analyze the effects of relative location demand on aggregate house prices.

2 The Cross-Section of Local Housing Demand

This section presents evidence that changing patterns of regional demand for housing played an important role for explaining the cross-section of house price growth from 2000-2006. We

⁶Autor, Dorn and Hanson (2017) discuss the fact that the manufacturing decline in general, and the China shock in particular, may have affected the cross section of housing prices.

begin in Section 2.1 by showing the strong relationship between population growth and housing demand. We also show that, cross-sectionally, house price growth is *not* well explained by cross-sectional changes in housing quality or size, and that population movements in turn are associated with exposure to the manufacturing sector. Section 2.2 presents evidence that similar patterns held for rents as for housing prices. This finding provides further evidence that cross-sectional house price changes reflected changes in relative demand. To show this, we create a new index of multifamily operating incomes using data from Commercial Mortgage-Backed Security (CMBS) records. Finally, in Section 2.4, we present evidence that allows us to reject an important alternative hypothesis - that regional patterns in house price changes are explained by regional changes in mortgage availability.

2.1 Population Growth and Housing Prices

This subsection presents evidence supporting the model’s key mechanism – that cross-sectional variation in housing prices was caused by demand to live in particular areas. To provide evidence for this mechanism, we present three stylized facts. 1) Almost all of the cross-sectional changes in housing units is due to population change. 2) Low-manufacturing areas and high-initial-rent areas had the greatest population and price increases. 3) These areas did not have greater increases in either the number of houses per capita or the quality of housing.

Figure 2 shows the relationship between population growth and housing unit growth from 2000 to 2006. The two variables are almost perfectly correlated: Population changes nearly one-for-one with housing units. Moreover, the R^2 in this relationship is 0.91– after accounting for population growth, very little of housing unit growth remains unexplained. Moreover, population growth and housing price growth are positively correlated, shown in Figure 3.

Next, we provide evidence that exposure to the manufacturing sector affected populations and prices in precisely the way one would expect from a regional demand shock.




Figure 2: Population Growth and Housing Units

The relationship between manufacturing and population is shown in left half of Figure 4. From 2000-2006 there was an overall net migration from areas with high to areas with low manufacturing. On average, a 10 percent higher fraction of the population employed in manufacturing corresponded to lower population growth of 0.007 log-points. A similar pattern holds for price changes, shown in the right half of Figure 4. Low-manufacturing areas had greater price increases as well.

Similarly, we see data consistent with an increase in demand for places with higher initial rents. Data for rents comes from the 2000 ACS, and we use the median rent by commuting zone.⁷ As can be seen in Figure 5, places with higher initial rents experienced larger increases in population and larger increases in house prices.

Along these two dimensions (manufacturing and initial rents), the increase in cross-sectional housing demand is driven by population changes, not changes in per-capita housing. This can be seen in Figure 6. The y-axis scales on these graphs were adjusted to be comparable to the graphs that showed the change in population. There is no statistically

⁷The public ACS data classifies geographies at the PUMA level, which we convert to commuting zones using the Missouri Census Data Center MABLE/Geocorr web application. This could lead to some mis-measurement, but we expect that to be small, as housing costs are typically thought to be fairly smooth geographically.



Figure 3: Population Growth and Housing Prices

or economically significant relationship between the increase in housing per capita and either the manufacturing share or the median rent. This null result is important because it emphasizes that population changes are a major driver of housing demand during this time period. In the next section, it will be the basis of assuming unit housing demand, regardless of location.⁸

2.2 Rents and Prices

Because regional demand for housing is the model's key mechanism, the cross-sectional predictions for housing prices hold for rental prices as well. Here we present new evidence that rents and house prices co-moved cross-sectionally.

To study the relationship between rents and housing prices, we construct a new regional index of multifamily net operating incomes using data from Commercial Mortgage Backed Securities. Net operating income is very closely related to property rents, as it measures the

⁸Along these lines, we have also investigated whether housing demand changes would have been caused by differential changes in the size of housing. Using data from the Census of Construction, we see no cross-sectional relationship between house price increases and the change in average square footage or lot size. Unfortunately, for these measures, our geographic scope is more limited, and we can only do the analysis at the Census Division level ($n = 9$). This result is consistent with the finding by Albouy and Zabek (2016): most of the increase in housing values is due to changes in land values, not to major changes in the dispersion of the size or amenities of houses.




Figure 4: Housing Demand and Manufacturing Share

difference between effective rental income and operating expenses from the perspective of landlords. We believe that our index has several advantages over other widely-used measures, such as the rent index provided by the BLS. Other series that have been used as proxies for rents, such as fair market rents, attempt to measure something different.

The source of our data is records of commercial mortgage-backed securities (CMBS) collected and provided by the firm Trepp. These data are used to price and track CMBS and have near-universal coverage of the CMBS market. Although CMBS have existed since the mid 1980s, they represented an insubstantial portion of the commercial mortgage market until the late 1990s. Black, Krainer and Nichols (2017) show that, as of 2017, 20 and 25 percent of all commercial mortgages are securitized, with a larger portion among Class A in the largest urban property markets. This may make the sample less representative of the entire rental market but may make it better for measuring the fundamental value of living in a city. Distortions in the rental market, such as rent control and insurance motives between renters and landlords, may mean that rents are not reflective of the fundamental value. A sample that skews towards more recently-built and high-occupancy buildings may be less




Figure 5: Housing Demand and 2000 Rents

subject to such biases.

The data sample is constructed as follows. We include only multifamily properties with mortgages originated in 1999 or later. We drop mortgages that are for more than one property or which appear as part of more than one security. For each property, we calculate average values of property's net operating income and occupancy rate by year. Operating incomes are deflated by the urban CPI. Property location information and construction year come from origination information. We calculate the property age from its construction year. Summary statistics are shown in the table below.

