

Chapter 6

Housing supply and housing affordability[☆]

Nathaniel Baum-Snow^a and Gilles Duranton^b

^a*Rotman School of Management, University of Toronto, Toronto, ON, Canada*, ^b*Wharton School, University of Pennsylvania, Philadelphia, PA, United States*

Contents

1 Introduction	354		
2 Facts	357		
2.1 Measuring housing prices and quantities	358	3.6 The productivity problem in construction	385
2.2 Housing prices	359		
2.3 Construction costs	362		
2.4 Accounting for changes in housing supply	363		
2.5 Housing quantities: new constructions and depreciation	364		
2.6 Changes in dwelling attributes	366		
2.7 Land use regulations and spatial patterns of housing development	368		
3 The economics of construction	370		
3.1 Modeling housing development	371	4 Beyond homogeneous housing services	390
3.2 Estimating housing production: traditional approaches	374	4.1 Housing durability	390
3.3 Estimating housing production: more recent work	378	4.2 Lumpy and differentiated housing	399
3.4 Estimating housing production: tall buildings	379	4.3 Filtering	406
3.5 Building height restrictions	383	4.4 Housing externalities	410
		5 Land use and regulatory constraints	415
		5.1 Bottom-up approaches and the estimation of local supply elasticities	416
		5.2 Unpacking land use regulations	424
		5.3 Estimating the consequences of land use regulation: border discontinuity approaches	427
		5.4 Urban growth boundaries	431
		5.5 Evidence from upzoning	432
		5.6 Property taxation and the political economy of land use regulation	434

☆ This work was supported by the Zell Lurie Center for Real Estate. It was in large part conducted while the second author visited CEMFI as part of the Maria de Maeztu Visitors Program. Their generous hospitality is gratefully acknowledged. We thank the two editors, Dave Donaldson and Steve Redding, our discussant Vernon Henderson, Tom Davidoff, David Green, Joe Gyourko, Albert Saiz, and seminar and conference participants for their helpful comments. We are also grateful to Minahil Fatima for her splendid research assistance. The views expressed here are those of the authors and not of any institution with which they may be associated. Replication data and code are available at <https://data.mendeley.com/datasets/kmps4d22p3/1> (DOI: 10.17632/kmps4d22p3.1).

5.7 Taxation and housing policy	438	7 Conclusions and future research	
6 Top-down approaches to housing supply	442	directions	447
		References	450

Abstract

The decline in housing affordability over recent decades has promoted an enhanced interest in housing supply. This chapter presents descriptive evidence about the evolution of US housing prices, quantities, and regulations since 1980, indicating that supply constraints appear to be increasingly binding. We then provide an overview of the various approaches used to model construction and land development for homogeneous and heterogeneous housing in static and dynamic contexts to understand housing supply. Our treatment incorporates empirical implementation and policy implications throughout. Finally, we provide an overview of quantitative evidence on the consequences of relaxing various types of supply constraints.

Keywords

Housing supply, Construction, Land use regulations

JEL classification

R31, R38

1 Introduction

Since 2000, housing costs have increased more rapidly than household incomes across the US and in many other countries around the world. Indeed, housing affordability is one of the great policy challenges of our time. A widely proposed remedy for improving housing affordability is to expand housing supply. Doing so effectively requires finding ways to reduce the costs of developing new housing and/or of maintaining existing housing. While there is little controversy about the link from reduced construction and maintenance costs to improved housing affordability, little is known about how to achieve this. Possible remedies include making more land available for development, reducing the costs of materials and labor, improving productivity in housing construction, and lowering regulatory limits. There is little consensus on which supply-side policies are most effective. Moreover, the targeting of different supply margins may have different distributional implications.

Developing a policy consensus requires an understanding of both some key facts and the various processes that may be generating the data underlying these facts. This chapter endeavors to provide both. We thus start with an overview of relevant facts about various housing markets in the US. We present information on housing prices and quantities for four broad market definitions. We also demonstrate the evolution of land use regulation and the density of housing development as a function of location within cities.

We then turn to various modeling approaches that have been used to rationalize these facts. We begin with a neoclassical treatment of housing construction,

in which developers combine land and capital into units of housing services. We demonstrate how this static framework can be adapted to estimate the production function for units of housing services conditional on construction and recover the associated housing supply elasticities at the intensive margin. We also consider the role of tall buildings in housing supply and the challenge of stagnant or declining productivity in the construction sector.

The static neoclassical framework is then extended in various directions, one at a time. We start with dynamic housing supply and real options, an active area of frontier research. We then turn to assignment models, in which new housing supply serves only select segments of the housing market in a commodity hierarchy. This leads to a natural process of newer units of higher quality inhabited by higher-income households filtering down the income distribution over time. The introduction of demand and housing unit heterogeneity naturally leads to a discussion of various externalities from new housing construction, including through the reduction of blight and the changing composition of neighborhood residents. Such externalities potentially justify roles for policy, among which we consider inclusionary zoning and spatially targeted development subsidies.

We then review land development constraints, including topography and land use restrictions. Increasingly restrictive land use regulations and declines in available land for development in good locations have resulted in sharply declining housing supply elasticities over time. The study of land use regulations faces the dual challenges of developing coherent and parsimonious measures and estimating their consequences in a well-identified way from contexts with strong external validity. This has led to a recent wave of research using boundary discontinuity empirical designs for the identification of impacts of various types of land use restrictions on property values while separating the capitalization of “own-lot” real option values, “external” density effects, and aggregate “supply” effects into property values. We highlight the role of quantitative models in facilitating the analysis of large-scale land use restrictions like urban growth boundaries.

Local jurisdictions typically determine land use regulations. In the US, these same jurisdictions also provide various local public services, including schools, parks, and police, levying property taxes to pay for them. We discuss how these joint decisions incentivize municipalities to enact exclusionary zoning, which limits negative fiscal externalities from owners of relatively low-value properties. Local jurisdictions use a variety of fiscal instruments beyond property taxation. We discuss the consequences of property transfer taxes, which disincentivize moving and are mostly capitalized into lower housing values. We then consider the consequences for the supply of housing of low-income housing policies, including tax credits for subsidized housing development, housing vouchers, and rent control.

Finally, we review general equilibrium modeling frameworks that articulate linkages across a system of housing markets and labor markets. Recent examples incorporate either within-region neighborhood and household heterogeneity or dynamics, urban growth, and endogenous land use regulation.

Such models have been quantified to isolate the motivations driving incumbent landowners to enact land use regulations and the welfare consequences associated with reducing the stringency of these regulations.

Central to these quantifications are three types of forces. The migration elasticity (alternatively viewed as the population supply elasticity) to each market is central for determining the housing cost and affordability consequences of reducing regulation. Heterogeneity in amenities across residential locations means that reducing regulation influences welfare by allowing more people to live in high-amenity locations. Finally, heterogeneity in productivity across work locations and the nature of agglomeration forces influence welfare through the determination of wages. Then, in some modeling frameworks, eliminating regulation also improves welfare by reducing commuting costs. On the other side, incumbent property owners have an incentive to regulate in order to maintain their property values since deregulation allows the amount of housing in the market to increase, thereby bidding down prices.

This chapter is certainly not the first attempt to summarize the literature on housing supply. Our overview reprises some of the ideas and results in Arnott's highly technical overview of the theory of housing markets in Volume 2 of this *Handbook* (Arnott, 1987) that have fallen out of recognition. We also reiterate some elements of Olsen's institutions and policy-focused overview of housing empirics in the same volume (Olsen, 1987). Twenty-five years ago, Denise Di-Pasquale (1999) lamented about the lack of research on housing supply. Her call went unheeded for many years, and it is only recently that research output on the topic increased, perhaps because of the growing housing affordability challenge. The literature on land use regulations prior to 2015 is masterfully discussed in the previous volume of this *Handbook* by Gyourko and Molloy (2015). We only repeat what is strictly necessary here and refer the reader back to this chapter for more details. We also de-emphasize issues surrounding the location of housing within cities and the earlier literature on housing durability, which are both reviewed by Duranton and Puga (2015). Finally, related discussions of the facts we present here appear in Molloy (2020) and Baum-Snow (2023).

As housing affordability is of such policy concern, we find it important to incorporate considerations of various policy instruments that may influence housing supply and affordability. We distribute such discussions throughout the chapter as they relate to the substantive issues under review. In Section 3, we consider the consequences of maximum floor area ratios, local building codes and permitting regimes, and unionization in the construction industry. In Section 4, we consider potential justifications for subsidized housing construction aimed at low-income tenants, inclusionary zoning, and coordinated subsidies for neighborhood economic development. In Section 5, we consider rent control, tax policy, and land use regulation, a topic to which we return in our review of quantitative general equilibrium approaches in Section 6. Our policy discussions are meant to illustrate the substantive issues we review. They are not meant to be exhaustive, nor do we aim to cover “housing policy” in any systematic way.

2 Facts

We begin by establishing a set of facts about housing affordability, housing costs, and housing quantities in the US since 1980, the first year with comprehensive microgeographic coverage in the Census. These facts motivate our review of the conceptual frameworks used to analyze housing supply and housing markets in the remaining sections. Undoubtedly, housing affordability has been declining. Rising construction costs and stricter land use regulations have resulted in lower rates of new construction, particularly in the most prosperous parts of the country. Geographical constraints and the increased scarcity of large tracts of undeveloped land in desirable locations also contribute to a weaker supply. While new constructions of larger units on smaller parcels are a crucial component of the supply of housing, it is not the only one. With fewer constructions, better maintenance and renovations are playing an increasingly important role in the supply of housing. We also highlight that despite weaker demographic growth, the demand for housing remains strong due to smaller households, greater incomes, and possibly more work from home and demand for home offices.

We present facts for four different types of locations. While these areas are probably too large for each to form a distinct housing market, our goal is to group locations with similar market conditions. We look separately at data for (i) small cities and rural areas, (ii) suburban areas, (iii) urban areas, and (iv) “superstar cities”.¹

To define these types of locations, we begin with the 2020 definition of Core Based Statistical Areas (CBSAs) and the 2000 definition of central cities for these CBSAs. “Small and Rural Counties” are all counties, either not part of a CBSA (2000 population: 18 million) or in the 846 CBSAs with a central city with fewer 100,000 inhabitants in 2000 (2000 population: 89 million). “Suburbs” are the regions of the 81 CBSAs with a central city of more than 100,000 that are outside of this central city (2000 population: 125 million). We assign the corresponding primary central city areas to “Central Cities” (2000 population: 49 million). Among Central Cities, we identify as “Superstar Cities” (Gyourko et al., 2013) those with at least a 200,000 dollar gap between price and construction cost at their urban fringe following Duranton and Puga (2023). These superstar cities are New York, San Francisco, Washington, Boston, Seattle, and San Diego (2000 population: 15 million).²

Of these four broad classes of locations, each has a unique profile of housing demand, built-up density, and regulatory environment. Small and Rural Counties have experienced the weakest demand growth, have the most land available

¹ Considering alternative geographies, including dividing the US into large regions, is also insightful. Due to length limits, we cannot look into other forms of heterogeneity here. However, our code and data are posted in this volume’s archive. Our code can be easily amended to consider alternative spatial groupings.

² In short: Small and Rural Counties, Suburbs, Central Cities partition the country, while Superstar Cities are a subset of Central Cities.

for development, and have the laxest regulatory environment. Suburban areas have faced the most robust demand growth and still have some land available for development, but also have portions with the most restrictive regulatory environments. Central cities generally have only moderate demand growth, but their high existing densities make new housing construction more costly. Finally, superstar cities have a combination of high demand growth, high construction costs, and strong restrictions on land development. It is not surprising that these are the locations with the most rapid increases in housing costs and unaffordability.

To put our facts in context, we first lay out a conceptual environment to ensure we understand our notations and measures of prices and quantities.

2.1 Measuring housing prices and quantities

Conceptually, we would like to measure housing services H_j , an index of observed and unobserved attributes. Housing services are what consumers consume and this index constitutes a broad and convenient measure of the quantity of housing supplied to a market j in each period. The flow of housing services from dwelling i in market j is h_{ij} , and the average flow of housing services per dwelling in market j is \bar{h}_j . Aggregating yields $H_j = \sum_{i \in j} h_{ij} = N_j \bar{h}_j$, where N_j is the number of dwellings in the market.³

Aggregating housing services at the dwelling level to the market level is only informative about the total supply of housing under stringent conditions. It requires that the unit price of housing services, P_j , be uniform across dwellings. Although dwellings may differ in size and quality, they must all offer housing services viewed as homogeneous among residents. As a result, dwellings vary only in the quantity of housing services they provide. Moreover, this aggregation should also be free from problems caused by the mainly indivisible nature of housing. Housing units may not be sensibly aggregated if they are too heterogeneous within a market area. Instead, it may be preferable to separately examine differentiated market segments for some questions (e.g., apartments versus single-family homes).

In practice, data limitations abound. We usually have records for sales prices and only a few characteristics of dwellings, including the number of bedrooms in most data sets and floorspace in some data sets. Then, the price of dwellings reflects the discounted sum of the value of future housing services, $P_{jt}h_{ijt}$, the product of the quantity of housing services and their unit price. Given these caveats, it is unclear how well the total number of dwellings or total floorspace in j captures the supply of housing H_j in this location, even after adjustment with a measure of observed quality.

³ We note that under some conditions, it is more convenient to express housing quantities as a vector of attributes rather than as a single index. One can view h_{ij} as the aggregation of such an underlying vector for dwelling i in market j . The weights used to aggregate up to the index of housing services would typically be estimated using hedonic analysis.

As the economic environment changes over time in j (or elsewhere), for example following increases in population or income on the demand side, prices of housing services P_j increase. In turn, higher prices lead to changes in the stock of housing. We can decompose changes in the housing stock into four sources: (i) new dwellings built on previously undeveloped land (“new developments”), (ii) new dwellings built on previously developed land (“redevelopment”), (iii) existing dwellings that fully depreciate or entirely disappear (“teardowns”), and (iv) existing dwellings that decay or are renovated (“renovations”). The first three forces are about changes in the number of dwelling units, and the last is about changes in the quality of dwelling units. Overall, the quantity of housing services in location j changes by ΔH_j . Understanding ‘housing supply’ is about understanding how ΔH_j responds to changes in economic conditions and the channels through which these adjustments occur. This section aims to provide some evidence about the evolution of housing prices and quantities without wedging ourselves to particular processes that generate the data.

2.2 Housing prices

Fig. 1 establishes our central motivating fact: housing affordability in the US has been declining since 2000 in all types of locations for both renters and home buyers, with some differences between types of location in the timing and magnitude of this decline. To show this, we present plots of average self-reported home values, gross rents, and household incomes in each location type for the 1980-2022 period. These quantities are calculated using household microdata from the 1980-2020 decennial censuses and the 2005-2022 annual American Community Surveys (ACS, Ruggles et al., 2024). Gross rents and home values are different measures of $P_j \bar{h}_j$, which additionally capitalize an expected future stream of (usually implicit) rents. These self-reported values are perhaps subject to biases from homeowner ignorance about the state of the housing market.

Each plot in Fig. 1 is indexed to 100 in Small and Rural Areas in 2000. We selected Small and Rural Counties as the reference location because this is where we see the lowest levels and growth rates of rents and home values, which, in turn, result from relatively weak demand growth and lax supply conditions. However, even in these areas, affordability declines after 2000. While average rents increased at an annualized rate of 5.0% between 2000 and 2022, nominal household incomes increased at an annualized rate of only 3.5%. The difference between these two figures is a rental “affordability gap” that increased on average by 1.5 percentage points per year (1.3 percentage points to 2018 before COVID). Despite much greater volatility, home values increased at an annual rate of 6.9% between 2000 and 2022, leading to a growth in affordability gaps for home values close to or greater than those for rents.

The story here is not just about lagging income growth. The growth of the consumer price index (CPI-U) in the economy was, on average, 3.2% per year, and the growth rate of income for rural households was similar to that of suburban incomes. Instead, the price of housing went up in Small and Rural Counties

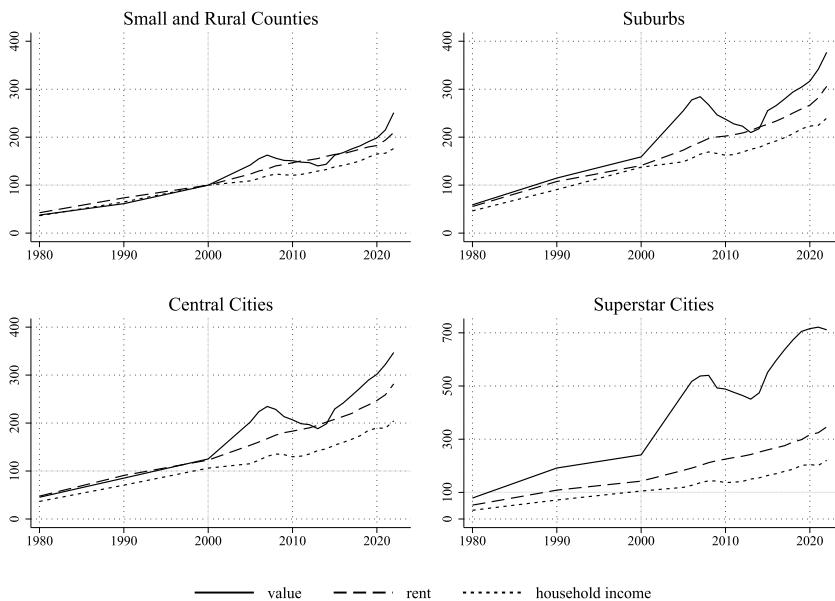


FIGURE 1 Housing affordability. *Notes:* This figure depicts descriptive evidence on self-reported nominal average log home values, log rents, and log household incomes for each indicated region over time. Plots are calculated using household level data from the 1980, 1990, and 2000 decennial censuses and the 2005-2022 ACS. County Groups in 1980 and Public Use Microdata Areas in other years are mapped to county and central city geographies using year 2000 population allocation factors. These allocation factors are constructed through spatial joins of 2000 census tract geographies. All values are indexed to 100 in Small and Rural Counties (top left panel) in 2000.

between 2000 and 2022. These patterns reverse the affordability increases that we observe for the 1980-2000 period when rental affordability gaps declined by two percentage points per year and value gaps declined by 0.5 percentage points per year in Small and Rural Counties.

These same qualitative patterns also exist in the other three types of locations that we study. In 2000, Suburbs started with household incomes, gross rents, and home values that were 37%, 41%, and 59% higher than in Small and Rural Counties. However, after 2000, these prices grew at very similar rates as in rural areas, except during the COVID-19 pandemic, when home values grew more rapidly in rural areas. The resulting growth in affordability gaps for Suburbs is very similar to that of Small and Rural Counties. However, the greater volatility in suburban home values is notable, perhaps reflecting lower housing supply elasticities (Davidoff, 2013). As incomes grew faster in suburbs during the 1980-2000 period, they experienced more rapidly declining home value-based affordability gaps than did rural areas in this period.

Central cities have larger shares of income spent on housing than do other regions, especially Superstar Cities. In 2000, incomes were 6% higher in central cities than in rural areas, but rents were 22% higher and self-reported home

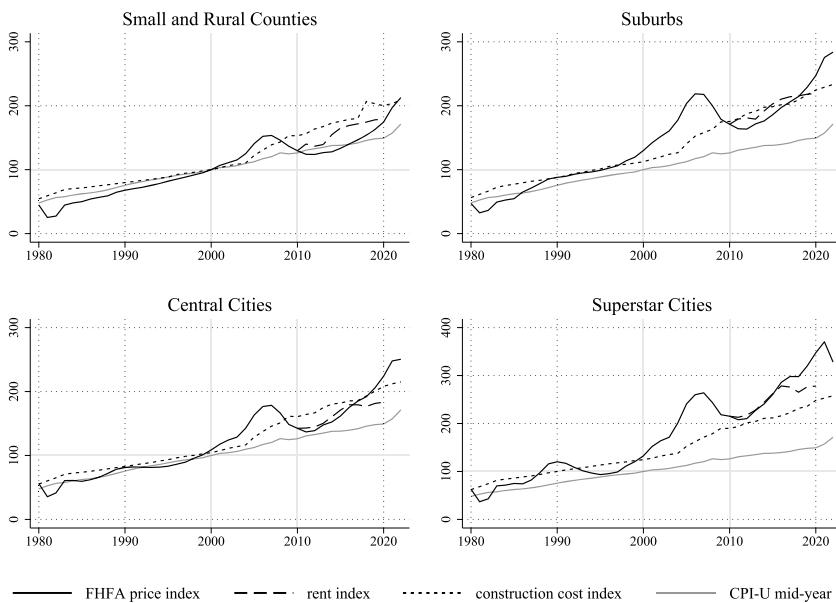


FIGURE 2 Indexed housing costs, homes prices, and rents. *Notes:* This figure depicts the average census tract level Federal Housing Finance Agency price index, zip code level Zillow Observed Rent Index, municipality level RSMeans Construction Cost Index for a typical single-family home, and national Consumer Price Index (CPI-U).

values 25% higher. In Superstar Cities, 2000 rents and home values were 42% and 141% higher than in rural areas. While the growth in the rental affordability gaps in Central Cities over 2000-2022 is similar to those in other regions, home value gaps increased faster, especially in Superstar Cities. In these locations, the affordability gap for home values increased by 5.3 percentage points per year from 2000 to 2018 (3.9 percentage points per year from 2000 to 2022). Like in other regions, this is a reversal of the rental affordability gains of the 1980-2000 period.

To confirm our findings from Fig. 1 regarding the ubiquitous increases in house values and rents, Fig. 2 shows the Federal Home Finance Agency home price index, which is a repeat sales index for single-family homes with prime mortgages, aggregated to census tracts.⁴ Since 1980, the evolution of this measure of indexed housing prices, arguably a measure of P_j , closely mirrors that of self-reported home values used in Fig. 1, a measure of $P_j \bar{h}_j$. Since self-reported home values could reflect an increase in dwelling quality, this similarity between the two implies that the decline in observed affordability is driven primarily by prices, P_j , and not by changes in the composition of the housing stock, \bar{h}_j .

⁴ For areas with the same coverage, the Federal Home Finance Agency and Case-Schiller house price indexes are very similar.

As a measure of indexed rents, with regrettably limited time coverage, we use the Zillow Observed Rent Index, a ZIP-Code level repeat listed rent index that started in 2010.⁵ The evolution of this index generally confirms our findings using self-reported rents in Fig. 1.

The affordability burden is generally growing less rapidly for renters. Rents consistently have lower growth rates and volatility than home values. With the US homeownership rate stable at 63%-68% since 1980, facilitating more renting may be one fruitful avenue to improve housing affordability. Forward-looking models, as in Molloy et al. (2022), can justify more rapid price than rent growth in supply-constrained regions like Superstar Cities, as rents are expected to grow only after housing supply constraints become even more binding.

2.3 Construction costs

To measure construction costs, we use the municipality-level RSMeans Construction Cost Index, which measures the cost of building a typical single-family home, excluding land and permitting costs. Hence, the differences between the Federal Home Finance Agency and RSMeans indices reflect, at least in part, land acquisition and regulatory costs. The national CPI-U is also graphed as a basis for comparison in Fig. 2.

Higher construction costs in Suburbs and Superstar Cities have existed since at least 1980. The main feature of Fig. 2 is that construction costs rose at about the same rate as the CPI-U until 2005 when they started growing about twice as fast in all four types of locations. Hence, the increase in construction costs is an essential part of the explanation of the patterns in Fig. 1 since the cost of structures represents 60% of US house values on average in 2019 (Davis et al., 2021).

To compare construction costs and house prices, we distinguish three periods. Between 1980 and 2000, construction costs increased at about the same pace as housing prices, or slightly less in Suburbs. As with affordability, we observe a first divergence between 2000 and 2018. The annualized growth in construction costs was 2.8 percentage points higher than the growth in housing prices in rural areas, 1.1 percentage points higher in suburbs, 0.4 percentage point higher in Central Cities, and 2.1 percentage points lower in Superstar Cities. However, these gaps closed or reversed after the onset of the COVID-19 pandemic following rapid growth in housing prices in all locations except Superstar Cities. Overall and acknowledging some complications associated with their cyclical behavior, housing prices have increased at about the same pace as construction costs in Small and Rural Counties, slightly faster than construction costs in Suburbs and Central Cities, and much faster in Superstar Cities.

Taken together, we find that land availability is a key source of supply constraints in the most expensive locations, such as Superstar Cities. In other locations, the increase in housing prices is driven chiefly by rising construction costs. This conclusion is, of course, consistent with the fact that in the most expensive locations, land represents a large share of the value of housing.

⁵ This index is constructed using almost the same methodology as Ambrose et al.'s (2015) repeat rent index calculated for a limited set of markets.

2.4 Accounting for changes in housing supply

Before exploring the evidence on quantities of housing supplied, we propose notations for the components of supply that we explore below. In the data, we observe measures of the stock of developed land in j , L_j , the stock of floorspace, S_j , and the number of dwellings, N_j . We focus on occupied dwellings so as to avoid overstating the housing stock to include units that are not habitable or in transition between occupants. Total housing is equal to the average quantity of housing services per unit of land $h_j \equiv \frac{H_j}{L_j}$ multiplied by the amount of developed land L_j . In turn, h_j can be decomposed further into housing services per unit of floorspace, $z_j \equiv \frac{H_j}{S_j}$ (typically unobserved) times average floorspace per dwelling $m_j \equiv \frac{S_j}{N_j}$ (typically observed). Also relevant is the average parcel size $l_j \equiv \frac{L_j}{N_j}$ (typically observed). Using these objects, we can decompose the housing stock in three convenient ways:

$$H_j = \frac{H_j}{L_j} \times L_j = \frac{H_j}{S_j} \times \frac{S_j}{N_j} \times N_j = \frac{H_j}{S_j} \times \frac{S_j}{N_j} \times \frac{N_j}{L_j} \times L_j \\ = h_j \times L_j \quad (1)$$

$$= z_j \times m_j \times N_j \quad (2)$$

$$= z_j \times m_j \times \frac{1}{l_j} \times L_j. \quad (3)$$

Eq. (1) distinguishes between the intensity of development per unit of land and the extensive margin of land development. Eq. (2) shows how dwelling units contribute to overall housing supply. Eq. (3) shows how land and parcel size contribute. Taking natural logarithms and time differencing equations (1)-(3) leads to

$$\Delta \log H_j = \Delta \log h_j + \Delta \log L_j \quad (4)$$

$$= \Delta \log z_j + \Delta \log m_j + \Delta \log N_j \quad (5)$$

$$= \Delta \log z_j + \Delta \log m_j - \Delta \log l_j + \Delta \log L_j. \quad (6)$$

Eqs. (6) states that the growth of the housing stock sums up the (mostly unobserved) growth in housing quality (leading to more units of housing services per unit of floorspace), the growth of dwelling size minus the growth in parcel size, and the growth of developed land. The last two terms decompose the total growth in dwellings shown in Eq. (5). Renovations and possibly new constructions, when they are of higher quality, increase housing quality z_j , while depreciation reduces it. Newly built housing and the redevelopment and renovation of existing dwellings affect floorspace per dwelling, m_j . New developments and redevelopments that impact the number of units in a parcel change the average size of the parcels. Finally, new developments also increase

the stock of developed land. Eqs. (5) and (6) inform our next three empirical exercises.

