

NBER WORKING PAPER SERIES

THE ECONOMIC IMPLICATIONS OF HOUSING SUPPLY

Edward Glaeser
Joseph Gyourko

Working Paper 23833
<http://www.nber.org/papers/w23833>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2017

Edward Glaeser is the Fred and Eleanor Glimp Professor of Economics, Harvard University, Cambridge, Massachusetts. He thanks the Taubman Center for State and Local Government at Harvard for financial support. Joseph Gyourko is the Martin Bucksbaum Professor of Real Estate, Finance and Business Economics and Public Policy, the Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania. He thanks the Research Sponsor Program of the Zell/Lurie Real Estate Center at Wharton for financial support. Their email addresses are eglaser@harvard.edu and gyourko@wharton.upenn.edu. The excellent research assistance of Yue Cao, Matt Davis and Xinyu Ma is much appreciated, but the authors remain responsible for any errors. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23833.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Edward Glaeser and Joseph Gyourko. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Economic Implications of Housing Supply
Edward Glaeser and Joseph Gyourko
NBER Working Paper No. 23833
September 2017
JEL No. D45,R12,R30

ABSTRACT

In this essay, we review the basic economics of housing supply and the functioning of US housing markets to better understand the distribution of home prices, household wealth and the spatial distribution of people across markets. We employ a cost-based approach to gauge whether a housing market is delivering appropriately priced units. Specifically, we investigate whether market prices (roughly) equal the costs of producing the housing unit. If so, the market is well-functioning in the sense that it efficiently delivers housing units at their production cost. Of course, poorer households still may have very high housing cost burdens that society may wish to address via transfers. But if housing prices are above this cost in a given area, then the housing market is not functioning well—and housing is too expensive for all households in the market, not just for poorer ones. The gap between price and production cost can be understood as a regulatory tax, which might be efficiently incorporating the negative externalities of new production, but typical estimates find that the implicit tax is far higher than most reasonable estimates of those externalities.

Edward Glaeser
Department of Economics
315A Littauer Center
Harvard University
Cambridge, MA 02138
and NBER
eglaeser@harvard.edu

Joseph Gyourko
University of Pennsylvania
The Wharton School of Business
3620 Locust Walk
1480 Steinberg-Dietrich Hall
Philadelphia, PA 19104-6302
and NBER
gyourko@wharton.upenn.edu

For most of US history, local economic booms were matched by local building booms. Into the 1960s, building was lightly regulated almost everywhere. Much housing was built in all high demand areas, including coastal California and New York City. However, between the 1960s and the 1990s, it became far more difficult to build in some areas with strong economic growth, especially those along the coasts. For example, there were 13,000 new housing units permitted in Manhattan in the single year of 1960 alone, which is nearly two-thirds of the 21,000 new units permitted throughout the decade of the 1990s (Glaeser, Gyourko and Saks 2005). Higher economic productivity in the San Francisco Bay Area, with its extensive restrictions on land use and building, now leads primarily to higher housing prices, rather than more homes and more workers (Ganong and Shoag 2013).

In this essay, we review the basic economics of housing supply and the functioning of US housing markets to better understand the distribution of home prices, household wealth and the spatial distribution of people across markets. We employ a cost-based approach to gauge whether a housing market is delivering appropriately priced units. Specifically, we investigate whether market prices (roughly) equal the costs of producing the housing unit. If so, the market is well-functioning in the sense that it efficiently delivers housing units at their production cost. Of course, poorer households still may have very high housing cost burdens that society may wish to address via transfers. But if housing prices are above this cost in a given area, then the housing market is not functioning well—and housing is too expensive for all households in the market, not just for poorer ones.¹ The gap between price and production cost can be understood as a regulatory tax, which might be efficiently incorporating the negative externalities of new production, but typical estimates find that the implicit tax is far higher than most reasonable estimates of those externalities.

We begin by discussing how to estimate the minimum profitable cost of production for a house in a lightly regulated housing market, where such costs are primarily determined by geography and characteristics of local markets for labor and materials. We can then classify US housing markets into three different groups. In lightly regulated housing markets with growing population and economies, like Atlanta, the supply curve for housing is relative flat. Thus, as demand for housing expands over time, the result is that competition in the home building industry holds the price of housing reasonably close to its minimum profitable production cost. In heavily regulated housing markets with growing economies, like the San Francisco Bay area, the supply curve for housing slopes up. As a result, additional demand for housing translates into prices that are substantially above the minimum profitable production cost, with rising land

¹ Policymakers often discuss housing prices through the prism of affordability: for example, many federal programs deem that housing is inappropriately expensive or unaffordable if the monetary costs of occupying your home exceed 30 percent of one's gross income. The social merits of this cutoff as a rough rule of thumb aside, economically, it lacks clarity about the extent to which the issue of housing affordability for a given area is due to a higher prevalence of households near or below the poverty level, or due to housing prices that are at relatively very high levels. It also fails to consider an implication of the standard special equilibrium model used in urban economics, which is that equalizing utility levels across space implies that housing costs will be a higher share of earnings in higher-wage locations (Rosen 1979; Roback 1982).

values driving up total costs. Finally, in a housing market like Detroit where the demand for housing declined sharply over time, the supply curve for housing has a kink at the existing level of housing because housing is durable and does not diminish quickly when demand falls. As a result, a reduction in demand leads to lower prices for housing and minimal new construction (Glaeser and Gyourko 2005).

The ratio of price-to-minimum profitable construction cost is akin to Tobin's q , the standard ratio of market value-to-firm replacement cost. Regulatory construction constraints can explain why this variant of q may be higher than one in some housing markets, just as capital adjustment costs can explain why q is higher than one in classical investment models (Hayashi 1982).

We then discuss two main effects of developments in housing prices: on patterns of household wealth and on the incentives for relocation to high-wage, high-productivity areas. Binding supply side restrictions shape the personal portfolios of millions of Americans, and much of the rise in the capital share can be attributed to rising rents on housing. However, only a small sliver of America is sitting on a large amount of housing wealth. We will argue that the rise in housing wealth is concentrated in the major coastal markets that have high prices relative to minimum production costs, and it is concentrated among the richest members of the older cohorts—that is, on those who already owned homes several decades ago, before binding constraints on new housing construction were imposed. In effect, the changes in housing wealth reflect a redistribution from buyers to a select group of sellers.

The restrictions on housing supply and corresponding high housing prices in certain areas is also a distortion that limits the movement of workers in areas with high productivity, high wages—and also high housing costs. Hsieh and Moretti (2017) have estimated that real GDP could be nearly 9 percent higher if there were plentiful new construction in just the three high productivity markets of New York, San Francisco and San Jose, so that people could move to equalize wages. We will discuss the basis for such estimates and show that there can be a fairly wide range of outcomes depending upon model and parameter assumptions. However, our analysis indicates that a lower bound cost of restrictive residential land use regulation is at least 2 percent of national output. If these regulatory distortions are efficiently internalizing negative externalities, then the benefit of increased aggregate output would also need to be weighed against the costs of local disamenities.

In the conclusion, we turn to some policy implications. The available evidence suggests, but does not definitively prove, that the implicit tax on development created by housing regulations is higher in many areas than any reasonable negative externalities associated with new construction. Consequently, there would appear to be welfare gains from reducing these restrictions. But in a democratic system where the rules for building and land use are largely determined by existing homeowners, development projects face a considerable disadvantage, especially since many of the potential beneficiaries of a new project do not have a place to live in the jurisdiction when possibilities for reducing regulation and expanding the supply of housing are debated.

Construction Costs and Regulations in Housing Markets

Variation Across Physical Geographies in the Cost of Supplying a Home

There is no reason to expect that the production costs of housing should be the same across markets, even if those places have similar levels of regulation. Geography will make housing more expensive to build in some areas than others. Bedrock makes it easier to build up (Rosenthal and Strange 2007). Steep ground makes it much more challenging to build (Saiz 2010). Bodies of water can limit land supply.

The flat cities of the American Midwest are close to the perfect physical environment for building, as is much of the Sunbelt region. Conversely, America's coastal cities are considerably more difficult geographical environments for builders. California cities often have significant changes in elevation within a single metropolitan area. Both east coast and west coast cities are limited in that they can only expand inland. All of America's oldest cities were built on major waterways because of the advantages of access to water-borne transportation. Consequently, the central business districts of markets such as Boston, New York, San Francisco and even Chicago are close to the waterfront. Developers in those places only have a semi-circle of land to develop. The island of Manhattan poses particularly unique challenges.

When supply of housing is relatively lightly regulated, as it is throughout much of the American Sunbelt and the interior of the country, construction seems to be close to a constant returns to scale technology. This relationship reflects the relative abundance of building materials such as wood, and less skilled construction workers.² Of course, construction costs do vary according to the physical geography of local building condition, but Gyourko and Saiz (2006) examine the heterogeneity of construction costs (discussed in more detail below) and find that the variance of such costs is much smaller than the heterogeneity of housing prices. This implies that we can talk sensibly about a single production cost.

Variations in Regulations on Land Use and Building

The United States is relatively unique in that land use is under local control, which accounts for the wide variation in regulation across communities. Many other countries, including the United Kingdom and France, have national planning agencies and guidelines set by their central governments. Local land use regulation in the United States, ranging from building code requirements to strict limits on the number of units delivered, also differs across markets and can affect construction costs associated with putting up the structure, as well as the underlying price of land.

² Taller buildings also display their own constant returns to scale, because the per square foot cost of building to seven stories is quite close to the per square foot cost of building 50 stories (R.S. Means, 2015). That said, the cost of building up is much higher than the cost of building low rise dwellings.

Modern land use regulation in the United States dates back at least to the 1910s, when the initial zoning laws were enacted to limit negative externalities from spillovers between different kinds of land users. While there are no consistent time series measures of the local residential land use regulatory environment, researchers generally agree that such regulation has proliferated across markets and become onerous in some places. The term NIMBYism (“not in my back yard”) dates back to Frieden (1979). The literature on this topic is now voluminous and Gyourko and Molloy (2015) provide a recent review.

There is no doubt that binding density restrictions affect supply. For example, the median Boston suburb has a minimum lot size over one acre—and larger minimum lot sizes are common. Unsurprisingly, minimum lot size is strongly negatively correlated with new building across communities in greater Boston (Glaeser and Ward 2009).

Restrictions often go far beyond minimum acreage or maximum height restrictions. Examples include laws that prohibit multi-family dwellings, stop development near wetlands (which are often loosely defined), and make it difficult to build across large swaths of historic neighborhoods. Since the 1972 “Friends of Mammoth” case, the California Environmental Quality Act (CEQA) has been interpreted to require an environmental impact review for “most proposals for physical development in California,” (<http://resources.ca.gov/ceqa/more/faq.html#when>). Environmental impact reviews may not prevent a project, but they will add time delays which increase development costs. Moreover, the environmental impact reviews only investigate the project’s impact on the local environment, and do not include the environmental benefits of building in California, where carbon emissions are low due to the mild climate, instead of less temperate Texas or Arizona (Glaeser and Kahn 2010).

The potential for a multi-year review process, which is not uncommon in many jurisdictions, is associated with higher project uncertainty, not just time delays. A project may be denied approval after many years of active planning. That risk also increases the expected costs to developers and deters new housing supply.