Table 1: Summary Statistics of CMBS Property Data

	Mean	Standard Deviation	10 th Percentile	Median	90 th Percentile
Net Operating Income (\$1000s)	725.8	1,246	85.75	406.9	1,659
Occupancy (Percent)	94.59	66.45	87.58	95.21	100
Loan-to-Value (Percent)	69.29	20.99	46.67	72.70	79.80
Appraised Value (\$1000s)	10,590	23,323	1,035	5,000	24,300
Units	155.0	230.9	20	108	320
Building Age (years)	36.91	25.97	11	32	76




Figure 6: Houses per capita and manufacturing, rents

There are three main advantages to using CMBS data as compared to BLS rents data. First, because it comes from an administrative data source, the CMBS data does not suffer from biases associated with the surveying process, as Crone, Nakamura and Voith (2010) document for the CPI rents series. Second, the CMBS data have broader geographic coverage than the BLS provides. Third, we believe net operating income of newer buildings is more likely to reflect the desirability of living in a specific location.

To create the income index from rents, we estimate annual changes in property-level net operating incomes using the following specification:

$$\text{Log}(\text{Income}) = \beta_t \text{Year}_t + \gamma_i \text{Property}_i + \delta \text{Log}(\text{Occupancy}_{it}) + \epsilon_{it}$$

The year coefficients β_t constitute the index. This “repeat incomes” index is akin to a repeat-sales house price index. We estimate this specification separately for each CZ. Our choice of how to include occupancy does not matter much, omitting it ($\delta = 0$) or calculating the index for income per occupant ($\delta = 1$) does not change the results.⁹

⁹Likewise, the index does not change significantly when weighting by the inverse time between observa-




Figure 7: Multifamily NOI and House Price Growth

Figure 7 shows the relationship between price growth and net operating income growth from 2000-2006 using the repeat income index. The slope of the fit line is large and positive, supporting the claim that positive fundamental shocks affected both prices and rents in the same cities.

Figure 8 verifies that prices and rents co-moved using the repeat-rent index constructed by Ambrose, Coulson and Yoshida (2015). Unlike the repeat-NOI index, Ambrose et al. (2015) measure rents directly. However, this index has the disadvantage that it is available at the CBSA-level and exists only for a handful of large CBSAs in the year 2000 because of data availability limitations (it is substantially expanded in later years). Nonetheless, it shows a similar relationship between prices and rents that we find using the NOI measure.

The relationship between rents and housing prices demonstrates that CZs with rising prices really did become more desirable to live in. This finding is important for distinguishing our model from models instead emphasizing changes to the possibility or desirability of home-ownership. If housing price changes were driven solely by changing characteristics of home-ownership, we would not expect such a relationship.

tions or when using FGLS to estimate optimal weights as a linear or quadratic function of time between observations. This is likely because the vast majority of properties provide new data every year, so weighting schemes are not as important as they are for repeat-sales indices.




Figure 8: Repeat Rent Index (RRI) and House Price Growth

As the NOI estimates control for both property and year fixed effects, it is not possible to control for property quality which changes as buildings age. Gordon and Van Goethem (2004) study the effects of changing quality in the BLS rents series and estimate a substantial downward bias as compared to a hedonic index. Therefore the within-property estimates we develop are suited only for understanding cross-sectional differences in NOI growth rather than changes in average NOI growth over this time. The average estimated NOI growth from 2000-2005 is only 2 percent in the repeat incomes index, but this estimate is downward-biased because building quality declines as buildings age. As long as one focuses on cross-sectional differences in log NOI, this bias is not important under the assumption that property depreciation, and hence bias, is the same across regions.¹⁰ In a simple hedonic model, we estimate a substantial increase in national NOI over the period we study, with the largest NOI increase 1999-2003 when prices increased by .11 log-points in total.¹¹

¹⁰This is not unique to rents; downward bias caused by property depreciation is a well-known problem for repeat-sales indexes (Nagaraja, Brown and Wachter, 2014) but this problem is more serious for rental properties than owner-occupied housing because of their faster depreciation rates (Shilling, Sirmans and Dombrow, 1991).

¹¹In the hedonic model, the vector of controls includes dummy variables for property age, the log(occupancy), and log(property square feet). To ensure that the estimates are not biased towards the cities where securitization is most popular, we weight the estimates by the inverse number of CMBS observations available, keeping only those cities and years where data from at least 30 properties are available. The estimates do not change significantly when interacting these control variables or using higher-order terms.

2.3 The Persistence of Shifts in Demand

As a final piece of evidence that cross-sectional house price changes reflect differential demands for location, we show that the relationship between house prices and manufacturing has lasted for many years after the crisis.

In Figure 9, we show the event study of manufacturing and rent's effect on house price changes since 2000. Specifically, we estimate the following regression, weighted by population:

$$\log(HousePriceIndex)_{i,t} = \beta_t ManufShare_{i,2000} + \alpha_t + \epsilon_{i,t}$$

where house prices are indexed to the year 2000. The left side of the figure plots the β_t 's. Similarly, the right-side of the figure shows the results for a similar regression, replacing manufacturing share with median rent.

The main takeaway is that the change in house prices with respect to these two variables, especially manufacturing, has persisted over a long time. In fact, the relationship between manufacturing and the house price change from 2000-2015 is about as strong as the relationship between manufacturing and the house price change from 2000-2006. For manufacturing this strong relationship never really went away, suggesting that this reflected a change in the desirability to live in those areas, not some factor specific to the housing boom itself. For rents, the relationship did seem much more temporary, but the effect has rebounded in recent years.

2.4 Housing Demand or Loan Supply?

Several theories of the housing boom argue that loan supply played an important role in explaining the housing boom. However, we argue that differential lending supply shocks cannot explain the cross-sectional price patterns. Areas with more manufacturing, or which had an elastic housing supply, had larger increases in loan application approval rates. Both of




Figure 9: The Persistence of House Prices

these patterns provide evidence that differences in mortgage demand, rather than mortgage supply, explains cross-sectional patterns. This finding justifies our focus on manufacturing as the source of cross-sectional demand variation which we argue led to a national increase in housing prices.