In Section 3, we will be interested in understanding both intensive and extensive margin factors driving the supply of housing services H_j .⁶ Eq. (4) shows that the stock of housing services grows both at the intensive margin, through the growth in housing quality and size, and at the extensive margin, through the growth of developed land. Then, it follows that

$$\frac{\Delta \log H_j}{\Delta \log P_j} = \frac{\Delta \log h_j}{\Delta \log P_j} + \frac{\Delta \log L_j}{\Delta \log P_j}. \quad (7)$$

The elasticity of housing services to housing price sums the corresponding elasticities of housing quality at the intensive margin and land development at the extensive margin.

2.5 Housing quantities: new constructions and depreciation

We begin our empirical exploration of housing quantities with Eq. (5) by looking at the evolution of occupied housing units N_j over time in different types of locations before turning to the evolution of other components of housing supply.

The stock of housing units changes due to the full depreciation of dwellings (teardowns) and new construction W_{jt} .

$$N_{jt+1} = N_{jt}(1 - \delta_{jt}) + W_{jt}, \quad (8)$$

which, after log differencing, can be written as,

$$\Delta \log N_{jt} \approx w_{jt} - \delta_{jt}. \quad (9)$$

That is, the growth rate in the number of dwellings, “units growth”, is approximately the construction rate $w_{jt} = \frac{W_{jt}}{N_{jt}}$ minus the dwelling units depreciation rate δ_{jt} . We emphasize that for the purpose of our descriptive analysis, this is an equilibrium dwelling unit depreciation rate that incorporates the possibility that endogenous maintenance can keep dwellings in the housing stock longer.

To measure trends in housing quantities, construction, and depreciation rates, we undertake a cohort analysis. We separately follow occupied dwellings that existed in the stock in 1980, 1990, 2000, or 2010.⁷ Fig. 3 reports the evolution of the stock of dwellings into each subsequent decade, which facilitates the joint visualization of changes in the overall stock due to depreciation and new

⁶ As H_j is unobserved, researchers often use floorspace S_j as a proxy. Even the stock of floorspace is empirically hard to obtain, as seen in Section 2. We often need to work with an even cruder proxy: the number of dwellings.

⁷ We exclude vacant dwellings from calculations. The national vacancy rates were 9% in 1980, 11% in 1990, 12% in 2000, 14% in 2010, and 11% in 2020, indicating a tightening of the housing market in the 2010-2020 period (US Census Bureau, Current Population Survey/Housing Vacancy Survey).

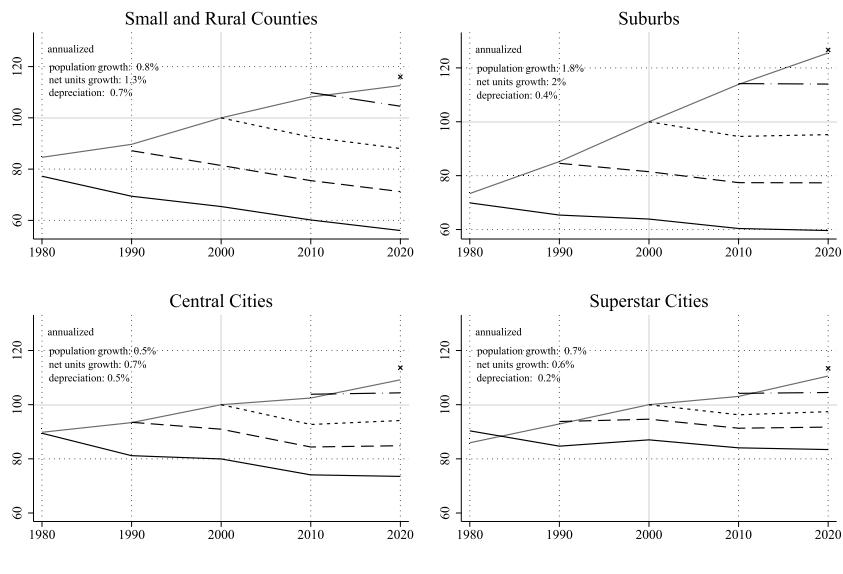


FIGURE 3 Housing quantities and depreciation rates. *Notes:* Downward sloping lines depict the number of housing units in the stock built prior to the indicated year that remain in the stock in the year on the horizontal axis. Stocks are measured using 100% count decadal census data. The upward sloping line is total population. All plots are indexed to the total housing stock or population in year 2000, which are assigned value 100.

construction. Our measure of depreciation is crude, as we can only measure it through tear-downs.

These data are constructed using decennial censuses and aggregate census tract data for 2008-12 and 2018-22 ACS, which are further aggregated into our four location types. To give a sense of demand conditions, we also report trends in aggregate populations. Therefore, by comparing trends in population and occupied housing units, trends in household sizes can also be inferred from the graph. The stocks of dwellings and populations are separately indexed to 100 in 2000 in each panel of Fig. 3 for the same four types of locations shown in Figs. 1 and 2.

During the 1980-2000 period, units growth, $\Delta \log N_{jt}$, equaled or exceeded population growth in all types of location, except Superstar Cities. In Small and Rural Counties, where land supply is the most elastic, the annualized unit growth of 1.3% exceeded population growth throughout the 1980-2020 period by 0.5 percentage points per year.

In Suburbs, which experienced much greater demand growth, units growth surpassed population growth by just 0.2 percentage points per year. Since 2000, these two growth rates have been nearly identical. We also note that these two growth rates were also markedly lower during the 2010-2020 period (see Baum-Snow, 2023, for more details). In Suburbs, as elsewhere, rising construction

costs and increasingly stringent land use regulations imply that the growth in the demand for housing is increasingly accommodated through rising home prices rather than increased quantities of dwellings.

Central Cities experienced weaker population growth that was also slightly outpaced by units growth by 0.2 percentage points per year, with a relatively high rate of units growth in 2010-2020 at 0.9% annualized, relative to 0.7% for the entire 1980-2020 period.

In Superstar Cities, population growth exceeded units growth during the 1980-2000 period.⁸ After 2000, units growth increased slightly in Superstar Cities to just offset population growth. Units growth has been declining in Suburbs and increasing in Central Cities in recent decades. However, these growth rates remain much lower in Central Cities, which also have a lower base stock of dwelling units. As a result, the vast majority of new dwellings are still built in suburban areas.

One way to cope with high demand and inelastic supply of new dwellings is to keep existing units in the housing stock longer (Baum-Snow and Han, 2024). In Small and Rural Counties, where construction is relatively easy, Fig. 3 shows approximately parallel downward-sloping lines for depreciation for all cohorts over all decades. In a typical year, 0.7% of the dwelling units depreciate out of the stock in these locations.⁹ Central cities also experienced relatively weak demand growth, manifested as average annualized population growth rates of only 0.5%. In this region, the depreciation rate was slightly lower at 0.5% per year, but with a marked slowdown in 2010-2020. Similar patterns of depreciation are seen in the suburbs. In Superstar Cities, where the pressure on the housing market is the greatest, the depreciation rates averaged only 0.2% per year. Relative to rural areas with a depreciation rate of 0.7% per year, Superstar Cities get an extra 0.5 percentage point in units growth from lower depreciation, about the same magnitude as units growth.

Together with the evidence on housing costs, Fig. 3 also strongly suggests that the decrease in average household size has put demand-side pressure on the housing market (Overman et al., 2008). The housing stock has grown at a higher rate than the population, yet housing prices and rents continue to rise above the rate of inflation.

2.6 Changes in dwelling attributes

Although there were no discernible differences between the evolution of dwelling values reported in Fig. 1 and that of the house price index reported in Fig. 2, suggesting little change in housing quality, the evolution of the composition of the housing stock deserves a more detailed look.

Fig. 4 plots the average number of rooms and bedrooms across dwelling units over time in each type of location. These two quantities are calculated

⁸ The ratio of the cost of replacement to the price of housing is at or below one for the majority of housing units in Small and Rural Counties in 2000 (Glaeser and Gyourko, 2018). When replacement costs are higher than house values, it is difficult to justify high volumes of new construction at market

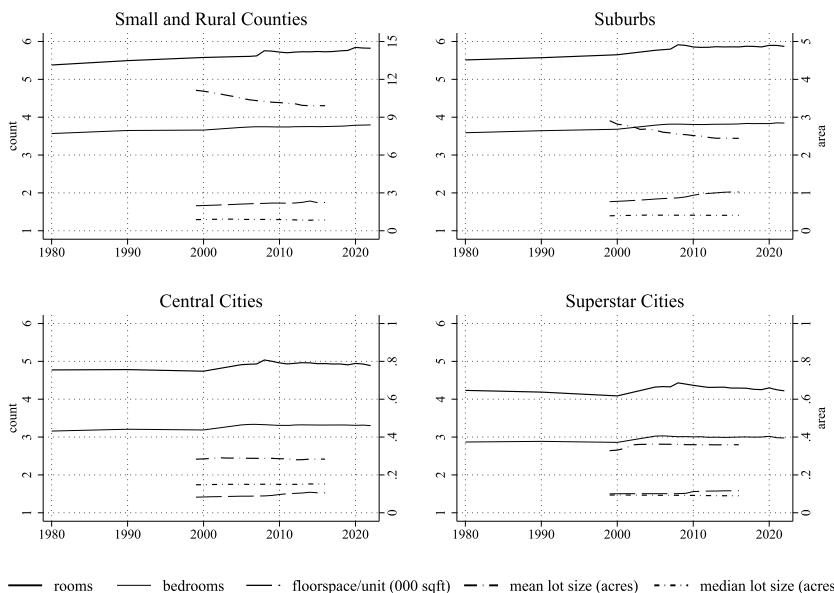


FIGURE 4 Attributes of housing units. *Notes:* Average rooms and bedrooms are calculated using decennial censuses and 2005–2022 ACS data. Mean floorspace per unit and lot size variables are calculated from census tract level data in the replication archive associated with Baum-Snow and Han (2024) (original source: Zillow ZTRAX). The sample includes all census tracts with a count of dwelling units in 2000 with information on floorspace and lot size that is within 10% of the number of occupied housing units reported in the 2000 census. This tract restriction retains 16%, 28%, 25%, and 21% of all tracts in Small and Rural Counties, Suburbs, Central Cities, and Superstar Cities, respectively. See Baum-Snow and Han (2024) for further discussion about sample constraints in the Zillow data. Median lot sizes are calculated as the median across tract average lot sizes in each year, weighted by the number of dwelling units in the tract.

from decennial censuses and the ACS following the same methodology as above. We view the average number of rooms and bedrooms as imperfect measures of housing services per dwelling, $h_j = m_j \times s_j$. We also report average square footage s_j for the 1999 to 2016 period only. This is computed from the Zillow data assembled by Baum-Snow and Han (2024), after restricting the sample to only include census tracts with year 2000 Zillow counts of dwellings that are within 10% of the 2000 census 100% count of occupied housing units and have information on floorspace and lot size.¹⁰

rates. Central Cities had an equally low ratio of replacement cost to price. Suburbs and Superstar Cities had much higher ratios.

⁹ Reported depreciation rates are calculated by averaging over all percentage reductions in the total initial stock across the four decades indicated in the figure, weighting by the stock at the beginning of the decade.

¹⁰ While this restriction has little effect on reported floorspace levels and trends, it improves the clarity of results about lot sizes.

Fig. 4 shows that in Small and Rural Counties and in Suburbs, where most new housing units have been built, the average number of rooms, bedrooms, and square footage have all increased over time. Despite demand pressures and rising housing prices, the quality of the typical housing unit has also increased slightly. This intensive margin component does contribute, albeit modestly, to the growth in overall housing supply.

A possible implication of larger housing units and smaller households is a possible mismatch between the two, as housing units are mostly indivisible. In particular, a “missing middle” of smaller single-family homes, townhouses, and larger apartments for middle-income families may be emerging as new homes are built disproportionately for the high-income segment of the market. We return to these issues below.

Fig. 4 also reports the evolution of mean and median lot sizes across properties in each type of location. We report both of these statistics to accommodate some imperfections in the lot size information in assessment data sets, including Zillow’s ZTRAX. In low-density areas, some of which are in our “Suburbs” region, many residential lots also include a large amount of non-residential land, including that in agriculture or completely undeveloped. For this reason, the lot size distribution for rural areas has a long right tail. In urban areas, multifamily buildings typically report the lot size of the entire building for each individual dwelling, greatly overstating average land per dwelling unit. This will artificially increase mean lot size in cities. As such, we also report medians of lot size distributions.

With these data limitations in mind, there remains useful information in the lot size distributions. While the size of housing units increased slightly, mean lot sizes l_j decreased rapidly between 1999 and 2016 in rural and suburban areas. This is from a mix of full abandonment of dwellings on very large rural lots and new housing developments on smaller, newly subdivided lots. This decline comes with the high net rates of new dwellings in these areas. Mean lot sizes in these areas remain nonetheless very large. In Small and Rural Counties, the average housing unit still sits on a 9.9-acre lot, down 1.2 acres since 1999, while the median unit sits on a slightly declining 0.9-acre lot. In Suburbs, the average lot is 2.4 acres, down almost 0.5 acres since 1999, but the median lot size is stable at 0.4 acres. Lot sizes in Central Cities are, of course, much smaller, with median lot sizes at 0.15 acres (and even smaller in Superstar Cities), with little discernible trend. The offsetting trends of bigger houses vs. smaller lots are consistent with the parallel evolutions of average dwelling prices and the house price index documented above.

2.7 Land use regulations and spatial patterns of housing development

To end our exploration of the data, Fig. 5 plots the density of housing units as a function of the distance to the central business district (CBD) across all

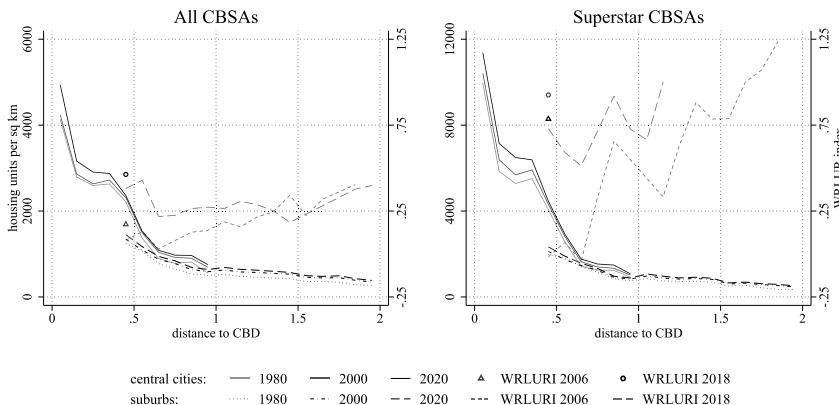


FIGURE 5 Density of housing units and land use regulations. *Notes:* Census tract level data from 1980, 2000, and 2020 is used to calculate average housing units per squared kilometer by distance to CBD separately for tracts in central city and suburban portions of CBSAs. Distance to CBD is indexed to be 1 for the central city census tract furthest from the CBD. Most CBSAs extend beyond index value 2, where the graphs cut off. Triangles and circles indicate the average Wharton Residential Land Use Regulation Index (WRLURI) for central city municipalities from 2006 and 2018, respectively. Remaining lines trace out averages in suburban municipalities as a function of indexed CBD distance.

CBSAs included in our Suburbs and Central Cities location types (left panel) and their Superstar Cities counterparts (right panel), again using tract-level data from decennial censuses and the ACS. The figure represents housing unit density plots for 1980, 2000, and 2020, separately in central cities and suburbs. The index on the horizontal axis scales the distance to the CBD to be between 0 and 1 for all census tracts within central cities. We cut the graph off at a distance index equal to 2, as the suburban density lines exhibit no discernible pattern of additional interest beyond this point.

Fig. 5 also represents the profiles of the Wharton Residential Land Use Regulation Index (WRLURI) of Gyourko et al. (2008) and Gyourko et al. (2021) for 2006 and 2018 on the right axis. This index, which aggregates many types of land use restrictions, covers most large municipalities and is discussed further in Section 5.2. These two vintage indices are independently standardized in each year, and are thus measured relative to a municipality with average regulation. Gyourko and Krimmel (2021) indicate that the average locale increased regulation between 2006 and 2018. We note that land use regulation tends to fall towards the periphery of CBSAs, beyond the right edge of the figure.

The declining rate of new housing construction in the suburbs seen in Fig. 3 is also evident in Fig. 5. The bottom line of the graph shows the density of suburban housing units in 1980. The next line up is its 2000 counterpart. The line above this is its 2020 counterpart, which is only slightly above the line for 2000, with a much bigger gap between 1980 and 2000. This reduction in the rate of growth of suburban density has coincided with increasingly stringent land use regulation in suburban jurisdictions, also represented on the graph (see also Gyourko and Krimmel, 2021). While lot sizes in suburban areas are declining, they

are still very large and provide ample opportunities for further densification. However, land use regulations impede densification in many suburban jurisdictions. One common type of regulation receiving a lot of attention in the literature is minimum lot size zoning, which explicitly limits housing unit densities, often with an exclusionary motive. We return to this impediment to construction later.

Since 2000, central cities have been more successful than suburban jurisdictions at facilitating densification. In the CBD distance range of 0.2-0.4, considerable densification occurred, especially over 2000-2020. As the amount of land in central cities is much smaller than in suburban areas, this increased city densification comes with fewer new dwelling units entering the aggregate city housing stock than entered the suburban housing stock through its small increases in unit density. We also note that the density of central cities exceeds that of suburbs at all distances in which they overlap, likely because this aspect of zoning rules in central cities is more lax, despite their overall increases in land use regulation, seen in Fig. 5 as the circles above the triangles.

Finally, we note that Superstar CBSAs are about twice as dense as all CBSAs and have stricter regulations. One result in the housing supply literature is that demand growth tends to beget more stringent land use regulations. This has been articulated in an overlapping generations model of urban development with agglomeration economies and endogenous regulation (Duranton and Puga, 2023) and empirically verified (Baum-Snow and Han, 2024). The reason is that regulations against densification make the housing supply less elastic. Then, demand growth is more strongly capitalized into housing values and thus benefits existing homeowners. Overall, there remains a lot of opportunity for increased densification. With housing unit densities roughly three times higher in central cities than in the next ring of suburbs, we see many well-functioning neighborhoods with high densities of dwelling units.

3 The economics of construction

In this section, we review the state of the literature on the production function for real estate. This treatment speaks primarily to the intensive margin of housing supply, the amount built conditional on development. We leave treatments of redevelopment and new developments for later sections.

Eq. (7) above states that the elasticity of housing services with respect to housing price sums the “building elasticity” $\epsilon_P^h(t)$ at the intensive margin with the “land development elasticity” $\epsilon_P^L(t)$ at the extensive margin. Generically, we can interpret both as “long-run” elasticities, as if each market area gets rebuilt from scratch in each long-run time period. These elasticities can be most naturally thought of in terms of cross-sectional comparisons of construction stocks and density of development across markets, given price differences. For example, if markets A and B both have $\epsilon_P^h(t) = 2$ and $\epsilon_P^L(t) = 0.5$ and $P_A = 1.1 P_B$, then each housing development in market A will provide 20% more housing services than in B, holding the amount of developed land fixed, and market A will

also have 5% more land developed than B . This results in 25% more housing services produced in market A than in B overall. We also want to make these elasticities a function of the initial time period t to capture the fact that changes in technology and regulation can result in long-run elasticities that change over time. For example, as market A is gradually built up, its development elasticity may fall from $\epsilon_p^L(t) = 0.5$ initially to zero when it is fully built up.

While long-run housing supply elasticities are important, we often wish to measure the total supply responses of the housing stock to changes in housing prices over a period of time, be it a year or a decade. Hence, the estimation of shorter-run housing elasticities will depend on both when they are estimated and over what length of time. Furthermore, the building elasticity, as measured over a specific time period, requires incorporating the redevelopment rate of the initial housing stock into the calculation. Markets with the same amount of redevelopment but higher initial stocks will mechanically have lower building elasticities.¹¹

Putting this all together, we have that the intensive margin elasticity as measured between periods t and t' equals the long-run elasticity for the period times the rate of redevelopment $R(t, t' | \mathbf{P})$. This rate is defined as the fraction of initial housing services that are redeveloped, which is typically the same as the fraction of initially developed land that is redeveloped. This redevelopment probability also depends on the price vector \mathbf{P} from period t into the future. Adding up,

$$\epsilon_p^H(t, t') = \epsilon_p^h(t, t') + \epsilon_p^L(t, t') = R(t, t' | \mathbf{P}) \times \epsilon_p^h(t) + \epsilon_p^L(t, t'). \quad (10)$$

In this expression, we maintain the short-run land development elasticity $\epsilon_p^L(t, t')$ as distinct from its long-run counterpart $\epsilon_p^L(t)$ to account for land development frictions that may take time to resolve or for the behavior of profit-maximizing builders waiting for the optimal time to redevelop.

We note some ambiguity in the notation of the elasticities ϵ_p^h and ϵ_p^L in Eq. (10) for two reasons. First, arguably, the entire sequence of current and future housing prices should matter for development decisions. We return to these issues in Section 4, where we take a fully dynamic approach. Second, in general, these elasticities additionally depend on market j , or more precisely, on market-specific factors that shape land availability. We return to these issues in Section 5. We study the building elasticity in the rest of this section on construction.

3.1 Modeling housing development

We follow Muth's (1969 and 1975) pioneering work and much of the literature, including Combes et al. (2021) and Baum-Snow and Han (2024), from whom

¹¹ The vintage of the initial building stock also affects the building elasticity since redevelopment typically results in larger per-dwelling upgrades for older properties (Brueckner and Rosenthal, 2009). We return to this below.

this exposition borrows. We view housing development as the result of a profit maximization problem by a competitive builder facing free entry. A housing production function, or its dual cost function, is at the core of this approach. This housing production function turns construction inputs, capital, labor, and land into housing. For simplicity, we lump together construction material (capital) and labor (which gets frozen into capital by the construction process) and only retain capital and land as factors of production. Housing is then produced according to the function $h(k, l)$, which we assume increases in both arguments and is strictly quasi-concave.

The profit of the builder of dwelling i is,

$$\pi_{ij} = P_j h(k, l_i) - r_j k - q_j(l_i) - c_{ij}, \quad (11)$$

where r_j is the price of capital in j , $q_j(l_i)$ is the price of a parcel of size l_i , and c_{ij} is a fixed cost of housing development for the parcel over which dwelling i sits. This fixed cost captures the physical cost of preparing a parcel for construction and all the regulatory taxes and fees needed to obtain a building permit (Glaeser and Gyourko, 2018; Duranton and Puga, 2023).

Housing development occurs in three stages. First, parcels are delineated. Second, the builder decides whether to pay the fixed cost c_{ij} to develop or redevelop a parcel. Third, the builder chooses how much capital to invest. We take the size of parcels as given for now.¹² In practice, parcel size is often predetermined (except for large subdivisions). The assembly and reorganization of the land are notoriously challenging (Brooks and Lutz, 2016). The size of the parcels is also highly regulated, as discussed in Section 5.¹³

Starting with the choice of capital investment, the builder's first-order condition for profit maximization is:

$$P_j \frac{\partial h(k, l_i)}{\partial k} = r_j. \quad (12)$$

¹² Builders in our setting resemble neoclassical firms, but some differences are worth noting. Parcel size (the land input) is typically assumed to be fixed from the perspective of individual builders. This leads to decreasing returns to scale and a naturally competitive market for real estate development, where the rents are dissipated into land prices. Because parcel size is exogenously given and the fact that builders only hold land parcels temporarily to develop and then sell housing, the role of the land input is quite different from that of the capital input in the neoclassical theory of the firm.

¹³ There are further complications associated with endogenous parcel size. When builders can choose the size of parcels in addition to capital to maximize profit, we can appeal to standard results from producer theory. If housing is produced with decreasing returns to scale, optimal lot size results from a trade-off between the decreasing returns to housing production on each parcel and its fixed cost of development. As we approach constant returns for all factors, the optimal size of parcels grows to infinity. Since, in practice, the production of single-family homes appears to be at or close to constant returns, this outcome seems implausible. A first resolution to this apparent contradiction is to appeal to the indivisible nature of housing. A huge house on a huge parcel is not a perfect substitute for ten smaller houses built on the same parcel. However, considering this feature would greatly complicate our modeling framework, and we reserve further discussion of it for later. A second possibility is that the returns to land are first increasing (it is hard to build on a tiny lot) and then decreasing (the tenth acre of a property likely provides little housing services). In a competitive equilibrium, builders operate at the point of constant returns.

This implicit equation uniquely determines the optimal capital investment k_i given the concavity of $h(\cdot, \cdot)$. Then, free entry implies that the profit from construction is dissipated into the price of the parcel,

$$q_{ij} \equiv q_j(l_i) = P_j h(k_i, l_i) - r_j k_i - c_{ij} = (1 - \alpha_i) P_j h(k_i, l_i) - c_{ij}, \quad (13)$$

where $\alpha_i \equiv \frac{\partial \log h(k_i, l_i)}{\partial \log k_i}$ is the elasticity of housing production with respect to capital. The last equality is obtained after substituting for r_j using the first-order condition (12) and simplifying.

Some comments are in order. First, a limitation of this approach is that it only considers the production of housing (services), and it is silent about how constructions are divided across different dwellings.

Second, because the price of parcels is generally bounded from below by their next best use $q_j l_i$, land development occurs only for parcels with fixed costs below a threshold $c_{ij} \leq \bar{c}_{ij} \equiv (1 - \alpha_i) P_j h(k_i, l_i) - q_j l_i$. More parcels are then developed as housing prices increase or the production technology improves. The distribution of fixed development costs, including regulatory costs, is central to determining housing development patterns at the extensive margin. At the intensive margin, applying the implicit function theorem to Eq. (12) also implies that higher prices lead to more capital investment, $\frac{\partial k_i}{\partial P_j} > 0$, and thus more intensive housing development $\frac{\partial h(k_i, l_i)}{\partial P_j} > 0$.