The plethora of restrictions on building makes it difficult to measure the overall strictness of the broader regulatory environment, but it is possible to describe the nature of different types of communities’ approaches to regulation. The Wharton Residential Land Use Regulation Index, based on surveys of local government officials, documents wide differences in the difficulty of obtaining building permits across metropolitan areas (Gyourko, Saiz and Summers 2008). The typical regulatory environment in their sample of 2,611 communities across 293 metropolitan areas can be described as follows: a) two entities are required to approve any project requiring a zoning change, so there are multiple opportunities for rejection; b) minimum lot size restrictions are omnipresent; c) “development exaction fee programs” also are now omnipresent; (d) the typical community exhibits about a six-month lag between submission of a permit request for a standard project and a decision on whether to approve it.

The one-third most highly regulated communities in the Gyourko, Saiz and Summers (2008) sample also share some additional traits. Local and state pressure groups are much more likely to be involved in the regulatory process in these communities. More than half the highly regulated places have at least one neighborhood with a one-acre (or more) minimum lot size rule; in contrast, only 5 percent of the one-third most lightly regulated communities had any neighborhood with a one acre minimum rule. Open space requirements, not just development exactions, are now common in highly regulated places. Finally, the most highly regulated places have project approval lags that average ten months in length, which is three times longer than in the least regulated one-third of communities. In another study, Saiz (2010) documents how both regulations and geography limit building and increase prices across space

A variety of models of local land use control embed the idea that not all local residents will share the same goals, so that the regulatory environment will be shaped by the incentives and influence of different actors in the political process. For example, Fischel (2001) emphasizes the role of existing homeowners, who have a strong incentive to protect what often is their most important asset. One obvious way to protect asset value is to restrict new supply. Theoretical analysis is much more challenging in a multi-community setting that permits Tiebout sorting and strategic interactions (for discussion, see Gyourko and Molloy 2015). In principle, regulation can be an efficient means of forcing developers to internalize negative externalities from construction. Moreover, the spatial heterogeneity in those regulations may reflect different external costs from construction, perhaps because of different local preferences.

The general conclusion of existing research is that local land use regulation reduces the elasticity of housing supply, and that this results in a smaller stock of housing, higher house prices, greater volatility of house prices and less volatility of new construction. Most results are consistent with these implications and we report additional evidence below. However, it has been a challenge in this literature to find convincing instruments or some form of experimental variation. Because empirical work in this area is cross sectional in nature, and it is subject to standard potential biases associated with omitted variables and reverse causality.³

What Does It Actually Cost to Supply Homes to the Market?

There are three components to the cost of delivering a unit of housing to the market: 1) the land (L) on which the housing unit sits; 2) construction costs (CC) associated with putting up structure itself; 3) the entrepreneurial profit (EP) needed to compensate the home builder. Thus, we define the “minimum profitable production cost” (MPPC) of a unit of housing as follows:

³ To understand the problem of finding experimental variation in this literature, consider the variation in difficulty of building across space generated by geographic variables of the type analyzed by Saiz (2010). In this setting, a location that is close to water increases housing demand, but also creates a more challenging geographical environment. More generally, home-building will occur in more challenging and costly locations only if they have something else going for it. Consequently, geography provides meaningful variation in the difficulty of building, but is not a valid instrument for housing supply in most situations.

$$\text{MPPC} = (\text{L} + \text{CC}) * \text{EP}.$$

Vacant land sales are rarely observed in the United States, so to estimate the value of a price of land, we use an industry rule of thumb based on an *ad hoc* survey of home builders that land values are no more than 20 percent of the sum of physical construction costs plus land in a relatively free market with few restrictions on building. We have used this metric in earlier research (Glaeser and Gyourko 2003, 2008), and it continues to be relevant and consistent with the data discussed below.

The gross profit margin on the builder's land and construction costs for a portfolio of homebuilders range from 9-11 percent per annum across the cycle. This implies gross margins of about 17 percent given the roughly 35-40 percent cost of operations for such companies. Hence, $\text{EP}=1.17$ in our calculations below.

Physical construction costs are more readily observable from the home building industry. We use R.S. Means Company data on physical construction costs as the foundation of our estimates of minimum profitable production cost. This firm provides and sells estimates of the cost of providing units of different qualities across more than 100 American housing markets. Their data have been used by us (and others) in previous research (for examples, see Glaeser and Gyourko 2003, 2005; Glaeser, Gyourko and Saks 2005; Gyourko and Saiz 2006).

The R.S. Means cost estimates cover material, labor, and equipment (but not land) for four different qualities of single family homes—economy, average, custom and luxury. Means reports costs per square foot and provides estimates for homes ranging in size from 600 ft² to 3,200 ft² of living area. Breakdowns are available by the number of stories in the house, and certain other characteristics (such as the presence of a basement). We focus on costs associated with a smaller, modest-quality, one-story home of economy quality described in R.S. Means Company publication, *Residential Cost Data 2015*.⁴ We choose this home because we believe it reflects the quality of the typical home (which is not new or very large) in most, if not all, markets. We have experimented with using this data with regard to homes of other quality characteristics and discuss possible biases below.

The first important stylized fact is that structure costs are modest for an economy-quality home. The interquartile range runs from \$72/sf² to \$86/ft², and the distribution is not fat-tailed. The 5th and 95th percentile values are \$68/ft² and \$95/ft², respectively. Thus, in cheaper markets physical construction costs associated with putting up a typical home with 2,000ft² of living space are about \$140,000 (approximately \$70 per foot); in the most expensive markets, the costs are about \$180,000 (approximately \$90 per foot).

⁴ Specifically, this is a one-story single family home, one full bathroom, one kitchen, asphalt roof shingles, hot air heat, gypsum wallboard interior finishes, mass produced from stock plans. The R.S. Means Company presumes that a given quality home is constructed in a common way across markets. It divides the home into a number of different tasks that require certain services, materials or labor. Means then surveys local suppliers and builders to determine the local price of those inputs. One-off construction of custom homes would be much more costly. See R.S. Means Company (2015) and Section 2 in Gyourko and Saiz (2006) for more on the underlying methodology.

A second noteworthy stylized fact is that real construction costs have not risen much over time. Measured in constant 2010 dollars, the cost was \$83 per square foot in 1980, had declined slightly to the mid-\$60s per square foot by the late 1990s and early 2000s, and then rose back to \$85 per square foot by 2015. This finding is consistent with much previous research and implies that rising real house prices cannot be explained by higher physical construction costs (for example, Davis and Heathcote 2004; Davis and Palumbo 2008; Gyourko and Molloy 2015).

These relatively constant physical production costs help us to understand the often-noted decline in total factor productivity in the construction sector.⁵ This decline does not seem to result from any change in building technology, but rather an increase in other costs associated with delivering housing, such as dealing with regulation.

Given the assumptions outlined above for costs of land and profits, minimum profitable production costs that take land and profit into account are nearly 50 percent higher than the R.S. Means physical construction cost numbers. This suggests that an efficient housing market should be able to supply economy-quality single-family housing with 2,000 ft² of living space for around \$200,000 in low construction cost markets and for little more than \$265,000 in the highest construction cost markets. The key factors that account for the cross-sectional variation in structure production costs in this data are the extent of unionization in the construction industry, the level of local wages in general, and difficult topography (Gyourko and Saiz 2006). For perspective, what R.S. Means calls the “average” quality home costs about 25 percent more than the economy home and the highest quality “luxury” home of the same size costs almost twice as much to construct as the economy home.

Comparing Minimum Profitable Production Cost and Actual Housing Prices

We can compute the ratios of house prices to the minimum profitable production cost using different data sources on home values. Most of our results below use self-reported house prices from the micro data in the biannual *American Housing Survey (AHS)* which runs from 1985-2013. It reports data on individual housing units and their occupants in 98 core-based statistical areas (CBSAs), which refers to a metropolitan area of one or more counties anchored by an urban center of at least 10,000 people, tied together by commuting patterns. These markets (which are listed in Appendix 1) contain approximately 75 percent of the urbanized population in the United States according to 2010 Census data and include virtually any market of significant size.⁶

⁵ For a recent story, see the August 25, 2017, article in *The Economist* at <https://www.economist.com/blogs/graphicdetail/2017/08/daily-chart-17>.

⁶ We cannot calculate a truly national ratio of housing price-to-minimum profitable production cost. Construction costs are not reported by R.S. Means for each market in the country, and no such data are available for rural areas either. Moreover, the *American Community Survey* does not report anything on housing unit size, which means that added assumptions need to be made if using its data to compare housing prices and costs. We did experiment with the median priced-unit from the 2014 *American Community Survey* in computing price-to-cost ratios like those discussed immediately below. Those findings are very similar in quality and quantity to those reported below using the *AHS*.

Some strengths of the *American Housing Survey* data are that it contains micro data, clearly identifies single-family detached units, and reports the square footage of living area. The latter is useful as it allows us to match units of different sizes with the appropriate construction cost in the R.S. Means Company data. Smaller units typically have higher costs per square foot. We do this for homes of 600, 800, 1000, 1200, 1400, 1600, 1800, 2000, 2400, 2800, 3200, 3600 and 4000+ square feet of living area. Specifically, if a house is reported to be less than or equal to 700 square feet of living area, this is matched to R.S. Means Company costs per square foot for a 600 square foot, economy-quality home.

Each single family home that includes data on living area is matched with cost data from R.S. Means and then grouped into one of four bins, based on the ratio of housing prices to minimum profitable production cost: 1) A ratio of 0.75 or less, which implies that market value of the house is at least 25 percent below our estimate of reproduction costs; 2) a ratio between 0.75 and 1.25, which we interpret to be the range within which prices are not materially different from minimum profitable production costs; 3) a ratio between 1.25 and 2; and 4) a ratio greater than 2, which implies that prices are more than double our estimate of production costs. We chose these four relatively wide bins because they are likely to be reasonably robust to the measurement error involved in the construction of our ratios.

These ratios are essentially the value of Tobin's q for housing. Just as in standard investment theory, a value of q below one implies that the capital would not be replaced if it were destroyed. Values of q above one must reflect some barrier to investment, which we believe is more likely to be regulation in the housing market rather than standard adjustment costs (Hayashi 1982). Values of q above one can also be a sign of market over-valuation, as in Las Vegas in 2005, but only in cases where land is abundant and regulations are few.

Table 1 reports our baseline results, which include data from 1985 to 2013. As of 2013, slightly less than three-quarters of all observations (73.6 percent) are priced near or below minimum profitable production costs, with more than half of them being valued more than 25 percent below. This leaves just over one-fourth (26.4 percent) living in expensive housing, with 10 percent of the underlying sample living in homes estimated to be more than double minimum profitable production costs. In a large swath of urban America—and especially if one focuses on the local housing markets in the bottom four-fifths of prices—the housing market is supplying units at quite reasonable prices, given all-in production costs.

Also, Table 1 shows that the housing cycle matters. For example, at the height of the last housing boom, the 2005 data indicate that more than one-half of all observations were at least 25 percent more expensive than minimum profitable production costs.⁷

⁷ We also experimented with different housing quality assumptions in computing minimum profitable production cost. Assuming the lowest quality that meets local building codes will result in misclassifying some observations as expensive, especially those living in elite suburbs. If we use the costs associated with what R.S. Means terms “average” quality (one above economy quality), the share of observations classified as expensive (that is, with a ratio over 125 percent) falls from 26 to 18 percent. Assuming the highest possible construction quality—the “luxury” homes in R.S. Means terminology—is required to dramatically lower the estimate of expensive homes. In that case, the share of observations valued at more than 125 percent of minimum profitable production cost falls

Given the inevitable measurement error arising from unobserved quality differences across households in the micro data, another way to examine the spatial distribution of housing prices is at the metropolitan-area level. We look at the ratio of the median housing price to the minimum profitable production cost in every housing market for which we have at least 25 individual observations.⁸ These results can reflect the status of the typical homeowner in a market, even though we are not using all the underlying data. Table 2 reports the findings.