We study the relationship between prices and loan supply in Table 2. The estimates in both tables come from long-differences panel regressions estimating specifications of the following form:

$$Approved_j = Region_{i(j)} + LoanControls_j + \mathbf{1}(Year_j = 2006) \times Manufacturing_{i(j)}$$

$$Approved_j = Region_{i(j)} + LoanControls_j + \mathbf{1}(Year_j = 2006) \times Elasticity_{i(j)}$$

where $Manufacturing_{i(j)}$ is the manufacturing share of the county the applicant is located in during the year 2000. The data is limited to loans from the years 2000 and 2006 in order to give this an interpretation as a long differences specification. The $LoanControls_j$ variables are control variables measured at the loan (and not regional) level. In the first two columns of Table 2, the independent variable is manufacturing share, and in Columns 3-4,

Table 2: Loan Demand and Manufacturing				
	(1)	(2)	(3)	(4)
	Approval Frac.	Approval Frac.	Approval Frac.	Approval Frac.
YR=2006	-0.0133* (0.00683)	-0.0135** (0.00628)	-0.168*** (0.0114)	-0.155*** (0.0103)
YR=2006×Manuf.	0.401*** (0.0233)	0.374*** (0.0222)		
YR=2006×Elast.			0.188*** (0.00841)	0.173*** (0.00767)
County FE	X	X	X	X
Individual Controls		X		X
N	1,427,379	1,356,189	1,449,814	1,377,386
R^2	0.126	0.141	0.126	0.141

Source: Approval rates data are the fraction of loans that are approved in HMDA at the county level. Individual controls are for ventiles of DTI interacted with main applicant gender. Data is limited to 2000 and 2006.

it is housing supply elasticity in the CZ.¹²

We see that lower manufacturing and greater elasticity – which we have argued correspond to lower housing demand – are both associated with a lower approved fraction of loans. This finding allows two possible interpretations, both of which are consistent with our model. A natural one is that fundamental regional shocks (to amenities or jobs) are the main source of changes in housing and mortgage demand. A second possible explanation is that there is regional variation in the slope of the mortgage demand curve which is correlated with housing supply elasticity. A uniform, national shock in mortgage supply would then lead to both lower approval rates and higher prices in inelastic areas.

By contrast, we reject the interpretation that the relative housing boom in low-manufacturing areas was due to a relatively greater increase in loan supply in these regions. If this alternative explanation held, we would low manufacturing areas to have higher increases in loan approval rates from 2000-2006.

¹²See Section 4.1 for how the elasticity variable is created.

3 Theory: The Geography Channel

In this section, we explain how changes in the relative preferences for location have an impact on the national price level of housing. We also derive simple closed-form expressions for the effect of two prominent shocks during this time period, and how they would affect the national house price index through the geography channel.

3.1 Setup

Consider a spatial model with many cities indexed by i . Agents are mobile between cities. Agents consume a tradable good and a non-tradable good, with the price of the non-tradable good normalized to one. Housing is produced locally using local land, which is in fixed supply. Housing depreciates at a constant rate δ .

Assume that agents each consume one unit of housing, regardless of location. While this assumption is obviously counterfactual, we showed in the previous section that local housing demand was strongly driven by population changes during this time period. By making this assumption, we are quantifying the increase in house prices purely due to changing locational demand. We are purposefully neglecting changes in demand relating to the size or quantity of housing, which has previously been the focus of the literature.

Finally assume that amount of construction has no effect on intermediate inputs of housing production. We will relax this assumption later.

Define σ_i to be the local elasticity of housing supply and $\bar{\sigma} \equiv E[\sigma_i]$ to be the population-weighted average elasticity of housing supply.

Lemma 1. *Consider any shock that changes the relative demand for location. Under the previous assumptions, the change in the initial-population-weighted aggregate house price index is given by*

$$\overline{d \log p^h} = - \frac{\text{Cov}(d \log p_i^h, \sigma_i)}{\bar{\sigma}}$$

where the covariance is initial-population-weighted.

This equation illustrates that if prices in inelastic places rise compared to prices in elastic places, prices must rise on average. The equation is simply a rearrangement of the housing market-clearing condition and the definition of elasticity. The only assumptions it uses are that each agent consumes one unit of housing, and that housing markets clear.

To set intuition, imagine a shift in population from a high-elasticity place to a low-elasticity place. In the high-elasticity place, there are fewer housing units needed, so prices go down, but not by much. In the low-elasticity place, there are more units needed, and prices rise by a lot. Prices rise by more in the low-elasticity area because it is relatively inelastic, and so, on average, they rise overall.

To make a bit more progress, we add some simple structure to the model. Suppose the agent has Cobb-Douglass utility over tradable and non-tradable goods:

$$\max_{c_{nt}, c_t} c_{nt}^\alpha c_t^{1-\alpha}$$

such that $w_i c_{nt} + c_t \leq w_i - r_i$. This problem gives us the indirect utility function $u = \frac{w_i - r_i}{w_i^\alpha}$

Further suppose there is a fixed amount of land within each city available to build new housing each period, and the housing production function is CES:

$$\hat{H}_i = \left(Z_i^{\frac{\epsilon-1}{\epsilon}} + A_K K_i^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$$

where $\hat{H}_i$ is the number of new houses, Z_i is the land used for construction, and K_i is the housing material input inputs. It can be shown that $\sigma_i = \epsilon \left((\hat{H}/Z)^{\frac{\epsilon-1}{\epsilon}} - 1 \right)$. Hence the elasticity is decreasing in the density of the city.

Lemma 2. *Under these assumptions, the effect on average house prices of a change in the relative demand for location is given by*

$$d \log p^h = -\frac{\text{Cov}(\sigma_i, d \log p_i^h)}{\bar{\sigma}} + \frac{\mathbb{E}[k_i(\sigma_i + \epsilon)]}{\gamma} \text{Cov} \left(\frac{k_i(\sigma_i + \epsilon)}{\mathbb{E}[k_i(\sigma_i + \epsilon)]} - \frac{\sigma_i}{\bar{\sigma}}, d \log p_i^h \right)$$

where k_i is an adjustment term for the capital-intensity of housing in the region. k_i is proportional to $(1 + \frac{\epsilon}{\sigma_i})^{\frac{\epsilon}{1-\epsilon}}$ and $\mathbb{E}k_i = 1$.

This term is larger than the expression from Lemma 1. The intuition behind this lemma is similar. In the less elastic areas, housing requires more capital to be built, so if relative house prices increase in inelastic areas, the total demand for housing capital rises. Depending on the elasticity of this supply, it will drive up the price of this capital, increasing the total price of housing overall. Critically this formula relies on estimates of γ and ϵ .