Third, although it is convenient to model land use restrictions as part of the fixed cost of building a dwelling, the regulation of housing development may have implications beyond this fixed cost. For instance, regulations often limit capital investment or certain forms of capital investment, such as building height or the maximum footprint per unit of land. Such constraints affect optimal capital investment in Eq. (12) and lead to more complex distortions when different forms of capital investment are available (Combes et al., 2021).

As a side note, it is also instructive, and as shown below in some cases useful, to consider the dual of the producer problem above and the variable cost function $C_j(h_i)$ for parcel i in location j . With fixed parcel size, the first-order condition for profit maximization and the chain rule imply that the marginal cost is given by $\frac{dC_j(h)}{dh} = r_j / \frac{\partial h(k, l_i)}{\partial k}$. In turn, after inserting this expression into the first-order condition (12) and using the resulting expression to substitute for P_j into the zero profit condition, we find that the price of parcels in Eq. (13) can equivalently be written as,

$$q_{ij} = C_j(h_i) \left(\frac{d \log C(h_i)}{d \log h_i} - 1 \right) - c_{ij}. \quad (14)$$

After accounting for the fixed development cost c_{ij} , the fraction of total variable construction cost $C_j(h_i)$ capitalized into the price of land parcels is increasing in the convexity of the cost function, which itself is increasing in the intensity of land in production given the exogenous parcel size l_i . Relative to Eq. (13), Eq. (14) is more general, as it accommodates multiple variable factors of production.

3.2 Estimating housing production: traditional approaches

The simplest production function for housing to consider is the following Cobb-Douglas aggregation of land and capital,

$$h = B_j k^\alpha l^\beta, \quad (15)$$

where B_j is the total factor productivity of builders in location j .

Before going deeper into estimating this production function and to understand what is at stake with this estimation, note that with exogenous parcel sizes and this production function, following Baum-Snow and Han (2024), the housing supply function of parcel i can be written as

$$h_i = B_j^{\frac{1}{1-\alpha}} (\alpha P_j / r)^{\frac{\alpha}{1-\alpha}} l_i^{\frac{\beta}{1-\alpha}}. \quad (16)$$

In turn, this expression can be aggregated across developed parcels within market j to get

$$H_j^s = B_j^{\frac{1}{1-\alpha}} \left(\frac{\alpha P_j}{r} \right)^{\frac{\alpha}{1-\alpha}} \sum_{i \in j} l_i^{\frac{\beta}{1-\alpha}}. \quad (17)$$

Hence, conditional on building, the elasticity of housing production relative to housing prices is $\epsilon_P^h \equiv \frac{\partial \log h}{\partial \log P} = \frac{\alpha}{1-\alpha}$ at both the parcel and the market levels. This is also true for the overall economy. This technological parameter estimated at the “micro” level is thus indicative of a “macro” elasticity of interest.

After taking logs in Eq. (15), we note that the production function for housing could potentially be estimated by regressing log housing on log capital and log land: $\log h_{ij} = \alpha \log k_i + \beta \log l_i + \log B_j$. Unfortunately, this regression cannot be estimated directly since units of housing services are not observed. Units of capital are also often not observed separately from their price. Instead, we can rely on the first-order condition for profit maximization above and estimate the following regression:

$$\log (P_j h_{ij}) = \alpha \log (r_j k_i) + \beta \log q_{ij} + (\log B_j + \epsilon_i), \quad (18)$$

where we add an idiosyncratic error term ϵ_i , perhaps reflecting some mismeasurement of the dependent variable since a perfect fit is unlikely to occur. This error term could also reflect some unobserved heterogeneity in construction costs across parcels. This regression is naturally estimated using a cross-section of houses. It can also be estimated for a cross-section of cities or neighborhoods.¹⁴

¹⁴ Alternatively, the cost share parameters can also be calibrated as part of a larger model as in Albouy (2009) who finds a land share of 0.23, a capital share of 0.15, and a labor share of 0.62. The main pitfall of this type of exercise is that it relies on a model that makes many more assumptions, including functional form assumptions, than the simple framework proposed above requires. In addition, such calibrations often rely on available aggregate figures that may not coincide with the concepts used in the model. See also Muth (1975) for further discussion.

Estimating the regression described by Eq. (18) is challenging regardless of the spatial scale. Productivity B_j is not observed and is treated as a residual. This generates a well-understood simultaneity problem. When the idiosyncratic error term ϵ_i reflects parcel-specific construction costs, a builder who observes this term will adjust their choice of input accordingly. A related issue is that efficiency units of housing cannot be observed separately from their price in the dependent variable. Thus, any demand shock affecting P_j is expected to be reflected in the choice of inputs. See Ackerberg et al. (2007) and Syverson (2011) for further discussion of this general problem that plagues the estimation of nearly all production functions.

Land values are also difficult to observe. Separate transactions for residential land parcels are rarely observed. There are two exceptions. Thorsnes (1997) observes land values for large undeveloped parcels purchased by home builders. These parcels are then subdivided and developed before being sold to incoming residents. Combes et al. (2021) use data from France, where a standard arrangement between home builders and buyers of new homes is for the future resident to buy land before contracting with the builder to construct a house paid in several installments. This arrangement is fiscally advantageous for the buyer who minimizes taxes on real estate transaction, while gradual payment helps the builder finance construction.

Land values can also be computed using land market transaction data from providers such as the CoStar Group. Because only transactions above a certain threshold are recorded, these transactions are spatially selected toward the urban fringe where large land parcels are available or toward trophy locations where land is the most valuable. The transactions may also select a specific type of house. Large parcels are generally subdivided before construction. New constructions in large new subdivisions are only a fraction of all constructions, perhaps a small one in some countries (Combes et al., 2021). Finally, these transactions are rare, raising sampling issues in small cities. To limit these problems, Haughwout et al. (2008), Albouy et al. (2018), and Davis et al. (2021) implement regression approaches to correct for the characteristics of the parcels, most notably for their location.¹⁵ Using the resulting predicted land value as an explanatory variable to estimate a regression such as Eq. (18) like Albouy and Ehrlich (2018) is potentially problematic because the implicit smoothing of the data can introduce classical measurement error in the regression and lead to coefficients biased toward zero.

¹⁵ There are important differences between these approaches. Albouy et al. (2018) estimate a ‘monocentric cone’ for each city with two key parameters: the price at the center and a distance gradient. Davis et al. (2021) rely instead on a ‘Kriging’ procedure where the prediction relies on the available information from the nearest neighbors. The former approach is perhaps more appropriate for city-level predictions, while the latter is likely to perform better for smaller spatial units such as ZIP Codes. It is also possible to predict the price of land as the difference between the value of a dwelling and the estimated cost of reconstructing the existing structure. Although this is a reasonable approach to generate aggregate land values (Davis and Palumbo, 2008), it would be problematic to introduce it into the regression derived from Eq. (18) as it would bias the estimation through a mechanical correlation between the value of a house and the value of its land.

Following Clapp (1979) and many others, the rest of the literature relies on appraised land values. Appraised land values are likely noisy and subject to various biases. For example, appraisers may compute land values by taking the difference between their estimated value of dwellings and the replacement cost of the structures before applying a discount to avoid disputes with the owners since these appraised values often form the property tax base.

A major shortcoming of Eq. (18) is that it imposes a unit elasticity of substitution between land and capital. This need not be the case, and the economic cost of regulations that push toward either high or low density depends crucially on the value of this elasticity. Rather than attempting to estimate a Cobb-Douglas functional form, much of the literature has focused instead on constant-elasticity-of-substitution (CES) production functions:

$$h = B_j \left(\nu k^{(\sigma-1)/\sigma} + (1 - \nu) l^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \quad (19)$$

where σ is the elasticity of substitution between land and capital, ν is a factor share parameter, and B_j remains a productivity shifter. Applying the CES functional form to the first-order condition for profit maximization in Eq. (12) and substituting the zero-profit conditions yields:

$$\log \left(\frac{k_i}{l_i} \right) = \sigma \log \left(\frac{\nu}{1 - \nu} \right) + \sigma \log \left(\frac{q_i}{l_i r_j} \right), \quad (20)$$

after taking logs and simplifying. Assuming a constant cost of capital ($r_j = r$) and some measurement error on the dependent variable, the previous equation leads to $\log \left(\frac{r k_i}{l_i} \right) = C + \sigma \log \frac{q_i}{l_i} + \epsilon_i$, where C is a constellation of parameters.

This regression of the log of the capital-to-land ratio on the unit price of land has been estimated many times in the literature, with a review by McDonald as early as 1981. Ahlfeldt and McMillen (2013) and Combes et al. (2021) provide more recent estimates. The range for the estimates of the elasticity of substitution between land and other inputs, when estimating some version of Eq. (20), is extremely wide, from less than 0.1 in Muth (1964) to slightly above one in McDonald (1979), with a mean perhaps around 0.5.

Sample heterogeneity is a first possible explanation for these differences in the estimated elasticity of substitution between land and capital. First, residential and commercial properties will likely differ in how they are built. There is also significant heterogeneity within each broad class: single-family houses differ greatly from large multifamily buildings within residential real estate; office towers require construction techniques different from big-box retail stores within commercial real estate, etc. Construction technologies also likely differ between countries (Yoshida, 2016). But even when they focus on US residential samples where single-family homes are likely to dominate, Ahlfeldt and McMillen (2013) estimate an elasticity of substitution between land and capital

of 0.95 in Allegheny County PA with one data set and of 0.43 and 0.60 for two other data sets in Chicago.

These differences may also be symptomatic of data measurement issues for the price of land. In the older literature, Thorsnes (1997), who used arguably better quality data than previously, estimates elasticities around 0.9, toward the upper end of the range. In Monte-Carlo simulations using synthetic data, Ahlfeldt and McMillen (2013) show that the estimation of Eq. (19) is sensitive to adding noise to the data. Even a moderate amount of mismeasurement for the price of land leads them to recover estimates of less than half the elasticity of substitution they used to generate the data. Combes et al. (2021) also show that when they estimate Eq. (19) after smoothing their land price data, the estimated elasticity of substitution about doubles from 0.49 to 0.99.

Another possibility is that even accurate land price data may be ‘overly noisy’. Parcel prices may reflect a range of phenomena extraneous to how Eq. (19) is derived. Variations in the price of land may occur because some sellers may be in a hurry to sell, because some buyers and sellers lie about the actual amount of the transaction to evade taxes, or because buyers and sellers of land parcels in different locations may hold different expectations about future rents, etc. Put differently, the data-generating process may be richer than the simple model used here (and in most of the literature). As a result, Eq. (19) may be misspecified, and naive estimates of σ in the regression corresponding to Eq. (19) will be biased downward.

Unobserved construction costs at the parcel level are also likely to bias the estimated elasticity of substitution downward. More challenging parcels for builders require additional capital investments and, as a result, face lower land prices to retain zero profit. More generally, using land prices as an explanatory variable in Eq. (20) is problematic given that the zero-profit condition (11) treats the price of the parcels as an outcome of the building process.

Another potential difficulty is that a CES functional form may not capture the complexities of the production of housing. Albouy and Ehrlich (2018) propose estimating an even more flexible translog cost function. They regress US metropolitan construction prices on construction costs, land prices, various terms of higher order imposed by the translog specification, and geographical and regulatory constraints measures. They instrument land prices and construction costs with local amenities and estimated productivity for tradable goods under the assumptions that (i) they affect the demand for housing and (ii) are otherwise uncorrelated with unobserved supply characteristics. While the first condition about the power of the instruments can be tested, the exclusion restriction may be more questionable. For example, natural amenities such as hills that make a location more attractive to residents could be correlated with unobserved housing productivity. However, the main limitation of the exercise is the size of the sample of metropolitan areas relative to the number of parameters to be estimated. This lack of power leads to imprecise estimates for the elasticity of substitution computed by Albouy and Ehrlich (2018).¹⁶

¹⁶ The coefficients on the role of land use regulation and geographic limits are much more precisely estimated. We return to them below.

3.3 Estimating housing production: more recent work

Following Epple et al. (2010) and Combes et al. (2021), more recent work has focused on developing new ways to separate prices from quantities to avoid biases arising from unobserved demand shocks when estimating the production function of housing.

We follow the exposition of Combes et al. (2021) and build on the builder's profit maximization approach derived above.¹⁷ The key is to note that the first-order condition in Eq. (12) and the free-entry condition in Eq. (13) both contain the price of housing and its quantity. Ignoring the fixed cost of development for now and assuming the same cost of capital everywhere, it is possible to insert the zero-profit condition (13) into the first-order condition (12) to eliminate the price of housing P_j . This yields the following partial differential equation:

$$\frac{\partial \log h_i}{\partial \log k_i} = \frac{r k_i}{r k_i + q(k_i, l_i)}, \quad (21)$$

after simplification. Eq. (21) shows that the elasticity of housing with respect to capital investment for a dwelling is equal to its share of capital in the total construction cost. All these quantities are measured at the private optimum of the builder, and we note the price of the land $q(k_i, l_i)$ to reflect this. In French construction costs data, the price of land and capital investment are observed separately. Hence, the cost share on the right-hand side of Eq. (21) is readily measured. For any parcel of size l_i , we can then compute the quantity of housing offered by dwelling i by integrating Eq. (21) over k :

$$\log h(k_i, l_i) = \int_{\underline{k}}^{k_i} \frac{r k}{r k + q(k_i, l_i)} d \log k + \log Z(l_i) \quad (22)$$

where $Z(l_i)$ is an integration constant for parcels of size l_i , and \underline{k} is the lowest level of capital investment observed in the data. In Eq. (22), housing production is calculated as the sum of marginal products. After a log transformation, the term to integrate is a simple cost share that is observed in the data. Then, after recovering h non-parametrically, it is possible to estimate how k maps into h . This is a partial identification of the housing production function.¹⁸

By construction, this approach removes any simultaneity between capital and housing production, working through the price of housing. However, some

¹⁷ The approach of Epple et al. (2010) is essentially the dual of that of Combes et al. (2021) and, empirically, it relies on different but equivalent observables. Combes et al. (2021) provide a detailed comparison of the two approaches within a unified framework.

¹⁸ Full identification is possible after duplicating the same approach to the first-order condition for land. Due to the zero-profit assumption, constant returns in housing production need to be assumed. Combes et al. (2021) econometrically reject constant returns in housing production, albeit in a way that is not economically meaningful. More importantly, in their data, most new constructions appear to be infills with severe constraints on dividing parcels. Taking parcels as given is arguably more appropriate than assuming some form of optimization in their delineation.

identification concerns remain. The most important of these regards factor heterogeneity, in particular the heterogeneity of land parcels.

The estimation of the quantity of housing in Eq. (22) is based on the notion that builders want to develop more expensive land with greater capital intensity since more expensive land reflects higher housing prices. Although Combes et al. (2019) provide evidence of a strong link between the price of housing and the price of land, land parcels also differ in how costly they are to develop due to, for example, their slope, how easy they are to excavate, any landscaping needed to avoid stormwater runoffs, etc. All else being equal, more difficult terrain will result in greater capital expenditure to build a given dwelling and a lower parcel price to satisfy the zero-profit condition.

To address this problem, Combes et al. (2021) use residualized values for k , q , and l , where these quantities are predicted only from variables that affect the demand for housing, but, arguably, not the unobserved development costs of parcels.¹⁹ These demand predictors of parcel prices include the city's population size and the distance to the center, conditional on other local characteristics that may reflect unobserved supply conditions.

The main findings of Combes et al. (2021) are an elasticity of housing production with respect to capital of 0.65 with little variation across house sizes. Consistent with this stability, they estimate an elasticity of substitution between capital and land close to one. Although they econometrically reject a Cobb-Douglas functional form (and more flexible CES and translog forms), Cobb-Douglas provides a good first approximation. The main heterogeneity in their results is a lower capital elasticity in larger cities (0.54 in Paris vs. 0.71 in the smallest cities). After discarding several possible explanations, Combes et al. (2021) conclude that these differences are caused by expectations of stronger rent growth in larger cities combined with the opportunity to redevelop housing in the future.

Importantly, as shown above, the elasticity of housing production relative to housing prices is $\epsilon_P^h = \frac{\alpha}{1-\alpha}$ with a Cobb-Douglas production function. Hence, an elasticity of housing production with respect to capital $\alpha \approx 0.65$ implies a supply elasticity at the intensive margin of $\epsilon_P^h = \frac{0.65}{1-0.65} \approx 1.86$ or around two.

3.4 Estimating housing production: tall buildings

The literature on housing production mainly focuses on recovering the mix of land and capital chosen by profit-maximizing developers to build single-family houses. This framework works well to understand the intensive margin of housing supply in recently developed residential neighborhoods in the US and, to some extent, Europe. It is also appropriate for the redevelopment of single-family houses. However, it is less obvious that it carries over seamlessly to tall buildings.

¹⁹ This last condition is an exclusion restriction. This approach to identification is in the same spirit as standard instrumental variables, but adapted to the specific case of Eq. (22).

The cumulative height of buildings taller than 55 meters in cities across the world is above 16,000 kilometers. This is equivalent to 42,000 times the Empire State Building. About one third of this stock is in the developing world, and about 90% of it was built after 1980. Tall buildings make up over 10% of total building volumes in cities with a population larger than one million (Ahlfeldt et al., 2023).

Thus, any analysis of aggregate housing supply in large cities requires a conceptual framework that can accommodate tall structures. To model tall buildings, it is helpful to focus directly on the choice of building height g by builders since data on building heights and volumes are readily available for the entire world on a consistent basis (Esch et al., 2024). This focus on building height is also useful when analyzing the consequences of ubiquitous land use regulations such as height limits and maximum floor-to-area ratios (FARS). While we consider the implications of FAR restrictions in the following sub-section, we delay a broader discussion of land use regulation to Section 5.

Since the height of a building is roughly proportional to floorspace and housing services per land area, the ideas developed in Section 3.2 above carry over to the analysis of tall buildings. For example, with Cobb-Douglas production for parcel i in location j and a common fixed parcel size of l_i as in Eq. (15), the variable cost function for housing services is,

$$C(h) = r_j \left[h / \left(B_j l_i^\beta \right) \right]^{1/\alpha}. \quad (23)$$

Then, the elasticity of the unit cost of housing, $C(h)/h$, with respect to the amount of housing built is $\theta \equiv \frac{1-\alpha}{\alpha}$, the inverse of the supply elasticity given zero profit. With height given by $g \equiv h/l$, the unit cost elasticity with respect to h , θ , is also the elasticity of unit costs with respect to height. $\theta > 0$ captures the idea that structural engineering, material, elevator, and building foundation costs are collectively convex in building height. Relative to that on single-family houses, the tall buildings literature thinks about making the same model more flexible off this baseline by allowing the cost elasticity of height to vary across locations and height, and in considering the consequences of regulations that limit building heights.

Cost data indicates that, while in general, costs are convex in height, there are regions of local increasing returns. Using data from RSMeans for apartment buildings in 50 US cities, Eriksen and Orlando (2022) trace the cost function for height. They find a declining marginal cost before discontinuous cost increases at 4 and 8 stories in buildings up to 12 stories tall. In addition, the marginal cost declines monotonically with floorplate size conditional on building height. For very tall buildings, the engineering evidence is that the “cost of height” elasticity θ is increasing in height.

Ahlfeldt and McMillen (2018) is the first study to explicitly estimate θ using building level data on construction costs and heights. In the spirit of the Cobb-Douglas production function in Eq. (15), θ can be estimated by regressing the

log ratio of capital to floorspace on the log of building height. Identification requires variation in heights that are driven by demand-side forces rather than by unobserved location-specific cost differences. For example, one may worry that building heights reflect in part unobserved regulatory burdens and thus cannot be taken as unrelated to developers' profit-maximizing choices of inputs and output. Measured construction costs per floorspace k/h could be lower in less regulated places that also have greater heights, thereby biasing estimates of the cost of height θ downwards.

Ahlfeldt and McMillen (2018) use microdata from Chicago on tall buildings from *Emporis* (now integrated into CoStar products) to measure building height g , construction cost k , and real estate services h (proxied by floorspace) for most tall buildings in Chicago over its history. Distance to the central business district is the demand side shifter used to generate variation in building height. This instrumentation strategy exploits insights from classical land use theory (Alonso, 1964; Muth, 1969; Mills, 1969) and its central result that better accessibility to the CBD gets capitalized into higher land values and more capital-intensive construction through demand side forces. Ultimately, Ahlfeldt and McMillen (2018) finds estimates of θ of 0.54 for commercial buildings and 0.61 for residential buildings of 20 floors. These values imply estimates of α for tall buildings of about 0.63, very similar to the capital share of 0.65 found in Combes et al. (2021) for single-family homes.²⁰

An essential feature of the construction of tall buildings is that the cost of building higher varies between locations mainly due to the depth of the bedrock, as discussed in length in Barr (2016) in the context of New York City. Bedrock that is too close to the surface requires costly blasting to establish foundations for tall buildings. Bedrock that is too deep requires costly engineering for the foundations to be stable.²¹ These features of the construction technology suggest a U-shaped relationship between θ and the depth of the bedrock.

Exploiting this insight, Ahlfeldt et al. (2023) specify a variable cost function for height per unit of land of the form

$$C_i(g) = v_i g^{1+\theta_i}. \quad (24)$$

This cost function is derived from a Cobb-Douglas production function in land and capital with exogenous parcel sizes as in Eq. (15). It assumes constant returns and that the ratio of floorspace to building height is constant but allows

²⁰ More broadly, Ahlfeldt and McMillen (2018) explore a CES technology in capital and land, in which, in addition to θ , they propose a parameter that defines the elasticity of the ratio of lot size to building footprint with respect to height: $\frac{d \log(l/f)}{d \log g} \equiv \lambda$, which they estimate to be 0.16. This can be viewed as a source of over-identification, as together with the definition of θ , a Cobb-Douglas technology in land and capital ($\sigma = 1$) is implied. Then, the elasticity of building height with respect to the unit price of land is $\frac{1}{1+\theta-\lambda}$.

²¹ Foundation stability is achieved through deeply bored piles that go all the way down to the bedrock, a "raft" floating in subsoil built below deep piles, or many small, very deeply bored piles that do not reach the bedrock.

the share of capital to vary between parcels. The specification of Eq. (24) is built to focus on how the cost elasticity of height, θ_i , varies with the depth of the bedrock. After taking logs for a chosen height g_i , this equation implies $\log(C_i/g_i) = \theta_i \log g_i + \log v_i$, which can be estimated as a regression of log cost of building per unit of floorspace on log building height interacted flexibly with bedrock depth. In turn, θ_i can be recovered as a function of bedrock depth.

Using data from *Emporis* for about 1,000 individual buildings taller than 55 meters, Ahlfeldt et al. (2023) estimate instrumental variables locally weighted regressions, smoothing across bedrock depths.²² As with Ahlfeldt and McMillen (2018), distance to the CBD instruments for building height. The resulting estimates of θ range from 0.2 at about a 17-meter bedrock depth to 0.8 at depths of both 0 and 30 meters.²³ These cost elasticities imply intensive margin supply elasticities ϵ_P^h of between 1.25 and 5.

Profit maximization by tall building developers to choose optimal heights incorporates the fact that both commercial and residential rents rise with height off the ground. Danton and Himbert (2018) and Ben-Shahar et al. (2022) estimate elasticities of rent with respect to height of about 0.03 for residential buildings while Liu et al. (2018) and Koster et al. (2013) estimate about 0.07 for office buildings. Developers' choices of optimal building heights depend on the interaction between the local cost elasticity of height, θ_i , and local real estate demand or idiosyncratic construction costs v_i . Lower cost elasticities of height imply taller building height choices in locations with stronger demand conditions and/or lower idiosyncratic construction costs.

For a large cross-section of world cities, Ahlfeldt et al. (2023) estimate the relationship between their aggregate stock of tall building and their population and built-up footprint. These estimates reflect how differences across cities in the cost of height θ affect building height. Differences in supply conditions, which drive the cost of height, generate variation in the stock of tall buildings, and, in turn, the stock of tall buildings affects the built areas of cities and their ability to accommodate residents. Ahlfeldt et al. (2023) find that the effects of bedrock depth (or, equivalently, θ) on heights are of larger magnitudes in bigger cities, where demand for floorspace is stronger. In turn, this variation can be used over the 1975-2015 period to infer an elasticity of average city population with respect to the stock of tall building heights of 0.13 across cities in the developing world. The analogous elasticity for built area is -0.16 .²⁴

²² A few earlier papers also use the idea that favorable bedrock depths are cost shifters for building heights and density. Rosenthal and Strange (2008) and Combes et al. (2010) use the depth of the bedrock as a source of plausibly exogenous variation in urban density to estimate the strength of local and city-level agglomeration forces.

²³ Most of the Emporis data with construction cost information is for steel construction buildings in the US. However, most of the new tall buildings built after 2010 are framed in concrete. As these buildings are heavier, concrete construction buildings typically require deeper foundations than steel-framed buildings of the same height. As a result, the optimal depth of the bedrock is probably a few meters deeper for most tall buildings built in recent years.

²⁴ Similar estimates are found by combining simple comparisons of 1975-2015 changes in log population, built area, and aggregate heights among cities over 1 million on favorable relative to unfavorable bedrock depths.

Ahlfeldt et al. (2023) use these reduced-form results and inferred city-specific gaps between actual and unconstrained heights to calibrate a monocentric land use model for each city worldwide. Model counterfactuals show that expanding real estate supply through relaxing existing building height limits improves renter welfare, largely by facilitating more compact cities, thereby reducing intra-city travel costs.

3.5 Building height restrictions

The result in Ahlfeldt et al. (2023) that tall buildings lead to more compact cities echoes earlier results in the small literature looking at the costs of restricting building height. In the context of a simple monocentric city, Bertaud and Brueckner (2005) were the first to formally prove that restrictions on the height of buildings lead to urban sprawl and higher commuting costs. Brueckner et al. (2017) also demonstrate the empirical importance of mechanism in the welfare costs of floor area ratio (FAR) restrictions in Indian cities.

The costs of building height restrictions can also be inferred from the difference between the price of floorspace and its marginal cost of construction. To see the logic, note that unconstrained developers build until the price of one unit of housing services is equal to its marginal cost: $P_j = C'(h)$. If no more than $h^{max} < h$ units of housing services can be built, the height constraint is binding and $P_j > C'(h^{max})$. Therefore, $P_j/C'(h^{max}) - 1 > 0$ can be viewed as reflecting the “regulatory tax” rate associated with the constraint h^{max} .