In 1985, over 90 percent of our metropolitan areas had median price-to-cost ratios less than or near 1. Only five (6.4 percent) had medians above 1.25 (and there were none where price was more than double production cost). This latter figure is only one-third of the 21.5 percent reported in Table 1 using all the micro data, but we do not find this surprising given the measurement error issue associated with unobserved unit quality, especially for homes located in the most desired suburbs. As of the middle of the 1980s, in only a handful of markets concentrated in California and Hawaii (none were east coast markets that year) was the typical home expensive relative to minimum profitable production cost. Based on earlier Census data, we presume that this distribution largely characterized housing markets before that point as well (Gyourko, Mayer, and Sinai 2013).

During the late 1980s boom in housing prices, median prices shifted up relative to construction costs. By 1991, the share of metropolitan areas with median value-to-cost ratios below 0.75 had fallen to 24 percent, but another 59 percent had values reasonably close to 1. The share of metropolitan areas with median price-to-cost ratios greater than 1.25 nearly tripled to just over 17 percent, with the Honolulu, Los Angeles, and San Francisco markets having prices more than double minimum profitable production costs.

The mid-1990s seems to have been a time of compression of metropolitan area prices, just as it was the only period in recent decades in which income inequality also declined. But between 1997 and 2007, median price-to-cost ratios in the most expensive markets rose dramatically. At the height of the boom in 2007, just over 48 percent of our metropolitan areas had median ratios with prices more than 25 percent above estimated reproduction costs, with well over one-third of those areas having price-to-cost ratios that were greater than two.

The years following the global financial crisis saw a distribution of median price-to-cost ratios that looked much like the early 1990s. By 2013, only three markets had median price-to-cost ratios above 2 — the same number as in 1991. Nearly 11 percent had ratios between 125 and 200 percent which is only slightly lower than the analogous share in 1991. Median price-to-cost ratios were less than 0.75 in one-third of markets in 2013, which is higher than the 24 percent in 1991. This implies that in a substantial fraction of urbanized America, it would not pay to

to just over 6 percent. Thus, our conclusion that the vast majority of homes are priced near or below their full social costs of replication is robust to virtually any assumption we could make.

⁸ This use of the median only for markets with 25 or more observations results in an unbalanced panel of markets, but the findings are not materially different if we restrict the data to the common set of metropolitan areas for which we have at least 25 observations each survey year dating back to 1985.

rebuild the typical home if it fell down today. Nominal prices have gone up in these areas since the late 1980s, but nominal construction costs have risen as well.

Perhaps the largest difference between 1985 and 2013 is that the share of metropolitan areas with median price-to-cost ratios above 1.25 has risen from 6.4 to 15.4 percent. There are a modest, but growing number of markets in America in which the typical owner is living in a home that is priced substantially above minimum profitable production costs. These markets include some of the nation's most productive labor markets, so they are important for our future.⁹

This gap between price and cost seems to reflect the influence of regulation, not the scarcity value arising from a purely physical or geographic limitation on the supply of land. For example, Glaeser, Gyourko and Saks (2005) show that the cost of Manhattan apartments are far higher than marginal construction costs, and more apartments can always be delivered by building up without using more land. This and other research we have done (Glaeser and Gyourko 2003) also finds that land is worth far more when it sits under a new home than when it extends the lot of an existing home, which is also most compatible with a view that the limitation is related to permits, not acreage per se.

It is possible that regulatory limits on construction are efficiently internalizing the negative externalities from construction, but the vast gap between price and construction cost in some coastal markets could only be justified by enormous construction externalities. Empirical investigations of the local costs and benefits of restricting building generally conclude that the negative externalities are not nearly large enough to justify the costs of regulation (Cheshire and Sheppard 2002; Glaeser, Gyourko and Saks 2005; Turner, Haughwout and van der Klaauw 2014). Glaeser and Ward (2009) also find that the impact of neighborhood level density on housing values in greater Boston is far too small to justify the current restrictions on new construction.

A Closer Look at Three Types of Markets: Detroit, Atlanta, and San Francisco

The housing market in Detroit-Warren-Dearborn, Michigan, is emblematic of a place in which home prices have been well under minimum profitable production costs for long periods of time. Graphically, this market is characterized in Figure 1. There is a kinked supply schedule of housing, with the vertical component reflecting the size of the current stock. The height of the supply schedule at that point is minimum profitable production costs. Prices in this type of market were pinned down by minimum profitable production cost in the past, when the market was growing (Glaeser and Gyourko 2005). This is reflected in the intersection of supply and demand, D_1 , which is on the horizontal part of the supply schedule.

Following a negative demand shock for the market (in this case, fierce foreign competition for the domestic auto industry that was concentrated in Detroit), demand dropped to D_2 and now

⁹ The three markets with ratios of median housing price to minimum profitable production cost above 2 are Los Angeles-Long Beach-Anaheim, CA, Oxnard-Thousand Oaks-Ventura, CA, and San Francisco-Oakland-Hayward, CA. Those with ratios between 1.25 and 2 are Baltimore-Columbia-Towson, MD, Boston-Cambridge-Newton, MA-NH, Denver-Aurora-Lakewood, CO, New York-Newark-Jersey City, NY-NY-PA, San Diego-Carlsbad, CA, Seattle-Tacoma-Bellevue, WA, and Washington-Arlington-Alexandria, DC-VA-MD-WV.

intersects supply on its vertical component. Prices are below the full production cost of new housing, because this intersection reflects the depreciated price of older housing. Most Americans, not just those in declining markets, do not live in new units. More than seven million occupied housing units were built before 1919, constituting approximately 6.2 percent of the occupied housing stock. Over 30 percent of occupied units in 2014 were built before 1960 and so were more than 50 years old. As shown in Panel A of Figure 2, the ratio of house prices-to-minimum profitable production cost in Detroit was well below 1 for much of the 1980s and 1990s, and then rose towards 1 during the recent long boom, before falling back after the bust ensued. Unsurprisingly, annual building permits in Detroit are not more than about 1 percent of market's 2000 housing stock in any year since 1985—and were near zero from 2007-2011, according to *American Housing Survey* data.

The Atlanta market is a canonical example of a local housing market in which supply is highly elastic and demand always is strong enough to keep prices at minimal profitable production cost. As Panel B of Figure 2 shows, new supply is highly volatile. Permitting intensity was running at 3 percent of market size in 1985; fell half by 1991 as the local economy declined; more than doubled to nearly 4.5 percent of market size by 2005; plummeted to below 0.5 percent of market size in the throes of the financial crisis by 2009; and it has only recently started to increase again. Amidst all this variation in new supply, the median owner's price-to-cost ratio never varies much from 1. This is consistent with a highly elastic supply side of the housing market, but one in which demand is intersecting it on the horizontal part of its schedule in Figure 1.

San Francisco represents the third type of housing market in which the price of housing is considerably above the minimum profitable production cost. In this situation, strict regulation of housing construction means developers in this type of market cannot bring on new supply even though it looks as if they could earn super-normal profits if they did. Unlike the graph in Figure 1, the supply schedule is upward sloping and demand is strong enough to intersect it well above where $P=MPPC$. Thus, shifts in the demand for housing affect price more than quantity. As Panel C of Figure 2 shows, the median house price in this market has been well above the minimum profitable production cost for the past three decades and reached dramatic heights at the peak of the last housing boom in 2005. However, permitting activity did not increase at all over the eight-year span from 1997-2005, even though the median price-to-cost ratio increased from below 2 to over 5. Although the ratio has fallen sharply from that peak, it remained a very high 2.84 as of 2013. The link between prices (relative to production costs) and new supply has been broken in this type of market.

San Francisco is a relatively high physical construction cost market, but that is not what makes its homes cost so much. The median housing unit in this market contained 1900 square feet, and the physical construction costs for this unit based on R.S. Means data were \$192,938, so the per square foot cost of the (presumed modest quality) structure was just over \$100 per square foot, which is one of the most expensive construction cost markets in the United States. Our earlier assumption that land is 20 percent of the physical-cost-plus-land total provides an estimated land price of \$48,235. Stated differently, that is what we think the underlying land would cost in a relatively unregulated residential development market. Add the builder's 17 percent gross

margin, and the minimum profitable production cost for this house is \$281,690. This compares with an actual price of the median house of \$800,000 (and thus a price-to-cost ratio of 2.84).

Clearly, San Francisco housing developers cannot actually earn super-normal profits on the margin. Instead, what makes San Francisco housing so expensive is the bidding up of land values. Our formula suggests that the land underlying this particular modest quality house cost about \$490,000—roughly 10 times the amount presumed for our underlying calculations of the minimum profitable production cost.

The time path of prices in the three cities is representative of a larger pattern: cities with inelastic housing supply generally experienced much more extreme price gyrations during the boom-bust cycles of the 1980s and 2000s. Glaeser, Gyourko and Saiz (2008) report that in the 1980s boom, mean price growth was 29 percent for most inelastic metropolitan areas and 3.4 percent for the most elastic metropolitan areas. During the 1996-2006 boom, mean real price growth was 93.9 percent in the most inelastic cities and 28.2 percent in the most elastic cities. The remarkable element in the 1996 to 2006 is that some relatively elastic cities, such as Phoenix and Las Vegas, still experienced extremely high price growth over a short time period, and equally sharp subsequent declines.

Additional Connections

Overall, more expensive housing markets tend to be both more regulated and have more inelastic supply sides. The correlation of median house price in 2013 with the Wharton Residential Land Use Restrictiveness Index (which has a bigger value the more restrictive the regulatory environment) is about 0.5. This is very similar to the magnitude of the correlation with Saiz's (2010) elasticity measure (although of the opposite sign because his measure declines in value the more inelastic is supply).

A broader look at our data also shows a clear connection between housing prices and new home construction activity. Figure 3 confirms that Atlanta, Detroit, and San Francisco are, indeed, representative of their market types. Price-to-MPPC ratios in 2013 are plotted against the magnitude of construction activity as reflected by the ratio of new units built between 2000-2013 to the 2000 stock. The modest negative slope that best fits that scatter plot of markets is driven by the following combination of facts: (a) among markets with high P/MPPC ratios of 1.5+, there was relatively little new home construction over this 13-year period (typically less than 15% in aggregate, or about 1% per annum on a compounded basis); in addition, there is little variation in permitting intensity among this group of the most expensive housing markets; (b) among markets with low P/MPPC ratios of 0.7 or less, there also was very little new home construction; building intensity in Detroit, Cleveland and Rochester is not much less than in Boston or New York City—for a very different reason, of course; in these lowest priced markets, developers cannot earn a normal profit given fundamental production costs; and (c) among the markets with P/MPPC ratios closer to one, there is a much wider range of building levels, depending upon the level of demand in each metropolitan area.

Our data also shows a marked increase in price dispersion across markets, with the right-tail of inflation-adjusted housing prices much longer now than it was three decades earlier. This is consistent with earlier research (for example, Gyourko, Mayer, and Sinai 2013).

Finally, there is a strong correlation between homeowner income and the degree of regulation in a market. Variation in the Wharton regulatory index or Saiz's (2010) elasticity can account for nearly 25 percent of the variation in the income of the owner of the median-priced home in 2014 based on *American Community Survey* data. Given the aforementioned positive correlation between house prices and the degree of regulatory constraint, it is not surprising to find higher-income people living in more expensive homes. Of course, no causal relation is implied from a simple, bivariate correlation. However, this does link to one of the most important new implications of inelastic supply sides in coastal markets—the potential impact on the distribution of wealth and on the geographic distribution of where people of different income levels are more likely to end up living. We now turn to these issues.