3.2 The Role of Manufacturing

In this section, make the simplifying assumption that $\alpha \rightarrow 1$, implying that utility is a transformation of $\log w_i - \log p_i^h$.

Suppose a shock hits that affects the productivity of manufacturing, taking the following form:

$$d \log w_i = \beta + \phi m_i \tag{1}$$

where m_i is the manufacturing concentration.¹³

Then, because of the second assumption, $d \log p_i^h = \beta' + \phi m_i$. Here, ϕ is the same as the effect on wages, but β' need not be.

Proposition 1. *If $\alpha \rightarrow 1$, $\gamma \rightarrow \infty$, and a shock takes the form of equation (1), then the change in house prices is given by:*

$$\overline{d \log p^h} = \phi \frac{\text{Cov}(m_i, \sigma_i)}{\bar{\sigma}}$$

where ϕ is the coefficient from the regression of $d \log p_i^h$ on m_i .

In the empirical section, we show that this covariance is positive and large. Therefore, if

¹³This form would fit the Autor, Dorn, and Hanson (2013) framework, or an exogenous shift in the relative productivity of manufacturing in any gravity model, combined with the assumption that $\delta = 0$, i.e. there is no mobility in equilibrium.

productivity of manufacturing falls, lowering the price of housing in such areas, the national house price index rises.

By assuming γ approaches infinity, we have shut down the channel that operates through intermediate housing materials. If $\gamma < \infty$, the effect would be larger. In fact, ϕm_i could be directly substituted in for p_i^h in the expression from Lemma 2. The assumption that $\alpha \rightarrow 1$ was made primarily to simplify the algebra. With smaller α 's, the term might be slightly smaller because wages and house prices in inelastic areas were higher to start with. However, this force is relatively small when taken to the data.

3.3 The Migration Margin in Response to Price-Rent Ratio

In this section, we consider the effect of a change in the price-rent ratio. We assume that the relationship between house prices and rents is given by the following:

$$p_i^h = \frac{r_i}{R}$$

where r_i is the rent in location i and R is the national price-rent ratio. A one-for-one relationship between rents and housing prices is not implausible given our empirical finding that rents and prices co-move closely.

Proposition 2. *If $\alpha = 0$ and $\gamma \rightarrow \infty$, then the effect of the price-rent ratio on house prices is given by*

$$\overline{d \log p^h} = - \frac{Cov\left(\frac{1}{r_i}, \sigma_i\right)}{E\left[\frac{\sigma_i}{r_i}\right]} d \log R$$

In other models, changes in the price-rent ratio affect house prices because rents are pinned down, usually by an outside option. This is not the case here, where house prices are pinned down by construction costs, and rents can move freely. This is clearly seen by considering the case where inverse-rent and elasticity are uncorrelated. In that case, the price-rent ratio have no effect on house prices at all.

The fact that the level of rents matter is a natural explanation for why certain regions experienced larger growth during this time period. As interest rates fell, areas that were more expensive to start with became relatively more affordable, encouraging movement in that direction. Here, the assumption that $\alpha = 0$ is not innocuous, if $\alpha \rightarrow 1$, this channel would not exist.

One of the main predictions of this channel is discussed in Aladangady (2014), that in response to monetary policy, inelastic areas have bigger fluctuations in house prices than elastic areas.

Again, that γ is zero is innocuous, a realistic calibration would have a larger effect on house prices, but the expression would be more convoluted.

4 The Magnitude of the Geography Channel

In this section, we put numbers on the expressions from the last section, arguing that the geography channel is large. First, we create a measure of elasticity that covers the entire country. We then use that measure of elasticity to compute the values of those expressions.

4.1 Constructing a measure of elasticity

In order to estimate the covariance between an economic shock affecting wages and the elasticity of housing supply, we must first have an estimate of housing supply elasticity. Existing measures of this elasticity are inadequate for our needs because they do not provide comprehensive coverage of the entire United States, mostly focusing exclusively on MSAs (e.g. Saiz, 2010). This is especially crucial because while the Saiz measures cover much of the population, it systematically undercovers areas that are high in manufacturing concentration. In this section, we construct a measure for commuting zones, providing full coverage for the United States.

One way to do this would be simply to plug in the population density of each CZ into

the formula from before. However, we do not wish to take the Z 's literally as the amount of land in each CZ, recognizing that in the real world, there is variation in the suitability and availability of land for housing development.¹⁴

Rather, we construct elasticities similar to the methodology of Saiz (2010) by directly estimating the effect of a change in housing units on house prices, projecting this relationship onto three measures associated with land availability: land use regulations, population density, and coastal areas.

Specifically, we estimate the following model for the years 1980-2000, using decadal data on housing units from the Census and house prices from the Federal Housing Finance Agency.

$$\Delta \log p_{i,t}^h = e^{\beta_0 + \beta_1 \text{WRLURI}_i + \beta_2 \text{WRLURI missing}_i + \beta_3 \log \text{pop density}_i + \beta_4 \text{coastal}_i} \Delta \log H_{i,t} + \alpha_t + \epsilon_{i,t} \quad (2)$$

where WRLURI is the Wharton Residential Land Use Regulatory Index, pop density is the population of a CZ divided by the total square miles in it, and coastal is a dummy variable for having a county defined as coastal by the National Oceanic and Atmospheric Administration.¹⁵ We use an exponential function to guarantee our elasticities are positive. The year fixed effect is meant to capture national changes in the cost of construction and inflation.

In our baseline specification, we estimate the model using GMM, imposing that each of the four variables in the exponential, $\Delta \log H_{i,t}$, and year dummies are orthogonal to the error term. The results are shown in Column (3) of Table 3. The identifying assumption is that these variables are as good as random. If some areas of the country had more extensions to their houses installed, this would violate this assumption because that would be an omitted variable that drove up house prices and drove down housing units, as more people could live

¹⁴For example, the formula from before would imply that the CZ with New York City in it would be less than one-one hundredth as elastic as the next most inelastic CZ.