In early work, Glaeser et al. (2005a) use hedonic analysis to compare the price per square foot of real estate to its construction cost and estimate the associated regulatory tax, driven by height limits and other restrictions. For Manhattan, they find that the price of floorspace is about twice its marginal cost. Cheshire and Hilber (2008) find even larger numbers for office markets in British cities.²⁵ Ben-Moshe and Genesove (2024) use a related frontier approach to trace out the marginal cost curve for height absent regulation for apartment buildings in Israel. This marginal cost curve is the lower envelope in sales in the price-height space. Comparing the distribution of all sales prices to these marginal costs, this paper finds a mean regulatory tax of 48% of the price.

The costs of building height restrictions can also be inferred from the capitalization of floorspace into land values. Following Brueckner et al. (2017), a constraint on FAR means that the profit-maximizing amount of capital employed is essentially capped at k^{max} given an exogenous parcel size, l_i . The marginal capitalization associated with an additional unit of capital after partially relaxing a binding FAR constraint is thus $P_j \frac{\partial h(k^{max}, l_i)}{\partial k} - r_j > 0$, where r_j is the cost of capital in market j . Importantly, the relaxation of FAR restrictions gets capitalized into land values with an elasticity that increases with the stringency of these

²⁵ The standard explanation for markups of price over marginal cost is market power. However, given that prices in most housing markets are determined by many small bilateral transactions, it seems unlikely that developers could exploit market power to that extent.

restrictions, reflecting diminishing marginal returns to capital in the production of housing. The capitalization of less stringent FAR restrictions is also increasing in the land share, an observation also made by Tan et al. (2020) in the context of a richer production technology that additionally allows for a quantity-quality trade-off in inputs producing housing services.

This theory suggests that regressing log land value on log FAR reveals the stringency of the FAR restrictions. All else being equal, if relaxing the FAR results in large increases in land values, this means that developers value the extra density and expect to build taller to exploit it. Using the procedure described above and data on land sales in Chinese cities, Brueckner et al. (2017) find that, on average, FAR restrictions result in building heights that are 62% of the free market level, with wide variation across cities and even within Beijing.²⁶ Evidence in Cai et al. (2017) that building above the legal FAR limits is common in high-demand areas (albeit coming at additional costs to developers) corroborates this evidence of binding FAR restrictions in most Chinese cities. Peng (2023) reaches a similar conclusion for New York City.

These empirical exercises face identification challenges. One can imagine unobservables that drive both the stringency of FAR restrictions and land values. For example, locations with higher real estate demand may be more likely to enact tighter FAR restrictions to maintain the character of the neighborhood. In this example, cost and regulation-driven estimates of $\frac{\partial \log q}{\partial \log \text{FAR}}$ would be understated by a simple OLS regression. To address this challenge, Brueckner and Singh (2020) and Tan et al. (2020) include fixed effects for small neighborhoods, with the idea that demand conditions are approximately constant within them. Cai et al. (2017) uses model structure and data on observable demand shifters to hold demand conditions constant.

An alternative approach to estimating the stringency of FAR regulations is to calculate the amount of bunching at the FAR limit \bar{F} . Assuming a smooth distribution of unregulated building heights around \bar{F} , Brueckner et al. (2024) uses a bunching estimator to recover the counterfactual distribution of building heights absent FAR restrictions in New York City. Assuming a Cobb-Douglas production technology with a capital share of 0.65, marginal costs around the associated kink in the cost function can be recovered. Brueckner et al. (2024) find that New York City statutory FAR restrictions are overcome at an additional marginal cost of floorspace of 11% for both $\bar{F} = 2$ and $\bar{F} = 0.5$. Eliminating all FAR restrictions would result in 10% more floorspace in the New York City housing stock affected by both $\bar{F} = 2$ and $\bar{F} = 0.5$, and 3-4% for intermediate \bar{F} , given recent demand conditions.

In conclusion, the tall building component of housing supply can be reasonably analyzed using the same production technology as that of single-family homes, with a capital share of 0.65. Evidence in the literature using such a framework indicates that relaxing restrictions on building tall remains a promising vehicle through which to accommodate rising housing demand in the world's largest cities.

²⁶ Brueckner and Singh (2020) carry out an analogous analysis using data from five large US cities and find that, among them, Chicago is the least height constrained.

3.6 The productivity problem in construction

We now turn to understanding how productivity in the construction industry changes over time. To start with an example, the Empire State Building was the tallest building in New York City in 2012 when it was surpassed by One World Trade Center.²⁷ The Empire State Building, which replaced the Waldorf-Astoria Hotel after its closure in May 1929, saw construction start on March 17, 1930. It opened 410 days later, on May 1, 1931, at a cost of 681 million dollars (2024 dollars). One World Trade Center, which is admittedly 100 meters taller and contains about 50% more floorspace than the Empire State Building, replaced the original World Trade Center after it was destroyed in the terrorist attacks of September 11, 2001. Construction for One World Trade Center began in April 2006. The building opened after 3,112 days on November 3, 2014, at a cost of 5.19 billion dollars. Put differently, building One World Trade Center took nearly eight times as long and cost about five times as much per square meter of floor area as building the Empire State Building.

Although perhaps extreme, this contrast between the Empire State Building and One World Trade Center is not an isolated case. Perhaps unsurprisingly to the casual observer, construction sites today look similar to those in 1950. Starting with Stokes (1981), Allen (1985), and more recently Sveikauskas et al. (2016), Garcia and Molloy (2023), Goolsbee and Syverson (2023), and D’Amico et al. (2024), economists have been puzzled by the poor apparent productivity performance of the construction industry measured through a wide range of metrics.

As summarized by Goolsbee and Syverson (2023) using data from the Bureau of Economic Analysis, an index of value added per worker between 1950 and 2020 increased by 290% in the US economy but decreased slightly in the construction industry. This mild decline for construction incorporates a strong increase between 1950 and 1970 followed by a staggering 40% decline since 1970. Although we only discuss evidence from the US, Goolsbee and Syverson (2023) show that other developed countries had a similar experience, though perhaps less extreme.²⁸ This poor productivity performance of the construction industry with a decrease in the value added per worker over time is not due to lack of investment or to deterioration in labor quality (Goolsbee and Syverson, 2023). Instead, it appears to be driven by total factor productivity.

Computing changes in labor productivity in an industry producing heterogeneous and durable goods tied to land is difficult. There are several measurement reasons why an observed decrease in productivity might be spurious. First, to compute value-added in construction, the value of land must be taken out. Although the price of land has increased dramatically, its share in construction

²⁷ Sources for this paragraph are https://en.wikipedia.org/wiki/Empire_State_Building, https://en.wikipedia.org/wiki/One_World_Trade_Center, and <https://data.bls.gov/cgi-bin/cpicalc.pl> for inflation adjustments, all consulted on August 8, 2024.

²⁸ Brooks and Liscow (2023) and Mehrotra et al. (2024) provide a diagnostic for the costs of providing infrastructure in the US that echoes many of the points we make here for housing.

appears stable, as noted by Garcia and Molloy (2023). Second, Garcia and Molloy (2023), Goolsbee and Syverson (2023), and D'Amico et al. (2024) provide convincing evidence that this decline in productivity is not an artifact of failing to account for faster unobserved quality growth in the construction industry's output relative to other industries since 1970. Similarly, there is no evidence of increased output mismeasurement, again relative to other industries since 1970. Although profit margins in the construction industry have increased, this increase has been modest, ruling out another source of possible bias in calculating productivity growth.

To avoid many of the complications associated with comparing aggregate values over time, Goolsbee and Syverson (2023) use information on the number of physical units produced and the labor employed to produce them to directly measure the physical productivity of labor in the industry over time. In the same spirit, Garcia and Molloy (2023) and D'Amico et al. (2024) use longitudinal information about the cost of building houses from RSMeans, a data provider and consulting firm that collects comprehensive cost data by task (or set of tasks, including whole houses) and geography. Although physical productivity does not decline as sharply as nominal productivity, the construction industry remains a clear laggard. Overall, the dismal productivity performance of the construction industry is, for the most part, something real and not a statistical illusion. Stokes (1981) already reached a similar conclusion many years ago.

Due to the size of the construction industry, its weak productivity performance has sizeable aggregate implications. Goolsbee and Syverson (2023) calculate that if construction labor productivity had grown by a modest 1% a year instead of declining by 1% a year since 1970, aggregate labor productivity (and to a large extent income per capita) would be about 10% higher today. Taking into account the linkages between industries in a general equilibrium framework, Foerster et al. (2022) find that changes in the productivity of the construction industry between 1950 and 1980 account for a full point of lower annual GDP growth between these dates.

Given this, it is crucial to understand the deeper causes of the poor productivity performance of the construction industry. Before diving deeper into this issue, looking at the variation in the cost of building in the US cross-section provides a useful first pass. Gyourko and Saiz (2006) exploit the richness of the RSMeans data by looking at the cost of construction per square foot for a standard 2,000 square-feet "economy quality" house across 140 US markets (cities). They first highlight substantial variation in construction costs between cities. In 2003, building in New York was 80% more expensive than building in the North Carolina markets of Raleigh, Charlotte, and Greensboro. Even when outliers are ignored, differences remain large, with an interquartile range of about 20% of the unit cost. When Gyourko and Saiz (2006) regress unit construction costs on the number of permits and controls, they estimate a tiny elasticity. All else being equal, places that build more do not experience higher building costs. This is consistent with a flat supply curve for housing and a very fragmented

industry. Gyourko and Saiz (2006) report that the number of establishments in the construction industry is nearly proportional to population size, with one construction establishment for every 1,200 residents. Unsurprisingly, they also estimate a strong association between construction costs and local wages.

Interestingly, Gyourko and Saiz (2006) also find a strong association between construction costs and the share of unionized construction workers. An interquartile range in unionization (22 percentage points) is associated with 11% higher construction costs, or more than half the interquartile variation in unit construction costs. Finally, they uncover a statistically significant association with some measures of construction regulations.

There is little doubt that, through building codes, city governments and construction unions promote inefficient and expensive work practices, mandate questionable improvements, and restrict cost savings. Building codes throughout the US also increase construction costs because they impose high compliance costs, including excessive fees, lead to administrative delays in construction, and entail extensive and expensive inspections. On the other hand, building codes are arguably needed, given the asymmetries of information between house-builders and residents and the risks such asymmetries entail for residents and their neighbors. Assessing the costs and benefits of building codes is thus inherently challenging. In their reviews of the literature, Listokin and Hattis (2005) and Dumm et al. (2011) conclude that excessive and inappropriate building codes raise housing costs by, perhaps, around 5%.

While we note that more up-to-date and better-identified evidence on the effects of building codes is needed, these effects are apparently smaller than those uncovered by Gyourko and Saiz (2006). These findings need not be inconsistent as Gyourko and Saiz (2006) measure an overall effect associated with unionization, of which inefficient building codes may be just one of the manifestations.²⁹ In any case, even the effects estimated by Gyourko and Saiz (2006) are somewhat modest. While remedying the excesses of building codes and high labor costs caused by unionization would be worthwhile, such small effects do not explain much of the poor productivity performance of the construction industry.

Perhaps more crucially, building codes may prevent innovations in building technologies from diffusing or even happening in the first place. Schmitz (2020) provides a detailed historical narrative that blames construction unions acting in conjunction with various levels of government for the failure of housing to “industrialize” with more housing being produced off-site. Although they are undoubtedly an obstacle, institutional blockages in the US are only one element among the production and standardization challenges faced by modular housing.³⁰ The industrialization of housing is also slow to develop in countries such

²⁹ While construction unions may push for make-work, a more benign interpretation may deserve consideration. Increasingly stringent construction codes, requiring, for instance, additional levels of insulation or resistance to fire for structures, increase construction costs. In a world of lower interest, more demanding specifications face lower user costs and might be economically justified. Mehrotra et al. (2024) propose a defense of this argument to explain the inflation of the cost of the development of new road infrastructure.

³⁰ For example, incompatible standards leave builders liable to being held up by their suppliers. In turn, this possible hold-up makes the funding of modular construction projects harder and more expensive.

as Sweden, which are welcoming of the technology. Sweden even actively supports modular housing because of long winters that impede construction and the availability of cheap timber (Bayliss and Bergin, 2020). Another limiting factor of the potential for modular housing is that the legacy stock of houses and buildings will likely remain mainly out of reach for such new technologies for many years. While a complete discussion of the mechanization and automation of construction is beyond our scope, we note that new construction technologies and materials also hold great promises for decarbonization, as they provide better insulation and rely less on concrete than standard construction technologies (see Fernandes Rocha et al., 2022, for a broad introduction).

Even if the importance of unions and city governments in preventing the emergence of radically new forms of construction is debatable, building codes prevent the diffusion of incremental innovations. In their classic work, Oster and Quigley (1977) show that four such incremental innovations regarding the intensity of material usage, the use of cheaper materials, and material pre-assembly were less likely to be permitted in jurisdictions where unionization rates were higher, where the chief planner was less educated, and where construction firms were smaller. Consistent with these micro findings, Goolsbee and Syverson (2023) note that, while the construction industry invests slightly more than the rest of the economy, it invests extremely little in intellectual property capital, which includes research and development and software. Productivity is more likely to be embedded into this kind of capital than in standard equipment and structure. At the very least, this lack of investment in intellectual property is a symptom of a broader malaise.

In a different explanation, D'Amico et al. (2024) highlight the effects of an interaction between societal change and the nature of construction projects. Unlike manufacturing facilities, construction sites are highly visible and immobile. A permit is also needed before construction can start. These features made construction sites prime targets of increased "citizen voice" (Brooks and Liscow, 2023; Glaeser et al., 2005b). Then, compliance and approval for a new construction are more manageable for smaller, highly customized, and less locally disruptive construction projects.³¹ In turn, these projects are more likely to be developed by small construction firms. As in most industries, smaller firms are less productive and invest less in productivity growth.

To support their argument, D'Amico et al. (2024) show that the construction industry remains remarkably fragmented relative to manufacturing. Although construction has experienced some consolidation since the Great Recession (Quintero, 2023), it pales in comparison to the consolidation of retail or even multifamily rentals (Calder-Wang and Kim, 2024). D'Amico et al. (2024) also show that the size of residential projects has declined, that small firms work disproportionately on small projects, that small firms are less productive in the

³¹ For large buildings in the office market, another possible strategy is instead to propose trophy buildings designed by famous architects, as underscored by Cheshire and Dericks (2020) in the case of Central London.

construction industry just like in most other sectors, and that cities with stricter land-use regulations have smaller construction firms. Interestingly, infrastructure construction has followed a similar path. Brooks and Liscow (2023) provide complementary evidence for this industry, which is consistent with what we observe for residential real estate.

Several further conjectures can be envisioned to explain the weak productivity performance of the construction industry. First, rather than institutional factors, it may have been available land that became scarce and limited new constructions. The evidence in Duranton and Puga (2023) suggests that this is undoubtedly the case for a number of coastal metropolitan areas. More systematic evidence on land scarcity, including the lack of large parcels to develop new subdivisions, is greatly needed. Second, an increasing share of residential construction work is about maintaining legacy houses, expanding them in some cases, or tearing them down and rebuilding them. Although increasingly required, maintenance and redevelopment are likely to appear as less productive in the data. In this case, it is the observed productivity growth in the construction industry prior to 1970 when construction work took place mainly at the extensive margin that may have been spurious, or at least unsustainable. Like the previous one, this conjecture remains to be verified of course. A further wrinkle to this argument is that lower-productivity maintenance is made more important by land use restrictions that prevent new constructions. Third, new constructions today are increasingly at the upper end of the market (Wang, 2022). Higher quality may be difficult to measure and may contribute to an apparent productivity decline.³²

At this stage, we draw the following conclusions about the economics of construction. The literature has made progress in our understanding of how factors in construction, land and capital (itself produced from labor and material), are combined to produce single-family homes and, in more recent research developments, tall buildings. There is an emerging consensus regarding how to frame the issues and some of the results.

The key problem for construction is the slow growth of the technology used to assemble factors, not the price and quality of labor and materials. Productivity in construction may be about a third of what it would be if technology in construction had followed the same improvements as manufacturing technologies since 1950. We note here an important disconnect between a strand of research

³² We keep in mind that productivity can be either an aggregate, like value added per worker as used by Goolsbee and Syverson (2023), or the cost of building a specified house using data from RS Means data as in several contributions listed above. The evolution of the former productivity measure is much worse than that of the latter. Our conjecture about maintenance may explain part of this gap. Depending on the exact (but unknown) methodology of RS Means, an increased share of maintenance jobs in construction may also be able to contribute (or not) to the evolution of the RS Means series. For instance, for a given task such as installing shingles on a roof, RS Means may average its costs calculations over all such jobs even if doing this is much cheaper on a new house than as a replacement of old shingles on an existing house.

interested in identifying construction production functions (and thus supply responses for floorspace) earlier in this section and the more aggregate approaches used to study changes in construction productivity.

Several explanations have been proposed to explain this dismal productivity performance. They include (i) the cumulative effects of restrictions imposed by unions and regulators, (ii) adverse effects of NIMBYism pushing toward low-productivity projects and low-productivity construction firms, and (iii) more ‘natural’ causes arising from the need to maintain and repair an aging stock of housing and other mundane explanations of the same type mentioned above. Despite its appeal, the explanatory power of the first of these explanations is unclear. There are clear productivity inefficiencies caused by building codes ‘excesses’ and building height restrictions. However, these inefficiencies appear moderate, if not modest, in light of the poor productivity performance of the construction industry. Results in the literature suggest that they each add perhaps around 10% or less to construction costs. Relaxing these restrictions is a worthwhile endeavor, be it only because it will increase the quantity of housing produced, even if it does not lead to dramatic productivity changes, at least in the short run. Turning to the second explanation, NIMBYism will also, of course, affect the volume of construction activity, which we explore in Section 5. The third type of explanation is still largely unexplored.

4 Beyond homogeneous housing services

The main limitations of the framework used in the previous section are that housing is: (i) durable, (ii) lumpy and differentiated, and (iii) subject to local spillovers. Here, we review the literature by considering each of these features. As these features can interact in interesting ways, we also consider the limited body of work that combines them.

4.1 Housing durability

Building housing is an irreversible decision. Housing also takes time to develop and only slowly depreciates over time.³³ As a result, the development problem is not only about solving how much to build, as explored above, but also about deciding whether and when to build, when to rebuild, and, before rebuilding, how much to invest in maintenance to slow down depreciation. These decisions depend on market conditions today and far into the future. In turn, the aggregation of these decisions will affect current and future market conditions. This future is uncertain and requires forming expectations about it. For any given

³³ We take these features as self-evident. Evidence for the irreversibility of housing is provided by Burchfield et al. (2006) for the US between the mid-1970s and the early 1990s and by Combes et al. (2024) for France between circa 1860 and 2020. They all show that developed land rarely returns to an undeveloped state. Evidence of slow depreciation is reported above in Fig. 3. Construction delays will be familiar to all readers and do not require further evidence here.

construction project, we also expect development, maintenance, and redevelopment decisions to be driven by unobserved idiosyncratic factors.

To further motivate the importance of housing durability, we return to the headline figures for the elasticity of housing supply of Section 3 with values of approximately two for single-family homes and between 1.2 and five for tall buildings. These are long-run elasticities that indicate how much floorspace is built after the decision to develop or redevelop is made. As indicated by Eq. (10), the long-run elasticity is multiplied by the probability of development to obtain the ‘short-run’ elasticity. In housing and construction, the long run is genuinely long. Hence, the ‘short-run’ elasticity, often measured for periods of up to 10 or 20 years, is usually the quantity that is relevant to policy. As we show below in Section 5.1, the estimated housing supply elasticities for, say, 10 years are often nearly an order of magnitude smaller than the long-run elasticities. The probability of development ultimately drives the magnitudes of supply elasticities that hold over the time intervals relevant to policy. Loosely speaking, the housing supply elasticities we care about are driven by the probability of development or redevelopment much more than by the long-run elasticities derived from the construction technology.

A comprehensive model of dynamic aggregate housing supply is still well beyond our reach, given the difficulty of modeling decisions in rich, dynamic contexts. Partial approaches have nevertheless provided fundamental insights regarding (i) the development and redevelopment of specific parcels and (ii) the broader implications of the decisions to develop or redevelop for urban expansion and densification. We note that the earlier literature on housing durability is covered by Duranton and Puga (2015). Another class of spatial dynamic models is expertly reviewed in this *volume* by Desmet and Parro (2025). The models they consider are more aggregated, sometimes covering the entire world, with a richer set of interactions through trade and migration, for instance.

The first major insight from thinking dynamically about housing supply is the existence of an option value of development or redevelopment, leading to delays. The existence of real options associated with irreversible investments has been known for a long time and was discussed in the context of real estate development by Titman (1985). To understand this core issue, consider a risk-neutral builder facing an irreversible development opportunity over two periods on a parcel of unit size, with the same Cobb-Douglas housing production function as in Eq. (15): $h = k^\alpha$ after normalizing the productivity parameter $B = 1$. There are no delays caused by the time to build. The rental price of housing in period 0 is known and equal to P_0 .³⁴ In period 0, the rental price of housing in period 1 is uncertain, with expected value $\mathbb{E}(P_1)$. The value of P_1 is realized at the beginning of period 1. To keep matters simple, we ignore discounting and assume a constant unit cost for capital. We also postpone our discussion of the effects of construction on the rental price of housing.

³⁴ In static settings, the ‘price’ (asset value) of a dwelling and its rental value are confounded.

The builder chooses whether to develop in period 0 ($d = 0$) and collect rental income in both periods or wait and develop in period 1 ($d = 1$). While delaying development until period 1 implies foregoing the collection of rental income for period 0, it also allows the builder to build with the knowledge of the realized value of P_1 rather than its expected value $\mathbb{E}(P_1)$. More formally, when developing in period 0, the expected profit of the builder is $\pi(d = 0) \equiv \pi(P_0, k_0) + \mathbb{E}(\pi(P_1, k_0)) = P_0 k_0^\alpha + \mathbb{E}(P_1) k_0^\alpha - k_0$. The first-order condition with respect to profit maximization yields the profit-maximizing investment in capital. After substituting for it into the profit function, we obtain,

$$\pi(d = 0) = (1 - \alpha) \alpha^{\alpha/(1-\alpha)} (P_0 + \mathbb{E}(P_1))^{1/(1-\alpha)}. \quad (25)$$

In the same way, choosing to delay and develop in period 1 yields a profit of $\pi(P_1) = (1 - \alpha) \alpha^{\alpha/(1-\alpha)} (P_1)^{1/(1-\alpha)}$. For a builder having to choose whether to build in period 0, the expected profit of delaying development until period 1 is

$$\pi(d = 1) \equiv \mathbb{E}(\pi(P_1)) = (1 - \alpha) \alpha^{\alpha/(1-\alpha)} \mathbb{E}((P_1)^{1/(1-\alpha)}). \quad (26)$$

From Eqs. (25) and (26), it is easy to see that delaying development is profitable if and only if

$$\pi(d = 1) > \pi(d = 0) \Leftrightarrow \mathbb{E}((P_1)^{1/(1-\alpha)}) > (P_0 + \mathbb{E}(P_1))^{1/(1-\alpha)}. \quad (27)$$

Because $0 < \alpha < 1$, we know from Jensen's inequality that this condition is satisfied when $P_0 = 0$. More generally, this inequality will also be satisfied when P_1 is large enough relative to P_0 or, less obviously, when the variance of P_1 is large enough. These two results suggest that delays are more likely when rental values are expected to grow and when uncertainty is higher.³⁵

To be clear, the value of waiting $\pi(d = 1) - \pi(d = 0)$ is driven by the convexity of the profit function on the left side of the inequality in Eq. (27). When the variance of P_1 increases, the profit of competitive builders who delay development increases (convexly) with P_1 . On the other hand, the value of building now on the right side of the inequality in Eq. (27) does not depend on the variance of P_1 . It depends only on its expected mean.³⁶

Although insightful, this framework may oversimplify matters and miss some key effects. This framework features only two periods but could easily be extended. Adding a fixed number of periods to the model above is conceptually

³⁵ When investment takes a long time to implement, for example, because of construction lags, greater uncertainty can promote more rapid investment instead (Bar-Ilan and Strange, 1996). When projects take time to build, the probability of extreme outcomes increases. When builders can abandon their project, perhaps at a cost, the opportunity cost of waiting increases with uncertainty. This effect can be strong enough to more than compensate for the increase in benefits from waiting under greater uncertainty.

³⁶ Although arguably less relevant in the context of real estate development, risk aversion provides another reason for delays by making outcomes with low prices in period 1 disproportionately costly.

straightforward. The burden is in the computations. As is well-known, infinite-horizon models are of a different nature because they can no longer be solved by backward induction, as above. They nonetheless seem appropriate to model housing development where decisions to develop or redevelop frequently occur while new constructions last for several decades and, in some cases, centuries.

Let us turn to an infinite-horizon extension of the model presented above. Market j is divided into L_{j0} parcels of unit size that are initially undeveloped. Over time, the number of vacant parcels evolves according to $L_{jt+1} = L_{jt} - d_{jt}$ where d_{jt} denotes the number of parcels developed at period t .

We consider the competitive builder who owns parcel i in market j . At each period, this builder chooses (i) whether to develop its parcel and (ii) how much housing h_{ijt} to supply if it gets developed, i.e. if $d_{ijt} = 1$. These decisions are irreversible. Undeveloped land does not generate revenue, while parcel i , if developed at time \bar{t} , generates a rental income of $P_{jt}h_{ij\bar{t}}$ for every period $t \geq \bar{t}$, which is discounted by a factor $0 < 1 - \delta < 1$.

Building on parcel i is instantaneous but involves a fixed development cost c_{ijt} , which is idiosyncratic. More specifically, and in line with much of the previous literature, we assume that c_{ijt} is i.i.d. and follows a Type 1 Extreme Value distribution (or logit) every period. Building also incurs a variable cost. With the same Cobb-Douglas production function as in Eq. (15): $h = k^\alpha$ and following Eq. (23), the variable cost function is $C_{jt}(h) = r_t h^{1/\alpha}$ after normalizing productivity $B_{jt} = 1$ to ease the exposition. Then, the stock of housing in j grows according to $H_{jt+1} = H_{jt} + h_{jt}$ where h_{jt} sums new housing produced across all newly developed parcels. Because parcels are all of the same size and face the same variable cost function: $h_{jt} \equiv \sum_i d_{ijt} h_{ijt} = d_{jt} h_{ijt}$.