The Impacts of Supply Restrictions: Household Wealth

If housing restrictions have helped cause the secular rise in coastal housing prices, and the enormous volatility of prices during boom-bust cycles, then they may help explain the movement in household wealth in the U.S. and elsewhere. Piketty (2014) estimates that the ratio of the US capital stock-to-GDP increased from 332 percent in 1970 to 410 percent in 2010, and that increases in the value of the housing stock accounts for 40 percent of this increase. Increases in housing capital account for 83 percent of the increase in the ratio of private capital-to-income between 1970 and 2010. As Rognlie (2015) has carefully documented, the net capital share increase in the post-World War II era due to housing was from 3 to 8 percent of domestic value added. La Cava (2016) argues that this increase in housing wealth in recent decades has largely been due to supply-constrained markets.

This growth in the stock of housing capital relative to GDP in recent decades is primarily about prices, not the physical supply of housing. Between 1973 and 2010, the average new home expanded from 1,660 square feet to 2,392 square feet, but this 44 percent increase is far less than the 100 percent increase in income over the same time period. Standard indices such as the S&P/Case-Shiller Index or the Federal Housing Finance Agency (FHFA) housing price index, which use repeat sales and other methods to control for changes in the quality of housing quality, still show impressive increases in prices in restricted markets, such as the 109 percent increase in real prices in greater San Francisco between 1991 and 2016. An admittedly back-of-the-envelope calculation detailed in Appendix 2 suggests that owners of even modest properties in San Francisco fortunate enough to have bought prior to the rise of restrictive building regulations have seen an increase in wealth of several hundred thousand dollars. This increase in wealth is due to higher costs of land, not higher costs of physical construction, and in turn, we believe that the higher cost of land has been driven by binding land use restrictions.

Yet housing wealth is different from other forms of wealth because rising prices both increase the financial value of an asset and the cost of living. An infinitely lived homeowner who has no intention of moving and is not credit constrained would be no better off if her home doubled

in value and no worse off if her home value declined. The asset value increase exactly offsets the rising cost of living (Sinai and Souleles, 2005). This logic explains why home-rich New Yorkers or Parisians may not feel privileged: if they want to continue living in their homes, sky-high housing values do them little good.

Ultimately, the source of high housing costs determines its impact on well-being and personal finances. For example, if higher housing prices reflect higher wages, then San Francisco may have become less affordable, but residents who have owned property for a time are also richer. This logic leads Moretti (2013) to conclude that nominal wage inequality overstates true inequality, because those with high incomes need to pay more for access to their well-paid labor markets. Conversely, Diamond (2016) argues that high housing prices in educated metropolitan areas reflect higher amenity values in those areas, which implies that real inequality is higher than earnings inequality. More generally, if higher housing prices reflect more amenities, then buyers are no worse off, but if they reflect a greater demand for the same amenities, then buyers' welfare has fallen.

In any event, the gains in housing wealth are not evenly distributed. When housing prices rise, those who already own housing are essentially hedged against a higher cost of housing (Sinai and Souleles 2005). Renters, conversely, experience the rising housing costs directly and become poorer in real terms.

Because home-owners tend to be older while renters are younger, the limited growth in housing supply has created an intergenerational transfer to currently older people who happened to have owned in the relatively small number of coastal markets that have seen land values increase substantially. On a per owner basis, the value of these wealth gains can be considerable, but the number of markets is relatively small and many are not particularly populous. Only 11 of our Core Based Statistical Areas have a housing-price-to-cost ratio above 1.25 in 2013. In total, they contained 58.8 million people and 22.9 million total housing units (according to American Community Survey data for 2014). More than half of this total for these markets consists of the 31 million people and 12 million housing units in the huge New York City and Los Angeles markets; in total, these areas contain only about 23 percent of total urban population. This relatively low share of the urban population should not be a surprise: after all, there are areas with strong constraints on building, and people cannot move to these cities without a place to live.

Table 3 presents data on net worth for six different pairs of age groups in 1983 and in 2013 from the *Survey of Consumer Finances* carried out by the Federal Reserve. The public use samples do not provide any geographical identifiers, but we focus here on facts about home equity. We report values for the 50th, 75th, 90th, 95th and 99th percentiles of the distribution. Given the aggregate sample size, there are 30-40 observations per percentile, and we report the average of those observations.

This table allows us to look at the same age cohort at two different times, three decades apart: for example, comparing housing wealth for 18-24 year-olds in 1983 and in 2013. For example, the 18-24 age group has little housing wealth in 1983, and less at each percentile level in 2013. For the intermediate age groups—25-34, 35-44, 45-54—housing wealth is lower in 2013 than in

1983 at the 50th and 75th percentiles, and either roughly the same or lower at the 90th percentile. However, housing wealth is somewhat higher for these groups at the 95th and 99th percentile in 2013. For the oldest age groups—55-64 and 65-74—housing wealth is up considerably at the 90th percentile and above, with the increases being especially notable in the oldest group. Many in these age groups established as homeowners 30-40 years earlier, and so were in the best position to benefit from a rise in housing prices. In short, the *Survey of Consumer Finance* shows sharp home wealth increases only among the richest members of the oldest cohorts. Given the potential magnitudes involved and the rising prices in many coastal markets since the latest data from 2013, these patterns seem likely to have continued.

The big winners from the reduction in housing supply are a small number of older Americans who bought when prices were much lower. Some of this wealth may be passed to the next generation as bequests. But much of the housing price appreciation has probably already vanished from the home equity line in housing balance sheets, and turned into consumption by retirees who have moved away from America's priciest areas. The *Survey of Consumer Finances* data show that home equity has risen much more slowly than aggregate housing wealth, because rising mortgage levels have offset rising home values. Younger Americans, in particular, are more likely to have paid for their homes using large mortgages than to have experienced large wealth increases.

Overall, these shifts in housing wealth seems to show that older groups in certain geographic areas getting most of the gains, but we have not established causality. More research is needed to identify causality, especially because non-housing wealth is skewing in at least somewhat similar ways among the same groups noted above.

Boom-bust housing cycles can be important redistributors of wealth, too. Pfeffer, Danziger and Schoeni (2013) document that the median household in the *Panel Study of Income Dynamics* lost more than 50 percent of its wealth between 2007 and 2011, and that 83 percent of that loss came from real estate. Wolff (2014) found that in 2010, 16.2 percent of homeowners under the age of 35 had negative home equity, but only 5.3 percent of homeowners between 55 and 64 had negative home equity.

Housing supply shapes these wealth transfers because it partially determines the extent of a housing convulsion. Glaeser, Gyourko and Saiz (2008) show that the 1980s housing boom and subsequent bust largely bypassed places with elastic housing supply. In those years, buyers seem to have recognized that where it was easy to build, housing prices would not remain above construction costs for long. Consequently, the transfers of wealth that occurred during that boom were located primarily in places with restricted supply. The boom of the 2000s also disproportionately impacted places with limited supply, yet there were some areas such as Phoenix and Las Vegas that experienced booms despite enjoying relatively elastic housing supply. Because it takes time to build, over-optimistic buyers can still bid prices up in such markets for a few years. Eventually, the glut of new building in Las Vegas generated one of the largest of America's housing busts. Nonetheless, Mian, Rao and Sufi (2013) show that wealth losses, and associated consumption declines, were higher in places where housing is less elastically supplied.

The Impact of Supply Restrictions: Urban Labor Markets and Productivity

Rising house prices represent a transfer from buyers to sellers, which is not itself obviously a welfare gain or loss. Yet constricted housing supply also generates a potentially profound distortion: people are unable to move into more desirable metropolitan areas. Hsieh and Moretti (2017) and Ganong and Shoag (2015) have raised the possibility that housing restrictions have led to a misallocation of labor that could have a serious adverse effect on US GDP. Given the large differences in productivity between Las Vegas and San Francisco, it seems virtually certain that America's GDP would rise if, for example, the San Francisco Bay Area built more housing, allowing more population to shift there from Las Vegas.

To better understand the possible GDP gains from eliminating land use controls, it is useful to make simplifying assumptions, some of which can bias the calculation in ways that are discussed at the end of this section. One such assumption is that there are no differences in negative externalities across locations. While there is little evidence to suggest that the negative effect of an extra home in a constrained area is worse than the negative effect of an extra home in an unconstrained area, if the externalities of construction were far worse in some places than others, then our estimates will overstate the benefits of deregulating housing markets.¹⁰ Another is that construction costs are the same everywhere, which as discussed above is a roughly plausible assumption. We will also ignore amenity differences, so an absence of regulation will tend to equalize housing costs and wages across space.

In this setting, the potential output benefits from reallocating a fixed amount of labor from low wage areas to high wage areas can be seen in Figure 4, which depicts demand curves for two areas. The horizontal axis shows population in the constrained area, and a higher population in that area causes wages to decline. Population in the unconstrained area is the remainder, and so more population in the constrained area means less in the unconstrained area, leading to the upward-sloping demand curve for labor shown here. In the absence of land use controls, prices equalize across the two areas, which is shown in the point in the middle of the figure where the two curves meet. When housing supply is restricted, the wage in the restricted area is higher than in the unrestricted area.

If we assume that the demand for housing comes only from local labor markets, then we can treat each of these lines as a transformation of the labor demand curve, which in turn reflects the marginal product of labor. The lost output from misallocation is then equal to the area under the higher line from the restricted population level to the level that causes the lines to meet. This difference represents a classic deadweight loss triangle. In addition, there is a rectangle that represents the transfer to the owners of land in the more expensive area.

Hsieh and Moretti (2017) offer a set of illustrative calculations that have received considerable attention. They use a Cobb-Douglas production function in which the share of labor is 0.65 and the share of fungible capital—which will move in response to shifts in labor between cities—is 0.25. In this framework, the elasticity of labor demand is -7.5. In their analysis, changing the

¹⁰ Glaeser and Ward (2009) show that if one assumes constant construction costs (a rough but reasonable assumption, as discussed earlier), then land values are maximized when the gap between the mark-up over construction costs relative to price is equal to the absolute value of the elasticity of price with respect to density. Glaeser and Ward find that this gap is roughly ten times larger than the elasticity.

housing supply regulation in just three highly constrained markets – New York, San Francisco and San Jose – to the median for the country results in a nearly 9% rise in aggregate GDP. This is achieved via massive shifts in employment location. Jobs in the New York market increase by 1,010 percent, with those in San Jose rising by 689 percent. Naturally, output is much higher in these markets, too. Wages in these areas do fall, but only by 25 percent in their model.

The Cobb-Douglas production function with fungible capital is an important driver of this result in which cities can grow enormously with relatively modest decreases in wage. Assumptions about the shape of the labor demand function also have a strong effect in shaping the conclusions about the welfare losses from distortions in labor supply. Cobb-Douglas production functions tend to deliver particularly elastic labor demand curves, especially when capital is also mobile. Consequently, they lead to the conclusion that even relatively small wage gaps will result in large population misallocations and welfare losses.