¹⁵The WRLURI is calculated for places. We took a population-weighted average for all places within a CZ. For CZs in which WRLURI is not available, we coded it as zero, and coded "WRLURI missing" as 1, so that these (primarily rural) areas would still have an elasticity calculated. Population density is the population-weighted average of county-level Census data. To classify coastal CZs, we used NOAA's definition of coastal counties. Any CZ that contained a coastal county was coded as coastal.

in each house. Another example that would violate this assumption is Hurricane Katrina, which destroyed a lot of the housing stock and directly affected the desirability of living in New Orleans. However, we think that these examples are rare, and do not introduce much bias in our regression, so we adopt this as our baseline.

Nonetheless, because of these concerns, we run two robustness checks, swapping out the moment condition that $E[\Delta \log H_{i,t} \epsilon_{i,t}] = 0$ for instruments. In column (1), we use a shift-share instrument, similar to Bartik (1991), based on two-digit SIC codes. In column (2), we simply use differences in log population. The second column should ease concerns such as changing housing types, and the first column covers a broader range of endogeneity concerns, such as events like Hurricane Katrina. In addition, we also estimate the model using non-linear least squares, a method which amounts to placing different weights on variables in our moment conditions (column 4).

The estimated coefficients are intuitive, as seen in column 3. Areas with more land regulation, higher population density, and coastal areas have house prices that move more when population changes. The negative point-estimate of β_2 also makes sense: it implies areas with no measured land-use regulatory index are similar to those with an index of about -0.99, which is the 10th percentile of the WRLURI distribution.

The results are quite similar using non-linear least squares or using population growth for the moment condition. For the estimation using the Bartik instruments, however, the results are much larger and noisier. Nonetheless they maintain the same signs for each of the coefficients. Interestingly, the model using Bartik instruments converges to where areas with no measure of WRLURI are perfectly elastic.¹⁶ This is not unreasonable given that these are almost entirely rural commuting zones.

To transform our estimates into elasticities, we construct each CZ's elasticity using the

¹⁶Or are e^{4335} times more elastic, anyway. This is the result of numerical approximation in the Stata gmm algorithm. Imposing a coefficient of $-\infty$ and re-running the regression gave the same estimates for the other coefficients.

Table 3: Estimates to construct elasticity

	(1) GMM, Bartik	(2) GMM, Pop Growth	(3) GMM, Housing Growth	(4) NLLS
β_0	-9.385 (17.89)	-2.234** (0.817)	-3.409** (1.131)	-3.597*** (0.946)
β_1 , WRLURI	0.979 (1.314)	0.198 (0.150)	0.362+ (0.188)	0.253 (0.189)
β_2 , WRLURI missing	-4355.1 (.)	-0.199 (0.413)	-0.358 (0.762)	-0.0579 (0.555)
β_3 , Log Pop Density	1.012 (0.956)	0.296* (0.149)	0.446* (0.206)	0.469** (0.161)
β_4 , Coastal	2.632 (11.45)	0.357+ (0.199)	0.522+ (0.291)	0.426+ (0.229)
N	918	922	922	922

Standard errors in parentheses

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < .001$

following formula and our preferred estimates from column (3):

$$\sigma_i = \exp(-\beta_0 - \beta_1 \text{WRLURI}_i - \beta_2 \text{WRLURI missing}_i - \beta_3 \log \text{pop density}_i - \beta_4 \text{coastal}_i)$$

The average elasticity we estimate is 2.93, in line with previous estimates. A map of the elasticities we estimate are presented in Figure 10.

Our measure is highly correlated with the estimates in Saiz (2010), at least for the geographic areas that he covers. In Figure 11, we present the comparison of estimated elasticities, for the areas of the country on which they overlap. Each dot represents a set of counties in the infimum of the MSA and CZ partitions, the largest partition that refines both. The correlation between the two measures is 0.62.




Figure 10: Map of Estimated Elasticities (bin cutoffs at the population-weighted 10th, 25th, 50th, 75th, and 90th percentiles)

4.2 The Geography Channel

We wish to see how much the geography channel contributed to the increase in house prices between 2000 and 2006. During this time period, real house prices increased 0.32 log-points.¹⁷ Now that we have estimated elasticities for each CZ in the country, we can plug these values into the previous formulas.

We use county-level house price data from the FHFA. In order to construct an index for each CZ, we take the population-weighted average increase by county. If a county does not have house price data, we do not include it in our averages. We look at total log-change between 2000 and 2006. In total, we have house price indices for 648 CZs, covering 99.8 percent of the total population.¹⁸

The geography channel of house price appreciation, as calculated using Lemma 1, $-\frac{Cov(\sigma_i, d \log p_i^h)}{\bar{\sigma}}$, explains an aggregate increase of 0.089 log-points, about 28 percent of the total increase in house prices.

Taking into consideration the increase in demand for housing construction materials increases this number. In Appendix B, we calibrate ϵ to 0.19. Obviously, it also depends significantly on γ . Unfortunately, we have not found a good way to discipline this parameter,

¹⁷We took the national FHFA data, and deflated by the CPI.

¹⁸In comparison, MSAs cover about 80 percent of the population.




Figure 11: Comparison of estimated elasticities with Saiz (2010)

so we present a range of estimates. As $\gamma \rightarrow \infty$, materials adds nothing to the previous estimate of 0.089 log-points. If, on the other hand, $\gamma = 1$, which seems plausible in the short-run, the number increases to 0.110 log-points.

To check the robustness of our results, we also use the Saiz (2010) elasticities instead of the ones created in Section 4.1. For this to be correct, the assumption would have to be that no one moved in or out of the MSAs covered by this measure. Nonetheless, because this measure of elasticity existed and was popular before our paper, it serves as a reasonable check that these results are not due to the choices we made in constructing our measure. When $\gamma \rightarrow \infty$, the Saiz elasticities suggest an increase of 0.072 log-points, in line with our number. Intuitively, it makes sense that our number is slightly larger because using Saiz's elasticities fails to capture any movement from non-metropolitan areas into MSAs.

4.2.1 The Role of Manufacturing

The decline of manufacturing played a large role in explaining this variation in house prices. Regressing the change in house prices on the manufacturing share gives an estimate of -1.57, with a confidence interval of [-1.91, -1.23].¹⁹ A one standard-deviation increase in

¹⁹The regression is population-weighted and uses robust standard errors.

manufacturing share, about 7.6 percentage points lowered local house prices by 0.118 log-points.