The evolution of housing in market j is described by L_{jt} for the number of vacant parcels and H_{jt} for the stock of housing, the actions of developers d_{ijt} and h_{ijt} , the idiosyncratic fixed cost of development c_{ijt} , and the rental price P_{jt} of housing and r_t of capital. The rental price of housing is exogenous and, as above, uncertain. More precisely, at period t , only P_{jt} is known with certainty. For $t + 1$ and beyond, developers form (rational) expectations. Given the fragmentation of the construction industry, this assumption of price-taking is reasonable in most cases. At the market level, new constructions affect the rental price of housing through their effect on demand. We return to this equilibrium effect in Section 6. To keep our notations simple in what follows, we denote the set of state variables by Ξ_{ijt} .

The ex-ante value of (undeveloped) parcel i is given by

$$V_{ijt} = \mathbb{E} \left(\max_{d \in \{0,1\}} v_{ijt}^d | \Xi_{ijt} \right), \quad (28)$$

where the value of remaining undeveloped is:

$$v_{ijt}^0 = (1 - \delta) \mathbb{E}(V_{ijt+1} | \Xi_{ijt}), \quad (29)$$

and the value of optimally developing this parcel is:

$$v_{ijt}^1 = \max_{h_{ijt}} \left\{ P_{jt} h_{ijt} - r_t h_{ijt}^{\frac{1}{\alpha}} - c_{ijt} + (1 - \delta) \mathbb{E}(\Pi_{ijt+1} | \Xi_{ijt}) h_{ijt} \right\}, \quad (30)$$

where $P_{jt} h_{ijt}$ is the rental income of the current period, $r_t h^{1/\alpha}$ is the variable cost of construction from above. In this expression, we also define Π_{ijt+1} , the value of housing per unit at period $t + 1$. It is the sum of the rental price of housing in $t + 1$, which is known at that point, and the expected value of the sum of discounted rental prices in $t + 2$ and after:

$$\Pi_{ijt+1} \equiv P_{ijt+1} + \mathbb{E} \left(\sum_{\tau=t+2}^{\infty} (1 - \delta)^{\tau - t - 1} P_{ij\tau} | \Xi_{ijt} \right). \quad (31)$$

These expressions call for a few comments. First, the value function in Eq. (28) is written from the perspective of a builder with a vacant parcel. This vacant parcel generates zero income. Once a parcel has been developed, there is no remaining decision, and the parcel will bring a rental income for each future period. Second, if the builder chooses to leave the parcel vacant, the associated conditional value function from Eq. (29) repeats the value of a vacant parcel at the next period. This next-period situation implies the application of the discount factor $1 - \delta$, knowledge of the realization of P_{jt+1} , and a new draw of the fixed cost of development, c_{ijt+1} . Third, if the builder chooses to develop the parcel in period t , the associated conditional value in Eq. (30) includes the rental revenue for the period, the cost of construction, and the expected value of the discounted sum of future rental incomes.

This model, which is closely related to Murphy (2018), Peng (2023), and Hsiao (2023), can now be solved for its two different margins: when to build and how much to build when building.³⁷ Starting with the latter, we can first insert Eq. (31) valued at period t into Eq. (30) and write:

$$v_{ijt}^1 = \max_{h_{ijt}} \left\{ \Pi_{ijt} h_{ijt} - r_t h_{ijt}^{\frac{1}{\alpha}} - c_{ijt} \right\}, \quad (32)$$

after simplification. This expression shows clearly that the investment decision (i.e., how much to build) if made in t , consists in maximizing the value of the

³⁷ Before them, Paciorek (2013) estimated a city-aggregate analog to this model to show that more stringent land use regulations can lead to more price volatility through a real options mechanism. Another strand of literature models the decision of a single builder facing prices that follow a Brownian motion, following the pioneering work of Capozza and Helsley (1990) (see also Lebret et al., 2025, for an example of recent developments). These models provide nice insights about individual decisions relying on price uncertainty. A key limitation is that they are hard to aggregate to the market level because feedbacks of new constructions into prices affect their assumed Brownian motion.

house built, $\Pi_{ijt}h_{ijt}$, minus the variable and fixed costs of development. Taking the first-order condition, we obtain:

$$h_{ijt} = \left(\frac{\alpha \Pi_{ijt}}{r_t} \right)^{\frac{\alpha}{1-\alpha}}, \quad (33)$$

which solves for the optimal investment after equating its marginal cost with its expected intertemporal marginal return. We note that this expression is the intertemporal analog of the optimal static investment in housing derived in Eq. (16).

Estimating this equation is conceptually straightforward and implies regressing the log of how much housing is built on its unit price. However, this estimation is subject to the complications discussed in Section 3 in the static case, including those that arise from the unobserved productivity term B_{jt} , which is here hidden because it was normalized to unity for the sake of the exposition. The intertemporal nature of (33) adds additional complications because, in practice, the price of a new house is observed only if sold after it is built and not when the building decision was made.

Working at the property level, Peng (2023) estimates a more complicated version of the simple framework we just described because she is looking at redevelopments, which, unlike the (simpler) decisions examined here, are not one-off decisions. Hence, her analog of Eq. (30) also contains an additional term associated with the effect of any current redevelopment on the cost of future redevelopments. This extra term is not observed. Then, using rent data, she only observed realized rents, not expected rents, at the time of the redevelopment decision. Using realized rents, in turn, introduces an expectation error, which leads to further complications in the estimation.³⁸

Also, working at the property level, Murphy (2018) implements an indirect approach. Because he considers that housing services are not observed, he cannot regress new supply on prices. Instead, he assumes that housing services result from house characteristics with a known functional form but with unknown parameters. He first regresses house prices on house characteristics using a standard hedonic approach. He then modifies Eq. (33) to obtain a marginal cost equation. This marginal cost can be reconstituted using the estimated coefficients of the house price regression and subsequently regressed on house characteristics to recover the analogs of $\alpha/(1-\alpha)$ in Eq. (33).

Turning to the decision to develop, this is a standard optimal stopping decision. Because the fixed cost of development follows a Type 1 Extreme Value distribution, the probability of development of parcel i , is given by:

$$\mathbf{P}(d_{ijt} = 1 | \Xi_{ijt}) = \frac{\exp\left(\frac{1}{\chi} v_{ijt}^1\right)}{\exp\left(\frac{1}{\chi} v_{ijt}^0\right) + \exp\left(\frac{1}{\chi} v_{ijt}^1\right)}, \quad (34)$$

³⁸ To keep this section of manageable length, we only briefly describe how this dynamic framework gets implemented empirically but do not enter the details of the estimation strategies and their limitations.

where χ is the scale parameter of the logit distribution of c_{ijt} and v_{ijt}^0 and v_{ijt}^1 are defined in Eqs. (29) and (30), respectively. In turn, this equation can be aggregated across all (remaining) vacant parcels in market j .

Eq. (34) recovers the responsiveness of new housing development with respect to the price of housing and the price of vacant land. In our decomposition of the elasticity of housing supply in Eq. (10), this loosely corresponds to the extensive supply elasticity $\epsilon_P^L(t, t+1)$. A key difference with the static models of Section 3.1, where the decision to develop depends only on the price of housing, is that, in a dynamic setting, this decision depends on the price of housing relative to the price of vacant land. According to Eq. (34), building on a vacant parcel is more likely when the value of a developed parcel increases relative to the value of vacant land, that is, when the rental value of housing at the current period P_{jt} is higher. Like in the two-period model, the opportunity cost of delaying development is the loss of rental value at the current period. Similarly, the opportunity benefit of delaying is a higher future revenue from building more (or less) if future rental values turn out to be larger (or lower) than expected. Getting a good realization of the idiosyncratic fixed cost of development is another factor of delay in our infinite-horizon framework that we did not consider in the two-period model above.

Estimating Eq. (34) (or its market-level analog) requires information about the development of parcels at period t . Then, this estimation requires the knowledge of the conditional value functions v_{ijt}^0 and v_{ijt}^1 . In our simple case, the optimal supply of housing h_{ijt} is known directly from the estimation of Eq. (33). It can be inserted into Eq. (32) so that v_{ijt}^1 depends only on the observed unit price of housing, the (previously estimated) parameters of the cost function, and the logit error term c_{ijt} . In Eq. (29), v_{ijt}^0 is the continuation value associated with vacant land. Following Kalouptsidi (2014), we note that this continuation value is just the price of a vacant parcel, which can be estimated or read directly when the data is available.

With neighborhood-level data, Hsiao (2023) employs a version of the strategy just described to estimate the responsiveness of housing supply to the price of housing and the price of vacant land. His estimation side-steps the direct estimation of Eq. (33) but accommodates it. Hsiao (2023) directly regresses a measure of new developments by neighborhood in Jakarta on local property prices. New constructions and property prices are likely to be simultaneously determined by unobserved local shocks, as already discussed. Hsiao (2023) instruments local property prices with local amenities viewed as demand shifters. We discuss additional demand shifters below in Section 5.1.

Peng (2023) uses a similar estimation strategy for redevelopment in New York City, where limits to the floor-to-area ratio changed over time and where building over the limit is possible at some additional costs, which are estimated. Murphy (2018) takes a slightly different two-step approach that does not rely on the values of developed and vacant land. Instead, he first estimates the probability of development and the transition probabilities for state variables from

the difference $v_{ijt}^1 - v_{ijt}^0$, that is, from the period profit rather than the values of land, vacant and developed. The output of this regression can then be used to estimate an analog of Eq. (34) for market j to retrieve the remaining structural parameters: $1 - \delta$ the discount rate, \bar{c}_{ijt} the mean fixed cost, and χ the scale parameter of its distribution, which measures its variance.³⁹

The payoff from this type of estimation is two-fold. First, a cost function can be retrieved with both the variable cost at the intensive margin and the fixed cost at the extensive margin. For single-family homes in Northern California, Murphy (2018) recovers variable costs per unit of floorspace close to the engineering estimates of RSMeans, with a similar evolution over time. His estimation also suggests large, highly dispersed fixed costs of development that increase over time. For multifamily and commercial properties in New York City, Peng (2023) estimates a convexity of the variable cost function in Eq. (32) of $1/\alpha = 1.47$, which implies $\alpha = 0.68$, a value close to those estimated in Section 3. She also estimates a marginal cost of construction close to the RSMeans estimates for buildings in the same class in New York City. Importantly, regulatory costs at the intensive margin represent about a quarter of the marginal cost of construction. These variable costs are nonetheless dwarfed by large and highly dispersed fixed costs of redevelopment.⁴⁰ We note that the results of Murphy (2018) and Peng (2023) are convergent despite looking at different classes of buildings and different markets. Their cost functions are close to those estimated by static approaches, but the fixed costs they uncover are considerable.

The second output of these studies is a short-run supply elasticity, which can be decomposed into a long-run elasticity supply at the intensive margin and a propensity to develop or redevelop at the extensive margin as in Eq. (10). Murphy (2018) shows a large wedge between the two elasticities, which appears to be driven by the builders' expectations of future prices. Peng (2023) also highlights slow construction responses. In her case, a dramatic relaxation in maximum floor-to-area ratios in parts of New York City in the early 21st century only led to an increase of about 0.5% in the housing stock after nearly 20 years since the reform started. Here again, these papers converge to indicate extremely long development or redevelopment lags resulting from the high fixed costs that they also highlight. The persistence of buildings with floor-to-area ratios much below what is authorized need not reflect any form of misallocation.

³⁹ Both Hsiao (2023) and Peng (2023) rely on asset values to estimate their version of Eq. (34). When the data is available, this approach only imposes minimal requirements regarding expectations since these are directly read in the price of vacant and developed parcels. Murphy (2018) uses a more involved strategy which differences out the value of vacant land to recover local time effects which can then be used to provide estimates of local values v_{ijt}^0 and v_{ijt}^1 used to estimate an aggregated version of Eq. (34) following Arcidiacono and Miller (2011). A full-solution estimation in the spirit Rust (1987) involves computing value functions using a (nested) fixed-point iteration for each guess of the parameters to be estimated. This type of approach is typically computationally prohibitive given the high dimensionality of the state space.

⁴⁰ Hsiao (2023) focuses on the effect of flooding on construction as one input into his investigation of the importance of moral hazard.

Instead, it may only reflect the normal operation of markets with durable goods or investments (Asker et al., 2014), at least in part.

This type of framework can be enriched to allow further decisions to be made. Adding exogenous depreciation to the housing investment decision above seems straightforward, but making depreciation endogenous adds another decision every period. Multiple investments are difficult to handle when they are additively separated because the first-order condition associated with Eq. (30) becomes highly non-linear. In a development context, Henderson et al. (2021) consider two building technologies for formal and informal housing and the transition between the two, depending on the location within a city (Nairobi). However, a key simplification they make is that informal housing is not durable.

The framework above could be extended in a similar spirit to study the densification of neighborhoods from heterogeneous single-family homes to multi-family apartments with land assembly. This channel used to be the main way cities developed at the intensive margin. Despite repeated calls for urban densification, it now accounts for only a marginal share of new housing supply (Baum-Snow and Han, 2024). Understanding the key barriers to urban densification is an important challenge. Given the salience of the irreversibility of housing development and its durability, there is little doubt that extensions of this framework will be useful in exploring a variety of issues associated with housing development.

Given the convexity of the cost of housing, the equilibrium of the builder's problem is generally unique. However, complementarities, either on the supply or demand side, can change this result. Due to the externalities arising from housing development or redevelopment described below, future rental prices may increase when supply expands. This occurs, for instance, when redevelopment leads to increased neighborhood quality and, thus, desirability. As a result, two equilibria are possible. In the first, no redevelopment occurs because builders do not expect other builders to redevelop. In the second equilibrium, redevelopment occurs because builders expect other builders to redevelop, making redevelopment worthwhile. A similar issue arises with the emergence of subcenters or entire cities (Henderson and Venables, 2009). Such complementarities are an important reason why coordination mechanisms like masterplans are potentially desirable. We return to this issue below.

Another key implication of the durability of housing is that it creates a fundamental asymmetry between situations with growing demand and situations with declining demand. With a growing demand in market j , new housing is developed, with some lags, following the mechanics we just described. This corresponds to a gradual adjustment of the supply curve rotating clockwise. When demand declines, the housing stock cannot be unbuilt. The supply curve is thus kinked with a vertical portion at the current level of stock. Hence, when demand declines, most of the adjustment occurs through prices. These downward price adjustments can be considerable as the price of unwanted properties can fall to zero or close to it. Only in the very long run will the housing stock adjust

through its depreciation. In turn, this argument implies that housing decline is often highly persistent following a large negative shock in a city or neighborhood. See Glaeser and Gyourko (2005) for the original articulation of the issues and supportive empirical evidence. Duranton and Puga (2014) provide further discussion.

4.2 Lumpy and differentiated housing

So far, we have assumed that a housing market is in equilibrium when the aggregate supply of housing services in this market is equal to its aggregate demand. This determines the unit price of housing, and, in turn, given this price, each household purchases a quantity of housing commensurate to its demand. Although useful for many purposes, this conceptualization of the housing market misses two critical features. First, housing is indivisible. Existing dwellings come as bundles of land and capital that are essentially fixed, at least in the short run. Second, housing is vertically differentiated, as these bundles differ in how much housing they offer. In simple words, some dwellings are nicer than others.⁴¹ Rather than setting the unit price of housing and splitting the aggregate stock of housing services between residents, the operations of the housing market set the price of each dwelling and allow each household to be assigned to a dwelling.⁴²

Beyond its greater realism, this conception of housing is appealing a priori because it allows us to consider the entire housing market without arbitrarily partitioning it into separate segments. It also generates some interesting and non-trivial insights. For example, as we show below, adding high-quality dwellings affects the market differently than adding low-quality dwellings. Finally, for policy, we often care about understanding the effects of providing housing units of a certain type (in addition to housing services) and their distributional implications. Decisions by housing developers about the number of units to build on a given parcel are challenging to handle conceptually, as they reflect an interaction between heterogeneous demand by different types of households for different locations with the construction decision. Assignment models seem well-suited to handle these issues.

Sweeney (1974) and Braid (1981) first proposed to model the housing market as an assignment process. See also Sattinger (1993) for an early introduction to assignment models with a focus on the labor market. While Määttänen and

⁴¹ We also expect housing to be horizontally differentiated. Such differentiation typically implies (unpriced) benefits from the diversity of the housing stock as it allows residents to find accommodation more suited to their tastes. Zhang (2022) proposes a new empirical framework to explore this issue and suggests that real wages in New York City may be up to 15% higher than what a naive hedonic regression indicates when the benefits from housing diversity are taken into account.

⁴² As it turns out below, we can no longer define a market as a set of dwellings that face the same unit price of housing since the latter varies across dwellings of different sizes. Instead, a housing market in an assignment model is a set of dwellings providing perfectly substitutable housing services, except for their size.

Terviö (2014) provide a more thorough treatment of the assignment equilibrium in the housing market in an exchange economy, we adapt Landvoigt et al.'s (2015) simplified model to keep the exposition simple.^{43,44}

Before turning to the analysis of the effects of changes in supply, we first consider the assignment of a set of households to a fixed stock of houses. Housing is vertically differentiated, and the quantity of housing services provided by a house, often referred to as its “quality”, is denoted by $h \in [\underline{h}, \bar{h}]$. There is a continuum of indivisible houses of unit mass described by its strictly increasing cumulative distribution function, $G(h)$.

There is also a continuum of households of unit mass. Households consume a composite numéraire good c and one house offering housing services h at price $P(h)$. They maximize their utility function $u(c, h)$, which is strictly increasing in both arguments and quasi-concave, subject to the budget constraint $P(h) + c = w$. Income, $w \in [\underline{w}, \bar{w}]$, differs across households. Its cumulative distribution function, $F(w)$, is strictly increasing.

The first-order conditions for utility maximization by households imply

$$MRS_{hc} \equiv \frac{\frac{\partial u(c, h)}{\partial h}}{\frac{\partial u(c, h)}{\partial c}} = \frac{\partial P(h)}{\partial h}, \quad (35)$$

as the marginal rate of substitution of housing for the numéraire good is equal to the marginal price of a dwelling offering more housing services, $\partial P(h)/\partial h$. We note that this condition is satisfied by households of different income levels along the distribution of dwellings. Thus, there is no reason for $\partial P(h)/\partial h$ to be constant and, consequently, for the price of a house to be proportional to the housing services it provides. This is the first major implication of housing indivisibility: the price of a house $P(h)$ is no longer necessarily equal to $P \times h$ and, instead, depends on the distributions of h and w .

To prove this last point rigorously, we first note that our model features positive assortative matching by h and w . With quasi-concave utility, the marginal rate of substitution of housing for the numéraire good decreases in housing. In turn, diminishing marginal rates of substitution in Eq. (35) imply that richer households have a higher willingness to pay for dwellings offering more housing. This can be derived formally by differentiating the marginal rate of substitution with respect to w after substituting in the budget constraint for c .⁴⁵

⁴³ A key difference with Määttänen and Terviö (2014) is that we consider a two-sided heterogeneity framework. In contrast, theirs is a one-sided heterogeneity exchange economy where households of heterogeneous income are initially allocated heterogeneous dwellings. As a result, the budgets that households face are endogenous and depend on the price of their dwelling in equilibrium. This perhaps more realistic setting comes at the cost of additional complexity.

⁴⁴ Models of vertically differentiated housing have also been used to understand house price dynamics following changes in the demand for housing and changes in the ability of households to obtain a mortgage to finance their house (Ortalo-Magné and Rady, 2006; Landvoigt et al., 2015; Nikolakoudis, 2024). Behrens et al. (2014) and Davis and Dingel (2020) have used assignment models to study the sorting of households across cities.

⁴⁵ This amounts to a standard single-crossing solution ensuring the perfect separation of households of increasing w across dwellings of increasing h .

Hence, in any equilibrium, richer households occupy bigger and more expensive dwellings.

This intuitive property is an important step towards solving the model since the assignment of dwellings to households can now be described by a strictly increasing matching function $h^*(w)$ for $w \in [\underline{w}, \bar{w}]$ with $h^*(\underline{w}) = \underline{h}$ and $h^*(\bar{w}) = \bar{h}$. In what follows, it is convenient to work with its inverse, the assignment of households to dwellings, $w^*(h)$. With positive assortative matching, the market clearing condition can be expressed as $F(w^*(h)) = G(h)$, as all households of rising income are assigned to the dwellings offering more housing services. In turn, this condition can be written as,

$$w^*(h) = F^{-1}(G(h)) . \quad (36)$$

Importantly, the assignment of households of higher incomes to larger dwellings only depends on their respective distributions and not on household preferences, unlike in standard models. Hence, “quantities” do not depend on preferences, but prices do. As made clear by Eq. (35), preferences determine the price of each dwelling. In turn, these prices make the allocation described by Eq. (36) sustainable in equilibrium. Intuitively, the household of income w that occupies dwelling h must pay enough to avoid being outbid by the household of income $w - \partial w$ who occupies dwelling $h - \partial h$. In turn, the price of dwelling $h + \partial h$ must be high enough to discourage household w from outbidding its current occupier of income $w + \partial w$. Hence, the price of dwelling $h + \partial h$ depends on the price of dwelling h , which in turn depends on the price of dwelling $h - \partial h$.

More formally, the price of dwelling h is given by

$$P(h) = P(\underline{h}) + \int_{\underline{h}}^h \frac{\partial P(\tilde{h})}{\partial \tilde{h}} dG(\tilde{h}) = \int_{\underline{w}}^{w^*(h)} MRS_{hc}(\tilde{w}) dF(\tilde{w}) , \quad (37)$$

where the second equality is obtained from Eq. (35) after normalizing $P(\underline{h}) = 0$ and changing the integration variable from dwellings to households. While we have not solved the equilibrium marginal rates of substitution yet, this equation delivers two important insights regarding the working of assignment models. First, the price of dwellings depends non-trivially on the distributions of houses $G(h)$ and income, $F(w)$. For example, dilating the distribution of incomes from $[\underline{w}, \bar{w}]$ to $[\underline{w}, \bar{w}']$ with $\bar{w}' > \bar{w}$ raises the marginal rate of substitution of housing for the numéraire good. In turn, this increase makes the price of dwellings increase faster with h , as the competition between the richest households for the best houses becomes more intense.

Second, an increase in the income of the richest households in the distribution does not affect the price of the dwellings of other (poorer) households. On the other hand, an increase in the income of the poorest households (which preserves their ranking for simplicity) increases the price of all dwellings, including the largest one. These price effects are sometimes referred to as directional price spillovers. They occur because of the additive nature of the price of dwellings

in Eq. (37). An increase in the income of the poorest households raises their marginal rate of substitution and, thus, how much they are willing to pay for dwellings with a larger h . As a result, richer households will need to pay more to retain their current dwelling, even if their income is unchanged. Simply put, demand-induced changes in the prices of dwellings ‘trickle up’ but do not ‘trickle down’.

To keep our derivations simple, we consider the case where $u(c, h) = c^{1-\eta} h^\eta$ and both income and housing services are uniformly distributed. From the Cobb-Douglas nature of preferences, we can rewrite Eq. (35) as

$$\frac{\partial P(h)}{\partial h} = \frac{\eta}{1-\eta} \frac{c}{h} = \frac{\eta}{1-\eta} \frac{w^*(h) - P(h)}{h}, \quad (38)$$

after using the budget constraint and the inverse matching function $w^*(h)$ to substitute for the consumption of the numéraire good. From the uniform nature of the distributions of w and h , it is easy to show that the assignment of households to dwellings is described by

$$w^*(h) = a_0 + a_1 h \quad \text{with } a_0 \equiv \frac{\bar{h}w - \underline{h}\bar{w}}{\bar{h} - \underline{h}}, \quad a_1 \equiv \frac{\bar{w} - \underline{w}}{\bar{h} - \underline{h}}. \quad (39)$$

Inserting Eq. (39) into Eq. (38) to substitute for $w^*(h)$ yields a first-order linear ordinary differential equation. Its solution is

$$P(h) = a_0 + (1-\eta)a_1 h - c_1 h^{-\frac{1-\eta}{\eta}} \quad (40)$$

where the integration constant c_1 is such that $c_1 = a_0 \underline{h}^{(1-\eta)/\eta} + (1-\eta)a_1 \underline{h}^{1/\eta}$ for $P(\underline{h}) = 0$ to be satisfied. The price of dwellings is therefore increasing and concave in h following the distributional assumptions made above.

Wang (2022) and Nathanson (2023) allow for a change in the supply of dwellings in this type of framework.⁴⁶ We consider the addition of a small discrete set of new dwellings of mass Δ to the housing stock. The housing services of these dwellings are distributed over the subsegment $[h_0, h_1]$ with $h_0 > \underline{h}$ and $h_1 < \bar{h}$. Assuming for now a fixed set of residents, this increase in housing supply leads households with income $w < w^*(h_1)$ in the original equilibrium to move to a larger dwelling. Hence, the increase in the stock of housing trickles down and benefits poorer households. To preserve positive assortative matching in the bottom part of the distribution, these moves to larger homes occur in a chain, with each household moving up by Δ ($= G(h'^*(w)) - G(h^*(w))$) in the distribution of dwellings. Now occupied by poorer households, the price of smaller dwellings declines. This price decline is magnified in equilibrium as

⁴⁶ Nathanson (2023) assumes a collection of discrete classes of dwellings with different levels of housing services rather than a continuum. The equilibrium is solved using boundary conditions between sets of dwellings with different levels of housing. Despite these differences, the key qualitative properties of the model appear similar.

poorer households now enjoy a higher level of utility, which reduces how much they are willing to pay for more housing services.

Thus, poorer households now occupy larger dwellings at a lower price. At the very bottom of the distribution, there is now a surplus of dwellings. The smallest occupied dwelling \underline{h}' is such that $G(\underline{h}') = \Delta$, the mass of new dwellings. Dwellings with $\underline{h} \leq h < \underline{h}'$ are now empty and end up with $P(h) = 0$. At the same time, for the dwellings with $h > h_1$, there is no change in their assignment. However, the price of these larger dwellings decreases. With households with lower income enjoying lower prices for $h < h_1$, the price pressure they put on larger dwellings weakens.

While everyone eventually benefits from the addition of new dwellings, two caveats are worth keeping in mind. First, we have considered so far that the set of households was fixed. If dwellings become cheaper in a market, they will attract newcomers. These newcomers will offset some of the benefits for residents of lower housing prices (Nathanson, 2023). The complication is that these newcomers will be selected in the income distribution depending on the utility they enjoy elsewhere, their moving costs, and the changes in the distribution of the price of local dwellings. At one extreme, if there is free mobility for all income levels at the initial equilibrium and if market j is small relative to the economy, adding to the stock of housing in j only implies an increase in population Δ , with income distributed over $[w^*(h_0), w^*(h_1)]$ matching exactly the distribution of new dwellings.