For example, empirical estimates of the link between wages and labor demand at the local level are often much lower than predictions from a Cobb-Douglas function. Beaudry, Green and Sand (2014) present city-level labor demand elasticities that seem matched to our needs. They find a city-level labor elasticity of -0.3, which suggests that the overall impact is 0.7 percent of GDP. Their city-industry level estimates are larger (-1.0) and those would imply a misallocation cost equal to about 2 percent GDP. Past demand elasticities have typically ranged from -0.25 to -1.0 (Hammermesh 1991). In addition, we have experimented with back-of-the-envelope estimates of these gains using linear demand functions for labor, rather than the curved demand functions implied by the Cobb-Douglas function. While the precise outcome depends on the parameters used, such calculations suggest that 2 percent of GDP may be an upper bound on the gains from reallocation of labor.¹¹

We view any gain which involves adding several percent to GDP as quite sizeable and worth pursuing. But clearly considerable work remains to be done in pinning down the likely size of the potential gains. This follow-up work might also keep in mind the likely biases from our simplifying assumptions. Appendix 3 contains a more detailed discussion of these issues and the potential biases in the different calculations.

Empirical estimate of the costs of labor misallocation also should be cognizant of the problem of omitted human capital. The average worker in Tulsa will not necessarily earn the average wages in Silicon Valley by moving to San Jose. Any misallocation calculation will typically increase with the variance in perceived productivities, and the noise created by unobserved human capital heterogeneity will generally cause an overestimate of misallocation costs.

Another issue is that if places with higher human capital-adjusted wages typically have more amenities, because cities are more likely to form only if an area is either productive or nice or both, then differences in the cost of housing will lead to an overestimate of the true differences in productivity. Conversely, there are some examples of large urban areas, like Orlando and Miami, which have lower-than-average wages and housing prices, but which have the amenity of Florida sunshine. Again, not taking that amenity into account will bias attempts to infer productivity from wage levels.

¹¹ For examples of these calculations, see the online Appendix available with this paper at <http://e-jep.org>.

On the other side, our calculations reflect only an estimate based on static factors. One might speculate that Silicon Valley and other high-productivity urban areas are about creativity, as well as high wages. So, more Silicon Valley residents could also mean more technological innovation and faster productivity growth. If agglomeration economies are important, and tend to increase with population size, then this will attenuate the downward impact of added population on earnings. We are ignoring the impact that higher output has on product demand, which is captured in Hsieh and Moretti (2017), which also pushes earnings and the benefits from better labor allocation upward.

Next, the reallocation of population implied in this analysis would mean that the overwhelming majority of cities would lose population, while a few such as New York and the San Francisco Bay Area would gain substantial numbers of workers. In some of our back-of-the-envelope calculations, the entire population of certain cities would depart! As discussed earlier in the paper, declines in local demand for housing, given the durability of housing, can easily cause housing prices to fall in those cities—which further complicates calculations about what reallocation of population and welfare gains might be possible as a result of less stringent limits on housing construction. Finally, we stress again that we have assumed away any benefits that regulation might create by reducing the negative externalities from construction, so that these estimates should be taken as suggestive, not definitive.

Conclusion

When housing supply is highly regulated in a certain area, housing prices are higher and population growth is smaller relative to the level of demand. While most of America has experienced little growth in housing wealth over the past 30 years, the older, richer buyers in America's most regulated areas have experienced significant increases in housing equity. The regulation of America's most productive places seems to have led labor to locate in places where wages and prices are lower, reducing America's overall economic output in the process.

Advocates of land use restrictions emphasize the negative externalities of building. Certainly, new construction can lead to more crowded schools and roads, and it is costly to create new infrastructure to lower congestion. Hence, the optimal tax on new building is positive, not zero. However, there is as yet no consensus about the overall welfare implications of heightened land use controls. Any model-based assessment inevitably relies on various assumptions about the different aspects of regulation and how they are valued in agents' utility functions.

Empirical investigations of the local costs and benefits of restricting building generally conclude that the negative externalities are not nearly large enough to justify the costs of regulation. Adding the costs from substitute building in other markets generally strengthens this conclusion, as Glaeser and Kahn (2010) show that America restricts building more in places that have lower carbon emissions per household. If California's restrictions induce more building in Texas and Arizona, then their net environmental could be negative in aggregate. If restrictions on building limit an efficient geographical reallocation of labor, then estimates based on local externalities would miss this effect, too.

If the welfare and output gains from reducing regulation of housing construction are large, then why don't we see more policy interventions to permit more building in markets such as San

Francisco? The great challenge facing attempts to loosen local housing restrictions is that existing homeowners do not want more affordable homes: they want the value of their asset to cost more, not less. They also may not like the idea that new housing will bring in more people, including those from different socio-economic groups.

There have been some attempts at the state level to soften severe local land use restrictions, but they have not been successful. Massachusetts is particularly instructive because it has used both top-down regulatory reform and incentives to encourage local building. Massachusetts Chapter 40B provides builders with a tool to bypass local rules. If developers are building enough formally-defined affordable units in unaffordable areas, they can bypass local zoning rules. Yet localities still are able to find tools to limit local construction, and the cost of providing price-controlled affordable units lowers the incentive for developers to build. It is difficult to assess the overall impact of 40B, especially since both builder and community often face incentives to avoid building “affordable” units. Standard game theoretic arguments suggest that 40B should never itself be used, but rather work primarily by changing the fallback option of the developer. Massachusetts has also tried to create stronger incentives for local building with Chapters 40R and 40S. These parts of their law allow for transfers to the localities themselves, so builders are not capturing all the benefits. Even so, the Boston market and other high cost areas in the state have not seen meaningful surges in new housing development.

This suggests that more fiscal resources will be needed to convince local residents to bear the costs arising from new development. On purely efficiency grounds, one could argue that the federal government provide sufficient resources, but the political economy of the median taxpayer in the nation effectively transferring resources to much wealthier residents of metropolitan areas like San Francisco seems challenging to say the least. However daunting the task, the potential benefits look to be large enough that economists and policymakers should keep trying to devise a workable policy intervention.

References

- Beaudry, Paul, David Green and Benjamin Sand. "In Search of Labor Demand", National Bureau of Economic Research Working Paper 20568, October 2014.
- Cheshire, Paul and Steven Sheppard. "The Welfare Economics of Land Use Planning", *Journal of Urban Economics*, Vol 52 (2002): 242-269.
- Davis, Morris and Jonathon Heathcote. "The Price and Quantity of Residential Land in the United States", Discussion Paper No. 2004-37, Board of Governors of the Federal Reserve System, Finance and Economics.
- Davis, Morris A. and Michael G. Palumbo. "The Price of Residential Land in Large US Cities", *Journal of Urban Economics*, Vol. 63 (2008): 352-384.
- Diamond, Rebecca. "The Determinants and Welfare Implications of U.S. Workers' Diverging Location Choices by Skill: 1980-2000". *American Economic Review*, Vol 106, no. 3 (2016): 479-524.
- Fishel, William. *The Homevoter Hypothesis*. Cambridge, MA: Harvard University Press, 2001.
- Frieden, Bernard. "A New Regulation Comes to Suburbia", *Public Interest*, No. 55 (Spring 1979): 15-27.
- Ganong Peter and Daniel Shoag. "Why Has Regional Income Divergence in the U.S. Declined?", Harvard Kennedy School Working Paper No. RWP12-028, March 28, 2013.
- Glaeser, Edward and Joseph Gyourko. "The Impact of Zoning on Housing Affordability", *FRBNY Economic Policy Review*, June 2003: 21-39.
- _____. "Urban Decline and Durable Housing", *Journal of Political Economy*, Vol 113, no. 2 (April 2005): 345-375.
- _____. *Rethinking Federal Housing Policy*. Washington, D.C.: The AEI Press, 2008.
- _____. "The Economic Implications of Housing Supply", Zell/Lurie Real Estate Center, Working Paper, January 4, 2017
- Glaeser, Edward, Joseph Gyourko and Albert Saiz. "Housing Supply and Housing Bubbles", *Journal of Urban Economics*, Vol. 64, no. 3 (2008): 198-217.
- Glaeser, Edward, Joseph Gyourko and Raven E. Saks. "Why is Manhattan So Expensive? Regulation and the Rise in Housing Prices", *Journal of Law and Economics*, Vol. 48 no. 2 (October 2005): 331-369.

- Glaeser, Edward and Matthew Kahn. “The Greenness of Cities: Carbon Dioxide Emissions and Urban Development”, *Journal of Urban Economics*, Vol. 67, no. 3 (May 2010): 404-418.
- Glaeser, Edward and Bryce Ward. “The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston”, *Journal of Urban Economics*, Vol. 65, no. 3 (2009): 265-278.
- Gyourko, Joseph, Christopher Mayer and Todd Sinai. “Superstar Cities”. *American Economic Journal-Economic Policy*, Vol. 5, no. 4 (2013): 167-199.
- Gyourko, Joseph and Raven Molloy. “Regulation and Housing Supply”. In *Handbook of Regional and Urban Economics*, Vol.5b, Edited by Gilles Duranton, J. Vernon Henderson and William Strange. Amsterdam: Elsevier, 2015.
- Gyourko, Joseph and Albert Saiz. “Construction Costs and the Supply of Housing Structure”, *Journal of Regional Science*, Vol. 46, no. 6 (October 2006): 627-660.
- Gyourko, Joseph, Albert Saiz and Anita A. Summers. “A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index”, *Urban Studies*, Vol. 45, no. 3 (2008): 693-721.
- Hammermesh, Daniel. “Labor Demand: What Do We Know? What Don’t We Know?”, National Bureau of Economic Research Working Paper 3890, November 1991.
- Hayashi, Fumio. "Tobin's Marginal q and Average q: A Neoclassical Interpretation", *Econometrica*, Vol. 15, no. 1 (January 1982): 213-224.
- Hsieh, Chang-Tai and Enrico Moretti. “Housing Constraints and Spatial Misallocation”, Working Paper, May 18, 2017.
- La Cava, Gianni. “Housing Prices, Mortgage Interest Rates and the Rising Share of Capital in the United States”, BIS Working Papers, No. 572, July 2016.
- Mian, Atif, Kamalesh Rao and Amir Sufi. “Household Balance Sheets, Consumption and the Economic Slump”, Chicago Booth Research Paper 13-42, 2013.
- Moretti, Enrico. “Real Wage Inequality”, *American Economic Journal: Applied Economics*, Vol. 5, no. 1 (January 2013): 65-103.
- Pfeffer, Fabian, Sheldon Danziger and Robert Schoeni. “Wealth Disparities before and after the Great Recession”. *The Annals of the American Academy of Political and Social Science*, 650, no. 1 (2013): 98-123.
- Piketty, Thomas. *Capital in the Twenty-First Century*. Cambridge, MA: Harvard University Press, 2014.

- Roback, Jennifer. “Wages, Rents and the Quality of Life”. *Journal of Political Economy*, Vol. 90 (December 1982): 1257-78.
- Rognlie, Matthew. “Deciphering the Fall and Rise in Net Capital Share: Accumulation or Scarcity?”, *Brookings Papers on Economic Activity*, Spring 2015, 1-54.
- Rosen, Sherwin. “Wage-Based Indexes of Urban Quality of Life.” In *Current Issues in Urban Economics*, edited by Peter Mieszkowski and Mahlon Straszheim, Baltimore: Johns Hopkins University Press, 1979.
- R.S. Means Company. 2015. RS Means electrical cost data. Kingston, MA: R.S. Means Co.
- Saiz, Albert (2010). “The Geographic Determinants of Housing Supply”, *Quarterly Journal of Economics*, Vol. 125, no. 3: 1253-1296.
- Sinai, Todd and Nicholas Souleles. “Owner-Occupied Housing as a Hedge Against Rent Risk”, *Quarterly Journal of Economics*, Vol. 210, no. 2 (2005): 763-789.
- Turner, Matthew A., Andrew Haughwout, Wilbert van der Klaauw. “Land Use Regulation and Welfare”, *Econometrica*, Vol. 82, no. 4 (July 2014): 1341-1403.
- Wolff, Edward. “Household Wealth Trends in the United States, 1962-2013: What Happened Over the Great Recession?”, National Bureau of Economic Research Working Paper No. w20733, 2014.