The relationship between manufacturing and elasticity is shown in Figure 12. As can be seen, there is a strong relationship between the two. The covariance of manufacturing and elasticity divided by the average elasticity is 0.018, which means that manufacturing alone increased the national house price index by .028 log-points. This is more than 30 percent of the total geography channel contribution.




Figure 12: Estimated Elasticity and Manufacturing Share by CZ

4.2.2 The Geography of Increasing Price-Rent Ratios

As outlined in the theory, changes in the price-rent ratio can have effects on patterns of mobility and aggregate house prices, even when agents have unit housing demand. In this section, we quantify how large these effects are on the aggregate. First, we need an estimate of the increase in price-rent ratios. As a proxy, we use the decline in 30-year mortgage rates from Freddie Mac, which declined about 0.32 log-points over this time period. Hence, we will do some calculations for an increase in the price-rent ratio of 0.32 log-points.²⁰

²⁰This is a smaller increase in price-to-rent ratios than one might get from other data sources. For example, Greenwald (2016) tries to match data that shows an increase of about 60 percent.




Figure 13: Inverse Rent versus Elasticity

The relationship between inverse rents and elasticity is shown in Figure 13. Not surprisingly, rents are higher in inelastic areas. The term multiplying the change in the price-rent ratio evaluates to -0.13 , meaning that a decline in interest rates will cause an increase in house prices as higher-priced areas become more affordable. Using our preferred measure of R , the total effect on national house prices from Proposition 2 is $.041$ log-points. This is of similar magnitude as the manufacturing shock from the last section. Note these are not directly comparable, since the two formulas make opposite assumptions about α .

Taking a step-back, we have shown that the geography channel explains between 25 and 30 percent of the increase in national house prices, and that much of this variation is explained by the decline of manufacturing and lower interest rates. Of course, these two forces are not sufficient to cause the entire change. However, plenty of other explanations for the desirability of inelastic locations exist: including better amenities²¹ or increased sorting.²² Furthermore, evidence exists to suggest that other forces may be amplified by the migration itself (Howard, 2017).

²¹Many papers, such as Albouy (2016) find coastal cities to have better amenities.

²²Diamond (2016) documents better-educated workers have sorted into specific cities, many of which are inelastic. The presence of these workers raises the amenities value. Gyourko, Mayer and Sinai (2013) presents a model where sorting increases the skewness of house prices and incomes.

4.2.3 Declining Cities and External Validity

It has been hypothesized and shown in the data (Glaeser and Gyourko, 2005; Notowidigdo, 2011) that cities with declining populations have lower housing-supply elasticities.²³ The time period of interest to us is a time where almost no CZs experienced a decline in housing units (see Figure 2). In fact, by 2000 population, only 1.4 percent of the country lived in a commuting zone with fewer housing units in it in 2006 than in 2000. More than a third of those people lived in New Orleans, where the decline in housing units was caused by Hurricane Katrina. Hence, we do not expect this force to have much relevance in this time period.

However, it could easily matter in other time periods, where there is not population growth or increases in national housing per capita. So it makes sense to consider what would happen with the geography channel at other times. Here, the first-order variation in elasticity is over whether a city is shrinking or not, swamping any effects of whether people prefer to live in coastal areas, highly land-use-regulated areas, or population dense areas.

Proposition 3 (Declining Cities). *If any city declining in housing becomes completely inelastic and aggregate housing is not increasing, then any change in the relative desirability of cities leads to a negative aggregate change in house prices.*

Increases in population, fluctuations in houses per capita, and housing depreciation all make this situation less likely. Nonetheless, this result provides a stark example of how the geography channel can lead to house price declines instead of increases, even if the relative desirability of coastal, highly-regulated, urban areas is increasing. The key difference between our results and this proposition is that aggregate housing was increasing significantly during this period, such that almost no cities were declining in their housing stock.

²³In our own data, we have also included a term in our elasticity estimation for a dummy of having declining housing units. The coefficient on that is 3.46 (with standard error 0.71), suggesting hugely more inelastic housing supply in declining cities, a factor of 30. Effectively, areas with housing unit declines are completely inelastic. Other coefficients did not change significantly. We strongly prefer the linear piecewise estimation that we are doing rather than a continuous specification as Notowidigdo (2011), because we believe the economics more strongly justify concavity around 0, but we do not believe that extrapolating that concavity to large house price changes is appropriate.

5 The Timing of the Housing Boom

Our theory makes precise predictions on the timing of the housing boom. Aggregate house prices should increase most when the cross-sectional covariance between local house price changes and elasticity is the most negative. In this section, we show that this covariance was most negative from 2000-2006, when the boom occurred. Furthermore, we estimate the size of this channel for a previous period, from 1994-2000, and show that the effect is economically insignificant.

In the previous section, we showed that there was a significant negative correlation between local house price changes and elasticity from 2000-2006. This did not occur in the period beforehand. Using the same formula from Lemma 1 over the period of 1994-2000, there would have been an increase in house prices of about 0.012 log-points, with most of it coming in the last few years.

The same point is illustrated with the blue line on the left side of Figure 14. In this figure, we regress changes in different local variables on house prices year-by-year. Hence, the point of the green line in the year 2000 represents the population-weighted regression coefficient of the change in log house prices from 1999-2000 on housing supply elasticity. Before 2000, these coefficients were close to zero, but between 2000 and 2006 they are strongly negative. Interestingly, during the bust years, they are significantly positive, suggesting the geography channel contributed some to the bust as well.

Also pictured on the left side of Figure 14 are the cross-sectional relationships of rents and population with elasticity. Interestingly, the movement of people seems to lead changes in house prices. The biggest population shift toward inelastic areas occurs between the years 1996 and 2003. This may represent some sluggishness in house price changes.²⁴ This is not a problem for our theory; if local house prices are sluggish to react to fundamental changes, our theory would simply predict that national house prices would be similarly

²⁴Interestingly, and in support of our theory, there was a similar movement of people towards inelastic areas before the 1980 housing boom. Consistent with the idea that housing supply has become more inelastic Ganong and Shoag (2017), there were larger changes in population and a smaller house price boom.




Figure 14: The year-by-year regression coefficient of house price, rent, and population growth on elasticity (left) and the regression coefficient of house price growth on manufacturing (right)

sluggish. We include the rents data to show that it exhibits similar timing patterns, although unfortunately, our data extends only back to 2000. The timing of rents reverses before the downturn in the housing market, which may be further evidence of housing sluggishness.