The second caveat is that new dwellings may be produced, not as net additions, but by upgrading existing dwellings over $h \in [h_{-1}, h_0]$. In this case, residents with income $w < w^*(h_{-1})$ do not experience any changes while richer households benefit from the upgrading of the stock of housing. When allowing for migrations in and out of the market, upgrading some dwellings will lead to the entry of new households and the exit (or homelessness) of some incumbents.

With a housing production function as in Eq. (15), builders face a trade-off between the convex marginal cost of construction and the fixed outlays associated with the acquisition of the parcel and the fixed cost to start construction. This cost function is represented on Fig. 6. The builder's profit maximization problem is essentially the same as in Section 3.1. The main difference is that the revenue from construction is no longer proportional to h as the builder faces a non-linear equilibrium price schedule. However, if $P(h)$ is concave in h as in the example derived above, there is still a well-defined level of housing services h^π per dwelling that maximizes profits. The difference between the price of the dwelling built and its costs is dissipated in the price of the land, as previously. Fig. 6 illustrates this situation.

Interestingly and perhaps surprisingly, all builders choose a similar level of housing services h^π for newly built dwellings. This result comes from the uniqueness of the solution to the profit-maximizing choice of dwelling size. This unique dwelling size occurs despite the heterogeneity of residents' income. In equilibrium, households reassign themselves across dwellings at no cost. Although we doubt this result holds in more realistic settings with frictions to

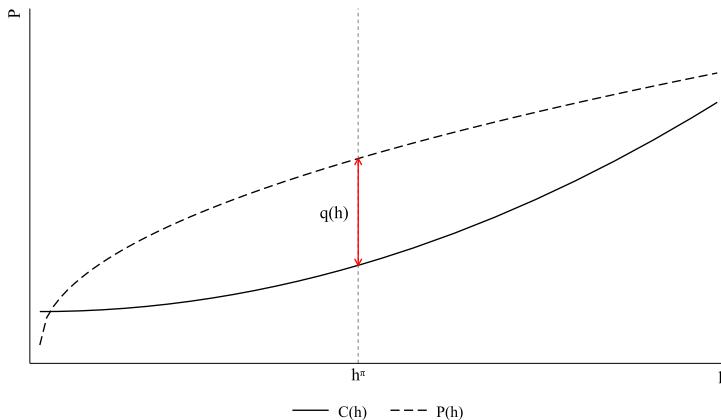


FIGURE 6 Supply decision in assignment models. *Notes:* The vertical axis is in numéraire and the horizontal axis is the quantity of housing services provided by a dwelling unit. The solid line is an example cost function $C(h)$, and the dashed line is an example price function anticipated by developers prior to construction. The difference between the price and the construction cost of a dwelling results being $q(h)$, the price of the land parcel. Following changes in fixed costs, $C(h)$ could be shifted up or down with no resulting change in h^π .

re-assignment, this result highlights a powerful force that pushes towards homogeneous neighborhoods: the variable cost of construction is convex.⁴⁷ Allowing for land parcels of different sizes or other differences in parcel characteristics would also lead to the construction of heterogeneous dwellings in equilibrium. These caveats aside, the homogeneity of new developments is consistent with casual observation in many markets.

Quite plausibly, new constructions may offer more housing services than existing dwellings ('bigger and nicer homes').⁴⁸ The effect of higher prices following an increase in demand is ambiguous. A proportional increase in $P(h)$ leads to a higher h^π . This would be illustrated by an anti-clockwise rotation of the price schedule in Fig. 6. However, changes in demand need not imply a proportional increase in the price of dwellings. For instance, with less binding credit constraints, demand may increase particularly strongly in the lower tail of the distribution, leading to strong price increases there. In turn, builders will also 'follow the demand' and may choose a lower h^π .

This approach can replicate two interesting features documented by Wang (2022). First, new constructions occur in markets such as Cleveland in the 2010s despite a general over-supply of housing and extremely low median prices.

⁴⁷ As long as it remains below a certain threshold, the fixed cost is irrelevant to the choice of h^π as it only affects land prices. Above this threshold, no profitable building occurs. Put differently, the fixed cost of development only affects the extensive margin of development. We return to this important point in Section 5.

⁴⁸ We do not model the dynamics of supply here, but the gradual decay of dwellings will push toward building new dwellings to be provided in the upper tail of the distribution – higher h^π – be it only to "replenish the stock". We develop this point in the next subsection.

These new constructions are at the upper end of the distribution to serve high-income households for whom several cheap houses are no substitute for a luxury one. She also documents that new constructions in San Diego during the boom of the early 2000s occurred for gradually smaller houses as the larger ones became unaffordable to all but the very richest households.

The empirical implementation of assignment models presents several challenges. While ‘testing’ for positive assortative matching in the housing market is likely foolish since it will never hold exactly, assessing how strongly it holds is a worthwhile empirical exploration, following Epple et al. (2020). The correlation between deciles of household incomes and deciles of house prices is strong, and the ordering is often perfect (Määttänen and Terviö, 2014; Wang, 2022). However, house prices are forward-looking and may not always reflect the user cost of housing well (Nathanson, 2023; Nikolakoudis, 2024), especially when comparing across markets. We also expect households to make their (indivisible) housing decision based on their permanent income and not their current income.

The trickling up of prices following income shocks at the bottom of the income distribution and the absence of trickling down after income shocks at the top is another interesting and specific property of assignment models. It is at the heart of the quantification of Landvoigt et al. (2015) and explored more broadly by Wang (2022) and Nikolakoudis (2024). While limited, the evidence about the asymmetric propagation of prices is encouraging. The trickling down of price changes after supply shocks is another interesting property of assignment models. It is explored by Wang (2022), Nathanson (2023), and Mense (2025) who provide supportive evidence. Understanding more precisely how prices adjust even if, as discussed below, we do not observe a full chain of moves would be interesting as well.

Consistent with the framework developed here, Handbury et al. (2024) argue that since the Great Financial Crisis, the growth in demand for rental housing has been particularly strong in the higher tiers of the distribution. Supply responded by expanding in those tiers, taming growth in rental prices for this segment of the market. At the same time, demand growth for rental housing in lower tiers remained unmet. Price growth was disproportionate for this segment of the market as demand grew, but supply did not respond. These findings beg two important questions. How much do rental prices at the lower end need to increase for new constructions to occur in this segment of the market? How much new construction is needed in the upper tier for prices to start declining in the lower tiers?

The existence of downward chains of moves between dwellings is another important prediction of assignment models. Obviously, the assumption of frictionless moves could only hold in the very long-run. As a result, a newly constructed dwelling with h units of housing services may result in vacancy in a dwelling with less than h , even if dwellings of even lower quality remain occupied. The empirical question is then the extent to which we observe moving

chains with households moving up the housing quality distribution after new constructions.

Using detailed address history data for 686 new large multifamily projects in the US, Mast (2023) tracks the former residences of current tenants and those who moved to replace them, etc. He documents moving chains that are fairly long and deep. A new market-rate building with 100 dwellings eventually leads 45 to 70 households to move out of below-median-income neighborhoods after six rounds of moves. While early moves are generally local, after six rounds, the chains reach a much wider housing market, essentially the metropolitan area. These moving chains are also relatively fast, as most of the reported effects occur within three years. Using data from Finland, Bratu et al. (2023) reach conclusions that are mostly similar to those of Mast (2023). The main difference is that moving chains in Helsinki go much deeper at each round, perhaps reflecting the more compressed nature of the Finnish housing stock. French and Gilbert (2024) also study vacancy chains in the US. Like Mast (2023), they find that moving chains deepen socially and broaden to the entire metropolitan area as the number of moves progresses. Unlike Mast (2023), they uncover short vacancy chains. It is unclear whether this difference is due to their much broader coverage, which may include dwellings in the lower segments of the distribution occupied by newly-formed households, or to difficulties in tracking movers and new housing.

The unobserved nature of housing services constrains empirical work on assignment models. Existing research mostly takes for granted the one-to-one mapping between $P(h)$ and h to work around h being unobserved. To quantify their model Landvoigt et al. (2015) had to take a stance on changes in the quality of houses on the market in 2000 and 2005. They do so using a repeat sales index, but this does not allow them to recover the underlying distribution of h . The price of a house is a linear function of its h in some cases, for instance, when the distributions of income and housing services are scaled versions of each other. These are, however, the situations where assignment models are the least interesting since they revert to the same properties as the models with divisible housing. In an innovative effort, Epple et al. (2020) explicitly treat housing services as a latent variable, which they estimate non-parametrically from repeated observations of the same market and a number of further identifying restrictions.

4.3 Filtering

Filtering links together the durability of housing and its vertical differentiation through depreciation. As dwelling i in market j ages, it decays, and the housing services it provides evolve according to $h_{ijt} = (1 - \delta) h_{ijt-1}$ with $\delta < 1$. In the simplest formulations, the depreciation rate, δ , is considered exogenous. Depreciation should nonetheless be viewed as a function of the owner's maintenance effort. In turn, this maintenance effort depends on housing prices in market j and the cost of maintenance.

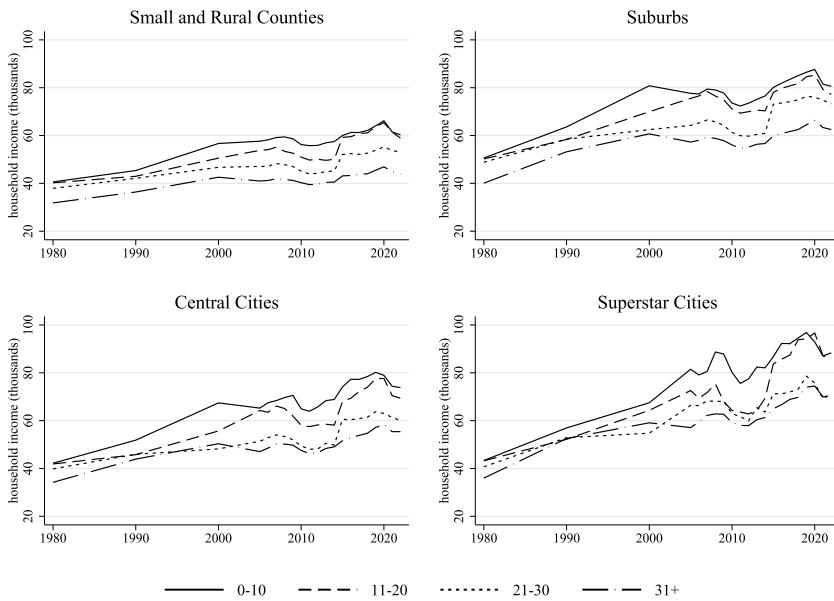


FIGURE 7 Mean household income by structure age. *Notes:* As in Fig. 1, plots use individual level census and ACS micro data aggregated to the household level. Household income is deflated to 2000 dollars using the CPI-U. For 2005-2022, indicated age ranges are only approximate due to limitations on built year information. For 2005-2014 on the horizontal axis, “0-10” covers structures built after 1999, “11-20” covers structures built in the 1990s, “21-30” covers structures built in the 1980s, and “31+” covers structures built prior to 1980. For 2015-2020 on the horizontal axis, “0-10” covers structures built in the prior 10 years, “11-20” covers structures built between 2000 and the indicated year minus 10, “21-30” covers structures built in the 1990s, and “31+” covers structures built prior to 1990. This reclassification results in upward jumps in 2015. For 2021-2022 on the horizontal axis, “0-10” covers structures built after 2009, “11-20” covers structures built 2000-2009, “21-30” covers structures built in the 1990s, and “31+” covers structures built prior to 1990.

These features are illustrated above in Fig. 3, which documents the slow disappearance of older dwellings from the stock of housing. Although the notion of housing decay is much broader than old dwellings taken out of the stock, the differences in exit rates documented by the various panels of the figure are striking: 0.7% per year in small and rural counties vs. 0.2% in superstar cities. See also Gyourko and Saiz (2004) for more systematic evidence about the endogeneity of housing investment and decay. While we discuss the strong correlation between household income and house prices above, there is also a strong correlation between the age of dwellings and the income of their residents, as illustrated by Fig. 7. See Brueckner and Rosenthal (2009) for further evidence on this. As a dwelling ages, not only does it decay, but it also hosts poorer households. This process is referred to as filtering.

Conceptually, in the absence of any other change, a constant rate of housing decay implies a leftward shift of the cumulative distribution function $G(h)$

of housing services: $G_t(h) = G_{t-1}(h/(1 - \delta))$.⁴⁹ Positive assortative matching remains, and residents occupy the same (decayed) dwelling from one period to the next. When all dwellings decay, the price schedule $P(h)$ adjusts following the marginal rates of substitution as per Eq. (35). In the special case where this marginal rate of substitution is constant, such as with Cobb-Douglas preferences, the price schedule shifts leftward over time in Fig. 6 with $P_t(h) = P_{t-1}(h/(1 - \delta))$.

Any steady state must then be such that new constructions only happen for the best houses offering \bar{h} units of housing to replenish the stock of housing from above and keep the quantity of housing and its distribution constant. This level of housing services \bar{h} for newly built houses should also correspond to what builders provide in Fig. 6, h^π . Hence, in the steady state of a simple assignment market with durable but decaying dwellings, the most profitable opportunities for builders are at the top of the market for the richest households. Newly built dwellings first serve this high-end market, after which they “filter” down the household income distribution as they age into lower-quality tiers.

While extremely stylized, this steady-state equilibrium is helpful to think about a number of possible perturbations affecting it. First, as households get richer over time, the richest households will be willing to pay more for dwellings that offer more housing services. Hence, the best dwelling will increase in quality over time: $\bar{h}_t > \bar{h}_{t-1}$. As the upper tail of the distribution of income shifts to the right, so will the supply of new dwellings, which will move toward providing ever-higher housing services to their wealthy residents. In turn, these dwellings will take longer to filter down.

It is also easy to understand how an increase in demand for middle- or lower-quality dwellings will not, in general, be directly served. As illustrated in Fig. 6, the price schedule in the body of the distribution or in the lower tail would need to increase a lot to make more affordable dwellings attractive to build. This situation is all the more difficult to reach since price increases for smaller dwellings are expected to percolate up and make new constructions at the upper end attractive again for builders. The more general point behind these two examples is that there are strong economic forces that push against building directly for poorer households. As a result, the supply of new housing in the lower tiers of the distribution may be highly inelastic, a fact noted by Handbury et al. (2024), among others.

As the supply of housing in the lower tiers of the market is mostly unresponsive to prices, it only expands when dwellings in the upper tiers decay. This lack of supply response implies that the housing stock may not match well with demand outside of the highest tiers of the market. For instance, there may be a missing middle of dwellings for the middle class, and this missing middle may persist for a long time.

⁴⁹ We assume here that a dwelling can decay all the way from \bar{h} to 0. In practice, multifamily apartment buildings are easier to reoccupy by poorer tenants, possibly after subdividing units. This is harder for high-end single-family houses.

Restrictive land use regulations may have a particularly perverse role in this type of model. By drying up the supply of housing, they end up excluding the poorest households from the market as the number of “viable” units declines (say, below a certain threshold for h). Allowing for maintenance to slow down this decline can exacerbate the problem. If new supply at the top is scarce, the wealthiest residents will try to slow down the decay of their dwellings (or even improve them). In turn, these actions slow down or completely stop the arrival of newly available, lower-quality housing for less well-off residents.⁵⁰

An oft-considered solution to this problem is to mandate new lower-quality dwellings for new constructions through inclusionary zoning (IZ). We discuss these policies below, but they seem to fail, as the lower cost of building at lower quality is dwarfed by much lower revenues for builders. Another approach is to make land use regulations less restrictive and reduce other impediments to new constructions. These new constructions happen in the upper tail of the quality distribution. In the short run, they allow lower-income residents to move up through moving chains. In the long run, the decay of these additional dwellings eventually filters down and increases the supply of lower-quality dwellings. The main limitations are that (i) moving chains may be too slow and shallow, and (ii) wealthy newcomers from different metropolitan areas may occupy new “luxury” housing.^{51,52}

We also note that housing decay also plays an important role in standard limitations of rent control policies Arnott (1995). With endogenous maintenance, landlords let rent-controlled housing quality decay until controlled rents correspond to their market rents. In this case, the lack of middle-quality housing no longer arises from a lack of filtering from the top but instead from too much filtering to the bottom.

At this stage, the question becomes how much filtering there is in the data. Despite a heated debate in the policy literature, very little is known about the rate at which housing “naturally” decays (say, with only minimal investment) and how maintenance and renovations affect decay. For instance, a 0.5% annual rate of decay over 50 years implies a decline in h for a dwelling of 22%. This would be arguably insufficient for filtering to be a major source of housing supply in the lower tiers of the housing market. On the other hand, a 2.5% decay would imply a decline in h of 70% over 50 years.⁵³ In Fig. 3, we crudely proxy filtering through the exit of dwellings from the housing stock. In path-breaking work, Rosenthal (2014) develops a “repeat-income” methodology to measure filtering.

⁵⁰ See Arnott and Braid (1997) for an early related formalization.

⁵¹ As we show in Section 6, having highly skilled workers moving to prosperous locations can generate large gains but does not solve the issue of insufficient housing for residents in the lower deciles of the income distribution.

⁵² A third option to provide affordable housing would be to greatly lower the variable cost of construction of such dwellings through, for instance, cheaper modular housing.

⁵³ Using hedonic methods, Francke and van de Minne (2017) find that the typical housing structure in the Netherlands loses 43% of its value over 50 years gross of maintenance.

This approach essentially consists of following dwellings and measuring the change in the income of their successive residents as these dwellings age.

We note that this metric does not capture the physical decay of dwellings *per se*, but instead captures changes in the assignment function of household income to dwellings. This is an analog to changes over time in the assignment function $w^*(h_{ijt}) - w^*(h_{ijt-1})$ in the model above. This metric is relative and captures both the decline of dwellings as they age and the possible movement of poorer households into higher-quality housing. Although not directly informative about changes in the housing stock, this is a relevant measure of filtering. A filtering rate of 1% corresponds to occupants of a dwelling being 1% poorer as the dwelling ages by one year.

Rosenthal (2014) estimates annual filtering rates of 0.5% for owner-occupied dwellings and 2.5% for rentals in the US between 1975 and 2011. For owner-occupied housing, the filtering rate is high during the first years of the life of a dwelling before stabilizing. It is even increasing again for older dwellings, acknowledging the selection of dwellings that pass the 50-year mark. Importantly, dwellings often switch from owner-occupier to rental status. Accounting for this, Rosenthal (2014) computes a 1.9% filtering rate overall, which corresponds to a 60% decline in household income over 50 years. This would make filtering an important source of housing supply for the lower tiers of the housing market.

However, as noted by Rosenthal (2014), filtering is locally highly sensitive to market conditions, and to housing prices in particular, as should be expected. This finding is confirmed by Liu et al. (2022), who find annual filtering rates that vary from -1.6% in Topeka, Kansas to a positive 0.7% in San Francisco. Spader (2024) extends Rosenthal's analysis to 2021 and finds that filtering essentially stopped in the US after 2011 when the housing market started to recover after the Great Financial Crisis while construction lagged. This finding suggests that filtering is highly responsive to prices and can only function when enough new housing gets built. In that case, filtering is a primary source of housing supply for the lower tiers of the housing market. When building activity slows down, however, filtering stops.

4.4 Housing externalities

Another key feature of housing is that the value of a dwelling depends on the number and quality of nearby dwellings and the people living inside them. In other words, there are housing externalities. These externalities can occur directly from the supply side. Re-painting a house can increase the value of nearby houses, as they become more appealing to potential buyers. In turn, higher house values in a neighborhood may encourage further maintenance and improvements. Housing externalities may also arise from the demand side. Nicer houses nearby may attract more desirable neighbors, in turn making the neighborhood more attractive and leading to further increases in supply. Under a slightly different causal chain, the arrival of more desirable neighbors may trigger improvements in the local housing stock, leading to further changes in the

composition of residents in a neighborhood and more changes to the local stock of housing, a process sometimes referred to as gentrification. Regardless of their exact mechanics, as housing externalities affect property values, they also affect the supply of housing. We also keep in mind that, because they affect migrations across locations, housing externalities can have effects far from where they take place. Investing in housing somewhere can lead to housing decay elsewhere.

Housing externalities may first occur at the construction stage. In his review, Brueckner (2000) distinguishes three market failures associated with the construction of new dwellings. First, new housing developments fail to account for the social value of open space. Open spaces at the outskirts of cities are enjoyed as recreational amenities by city dwellers. Open spaces are also desirable locations for new constructions. In turn, new constructions make open spaces less open and reduce their amenity values. This negative externality of urban sprawl is nonetheless more subtle than this simple story. Turner (2005) shows that preferences for open spaces nearby will lead to too little open space in equilibrium, that is, too much urban expansion. These preferences also imply an overly dense and populated urban core. The reason for this surprising result is that preferences for open spaces cut both ways. Some residents will rush to occupy peripheral urban spaces, but by doing so, they will destroy the appeal of these locations, and thus many will elect to live in the (oversized) core. More generally, spatial externalities associated with the construction or renovation of housing often generate subtle effects in equilibrium.

The second market failure associated with new constructions is the failure to account for the traffic congestion they generate. Beyond congestion, driving also generates pollution, noise, and accidents. Poor parking pricing also generates too much cruising for parking, and before that, it also distorts incentives to drive to some locations. In short, residents do not pay the full social cost of their commutes and other errands. As a result, city residents will tend to live in overly remote locations. Cities will thus spread out inefficiently. There is an empirical literature documenting the patterns of housing development with notable contributions by Burchfield et al. (2006) and a lively literature in urban planning and remote sensing (Artmann et al., 2019, provide a recent review). However, drawing convincing welfare estimates about the cost of urban sprawl and sound policy conclusions is still an open (and highly disputed) issue.

The last market failure associated with new constructions highlighted by Brueckner (2000) is the failure to account for the cost of providing infrastructure and services to new development. In the US, this failure has been documented by Altshuler and Gómez-Ibáñez (2000) among others. This market failure also leads to the development of new housing overly far from city centers and at suboptimally low levels of density.⁵⁴

⁵⁴ One could also add externalities associated with local population density and other measures of urban form, either in terms of production or consumption. With such externalities, the equilibrium building stock will not be optimal, but the direction and magnitude of the inefficiencies differ across locations and are sensitive to functional forms. See Duranton and Puga (2015) for further discussion. We discuss the effects of neighborhoods and the composition of their housing stock below.

For redevelopment and maintenance, another set of externalities becomes more salient. In their analysis of the local spillover effects of new Low Income Housing Tax Credit (LIHTC) housing, Diamond and McQuade (2019) estimate that a new (multifamily) development in a poor neighborhood increases nearby house prices by 6.5% while lowering crime rates and attracting diverse populations. In a rich neighborhood, a new LIHTC development reduces nearby house prices by 2.5% and attracts poorer residents. At least in poor neighborhoods, some local externality must be at play for a new housing development to raise the value of other dwellings despite an expansion of supply. It is harder to tell whether this externality is associated with the new building itself or with the residents living in it.

In a radically different setting, Hornbeck and Keniston (2017) document the reconstruction of Boston after the great fire of 1872. They show that, after reconstruction, property prices are significantly higher for burned plots and nearby unburned plots, as new constructions and building improvements encouraged further improvements nearby. Spillovers of a similar sort are also documented in Davidoff et al. (2022), which we discuss further below, and González-Pampillón (2022).

In a setting where it is easier to distinguish between externalities arising from better housing with those coming from more desirable neighbors, Rossi-Hansberg et al. (2010) report external land value gains of two to six dollars per dollar invested in housing renovation in a neighborhood improvement program in Richmond (VA). These positive effects decline by half over 300 meters. The counterpart of these results is that we also expect the contagion of blight that leads to low maintenance, more abandonment, and underutilized land even in otherwise high-cost areas (Owens et al., 2020).⁵⁵

There are also alleged externalities associated with bringing together housing units of different levels of quality within the same subdivision or the same multifamily complex. This practice is often referred to as inclusionary zoning (IZ). Before discussing these externalities, some background about IZ and some understanding of its mechanics are in order. Stacy et al. (2023) report the existence of IZ laws in 886 jurisdictions in 25 states and the District of Columbia in the US in 2016. These restrictions are usually hard constraints on new developments. They mandate a share of “affordable” units to be sold below market prices in all new developments. In some cases, developers have the option to opt in and receive fiscal incentives for doing so (See Stacy et al., 2023, for further descriptions and details).

There is little doubt that these mandates act as a tax on the supply of new housing. Returning to Fig. 6 and our simple assignment framework, the builder would like to build dwellings offering h^π units of housing services per unit of land in the absence of restrictions. An IZ mandate will constrain the builder to

⁵⁵ A related literature focuses on the effects of house foreclosures during the Great Financial Crisis. For instance, Gerardi et al. (2015) find that a foreclosed home lowers the price of non-distressed houses within a tenth of a mile by about one percent.

devote a fraction s^{IZ} of the new developments to be ‘affordable’ dwellings with each $h^{IZ} < h^\pi$ units of housing services to be sold at a set price P^{IZ} . The cost for the builder is twofold. First, an IZ forces the builder to build low-quality units. This is inefficient as land is not put to its highest and best use. Second, these dwellings are also generally sold at a set (or negotiated) price P^{IZ} , which is usually well below the corresponding market rate for such dwellings $P(h^\pi)$. While the builder may make a loss on affordable dwellings, $P^{IZ} < C(h^{IZ})$, a multi-dwelling project will remain profitable provided it remains profitable overall. If a housing project gets developed, the main cost is thus perhaps the small inefficiency associated with building low-quality units. Everything else is a transfer from landowners to the buyers of affordable housing units as we expect the reduced profits of the builder to be capitalized in the value of land.

The main issue is whether such mandates reduce the supply of new housing as the implicit tax makes marginal projects unprofitable so that landowners will no longer sell their parcel for development (or redevelopment) or sell it for another use. Unfortunately, evaluating the effects of IZ policies on housing supply is made difficult by the existence of multiple offsets, which are often adjusted after the introduction of an IZ mandate. Among others in the literature, Schuetz et al. (2011) find evidence of an adverse effect on housing supply for an IZ mandate in Boston but not for San Francisco. Krimmel and Wang (2024) use a boundary discontinuity approach to show that, in some parts of Seattle, allowing for upzoning combined with an IZ requirement led to a large decline in new construction in the treated parts. Less construction should lead to higher prices, a recurring finding in the literature. More subtly, IZ aims to change the composition of what is built and should thus affect the entire price schedule $P(h)$ in Fig. 6 through equilibrium effects. With relatively more units in the bulk of the distribution of housing services, we expect to see a decline in (relative) prices in this segment of the market, and following the assignment logic of Section 4.2, a decline in prices of all dwellings offering more housing services as well.