Figure 1: Kinked Supply Schedule from Durable Housing and Urban Decline

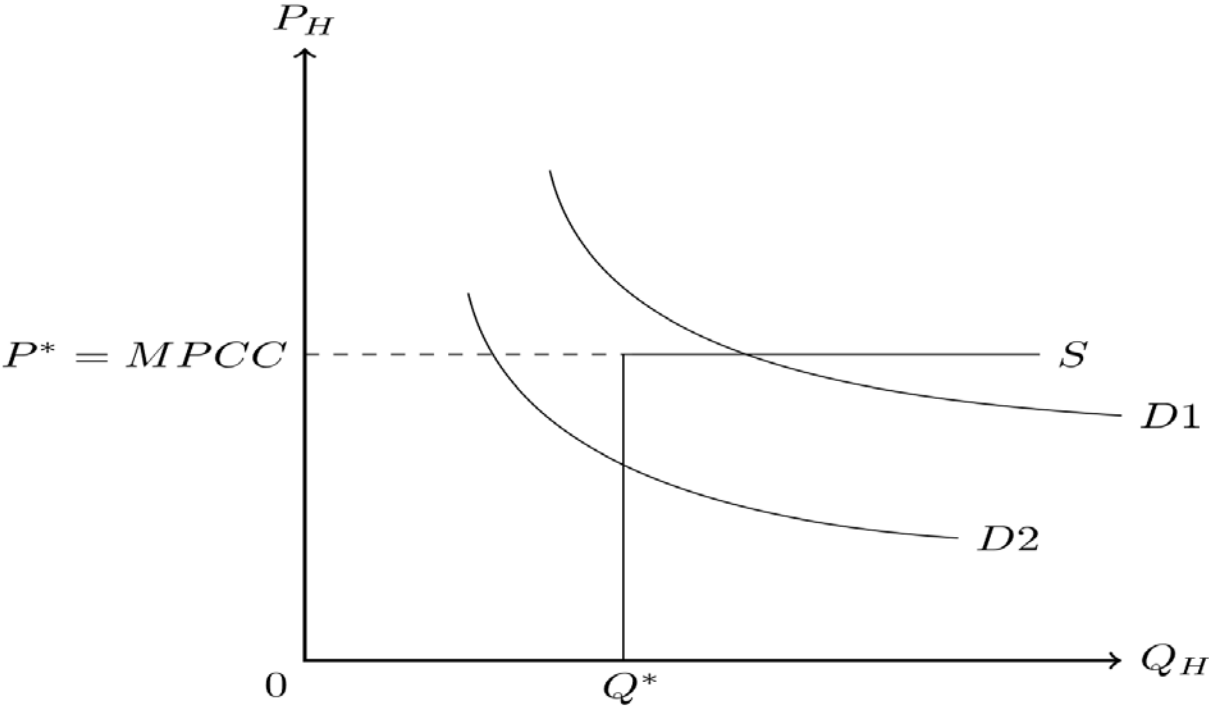
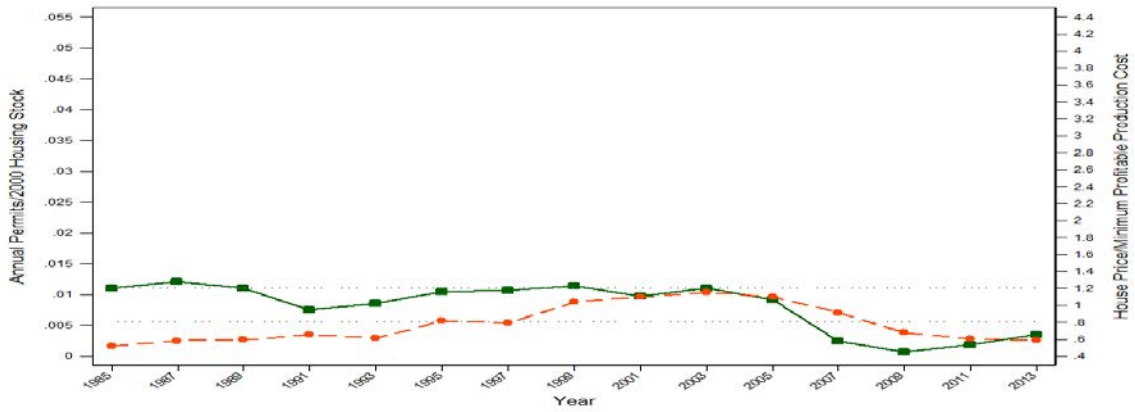
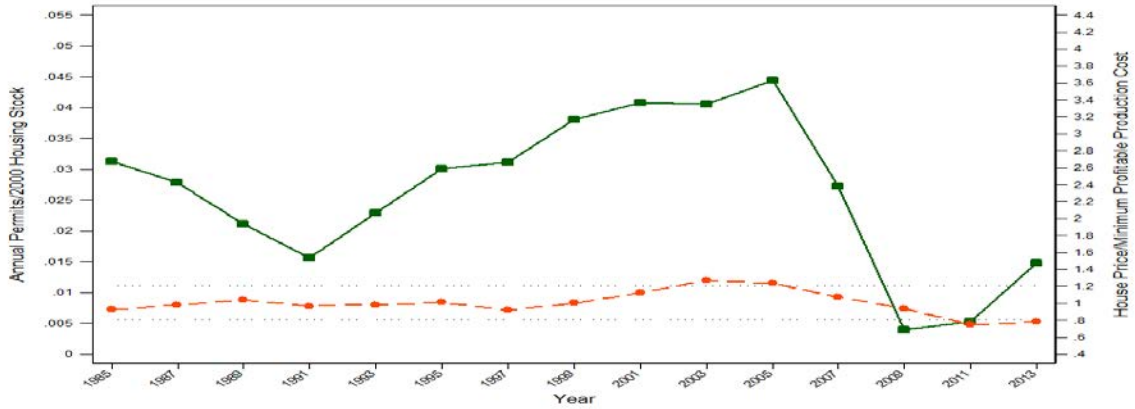


Figure 2: New Housing Supply and House Prices (Relative to Costs)

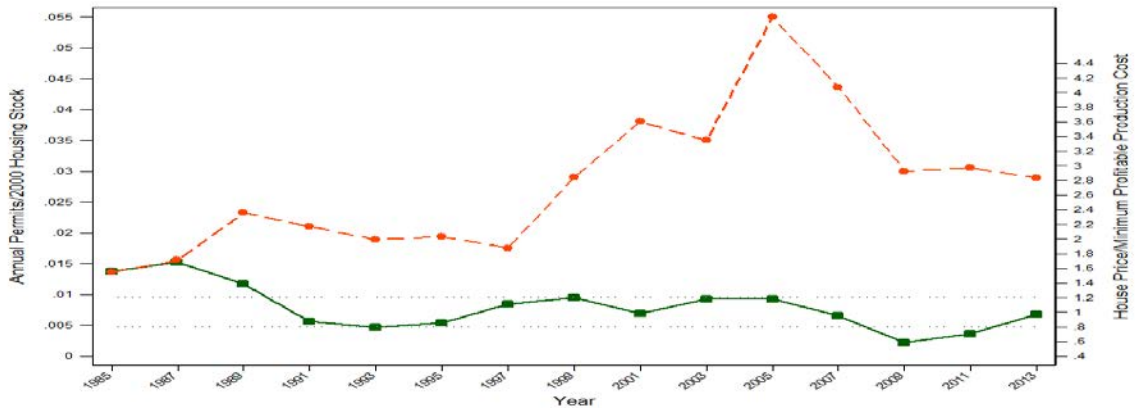
Panel A (Declining Market): Detroit-Warren-Dearborn, MI



Panel B (Growing, Elastically-Supplied Market): Atlanta-Sandy Springs-Roswell, GA



Panel C (Growing, Inelastically-Supplied Market): San Francisco-Oakland-Hayward, CA



—■— Annual Permits/2000 Housing Stock
- - -●- - - Median MSA House Price/Minimum Profitable Production Cost

House price are AHS micro data
 MPPC are for 1800 sf, economy class, 20% land share and 17% gross margin homes