The right side of Figure 14 does a similar exercise to assess whether the theory predicts similar timing for the contribution of the decline in manufacturing. Of course, the location of manufacturing has not shifted significantly over time, but its effect on house prices has. Here, we run the regression of the change in local house prices on manufacturing share year-by-year, and plot each coefficient. From 2000-2006, the coefficients are significantly negative, and combined with the fact that manufacturing is concentrated in inelastic areas, this would predict rising house prices exactly during those years. Before 2000, there was even a positive relationship, though much less economically significant.

6 Conclusion

Much of the literature on the housing boom has focused on why the demand for the quantity of housing or home-ownership increased, but little has focused on the location of that housing. This paper demonstrates that considering this geography can explain a significant fraction

of the overall boom.

The result also questions the appropriateness of testing theories using the cross-section of housing without explicitly considering the role of mobility. At the very least, it demonstrates that local housing markets cannot be viewed in isolation.

Our results have abstracted from several important facts, including the dynamics of the boom. We believe our results on rents and house prices could be useful in calibrating a model of expectations that could help explain the overshooting and subsequent housing bust.

References

- Aladangady, Aditya**, “Homeowner balance sheets and monetary policy,” 2014.
- Albouy, David**, “What are cities worth? Land rents, local productivity, and the total value of amenities,” *Review of Economics and Statistics*, 2016, 98 (3), 477–487.
- **and Mike Zabek**, “Housing inequality,” National Bureau of Economic Research Working Paper 2016.
- Ambrose, Brent W, N Edward Coulson, and Jiro Yoshida**, “The repeat rent index,” *Review of Economics and Statistics*, 2015, 97 (5), 939–950.
- Autor, David, David Dorn, and Gordon H Hanson**, “The China syndrome: Local labor market effects of import competition in the United States,” *The American Economic Review*, 2013, 103 (6), 2121–2168.
- , – , **and Gordon Hanson**, “Response to Robert Feenstra, Hong Ma, and Yuan Xus Comment on Autor, Dorn, and Hanson (AER 2013),” *Working Paper*, 2017.
- Bartik, Timothy J**, *Boon or Boondoggle? The debate over state and local economic development policies*, WE Upjohn Institute for Employment Research, 1991.
- Black, Lamont, John Krainer, and Joseph Nichols**, “From origination to renegotiation: A comparison of portfolio and securitized commercial real estate loans,” *The Journal of Real Estate Finance and Economics*, 2017, 55 (1), 1–31.
- Charles, Kerwin Kofi, Erik Hurst, and Matthew J Notowidigdo**, “The masking of the decline in manufacturing employment by the housing bubble,” *The Journal of Economic Perspectives*, 2016, 30 (2), 179–200.
- Crone, Theodore M, Leonard I Nakamura, and Richard Voith**, “Rents have been rising, not falling, in the postwar period,” *The Review of Economics and Statistics*, 2010, 92 (3), 628–642.
- Davidoff, Thomas**, “Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors,” *Critical Finance Review*, 2016, 5 (2), 177–206.
- Davis, Morris, Jonas Fisher, and Marcelo Veracierto**, “Gross Migration, Housing and Urban Population Dynamics,” in “2014 Meeting Papers” Society for Economic Dynamics 2014.
- DeFusco, Anthony A, Charles G Nathanson, and Eric Zwick**, “Speculative dynamics of prices and volume,” Technical Report, National Bureau of Economic Research 2017.
- Diamond, Rebecca**, “The determinants and welfare implications of US workers’ diverging location choices by skill: 1980–2000,” *The American Economic Review*, 2016, 106 (3), 479–524.

- Foote, Christopher L., Kristopher S. Gerardi, and Paul S. Willen**, “Why did so many people make so many ex post bad decisions? The causes of the foreclosure crisis,” NBER Working Paper 2012.
- Ganong, Peter and Daniel Shoag**, “Why Has Regional Income Convergence in the US Declined?,” *Journal of Urban Economics*, 2017.
- Garriga, Carlos, Aaron Hedlund, Yang Tang, and Ping Wang**, “Rural-urban migration, structural transformation, and housing markets in China,” National Bureau of Economic Research Working Paper 2017.
- Glaeser, Edward L and Charles G Nathanson**, “An extrapolative model of house price dynamics,” *Journal of Financial Economics*, 2017.
- **and Joseph Gyourko**, “Urban decline and durable housing,” *Journal of political economy*, 2005, 113 (2), 345–375.
- Gordon, Robert J and Todd Van Goethem**, “A Century of Downward Bias in the Most Important Component of the CPI: The Case of Rental Shelter, 1914-2003,” *Hard-to-Measure Goods and Services: Essays in Memory of Zvi Griliches, Studies in Income and Wealth*, 2004, (67).
- Greenwald, Daniel L.**, “The mortgage credit channel of macroeconomic transmission,” 2016.
- Gyourko, Joseph, Christopher Mayer, and Todd Sinai**, “Superstar cities,” *American Economic Journal: Economic Policy*, 2013, 5 (4), 167–199.
- Herkenhoff, Kyle F., Lee E. Ohanian, and Edward C. Prescott**, “Tarnishing the Golden and Empire States: Land-Use Restrictions and the U.S. Economic Slowdown,” NBER Working Paper 2017.
- Howard, Greg**, “The Migration Accelerator: Labor Mobility, Housing, and Aggregate Demand,” 2017.
- Hsieh, Chang-Tai and Enrico Moretti**, “Housing Constraints and Spatial Misallocation,” NBER Working Paper 2017.
- Kaplan, Greg, Kurt Mitman, and Gianluca Violante**, “Consumption and House Prices in the Great Recession: Model Meets Evidence,” 2017.
- Liebersohn, Carl Jack**, “Housing Demand, Regional House Prices and Consumption,” 2017.
- Mayer, Christopher and Todd Sinai**, “US house price dynamics and behavioral finance,” *Policy Making Insights from Behavioral Economics. Boston, Mass: Federal Reserve Bank of Boston*, 2009.