With these difficulties in mind, Soltas (2024b) takes a less direct route and estimates instead the net fiscal cost per dwelling of a New York City’s voluntary IZ mandate using builders’ responses to several changes in the fiscal exemption granted for new developments with affordable units. He estimates that the fiscal cost of the marginal dwelling under this program is 1.6 million dollars, which is many times the cost of other forms of housing assistance. There is, however, a lot of heterogeneity across neighborhoods, with costs of more than two million dollars in parts of Manhattan and less than one-tenth that in many parts of Bronx, Queens, or Staten Island.

These estimates at the lower end of the fiscal cost of new affordable housing are about the same magnitude as the lifetime undiscounted benefits for a child moving from a low- to a high-opportunity neighborhood as reported in Bergman et al. (2024) and building on results from Chetty et al. (2018).⁵⁶ While

⁵⁶ This comparison should not be read as part of a cost-benefit calculation, however crude. The benefits from growing up in a high-opportunity neighborhood are not discounted, and it is not clear that IZ dwellings are occupied by residents who would otherwise live in a low-opportunity neighborhood.

a complete discussion of this type of externality is beyond our scope, we note that we should expect some equilibrium effects from changes in neighborhood characteristics both where IZ residents live and where they are from (or where they would otherwise live). Fogli et al. (2024) and Eckert and Kleineberg (2021) examine these issues in a full general equilibrium setting. Unsurprisingly, they find reduced benefits from moving to high-opportunity neighborhoods in general equilibrium.

Property developers also complain heavily about IZ mandates and point to a direct negative effect of the presence of below-market-rate dwellings on the willingness to pay of market-rate buyers. We are not aware of any research about this issue and do not know if the results of Diamond et al. (2019) about externalities across buildings also apply within buildings. This question is for future research to answer.

The presence of externalities complicates the estimation of the price effects associated with an expansion of the housing stock. With no externality, any increase in the housing stock in a market leads to lower prices in that market and markets that are gross substitutes. With externalities, prices will also reflect how changes in the housing stock make neighborhoods more or less desirable. Housing prices in a market play two roles: they reflect the scarcity of housing in this market and its desirability. The externalities we discuss here are expected to have a spatial decay, but so will supply effects in all likelihood. Hence, it will be challenging to distinguish between the price changes caused by supply effects and price changes caused by external effects. The examples above are cases where increasing the stock of housing increases prices as positive externalities from new constructions dominate the effects of increased supply. Positive externalities making a location more appealing are arguably a good thing, with a caveat about distributional effects. It is also the case that research has focused on situations where improving housing supply led to higher prices to provide evidence about externalities (and likely misses price changes in the market where new residents are coming from).

The literature that examines the effects of new construction on housing prices generally finds negative effects. In a thorough study of how new apartment buildings affect housing rents in nearby apartment buildings, Asquith et al. (2023) find that a new apartment building lowers nearby rents in similar buildings by 5 to 7%, overwhelming any positive externality. Li (2021) finds a similar effect for new residential high-rises in New York City with an elasticity of the price close to -0.1 with evidence of positive spillover effects in the form of new restaurant openings, for instance.⁵⁷ The findings of Li (2021) also underscore that the popular perception that new constructions increase rents is likely driven by developers (rationally) choosing to build where and when prices increase. Pennington (2021) finds again similar results for San Francisco, where new buildings, built (plausibly exogenously) following fire damage, lead

⁵⁷ In a different context, Almagro and Domínguez-Iino (2025) provide a detailed estimation of the resorting and changes in amenities following a supply shock.

to a 1.2-2.3% percent decline in nearby rents. Rents decline despite improving amenities, making (poor) incumbent residents less likely to relocate to another neighborhood. Finally, using German data, Mense (2025) also estimates an elasticity of prices with respect to quantities of about -0.2 . This is lower than an elasticity of about -0.5 that we might expect by reverting the supply elasticity found for construction of about 2 for single-family homes in Section 3. To explain this gap, we note that the supply elasticity for multifamily housing may be higher, that new constructions generate positive spillovers, and that price effects imperfectly percolate across tiers of the housing market.

To assess the full effect of changes in housing supply in a market on prices and welfare everywhere, a complete general equilibrium framework is needed to model the direct price effects, the external effects, the effects on other tiers of the same geographic market, and other markets. Owens et al. (2020) show that the uncoordinated shrinking of Detroit is highly inefficient. Some areas with strong fundamentals may get trapped in low equilibrium following a small negative shock and the external effects of the blight that ensue. Builders and residents are unable to coordinate their actions. No resident wants to be first to move into a vacant neighborhood. No developer wants to be first to invest in a vacant neighborhood. Guerrieri et al. (2013) also develop a general equilibrium framework to understand gentrification. As richer neighborhoods no longer expand their supply of housing, the arrival of rich newcomers leads to the redevelopment of low-quality housing. Because the presence of rich neighbors nearby is desirable to these newcomers, these redevelopments will take place close to rich neighborhoods and will be spatially highly concentrated, leading to what is commonly known as gentrification.

In mild cases, the externalities discussed here can lead to inefficient equilibria with, for instance, neighborhoods suffering from insufficient housing maintenance. When these externalities are hyperlocal, they are increasingly handled by homeowners associations (HOAs) in the US. Clarke and Freedman (2019) estimate that being part of an HOA leads to a 4% higher price for a house. When these externalities are stronger, they lead to the existence of multiple equilibria and coordination failures. Rossi-Hansberg et al. (2010) argue that the housing improvement program they study in Richmond (VA) was designed with some awareness of the need to concentrate the intervention spatially. Similarly, Wheeler (2022) argues that one objective behind the concentration of investments into Opportunity Zones was to resolve coordination failures in blighted neighborhoods.

5 Land use and regulatory constraints

The literature explored in the previous two sections is mostly about the intensive margin of supply, that is, how the quantity of housing services built responds to demand conditions conditional on development. In this section, we dig deeper

into the land development margin. More formally, and returning to the decomposition of housing supply elasticities in Eq. (10), this section focuses on the elasticity of land development with respect to the price of housing between periods t and t' , $\epsilon_p^L(t, t')$.⁵⁸

Land availability has received a lot of attention in the literature as a central driver of variation in local housing supply elasticities. In addition to topography and natural features, a central determinant of land availability is land use regulation. For this reason, our discussion of the role of regulatory constraints in this section naturally pairs with our review of the literature on land development. A starting point is the excellent prior handbook chapter on land use regulation (Gyourko and Molloy, 2015), which we only briefly reprise. Since this chapter came out, there is extensive evidence, reviewed in Gyourko and Krimmel (2021) and shown in Fig. 5 above, that land use regulations in the US have only become stricter between 2006 and 2018. More restrictive regulation and declines in land availability for development have resulted in generally declining housing supply elasticities over time. Land use planning restrictions have proven very difficult to change once “zoning straitjackets” are established, locking in initial development patterns for decades (Ellickson, 2020).

5.1 Bottom-up approaches and the estimation of local supply elasticities

To understand the determinants of the elasticity of land development $\epsilon_p^L(t, t')$, consider a market j platted with lots of the same size, normalized to 1.⁵⁹ Within that market, the fixed cost of land development differs across lots and is described by its probability distribution function, $f_j(x)$, and its associated cumulative distribution function, $F_j(x)$. As discussed in Section 3.1, under perfect competition, land development occurs for all lots with fixed costs below the cutoff $\bar{c}_j(P_j)$ at which the parcel price is equal to the price associated with the lot's next best non-urban use \underline{q}_j . This cutoff is the variable profit given optimal development minus the opportunity cost of the marginal lot:

$$\bar{c}_j(P_j) = \left\{ \frac{d \log C_j[h_j^s(P_j)]}{d \log h} - 1 \right\} C_j[h_j^s(P_j)] - \underline{q}_j. \quad (41)$$

Importantly, we express the cost function as depending indirectly on the price P_j . The resulting fraction of land developed in market j is $F_j(\bar{c}_j)$.

⁵⁸ This is the percentage increase in developed land off of base development in period t for a percentage increase in the local housing services price index up to period t' . Because of the irreversibility of land development for urban uses and associated kinked supply, we tend to think only of additional land developed given price increases. As discussed in Section 4.1, land will typically not leave its developed state in response to price declines.

⁵⁹ The analysis below goes through with variable lot sizes but is messier. The imposition of minimum lot sizes discussed below is an example of a mechanism through which fixed lot sizes may occur.

Unsurprisingly, an increase in the price of housing P_j results in more land development since the higher marginal revenue allows builders to afford higher fixed costs. The amount of extra land developed depends on the thickness of the tail of the fixed cost distribution on the margin. A thicker tail, formally a larger $f_j(\bar{c}_j)/F_j(\bar{c}_j)$ leads to more land development given an increase in P_j . In turn, the thicker the tail, the higher the elasticity of land development.

Baum-Snow and Han (2024) consider the same Cobb-Douglas production function of housing with capital output elasticity α as in Eq. (15), an opportunity cost of land q_j normalized to 0, and Fréchet distributed fixed costs with distribution function $F_j(x) = \exp(-\Gamma_j x^{-\lambda})$. Each market j has its own (inverse) scale parameter $\Gamma_j > 0$, but the shape parameter $\lambda > 1$ is common to all markets.

The land development elasticity of market j then takes the form

$$\begin{aligned}\epsilon_p^L &= \frac{d \log F_j(\bar{c}_j(P_j))}{d \log P_j} = \left\{ \frac{d \log C_j[h_j(P_j)]}{d \log h} - 1 \right\} \epsilon_p^h \frac{f_j(\bar{c}_j(P_j))}{F_j(\bar{c}_j(P_j))} P_j h_j^s(P_j) \\ &= (1 - \alpha)^{-1 - \lambda} \lambda \rho_j^{-\lambda} P_j^{-\frac{\lambda}{1 - \alpha}} \Gamma_j,\end{aligned}\quad (42)$$

where the equalities on the first line hold for any regular production function. With Cobb-Douglas, the two terms that contain elasticities cancel out to give the expression on the second line, where ρ_j is a cluster of parameters that is increasing in the productivity of housing B_j .

As the price P_j of housing or its productivity B_j increases, there is a depletion effect of moving further into the thinning tail of the fixed cost distribution, thereby reducing the elasticity of land supply. An empirically observable proxy used by Baum-Snow and Han (2024) is the fraction of already developed land in the market.⁶⁰ In contrast, as the Fréchet (inverse) scale parameter Γ_j rises, the fixed cost distribution for market j has a higher mean and variance, thereby thickening the tail of the fixed cost distribution for any given fixed cost. Possible empirical proxies for Γ_j include topographical features, such as slope and ruggedness, and observed land use regulations, measured, for example, by the 2006 Wharton Residential Land Use Regulation Index (Baum-Snow and Han, 2024).

In Eq. (42), differences in land supply elasticities across markets follow from differences in prices and differences in the distributions of fixed costs, even given a common underlying housing production technology. Saiz (2010) was the first to formalize the idea that land availability can be a key force regulating housing supply elasticities. Saiz (2010) develops a motivating monocentric model in which the cost of acquiring land for development in a city increases with the unavailability of land in its metropolitan area, thereby generating reduced housing supply elasticities. This land constraint mechanism for

⁶⁰ Cosman et al. (2018) proposes a dynamic monocentric model with variable marginal land availability with some of the same intuition.

generating variation in land supply elasticities is fundamentally the same as that studied through Eq. (42).

Saiz (2010) proxies for land unavailability with the fraction of each metropolitan area that has a steep slope of over 15% or is covered by water within 50 kilometers of the central business district, since it is not possible (or extremely hard) to build in such areas. Constraints imposed by geography have also been used in various ways as direct proxies for metropolitan-level supply elasticities, with the idea that housing-demand shocks in more constrained metro areas will result in more rapid price increases due to less elastic housing supply (e.g. Diamond, 2016).

Estimating housing supply elasticities comes down to relating housing prices and quantities using housing demand shocks as sources of identifying variation while also allowing for land availability and regulation to vary flexibly across markets. The estimation is typically performed for one long first difference. Commensurate with various decompositions of changes in housing quantities discussed in Section 2, Baum-Snow and Han (2024) examine the 2000-2010 supply responses of total floorspace and total dwelling units to changes in prices. These responses are decomposed into additive components that come from flows of new constructions on already developed land (redevelopment) ($\Delta S_j^R > 0, \Delta N_j^R > 0$), other new constructions ($\Delta S_j^U > 0, \Delta N_j^U > 0$), tear-downs ($\Delta S_j^T < 0, \Delta N_j^T < 0$), and renovations ($\Delta S_j^E, \Delta N_j^E$), respectively. For floorspace, this decomposition can be written as

$$\Delta \log S_j \equiv \frac{\Delta S_j}{S_j} = \frac{\Delta S_j^R}{S_j} + \frac{\Delta S_j^U}{S_j} + \frac{\Delta S_j^T}{S_j} + \frac{\Delta S_j^E}{S_j}. \quad (43)$$

The difference between the floorspace elasticity, $\Delta \log S_j / \Delta \log P_j$, and the land development elasticity, $\Delta \log L_j / \Delta \log P_j$, yields the intensive margin component of the total housing supply elasticity, $\epsilon_p^h(t, t')$ in Eq. (10).

Baum-Snow and Han (2024) use microdata on local property assessments, compiled by Zillow in the ZTRAX data set, to measure housing quantities in Eq. (43) and to construct tract-level housing price indices. For each component of housing supply, including those on both sides of Eq. (43), they estimate regression equations of the form

$$\Delta Q_j^s = \gamma(Z_{j2000}) \Delta \log P_j + X_{j2000} \phi + u_j, \quad (44)$$

where ΔQ_j^s is a housing quantity change in market j , including the growth rate in floorspace and its components, the growth rate in dwelling units and its components, and the growth rate in developed land. The vector of variables Z_{j2000} , which is interacted with the changes in housing prices, ΔP_j , contains variables that are either time-invariant (including the distance to the CBD and the local topography) or measured for the base year (the fraction of land developed). The vector X_{j2000} controls for potential supply shifters that should include variables

in Z . The estimated vector of coefficients $\gamma(Z)$ is the supply elasticity function. The simplest specification for the interaction is linear in Z . More flexible specifications indexing parameters by one of two latent classes can also be considered.⁶¹

The primary identification challenge is to find instruments that affect $\Delta \log P_j$ by shifting housing demand while remaining uncorrelated with the error term u_j . That is, these instruments should be orthogonal to the cost of construction and land availability, conditional on controls X . To help identification, Baum-Snow and Han (2024) include metro area fixed effects in the specification. These controls require the identifying variation to come from comparisons between census tracts in the same metro area conditional on the distance to the CBD, topography, and the fraction of developed land.⁶²

A common approach is to use Bartik-style shift-shares as instruments. In the simplest case, the growth of the employment of the national industry, interacted with the shares of local employment by industry, likely predicts the growth in the local demand for housing. However, such a predictor, when computed for metropolitan areas, is collinear with metro area fixed effects and, in their absence, it is unlikely to be orthogonal to the unobserved drivers of land or housing development in regression (44). Rather, Baum-Snow and Han (2024) and Büchler et al. (2021) use industry shift-share shocks in potential commuting destinations filtered through the commuting time matrix. The idea is that labor demand shocks in labor markets outside of market j generate shocks to housing demand in market j if the former can be accessed from market j through commuting. This approach can be justified with a quantitative spatial model. These housing demand shifters are then interacted with Z to form a full set of instruments.⁶³

In separate reduced-form price and quantity growth regressions, Gorback and Keys (2023) uses foreign investment demand shift-shares across countries of origin as exogenous demand shocks, with time variation driven by changes in foreign buyer taxes in countries outside the US. In the reverse regression of (44), Saiz (2010) uses such industry shift-shares along with immigration flows as instruments for quantity growth.

⁶¹ The most flexible specification has $\gamma_{ir} = \Lambda(Z_r^1)Z_{ir}^2\mu_1 + (1 - \Lambda(Z_r^1))Z_{ir}^2\mu_2$, where $\Lambda(Z_r^1)$ is the logit probability of being in latent class 1, which depends on three variables in Z_r^2 measured at the metro region r level: the 2006 metro area Wharton Index (WRLURI), the fraction of land within 50 kilometers of the CBD that was developed in 2001, and the fraction of space unavailable for development due to steep slopes or water coverage within 50 kilometers of the CBD. Z_{ir}^2 are tract-level measures of CBD distance, the fraction of flat land area, and a quadratic of the fraction of land area developed in 2001.

⁶² Saiz (2010) estimates the inverse of regression (44), which delivers estimates of inverse-supply elasticities for the number of occupied dwellings. Except potentially for the calculation of standard errors, the estimation of Eq. (44) or its inverse are mathematically equivalent, but they differ econometrically.

⁶³ One advantage of the Bartik style approach is that “placebo” checks can be performed by verifying that instruments do not predict changes in quantities or prices prior to the period of study.

For empirical settings in which regulation explicitly enters as a supply factor, another identification challenge is to find exogenous variation in such regulation. For this, Saiz (2010) uses the share of the population that belongs to non-traditional Christian denominations, with the idea that this group is more averse to government regulation, and Hilber and Vermeulen (2016) uses a discrete national policy change that lifted some limits on construction as sources of identifying variation.

Three further issues need to be discussed. First, given the irreversibility of housing investment (Glaeser and Gyourko, 2005), inward housing demand shocks would tend to identify very small housing supply elasticities, as we expect the supply curve to be vertical below the current level of supply. Hence, the identification of supply elasticities requires isolating variation in price growth in Eq. (44) that is different for inward and outward housing demand shocks.

Second, the theory developed throughout this chapter has treated a “market” j as a set of dwellings that all face the same demand conditions. In equilibrium, they are perfect demand substitutes, and they face the same equilibrium price per unit of housing services P_j . For this condition to hold, markets must arguably be small. For this reason, Baum-Snow and Han (2024) use US census tracts and Büchler et al. (2021) use 2×2 -kilometer tiles in Switzerland as spatial units of analysis. Under certain aggregation assumptions discussed below, one can rationalize the use of broader spatial units that include areas with different demand conditions. Other work estimating housing supply elasticities uses metropolitan areas (Saiz, 2010; Aastveit et al., 2023), ZIP codes (Gorback and Keys, 2023), and English local planning authority regions (Hilber and Vermeulen, 2016) as spatial units.

Third, satellite-based measures of land development are imperfect. The US National Land Cover Database codes each 30×30 -meter pixel into one of four urbanized categories or one of various unurbanized categories. Urbanized categories are assigned on the basis of the ranges of the percentage of each pixel covered in impervious surfaces. As a result, small changes in land development are usually not observed in these data. This may cause studies using this data set to underestimate land development elasticities. As some redevelopment occurs on pixels with low levels of initial urbanization, estimated redevelopment elasticities could also be understated.

All of the studies mentioned above find that more land available for development and less stringent regulations predict more elastic housing supply. Baum-Snow and Han (2024) find that these attributes matter both at the census tract and metropolitan area levels, so that all markets in more highly regulated and land-constrained metropolitan areas are themselves more supply-constrained, even conditional on land unavailability at the tract level.

Housing supply elasticities have fallen over time. For example, across 237 metro areas, Saiz (2010) finds an average supply elasticity for the number of occupied dwellings of 2.6 (1.6 weighted by the size of metropolitan areas) for the 1970-2000 period. Using Saiz’ methods and price and quantity measures for the

2000-2010 period instead, this number falls to 1.3 (1.1 weighted by metro size) (Baum-Snow and Han, 2024). Using a panel of metropolitan areas, Aastveit et al. (2023) finds declining supply elasticities over time of similar magnitudes. For these same metro areas, Baum-Snow and Han (2024) finds a supply elasticity for dwelling units of 0.5 for the 2000-2010 period (0.3 weighted by metro size). For the 2010-2020 period, Gorback and Keys (2023) find elasticities that are smaller by another 10%. Rapidly declining rates of new housing starts, as seen in Figs. 3 and 5 above, also fit this evidence of declining estimated supply elasticities. A key open question in the literature is to what extent these declines in housing supply elasticities can be attributed to reductions in land available for development (or, equivalently, the selection effect of moving further into the tails of local fixed development cost distributions), increases in land use regulation, and/or something else. In fact, both land unavailability and land use regulation have increased over time.

Returning to the decompositions in Eqs. (10) and (43), Baum-Snow and Han (2024) finds supply elasticities for average floorspace and the number of dwellings of 0.5 and 0.3, respectively, across all urban neighborhoods in the US. New constructions represent 69% of the supply response for average floorspace and 54% for the number of dwellings. In both cases, the remainder is approximately equally split between reduced teardowns and expansions.

The average estimated elasticity of land development, $\epsilon_P^L(2000, 2010)$ is about 0.1, and the intensive margin elasticity $\epsilon_P^h(2000, 2010)$, computed using Eq. (10), is thus up to 0.4. The average estimated redevelopment elasticity is very low at 0.03. This tiny elasticity is consistent with the very small housing investment response to rent decontrol in Cambridge (MA) in 1995 documented in Autor et al. (2014). It is also consistent with the small changes in densification in New York City following changes in maximum floor-to-area ratios studied by Peng (2023) and discussed above.

Another striking feature of supply elasticities is that they differ across space within metropolitan areas, actually even more than they vary across metropolitan areas. As redevelopment is estimated to be very costly (as discussed in Section 4.1), most housing supply expansions must come on undeveloped land. Moreover, as seen in Fig. 5, inner suburbs tend to have quite restrictive land use regulations. The result is low floorspace supply elasticities of about 0.1 near CBDs, where available land is scarce. These supply elasticities then rise to 0.6 at 25% of the way to the edges of metropolitan areas before inching up further and reaching 0.75 at urban peripheries, where loosely regulated land remains available for development. Similar profiles are seen for the number of dwellings and new constructions.

A difference among the studies discussed above is the time length over which elasticities are measured. Given how slow the redevelopment process is, ten years may not be enough to be considered “long-run”. To go deeper, note that, commensurate with Eq. (8), a steady state requires:

$$H = (1 - \delta_T)H + A P^{\epsilon_P^H(t, t+T)}, \quad (45)$$

where δ_T is the T year depreciation rate and A is a supply shifter. We then define the long-run supply elasticity from the comparison of two different steady states with price difference $\Delta \log P$.⁶⁴ After solving for the steady state we find the long-run elasticity to be,

$$\frac{\Delta \log H}{\Delta \log P} = \frac{\epsilon_P^H(t, t+T)}{1 - (1 - \delta_T)}. \quad (46)$$

Glaeser and Gyourko (2005) find an annual depreciation rate of 3.5%. Using supply elasticity estimates over 10 years, we have $1 - \delta_{10} = (1 - 0.035)^{10} = 0.7$. So, a long-run elasticity is 3.33 times a 10-year elasticity and, by the same logic, 1.5 times a 30-year elasticity. After accounting for the length of time, the relative magnitudes of Saiz-type elasticities discussed above are quite similar for the 1970-2000 and 2000-2010 periods in elastic metropolitan areas but have fallen quickly in the relatively inelastic larger metropolitan areas.

Spatial aggregation can also matter. Metropolitan areas contain less supply-elastic central areas and more elastic suburbs. If the identifying variation for a metropolitan area is primarily from suburban submarkets, where more land is available for development, the estimated supply elasticity is biased upward. To be more precise, consider regional market r , which aggregates a number of submarket j . The log change of housing supplied in this market can be decomposed as $\Delta \log H_r = \sum_{j \in r} \omega_j \Delta \log H_j$, where ω_j are submarket weights that add to one. After replacing for each submarket the change in quantity $\Delta \log H_j$ by the equilibrium response from the change in price, $\epsilon_{Pj}^H \Delta \log P_j$, we can compute the supply elasticity for the aggregated regional market r as,

$$\epsilon_{Pr}^H \equiv \frac{\sum_{j \in r} \omega_j \epsilon_{Pj}^H \Delta \log P_j}{\sum_{j \in r} \omega_j \Delta \log P_j}. \quad (47)$$

It is standard to specify $\omega_j = H_j/H_r$. If markets are linked in a residential demand system, demand shocks are correlated across markets. In turn, this means that the aggregation of supply elasticities across markets depends crucially on the nature of the underlying demand shock.

If markets are close substitutes, aggregate elasticities have a natural interpretation as average market elasticities. In the edge case of perfect substitutes, any price change must be the same in all markets. From Eq. (47), the aggregate elasticity is thus the weighted average of elasticities across markets in r . If markets are segmented, the aggregation is more complicated. The nature of the demand shock matters, and demand shocks that disproportionately hit more

⁶⁴ To help calibrate a model predicting future rent and price growth across housing markets after the rise of work from home, Howard et al. (2023) develop these ideas to construct long-run supply elasticities for neighborhoods across the US using decadal estimates from Baum-Snow and Han (2024).

elastic market segments will result in elasticity estimates that are greater than the aggregate elasticity.⁶⁵

Market segmentation also occurs between the owner-occupied and rental markets. Greenwald and Guren (2021) demonstrate that the growth of housing demand driven by shocks to the supply of mortgage credit mostly affects the owner-occupied segment. In turn, these shocks increase the price-to-rent ratio of housing with little effect on the homeownership rate. That is, low aggregate housing supply elasticities can be understood in part through the limited conversion of rental units to owner-occupancy. Consistent with our discussion in Section 4.3, housing units do not easily “filter up”.⁶⁶

After estimating supply elasticities in many markets, it is tempting to use them as instruments to predict housing prices. Davidoff (2015) provides evidence that, indeed, supply elasticities do predict price levels and growth. However, he also shows that supply elasticities are also correlated with the demand for housing. More supply-constrained markets tend to have stronger demand, perhaps because natural features that constrain supply also tend to be positive amenities. Moreover, demand growth tends to beget more stringent land use regulations. Hence, supply elasticities are not good instruments for prices in settings in which the goal is to estimate parameters governing housing or residential demand.

To avoid this problem, Lutz and Sand (2023) propose the following alternative instrument for house price growth. They train the XGBoost machine learning algorithm for house price growth at the ZIP code level using predictive features built from satellite-based measures of land unavailability in and around each ZIP code and its CBSA. This model is trained on 80% of ZIP codes and is used to predict price growth in the remaining 20%. Repeating the process five times delivers out-of-sample predictions for all areas, which are then bootstrapped four more times. The resulting average predicted price growth at the ZIP code level can then be used as an instrument for actual price growth. Lutz and Sand (2023) provide evidence that this instrument is not correlated with local demand factors.