Figure 3: Price-to-Cost Ratios and Permitting Intensity, 2000-2013

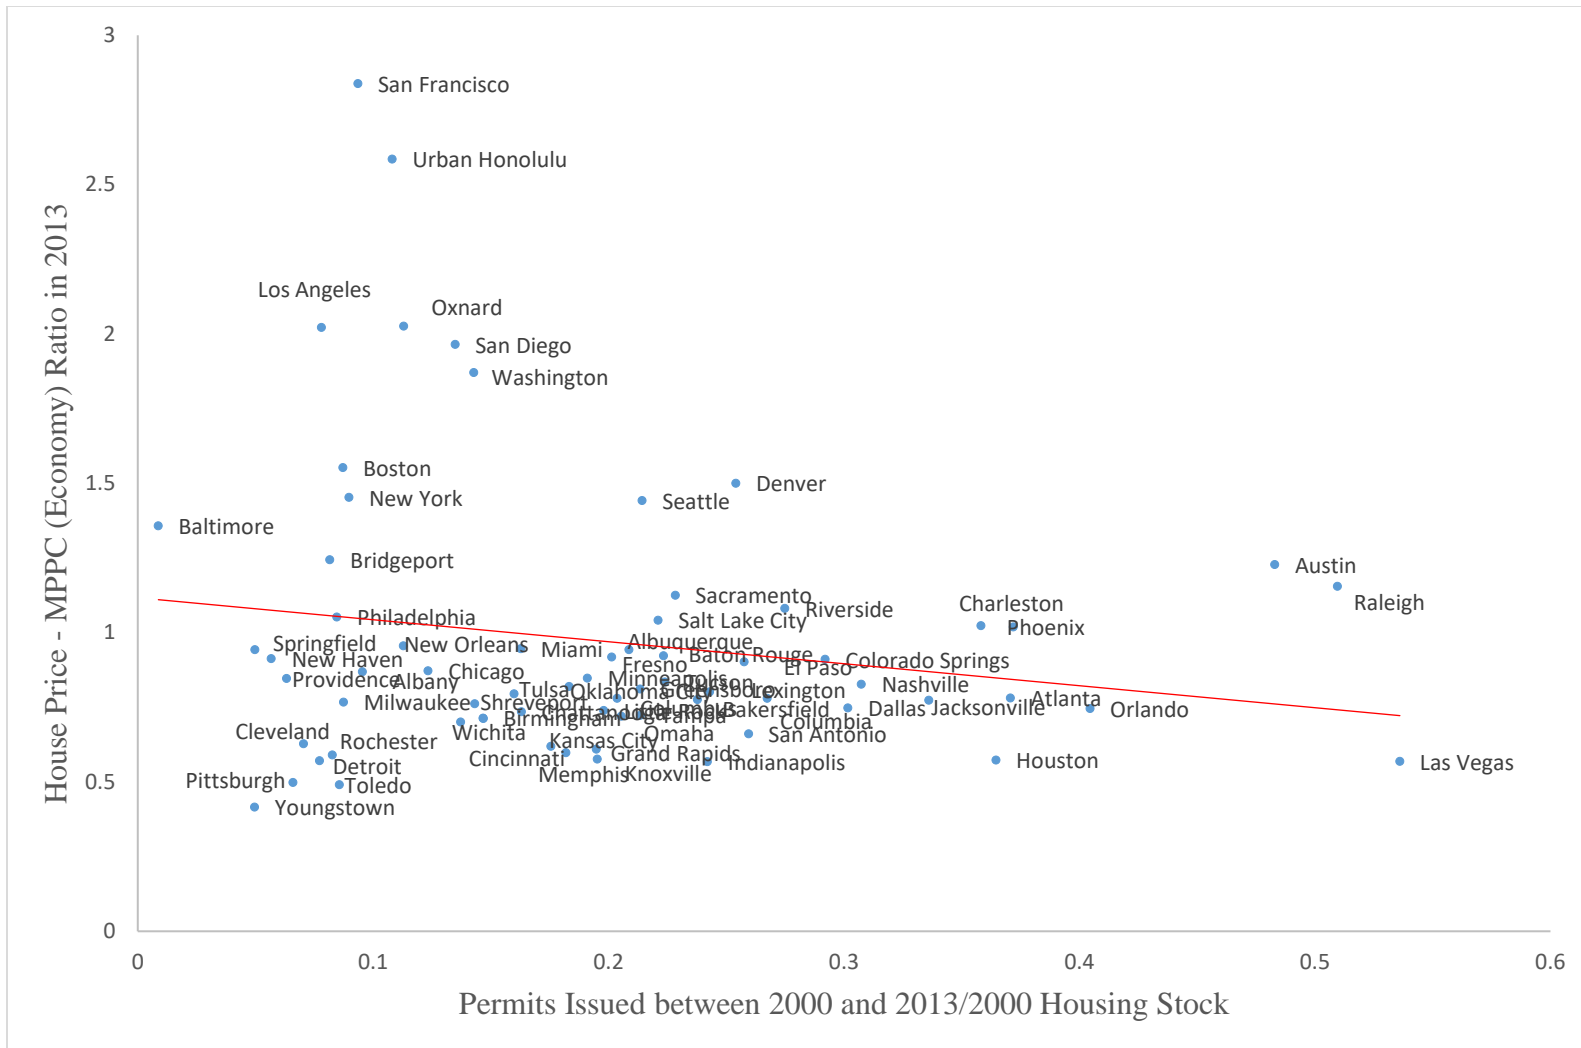


Figure 4: Welfare Consequences of Restricting Development in a Productive Market

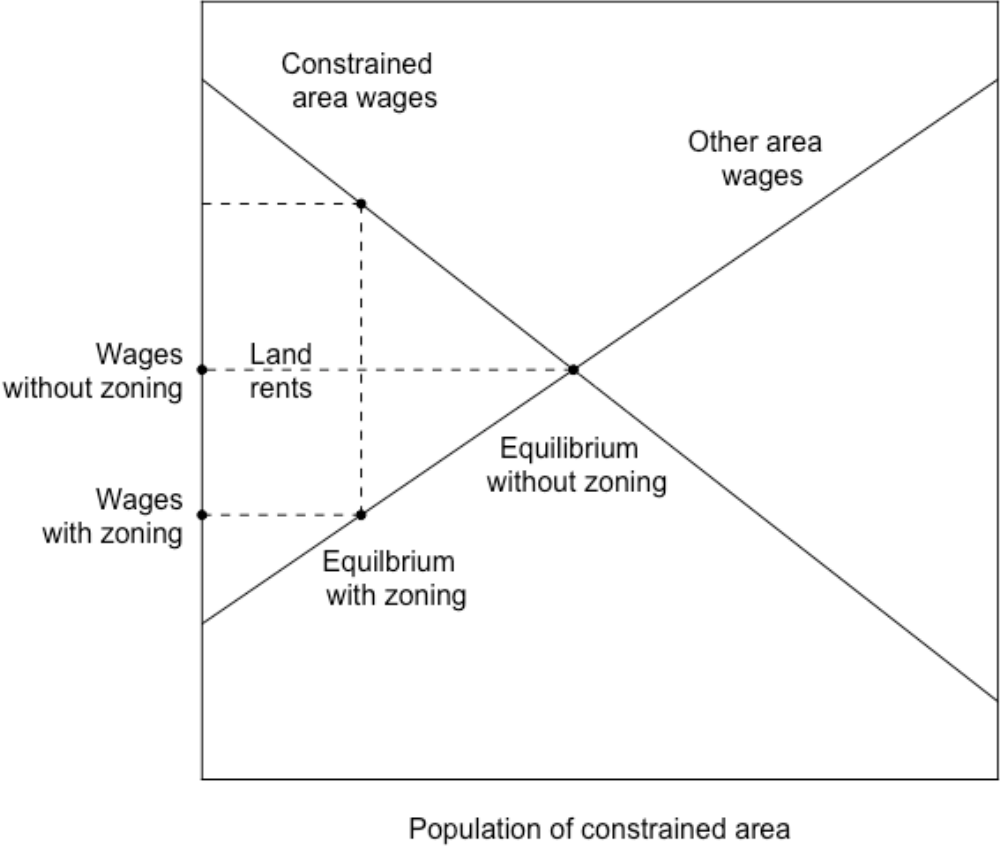


Table 1: House Price-to-Minimum Profitable Production Cost Ratio (HP/MPPC)
(Micro Data)

Year	P/MPPC ≤ 0.75	$0.75 < \text{P/MPPC} \leq 1.25$	$1.25 < \text{P/MPPC} \leq 2$	P/MPPC > 2
1985	38.0%	40.5%	17.9%	3.6%
1987	33.4%	38.3%	21.7%	6.6%
1989	31.8%	34.6%	20.3%	13.3%
1991	31.1%	35.3%	22.5%	11.1%
1993	31.8%	36.1%	23.6%	8.5%
1995	27.4%	37.7%	26.5%	8.4%
1997	31.5%	40.0%	23.0%	5.5%
1999	22.0%	40.1%	26.2%	11.8%
2001	19.4%	38.2%	25.2%	17.1%
2003	16.2%	32.1%	25.9%	25.9%
2005	18.0%	28.7%	25.3%	28.0%
2007	19.9%	28.1%	24.0%	28.0%
2009	31.4%	33.9%	21.6%	13.1%
2011	37.4%	35.4%	16.0%	11.2%
2013	40.3%	33.3%	16.2%	10.2%

Source: Authors' calculation uses American Housing Survey and R.S. Means Company data. See the text for details.

Table 2: House Price-to-Minimum Profitable Production Cost Ratio (HP/MPPC)
(Median Values, by CBSA)

Year	No. MSA	P/MPPC ≤ 0.75	$0.75 < P/MPPC \leq 1.25$	$1.25 < P/MPPC \leq 2$	P/MPPC > 2
1985	78	37.2%	56.4%	6.4%	0.0%
1987	72	29.2%	56.9%	13.9%	0.0%
1989	78	34.6%	50.0%	10.3%	5.1%
1991	70	24.3%	58.6%	12.9%	4.3%
1993	78	25.6%	61.5%	11.5%	1.3%
1995	69	18.8%	68.1%	10.1%	2.9%
1997	65	13.8%	72.3%	13.8%	0.0%
1999	68	7.4%	75.0%	14.7%	2.9%
2001	67	4.5%	70.1%	17.9%	7.5%
2003	69	4.3%	62.3%	23.2%	10.1%
2005	65	10.8%	44.6%	27.7%	16.9%
2007	64	10.9%	40.6%	29.7%	18.8%
2009	62	25.8%	50.0%	21.0%	3.2%
2011	64	28.1%	51.6%	15.6%	4.7%
2013	65	33.8%	50.8%	10.8%	4.6%

Source: Authors' calculation uses American Housing Survey and R.S. Means Company data. See the text for details.

Table 3: Housing Net Worth – 30 Year Changes (\$2013)				
Percentile	1983		2013	
	18-24 year olds	45-54 year olds	18-24 year olds	45-54 year olds
50	\$0	\$87,120	\$0	\$30,000
75	\$0	\$152,159	\$0	\$109,000
90	\$24,803	\$248,818	\$5,500	\$250,000
95	\$47,488	\$353,190	\$43,000	\$400,000
99	\$141,808	\$862,359	\$95,000	\$1,000,000
Percentile	25-34 year olds	55-64 year olds	25-34 year olds	55-64 year olds
50	\$0	\$94,184	\$0	\$60,000
75	\$45,352	\$161,886	\$21,000	\$167,000
90	\$91,827	\$255,361	\$74,000	\$350,000
95	\$123,135	\$353,190	\$140,000	\$543,000
99	\$230,751	\$760,380	\$256,000	\$1,500,000
Percentile	35-44 year olds	65-74 year olds	35-44 year olds	65-74 year olds
50	\$55,799	\$82,411	\$6,000	\$100,000
75	\$118,660	\$150,136	\$58,200	\$225,000
90	\$180,763	\$279,972	\$168,000	\$440,000
95	\$247,349	\$426,936	\$300,000	\$701,000
99	\$531,198	\$941,840	\$1,025,000	\$2,000,000

Notes: Data compiled from the 1983 and 2013 *Survey of Consumer Finances* using publicly available samples.

Appendix 1: CBSA List

CBSA Code	CBSA Name	CBSA Code	CBSA Name
10580	Albany-Schenectady-Troy, NY	33340	Milwaukee-Waukesha-West Allis, WI
10740	Albuquerque, NM	33460	Minneapolis-St. Paul-Bloomington, MN-WI
10900	Allentown-Bethlehem-Easton, PA-NJ	33660	Mobile, AL
12060	Atlanta-Sandy Springs-Roswell, GA	34980	Nashville-Davidson--Murfreeseboro--Franklin, TN
12420	Austin-Round Rock, TX	35300	New Haven-Milford, CT
12540	Bakersfield, CA	35380	New Orleans-Metairie, LA
12580	Baltimore-Columbia-Towson, MD	35620	New York-Newark-Jersey City, NY-NJ-PA
12940	Baton Rouge, LA	36420	Oklahoma City, OK
13820	Birmingham-Hoover, AL	36540	Omaha-Council Bluffs, NE-IA
14460	Boston-Cambridge-Newton, MA-NH	36740	Orlando-Kissimmee-Sanford, FL
14860	Bridgeport-Stamford-Norwalk, CT	37100	Oxnard-Thousand Oaks-Ventura, CA
15940	Canton-Massillon, OH	37900	Peoria, IL
16700	Charleston-North Charleston, SC	37980	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD
16860	Chattanooga, TN-GA	38060	Phoenix-Mesa-Scottsdale, AZ
16980	Chicago-Naperville-Elgin, IL-IN-WI	38300	Pittsburgh, PA
17140	Cincinnati, OH-KY-IN	39300	Providence-Warwick, RI-MA
17460	Cleveland-Elyria, OH	39580	Raleigh, NC
17820	Colorado Springs, CO	40140	Riverside-San Bernardino-Ontario, CA
17900	Columbia, SC	40380	Rochester, NY
18140	Columbus, OH	40900	Sacramento--Roseville--Arden-Arcade, CA
19100	Dallas-Fort Worth-Arlington, TX	41620	Salt Lake City, UT
19340	Davenport-Moline-Rock Island, IA-IL	41700	San Antonio-New Braunfels, TX
19740	Denver-Aurora-Lakewood, CO	41740	San Diego-Carlsbad, CA
19820	Detroit-Warren-Dearborn, MI	41860	San Francisco-Oakland-Hayward, CA
21340	El Paso, TX	42540	Scranton--Wilkes-Barre--Hazleton, PA
22420	Flint, MI	42660	Seattle-Tacoma-Bellevue, WA
23420	Fresno, CA	43340	Shreveport-Bossier City, LA
24340	Grand Rapids-Wyoming, MI	44060	Spokane-Spokane Valley, WA
24660	Greensboro-High Point, NC	44140	Springfield, MA
26420	Houston-The Woodlands-Sugar Land, TX	45060	Syracuse, NY
26900	Indianapolis-Carmel-Anderson, IN	45300	Tampa-St. Petersburg-Clearwater, FL
27260	Jacksonville, FL	45780	Toledo, OH
28140	Kansas City, MO-KS	46060	Tucson, AZ
28940	Knoxville, TN	46140	Tulsa, OK
29820	Las Vegas-Henderson-Paradise, NV	46520	Urban Honolulu, HI
30460	Lexington-Fayette, KY	46700	Vallejo-Fairfield, CA
30780	Little Rock-North Little Rock-Conway, AR	47900	Washington-Arlington-Alexandria, DC-VA-MD-WV
31080	Los Angeles-Long Beach-Anaheim, CA	48620	Wichita, KS
32820	Memphis, TN-MS-AR	49340	Worcester, MA-CT
33100	Miami-Fort Lauderdale-West Palm Beach, FL	49660	Youngstown-Warren-Boardman, OH-PA