- Mian, Atif and Amir Sufi**, “The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis,” *The Quarterly Journal of Economics*, 2009, 124 (4), 1449–1496.
- Nagaraja, Chaitra, Lawrence Brown, and Susan Wachter**, “Repeat sales house price index methodology,” *Journal of Real Estate Literature*, 2014, 22 (1), 23–46.
- Nieuwerburgh, Stijn Van and Pierre-Olivier Weill**, “Why has house price dispersion gone up?,” *The Review of Economic Studies*, 2010, 77 (4), 1567–1606.
- Notowidigdo, Matthew J**, “The incidence of local labor demand shocks,” National Bureau of Economic Research Working Paper 2011.
- Roback, Jennifer**, “Wages, rents, and the quality of life,” *Journal of political economy*, 1982, 90 (6), 1257–1278.
- Rosen, Sherwin**, “Wage-based indexes of urban quality of life,” *Current issues in urban economics*, 1979, 3, 490.
- Saiz, Albert**, “The geographic determinants of housing supply,” *The Quarterly Journal of Economics*, 2010, 125 (3), 1253–1296.
- Shilling, James D, CF Sirmans, and Jonathan F Dombrow**, “Measuring depreciation in single-family rental and owner-occupied housing,” *Journal of Housing Economics*, 1991, 1 (4), 368–383.

A Proofs

A.1 Proof of Lemma 1

Agents demand one unit of housing, implies that the total amount of housing is fixed by the population. Therefore,

$$\sum_i N_i d \log \hat{H}_i = 0$$

By the definition of housing supply elasticity, $d \log \hat{H}_i = \sigma_i d \log p_i^h$, so

$$\sum_i N_i \sigma_i d \log p_i^h = 0$$

The result is an algebraic rearrangement, solving for $\sum_i \frac{N_i}{N} d \log p_i^h$. □

A.2 Proof of Lemma 2

The two market-clearing conditions are that

$$\sum_i N_i d \log \hat{H}_i = 0$$

$$\frac{1}{\sum_i k_i} \sum_i k_i d \log k_i = \gamma d \log r^k$$

Under the CES housing supply assumption,

$$d \log \hat{H}_i = \sigma_i (d \log p_i^h - d \log r^k)$$

$$d \log k_i = (\sigma_i + \alpha) (d \log p_i^h - d \log r^k)$$

After simple algebra, the following two equations relate $d \log r^k$ and $\overline{d \log p^h}$.

$$d \log r^k = \frac{Cov(\sigma_i, d \log p_i^h)}{\bar{\sigma}} + \overline{d \log p^h}$$

$$(\gamma + \alpha + \mathbb{E}[\kappa_i \sigma_i]) d \log r^k = Cov(\kappa_i (\sigma_i + \alpha), d \log p_i^h) + (\alpha + \mathbb{E}[\kappa_i \sigma_i]) \overline{d \log p^h}$$

where $\kappa_i = \frac{k_i}{K} \frac{N}{N_i}$. Plugging in r^k leads to

$$\gamma d \log p^h = -(\gamma + \alpha + \mathbb{E}[\kappa_i \sigma_i]) \frac{Cov(\sigma_i, d \log p_i^h)}{\bar{\sigma}} + Cov(\kappa_i (\sigma_i + \alpha), d \log p_i^h)$$

This can be decomposed into the original term and the contribution from additional

$$d \log p^h = -\frac{Cov(\sigma_i, d \log p_i^h)}{\bar{\sigma}} + \frac{\mathbb{E}[\kappa_i (\sigma_i + \alpha)]}{\gamma} Cov\left(\frac{\kappa_i (\sigma_i + \alpha)}{\mathbb{E}[\kappa_i (\sigma_i + \alpha)]} - \frac{\sigma_i}{\bar{\sigma}}, d \log p_i^h\right)$$

The last step is to show that κ_i satisfies the stated properties. Simply note that $\kappa_i \propto \frac{k_i}{N_i} \propto \frac{k_i}{\hat{H}_i} = \left(\frac{\sigma_i + \alpha}{\alpha}\right)^{\frac{\alpha}{\alpha-1}}$, and that $\mathbb{E}\left[\frac{N}{N_i} \frac{k_i}{K}\right] = \sum \frac{k_i}{K} = 1$. $\square$

A.3 Proof of Proposition 1

Our assumptions imply that $d \log p_i^h = \phi m_i + \beta$. Plugging this into Lemma 1 and factoring the ϕ from the covariance gives the result. $\square$

A.4 Proof of Proposition 2

When wages are held constant, the following equation must hold for all i :

$$d \log u = \frac{r_i}{u} (d \log p_i^h + d \log R)$$

It is still the case that $\sum_i N_i \sigma_i d \log p_i^h = 0$, and $d \log u$ is constant across cities. Doing algebra,

$$d \log u = d \log R \frac{\sum_i N_i \sigma_i}{\sum_i N_i \sigma_i \frac{u}{r_i}}$$

Plugging this in and solving for $\sum_i \frac{N_i}{N} d \log p_i^h$ gives:

$$\overline{d \log p_i^h} = \frac{Cov(\frac{1}{r_i}, \sigma_i)}{E \left[\frac{\sigma_i}{r_i} \right]}$$

$\square$

A.5 Proof of Proposition 3

Suppose house prices anywhere increased. Then housing supply in that area increases, and housing supply somewhere else must decrease, by the assumption on the aggregate. But housing supply cannot decrease anywhere because it is completely inelastic in declining areas. Therefore, house prices cannot increase anywhere. If there is a change in the relative desirability, and house prices cannot increase in the relatively more desirable location, they must fall in the relatively less desirable location. Therefore, on average, house prices fall. $\square$




Figure 15: Land Shares and House Prices

B Calibration

B.1 ϵ

We made the assumption that housing was produced CES.

$$\hat{H}_i = \left(Z_i^{\frac{\epsilon-1}{\epsilon}} + A_K K_i^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$$

If everyone faces the same price of capital, it can be derived that

$$\log(1 - \text{Land Share}_i) = \beta + (\epsilon - 1) \log p_i^h$$

To get a value for ϵ , we used the Lincoln Land Institute's data from 2000 on Land Share and Home Value. The results are in Figure 15. The regression gives a value of -0.81 , implying ϵ is approximately 0.19.