Appendix 2: Imputing the Rise in Land Values for a 1985 Buyer in San Francisco

A number of assumptions have to be made to impute the capital gain on homes bought three decades ago in a market such as San Francisco. A given quality unit has to be defined as the starting point for such a calculation. For example, the unit underlying the median 1985 HP/MPPC value of 1.55 in San Francisco contained 1,300 square feet of living space and was reported to be worth \$150,000 (in 1985 dollars; \$324,000 in 2013 dollars given the 116% increase in the general urban price level between 1985-2013). Given our knowledge of construction costs and presuming a 17% gross builder's margin, we can impute a nominal raw land value of \$66,284 for this unit (\$143,279 in 2013 dollars) using equation (1) above. If we further presume that this owner kept the home and experienced the same 98% real increase reported for the median home in this market, the underlying land increased in value by just over \$272,000 to about \$416,000 (in 2013 dollars).

The details behind that calculation are as follows. The 98% real appreciation on the \$324,000 value of the home in 2013 dollars yields a value of \$641,520 in 2013. If we conservatively subtract the real value of construction costs times the builder's 17% gross margin, that leaves \$486,416 in value. Presuming that 17% of that remainder somehow gets captured by a builder still leaves a land price of \$415,742 in 2013. Subtracting off the \$143,279 that the owner paid in 2013 dollars in 1985 yields the gain of \$272,463. One could argue the gain is higher, as there is no 'builder' involved if the owner simply kept the property. One could argue over various assumptions and our point is not to provide a precise dollar figure. Rather, it is to show that a readily defensible, back-of-the-envelope calculation indicates that owners of modest properties in San Francisco in 1985 have seen more than a quarter million dollars of wealth come their way over the past three decades from land value appreciation that we believe is driven by binding land use restrictions. This is a near tripling of real land value for a long-term owner of a very modest house in San Francisco over the past three decades.

Appendix 3: Speculative Calculations of Welfare Losses from Land Use Restrictions

Our discussion of possible G.D.P. gains from eliminating land use controls assumes away construction cost differences across space, as well as congestion externalities and the like. We will also ignore amenity differences, so an absence of regulation means that housing costs will be equal and hence wages will also be equal across space.

The basic algebra of misallocation costs can be seen by assuming that wages reflect the marginal productivity of labor in each location i : $F_L^i(L_i)$. The output gain from reallocating Δ individuals from place B to place A, for any two locations is $\int_0^\Delta F_L^A(L_A + z)dz - \int_{-\Delta}^0 F_L^B(L_A + z)dz$. If we use a linear approximation $F_L^i(L_i + z) = F_L^i(L_i) \left(1 - \frac{\alpha z}{L_i}\right)$, then to equalize wages between B and A, there must be a change in population of $\Delta = \frac{L_A L_B (F_L^A(L_A) - F_L^B(L_B))}{\alpha (F_L^A(L_A) L_B + F_L^B(L_B) L_A)}$. The total output impact of the change is $\frac{F_L^A(L_A) - F_L^B(L_B)}{2} \Delta$, or the traditional welfare triangle of 0.5 times the gap in wages times the predicted population movement to eliminate misallocation. Everything needed for this calculation is observable directly from the data, except for α , which represents the inverse elasticity of labor demand.

To understand just how big the possible range of welfare gains could be, assume that $F_L^A(L_A) = 1.5 \cdot F_L^B(L_B)$, and that $L_A = .5L_B$. In that case, Δ must equal $\frac{L_B}{8\alpha}$ and the welfare gains equals the current earnings in area B, times $\frac{1}{32\alpha}$. Consequently, if α equals one, then this benefit would equal no more than 1/32 of total payroll in the lower paying area, which is significant but not massive.

There are functional forms that would deliver far higher welfare gains. Following Hsieh and Moretti (2017), assume instead that $F_L^i(L_i + z) = F_L^i(L_i) \left(\frac{L_i + z}{L_i}\right)^{\frac{\gamma + \eta - 1}{1 - \eta}}$, where γ represents the share of labor in a Cobb-Douglas production function (assumed to be .65) and η represents the share of fungible capital (assumed to be .25), which will move in response to labor. In that case, a 50 percent initial wage gap can only be closed if 87 percent of the population of the less productive area moves to the more productive area. Assuming that output is

$F_L^i(L_i) \left(\frac{L_i + z}{L_i}\right)^{\frac{\gamma + \eta - 1}{1 - \eta}} \frac{L_i + z}{\gamma}$, then the increase in output is 40 percent of output in the initially less productive place. This Cobb-Douglas formulation produces a value of $-\frac{\partial \text{Log}(Wage)}{\partial \text{Log}(Labor)}$ of 0.13, and if this were the value of α in our linear model, the welfare gains would rise to 1/4 of payroll in the lower paying area.

The Cobb-Douglas structure with fungible capital implies that cities can grow enormously with only modest decreases in wages. Perhaps, this is true. Agglomeration economies would only further attenuate the downward impact of added population on earnings. Yet, as we will shortly discuss, the empirical literature on local labor demand tends to find that labor demand is far less responsive to wages than this Cobb-Douglas model would imply.

Before proceeding with our main calibration, it is worth stressing that any spatial allocation exercise must face the problem of omitted human capital. Any misallocation calculation will typically increase with the variance in perceived productivities, and the noise created by unobserved human capital heterogeneity will generally cause an overestimate of misallocation costs.

Housing costs can themselves be used to assess the heterogeneity in human-capital adjusted wages. If places with higher human capital-adjusted wages typically have lower amenities, because cities are more likely to form only if an area is either productive or nice or both, then these cost of living differences may underestimate the true heterogeneity of productivity. If more productive people live in places with more amenities, then housing differences will also overestimate true productivity heterogeneity.

For our exercise, we will treat differences in payroll per worker as the true differences in the marginal product of labor, but we recognize that this is likely to lead to an overestimate of the true gains from reallocating labor. Using our linear approximation, if we have a large number of areas, with initial populations L_i and initial wage levels $F_L^i(L_i)$ and we move their populations to the point where their wages are equal to a constant $\hat{w} = \sum_i L_i / \sum_i (L_i / F_L^i(L_i))$, then the total gains from reallocation equal:

$$\text{Gain from Reallocation} = \frac{1}{2\alpha} \sum_i L_i (F_L^i(L_i) - \hat{w})$$

Our linear approximation means that when labor moves from less productive places to more productive places of equal size, then the average wages will fall since the marginal product curve is steeper in the more productive place. Equalizing wages will generate a reduction in the total wage bill and the output gain from reallocation will be proportionate to this total wage bill reduction. This wage reduction is a feature of our approximation, not a general feature of reallocation models. Still, this calculation suggests that the elimination of land use barriers would primarily redistribute from land owners to employers (and ultimately to customers).

Using the 2014 *County Business Patterns*, we can gauge the magnitudes of this quantity if we treat annual payroll per workers as synonymous with wage. We restricted our analysis to the 266 metropolitan areas with more than 50,000 workers. Assuming that α is low enough so that all areas maintain a positive population, equalizing wages would involve a total movement of $\frac{1}{\alpha}$ times 8 million workers, or about $\frac{1}{\alpha}$ times 8 percent of the employees in that sample. The largest gainer would be New York City (an extra 2.2 million times $\frac{1}{\alpha}$ workers). The overwhelming majority of cities would lose population, because they have current wages that are below the equalizing wage of \$49,000. Cities such as Orlando and Miami would lose particularly large numbers of workers, because they are large and relatively low wage. Since these areas may benefit from high amenities, this illustrates a shortcoming of our approach.

The total output gain would be $\frac{1}{\alpha}$ times 109 billion dollars, relative to a total payroll of 5.1 trillion in this sample. If we follow Hsieh and Moretti (2017) and assume that payroll is 65 percent of total output, this gain would represent $\frac{1}{\alpha}$ times 2.12 percent of total output. The obvious empirical necessity in this calculation is an estimate of α , the *inverse* elasticity of demand for labor.

The Cobb-Douglas formulation used by Hsieh and Moretti (2017) implies a value of $\frac{1}{\alpha}$ of 7.5. This produces in our calculations, as in theirs, a large misallocation effect. Our calculations suggests reallocation could increase total output by over 15 percent of G.D.P., but this would be reduced somewhat since some metropolitan areas would hit their lower bound of zero population.

Yet, the relatively large reaction of employment to wages implied by their Cobb-Douglas formulation is somewhat at odds with the empirical estimates of the link between wages and labor demand. For example, Beaudry, Green and Sand (2014) present city-level labor demand elasticities that seem matched to our needs. They find that a city-level labor elasticity of -0.3, which suggests that the overall impact is 0.7 percent of G.D.P. Their city-industry level estimates are larger (-1) and those would imply a misallocation cost equal to about 2 percent G.D.P. Past demand elasticities have typically ranged from -0.25 to -1.0, which suggests that two percent may be an upper bound on the gains from reallocation.

Labor demand elasticities are so important for these calculations because they determine how quickly an influx of labor into New York City would cause New York wages to fall to the national average. The Cobb-Douglas assumptions mean that an area with wages that are 50 percent above the national norm could see its employment increase 20 fold before wages fell to the national norm. In our formulation, if $\alpha = 1$, then a mere 1/3 increase in population will drop wages to the national norm.

We have nothing to add to discussions about labor demand elasticities at the local level. As 2 percent of G.D.P. is itself a large, we believe that these exercises illustrate that the benefits of reducing local land regulations may be sizable. If local labor demand is quite elastic, then Hsieh and Moretti (2017) may be right, and the output gains may be far larger.

Amenity differences and heterogeneity in building costs will tend to reduce this figure, but our calculations reflect only an estimate based on entirely static factors. It is quite possible that Silicon Valley is about creativity as well as high wages, and more Silicon Valley residents could also mean more technological innovation and faster productivity growth. Such hypotheses are quite speculative, but it is possible that the longer term costs of keeping people away from the most dynamic parts of the U.S. economy will prove higher than our short-term calculation.