Supply elasticity estimates are commonly used to make predictions about the home price responses to specific demand shocks. Using recently estimated supply elasticities, Davis et al. (2024) find that the growth in housing demand

⁶⁵ Imagine two equally-sized fully demand-segmented neighborhoods *A* and *B* with supply elasticities of 1 and 2, respectively. Neighborhood *A* is richer, while neighborhood *B* is poorer with more migrants. An employment shift-share may increase housing prices by 1% in both markets. In response, supply increases by 1% in *A* and 2% in *B*. In this case, the aggregate supply elasticity is 1.5. A second source of identification, such as an immigration shock, may boost prices by 2% in market *B* only, where it leads to a 4% growth in housing supply. In this case, the average price increase is 1%, and the average quantity increase is 2%, resulting in an estimated aggregate supply elasticity of 2.

⁶⁶ Diamond and McQuade (2019) examine the end of rent control in San Francisco in 1994. They find that rental supply was constrained by about 15% as a result of rent control. Some of the ‘disappearing’ rental units nonetheless “filtered up” due to building conversion to owner-occupancy.

associated with the rise of remote work accounts for a large fraction of house price growth after the COVID-19 pandemic, Couture et al. (2024) characterize the welfare costs of gentrification for incumbent renters, and Favilukis and Van Nieuwerburgh (2021) assess the importance of growth in foreign investor demand in the overall growth in the price of housing. Conclusions in these and other studies depend crucially on the flexibility to allow supply elasticities to be heterogeneous within and/or between different metropolitan areas.

By definition, housing supply elasticities are aggregate quantities that measure changes over many types of housing markets with widely differing characteristics. The identification of these elasticities often relies on broad-based exogenous variation in housing demand shocks. In turn, these elasticities are determined by both geography and regulations. Separating between both types of effects remains challenging, given that areas with less land available for development tend to have stricter land use regulations. As a result, much remains to be learned from a closer look at how different specific types of regulation influence housing costs and local amenities. The remainder of this section reviews the recent literature on the consequences of land use regulations while incorporating conceptualizations of and evidence of motivations behind their enactment.

5.2 Unpacking land use regulations

A central challenge for research on regulation is that land use and real estate development restrictions take many forms. Each municipality's zoning code is different and requires a unique process to navigate. Even with "as-of-right" developments that do not require zoning variance, developers must typically work closely with city planners to stay within rules that differ between jurisdictions. In many jurisdictions, this involves going through an uncertain, time-consuming, and costly process of permitting, inspection, and environmental reviews. In addition, zoning codes typically prescribe the land available for each specific type of development through zoning designations, and the intensity of development through density restrictions.⁶⁷ Then, the fraction of land or properties that require a more arduous process to develop or redevelop varies within each jurisdiction. Additional restrictions, such as requiring historic preservation, banning alcohol sales, or preserving open space, also limit developer flexibility. Development projects that require zoning variances typically must first undergo a series of uncertain reviews that require broad public support to pass, further influencing the time, cost, and probability of success. In all of these ways, stricter regulatory environments raise developers' costs, thereby restricting housing supply and reducing its price elasticity.

Given the many forms of regulations governing land use and real estate development, building succinct and informative summary measures of them is a

⁶⁷ Shertzer et al. (2018) and McMillen and McDonald (2002) demonstrate that most aspects of Chicago's original 1923 zoning code and land use planning regime persist to today, durably influencing land use and housing values.

challenge. As already mentioned, the Wharton Residential Land Use and Regulation Index (WRLURI) is the most widely used summary measure encompassing the difficulty of development and restrictiveness of land use regulations.⁶⁸ The 2006 version of the WRLURI (Gyourko et al., 2008) is derived from a 15 multi-part question survey sent to the 6,896 US municipalities with available contact information listed by the International City Managers Association. The 2,649 responding municipalities tend to be larger, including 62% of the 241 municipalities with a population over 100,000. Only 28% of the 1,969 municipalities with a population under 5,000 responded.

Eleven sub-indices are calculated to capture the restrictiveness due to approval delays, local political pressure, state political involvement, density restrictions, local project approval, open space requirements, local assembly approval, supply restrictions, exactions, state court involvement, and local zoning approval.⁶⁹ For the 2,611 municipalities with sufficient data, these sub-indices are then combined to compute the WRLURI as their first factor using factor analysis. The WRLURI is highly correlated with the sum of its standardized component indices. Factor loadings on sub-indices are in the same order as listed above, from largest positive to slightly negative. The WRLURI is standardized to have a mean of 0 and standard deviation of 1 across the jurisdictions with enough data to compute the WRLURI.

Characterizing the national or metropolitan distribution of regulation requires the use of weights reflecting the probability of being sampled. These weights are constructed as the inverse predicted probability of being in the data from a logit regression of survey return on various municipality characteristics, estimated using the universe of census-designated places. Weighting can be used to correct for the fact that larger and more regulated locations were more likely to be sampled. Indeed, rural municipalities have an average WRLURI of -0.46 ; unincorporated areas, which tend to have the laxest land use regulation, were not surveyed. More regulated municipalities tend to have higher median incomes and housing values and are better educated. The most regulated municipalities also tend to have lower population densities.

The 2018 WRLURI is constructed using a similar set of sub-indexes, with the addition of an affordable housing sub-index, based on survey responses from 2,472 mostly suburban municipalities (Gyourko and Krimmel, 2021). Factor loadings on each sub-index are somewhat different for 2018, with all of them positive. As the resulting 2018 WRLURI is also standardized, the two indices for 2006 and 2018 are not directly comparable, though metropolitan area rankings remain quite stable. Gyourko and Krimmel (2021)'s comparisons of common

⁶⁸ See Jackson (2018) for a discussion of other land use regulation indices that have been used for the US.

⁶⁹ So as to be comparable across municipalities, these questions ask city planners to rate various aspects of the zoning and approval process rather than for details specific to their jurisdictions. For example, one set of questions takes the form “On a scale of 1 to 5, please rate the importance of each of the following factors in regulating the rate of residential development in your community”. Two of the eleven factors subsequently listed are “supply of land” and “density restrictions”.

questions across the two surveys reveal that regulation is stable or increasing over time, with only a few lightly regulated areas as of 2006 that reduced regulation.

Multiple efforts are currently underway to develop further disaggregated measures of land use regulation. One dimension of disaggregation is in expanding the list of variables used to measure the attributes of a jurisdiction's zoning code. A further dimension of disaggregation is the partitioning of space within jurisdictions into zoning districts, developing a set of measures for each district.

Following the first approach, The National Zoning Atlas project (Bronin et al., 2023) is in the process of painstakingly coding up a set of consistent variables by carefully reading zoning codes from multiple jurisdictions, with teams of researchers assigned to read municipalities' zoning codes in different regions around the US toward this goal. This effort nationalizes the prior Pioneer Institute initiative in 2004 to consistently code the zoning codes of 187 municipalities in the Greater Boston area (Glaeser and Ward, 2009).

Such manually coded data sets can be used as training data in the development of large language models (artificial intelligence) to automate the reading of zoning codes. The advantage is scalability, though it potentially comes with some loss of accuracy. For example, Bartik et al. (2024) uses a large language model trained on the Pioneer Institute data and verified against elements of the 2018 WRLURI to predict key aspects of the zoning code for 25% of US municipalities, covering 63% of the US population. The analysis feeds each municipality's zoning code into a trained large language model (OpenAI GPT-4) and asks the same survey questions as in the Pioneer Institute study. This study then relates local supply growth to aspects of local zoning codes.

Echoing the evidence in Baum-Snow (2023) and Burchfield et al. (2006), Bartik et al. (2024) find that a disproportionate amount of new housing development is greenfield development in unincorporated exurbs. Incorporated locales restrict densities through minimum lot size restrictions and bans on multifamily housing at very high rates. Two-thirds of incorporated locales have minimum lot sizes of at least 5,000 square feet and only 31% of the land area is zoned for multifamily housing. Using the AI-generated answers to questions about the zoning code as inputs, Bartik et al. (2024) undertake a revealing principal components analysis to draw out central attributes of zoning codes. The first principal component reflects regulatory complexity, which increases developers' fixed costs. This attribute of zoning codes is more prevalent in higher-density areas of central municipalities. The second principal component encompasses policies that restrict development or exclusionary zoning. This attribute is more prevalent in high-income suburban areas.⁷⁰

⁷⁰ In related work, Shanks (2021) uses the Latent Dirichlet Allocation unsupervised learning procedure to assign each municipality in Massachusetts zoning code probabilities of membership in various clusters generated based on the prevalence of various words in the zoning code, with three clusters chosen ex-ante. Using simpler keyword-based natural language processing on zoning codes in most 2006 WRLURI municipalities to roughly replicate WRLURI sub-indices, Mleczko and Desmond (2023) also construct an alternative zoning index.

As real estate development restrictions are so heterogeneous, many studies focus on recovering the consequences of one common type of zoning restriction. In addition to maximum floor-to-area ratios explored in Section 3.5, which are prevalent around the world, one very common regulation in the US is minimum lot size (MLS) zoning. Zabel and Dalton (2011) pioneered the idea of automating the measurement of local MLS restrictions using the modes of observed lot size distributions or equivalently by isolating the lot size at which the distribution function in each jurisdiction exhibits a structural break. This study shows a positive relationship between these inferred MLS and property prices in Eastern Massachusetts. Song (2024) hones this method to characterize the distribution of MLS restrictions in zoning districts across the US. She finds that for the 16,217 municipalities for which MLS regulations can be inferred, the mean MLS is 0.37 acres (16,000 sq ft), and 75% of municipalities impose a MLS of at least one acre in at least one zoning district or census block group.

One fruitful way to summarize the stringency of MLS constraints given demand conditions, without even having to measure minimum lot sizes directly, is through the calculation of “regulatory taxes” (Glaeser and Gyourko, 2018). The idea is to compare the hedonic prices of a marginal unit of land as viewed by builders and homeowners. If the value of a unit of land to a homeowner is below its value to builders, we would expect the homeowner to subdivide and sell off part of her lot through a standard no-arbitrage argument. As a result, the greater the gap between the price of vacant and built-up land, *ceteris paribus*, the greater the regulatory distortion that prevents this arbitrage.

Using data on vacant land sales for development into single-family homes from CoStar and property transactions data from CoreLogic, Gyourko and Krimmel (2021) carry out a comprehensive analysis of regulatory taxes for 24 large US metropolitan areas. The study uses standard hedonic methods to recover homeowners’ implicit marginal willingness to pay for land by US metropolitan area. The study period is 2013-2018. Their results indicate regulatory taxes of 400,000 dollars per half-acre in San Francisco, 150,000-200,000 dollars per half-acre in New York, Los Angeles, and Seattle, and smaller numbers for other markets studied. Relatively elastic housing supply cities, including Atlanta, Charlotte, Dallas, and Cincinnati, are found to have very small regulatory taxes.⁷¹

5.3 Estimating the consequences of land use regulation: border discontinuity approaches

Land use restrictions generate winners and losers. Restrictive land use regulations reduce the option value of development by limiting the production of housing services from a parcel. However, restrictive regulations also protect

⁷¹ In the same spirit is Duranton and Puga’s 2023 identification of highly regulated metro areas by looking at property values at urban fringes. At the upper end, their estimated values are close to those of Gyourko and Krimmel (2021).

against nearby developments and their potential external costs. This external effect follows the classic compensating differentials logic for local amenities, as in Roback (1982). In addition, restrictive regulations reduce aggregate supply and thus also increase property values through this channel.

Because of their effect on property values, land use regulations are likely to reflect voter interests and, therefore, be endogenous to demand conditions. As a result, credible studies of the impacts of land use regulations require the use of identifying variation in land use regulation that is orthogonal to demand conditions. To resolve this endogeneity problem, one arguably credible identification strategy is to make comparisons across zoning districts at small spatial scales for which demand conditions are the same, a boundary discontinuity approach.

In an important paper, Turner et al. (2014) exploit discontinuous changes in land use regulations across municipal boundaries to recover the effects of these regulations on the values of undeveloped lots. When crossing a boundary between two jurisdictions, the regulation changes but, right at the boundary, the existing or expected built environment is approximately the same on either side. Hence, the discontinuity in land values at the boundary primarily reflects how stricter land use regulations reduce the option value of development of a parcel, which they call the “own-lot effect”. As developers are assumed to be competitive, the full foregone value associated with restricting development gets capitalized into reduced land values rather than reduced developer profits, following the logic of Eq. (13) derived in Section 3.1.

While the “external effect” resulting from the regulation of nearby properties on land values is approximately the same at the boundary, exposure to stricter or less strict regulations changes as we move inside each jurisdiction. Hence, the comparison of land values within a jurisdiction between parcels close to the boundary and those far from the boundary, which are subject to the same own-lot effect, can thus isolate the external effect of land use regulations.

Conceptually, stricter land use regulations can have ambiguous effects on land values. They are expected to reduce the option value of development (negative own-lot effect) but increase amenity values from regulations on other parcels (positive external effects). Through the “supply effect”, stricter land use regulations also contribute positively to land values. Using data on vacant land sales mostly from peripheries of US metropolitan areas, Turner et al. (2014) estimate an own-lot effect of minus one-third of value per standard deviation increase in the 2006 WRLURI and an imprecise and much smaller external effect. With negligible estimated supply effects, the conclusion is that land use regulations reduce land values. Of course, the magnitude of the supply effect in practice depends on the spatial reach of land use regulation within the market.

Negative estimates for own-lot effects provide direct evidence that MLS zoning constrains development. Using Eq. (16) with Cobb-Douglas housing production, the elasticity of housing services and property value with respect to lot size is $\frac{\beta}{1-\alpha} \approx 1$, reflecting approximately constant returns to scale. Holding the supply and external effects constant, the elasticity of parcel value per unit of land

with respect to lot size would thus be approximately 0, given small fixed development costs and no constraints on development. A negative estimated own-lot effect thus means that capital intensity, or the quantity of structure built per unit of land, is constrained to decline in lot size. This decline likely reflects developers' optimal responses to prohibiting the construction of additional dwelling units on larger lots.

Given this negative own-lot capitalization effect, it is worth considering why minimum lot size MLS regulations exist. One possibility is that the supply effect is large enough that laxer land use regulation drives down land values through general equilibrium forces, thereby incentivizing property owners to support them. This logic is considered in Section 6. In addition, negative externalities may be large enough in some circumstances to outweigh the increase in real option values from the own-lot effect as regulations are relaxed. There is evidence, discussed below, that negative external effects of dwelling unit density may be larger and own-lot effects smaller in already built-up neighborhoods than in the undeveloped areas studied in Turner et al. (2014). Costly redevelopment limits the option value of already developed parcels. Moreover, the external effects of increased development in already built-up areas are current, whereas those in unbuilt areas mostly reflect future spillovers that are discounted back to the present.

More recent studies use boundary discontinuity approaches to estimate own-lot, external, and supply effects for developed properties rather than vacant land. The use of developed properties has the advantage of covering a much wider and more representative range of contexts and locations. There is also the disadvantage of more challenging identification given potential unobserved heterogeneity in developed properties. Regressions like the following, using a sample of home transactions of properties close to zoning boundaries with boundary fixed effects, underlie this evidence.

$$\log(P_{jt}h_{ijt}) = \kappa^{\text{MLS}} \log \text{MLS}_{jt} + \rho_{b(i)} + \lambda_{m(j)t} + u_{ijt}. \quad (48)$$

Eq. (48) describes a regression in which the dependent variable is the log of the transaction price of home i in zoning district j at time t , $\log(P_{jt}h_{ijt})$. The explanatory variable of interest is the log of minimum lot size, $\log \text{MLS}_{jt}$. The estimated coefficient, κ^{MLS} , captures the impact of the regulation on the total value of housing services on each lot. The regression further conditions out boundary fixed effects, $\rho_{b(i)}$, and municipality-time fixed effects, $\lambda_{m(j)t}$. Boundary fixed effects control for location and local neighborhood amenities. Boundary distance and/or local demographic controls are sometimes also included to improve accounting for such factors. Municipality-time fixed effects $\lambda_{m(j)t}$ control for local public goods and finances. Some studies must leave these out, as they only have more aggregated data with one observation per jurisdiction.

Using regressions in the spirit of Eq. (48), Song (2024), Kulka et al. (2022), and Gyourko and McCulloch (2024) all find that more binding MLS regulations

increase average dwelling unit prices and rents. Song (2024) finds that doubling the MLS increases dwelling prices by 14% and rents by 9% using data from across the US. Using data from Massachusetts, Kulka et al. (2022) find similar estimates. In addition, they find that MLS increases dwelling unit size, with an estimated floorspace elasticity with respect to MLS of about 0.25, and reduces dwelling unit density. Looking across municipal borders, Gyourko and McCulloch (2024) find price elasticities of MLS of about 0.27 for single-family homes.⁷² If developers were unconstrained, the price elasticity would be one.

Positive estimates of κ^{MLS} reflect a combination of differences in prices and quantities of housing services. By adding property characteristics and lot size to Eq. (48), Song (2024) finds that the coefficient on log MLS declines by about three-quarters, indicating that quantities account for at least three-quarters of the effect. Residual price effects could occur due to (likely small) differences in local amenities across zoning district boundaries or unobserved housing quality differences.

This evidence of positive property value and housing quantity effects along with Turner et al.'s (2014) evidence of negative own-lot effects for land echoes earlier cross-sectional results in Ihlanfeldt (2007), finding own-let, external, and dwelling size responses to stricter land use regulation using data from 100 Florida cities. Been et al. (2016) analyze the impact of historic preservation districts in New York City. They uncover positive property transaction price responses within these districts and just outside for outer boroughs where the option values of redevelopment are lower because of lower demand. These districts are local amenities but they also restrict redevelopment.⁷³

The main goal in Gyourko and McCulloch (2024) is to estimate the external effects of density, or the willingness to pay for lower density by different types of households. They propose a hedonic regression like in Eq. (48) that additionally controls for own-lot property characteristics and replaces $\log \text{MLS}_{jt}$ with a flexible function of dwelling density within 500 meters of the focal dwelling unit i . Differentiating the estimated dwelling density function reveals that residents are willing to pay a mean of 9,500 dollars to avoid an additional dwelling unit per two acres of surrounding land, with this amount increasing in magnitude with neighborhood income. Indeed, heterogeneity in the marginal willingness to pay for low density really matters for understanding motivations to limit density. Sahn (2021) and Cui (2024) provide evidence that many MLS restrictions were enacted with exclusionary racial motives, which may also reflect exclusionary income-based motives. This is consistent with the pervasive evidence of negative external effects from neighbors with lower income and education and of

⁷² More precisely, across boundaries for 2018 WRLURI municipalities, Gyourko and McCulloch (2024) estimate average differences in MLS of about 3,000 squared feet on a mean of 12,574 and in sales prices of about 30,000 dollars on a mean of 469,000 dollars.

⁷³ In a similar empirical setting, Kahn et al. (2010) and Severen and Plantinga (2018) study the California Coastal Act of 1976, which restricts additional property development on a strip of land adjacent to the Pacific coast. Comparing property sales and rents across the boundary that defines the restricted region, these papers both find evidence that this regulation improves local amenities.

other races (Bayer et al., 2007; Diamond et al., 2019; Almagro et al., 2023). Residents' concerns about negative externalities can influence regulations, which in turn influence supply conditions. Kulka (2019) demonstrates that higher-income households sort into areas with larger minimum lot sizes, also consistent with an exclusionary motive for their enactment.

There is less work using boundary discontinuity research designs to measure the supply effect. Doing so requires combining estimates of the impacts of regulation on housing quantities with a model that has sufficient structure to allow for recovery of welfare consequences. Anagol et al. (2025) carries out such an exercise, using the 2016 relaxation of density constraints in various neighborhoods in São Paulo, Brazil, as a source of exogenous variation in supply. São Paulo is a good setting for this as strong demand means that many neighborhoods face binding constraints regarding maximal allowed construction. In areas where regulations were relaxed, there was a large supply response. As areas experiencing no change in regulation were already mostly built to the maximum density and experienced no change in construction with the implementation of this policy, there is little possibility that new construction in newly deregulated areas came from substitution across zoning district boundaries. Through the lens of a quantitative spatial equilibrium model of São Paulo's housing markets estimated in part using the identifying variation from boundary discontinuities, the authors find large housing wealth transfers from current to future homeowners, with a net welfare gain of 0.76%. This loss in value for current homeowners, a theme we return to in Section 6, helps explain the opposition of homeowners to relaxing land use and density restrictions.

5.4 Urban growth boundaries

Urban growth boundaries (UGBs), also known as green belts, are a common restriction on land use. They prevent any new development beyond a specified urban fringe. Many English cities, Toronto, Vancouver, Portland (OR), Minneapolis-St Paul, Miami, Seattle, and most Chinese cities, impose some form of UGB. As the elasticities for redevelopment and infill are very low, the increased scarcity of large tracts of undeveloped land in desirable locations in markets with binding UGBs contributes to their weaker supply responses to demand growth and higher housing costs.

In early work, Cunningham (2007) provides evidence that the area around Seattle's UGB, imposed in 1995, experienced less development on net. However, since the UGB was expected up to five years in advance, it also initially reduced the real option value of waiting to develop land beyond the boundary (as per our discussion in Section 4.1). While the Seattle UGB eventually reduced new developments beyond the boundary, it nonetheless accelerated them during the period between its announcement and its implementation.

UGBs can increase local amenities by preserving open space and reducing negative traffic externalities, given that higher-density cities have less aggregate

travel (Ewing and Cervero, 2010; Duranton and Turner, 2018).⁷⁴ As each UGB affects the equilibrium of its entire housing market, quantitative spatial equilibrium models, similar to those discussed in Redding (2025) in this volume, can be used to perform welfare analysis.

Using such a model, Koster (2024) reports quantitative evidence that in the English case, UGBs have large positive external effects. Using various identification strategies, he finds that being in a greenbelt raises house prices by about 12%, which is naturally interpreted as reflecting the amenity advantages of access to open space. In addition, the housing stock is much lower in greenbelts. A counterfactual of this quantitative spatial model developed and estimated for England in which all greenbelts are removed reveals that improved housing affordability would not be large enough to outweigh the associated reduction in consumer amenities. Central to this conclusion is that the capital share in housing production (α from Section 3) is greater than 0.55, which means that existing residentially zoned land can be built more densely at a sufficiently low cost. Another caveat is that Koster (2024) estimates an extremely high distaste by residents for commuting, which limits the benefits of urban expansion. In contrast, the amenity value of greenbelts decays only slowly with distance. This study also finds little role for productivity effects of UGBs.

Yu (2024) draws the opposite conclusion about the welfare consequences of a similar type of land use restriction that was imposed on Chinese cities in 1999. This restriction requires cities to maintain a constant stock of farmland nearby. As a result, each reclassification of land from rural to urban, which facilitates urban expansion, requires converting an equal amount of nominally urbanized land to farmland. While these are not strictly speaking UGBs, this restriction imposes a similar type of development constraint. This constraint is useful for empirical work because it varies across cities by the prevalence of rugged land at the urban fringe. Difference-in-differences evidence demonstrates that these development constraints reduce city population and GDP, raise prices, reduce productivity, and ultimately, through the lens of a quantitative model, reduce the welfare of workers. They also lead to aggregate land misallocation, a cost of land use restrictions to which we return in the following section.

5.5 Evidence from upzoning

Commensurate with the boundary discontinuity evidence discussed above, a number of additional studies use cross-sectional or panel variation across US jurisdictions to provide evidence that stricter land use regulation reduces new construction and/or increases housing costs (Mayer and Somerville, 2000; Quigley

⁷⁴ There is a long-standing debate in urban economics about the role of land use restrictions, and in particular UGBs as a second-best instrument to mitigate the effects of traffic congestion (Pines and Sadka, 1985; Anas and Rhee, 2006; Anas and Pines, 2008). Brueckner (2007)'s calibrations of a monocentric model with endogenous congestion indicate that a UGB is a vastly inferior second-best alternative to optimal congestion pricing for internalizing congestion externalities.

and Raphael, 2005; Kok et al., 2014; Jackson, 2016). Episodes of large-scale upzoning in individual cities provide additional evidence to corroborate these observational studies.

An empirical challenge in looking at such large-scale events is that general equilibrium effects generate spillovers from treated to potential control areas. As a result, comparison regions within the same city or metro area do not really exist. As upzoning typically happens in unique ways, it is also difficult to find valid counterfactual areas in other housing markets. This leaves a role for the use of set identification restrictions and/or quantitative modeling to assess upzoning's impacts. Moreover, construction and price responses to upzoning depend, of course, on the details. Changes in zoning regimes that give more development flexibility will have more significant impacts.

Greenaway-McGrevy and Phillips (2023) examines the impacts of the 2016 upzoning of three-quarters of the residential land in Auckland, New Zealand. While various changes in land use restrictions were implemented, a binary upzoning treatment is defined from the increases in the maximum allowed floor-to-area ratio. This study defines this treatment to be assigned at the New Zealand Statistical Area level, similar to US census tracts. The number of new building permits requested as a result of the upzoning represented 4% of the aggregate dwelling stock. It is implausible that this response was driven by substitution from areas that were not upzoned to upzoned areas. This study follows evidence in Greenaway-McGrevy et al. (2021) that upzoning primarily increased the values of underdeveloped properties, for which the option value of redevelopment increased the most.⁷⁵

Various US cities and states have recently become more involved in requiring municipalities to allow denser housing development. For example, Massachusetts' 2021 "MBTA Communities Law" compels municipalities who have not done so to establish zoning districts with multi-family housing development allowed "as of right" with densities of at least 15 units per acre within 0.5 miles of a transit station. Municipalities that do not comply lose state infrastructure funding. Nonetheless, the high-income communities of Milton and Needham have voted to reject the rezoning proposals that meet this upzoning requirement.⁷⁶

⁷⁵ Since 2009, the city of Minneapolis has undertaken several large-scale upzoning initiatives, including in single-family home neighborhoods and eliminating minimum parking requirements. While identification is more difficult in this Minneapolis case, evidence is a large development response, less rapid rent growth, and more rapid price growth for single-family homes, especially in lower-income neighborhoods. The latter price growth likely reflects the own-lot effect and the increased values of redevelopment options (Kuhlmann, 2021; Liang et al., 2024).

⁷⁶ <https://www.mass.gov/info-details/mbta-communities-law-qa>. The Canadian province of Ontario is one of the few jurisdictions in North America in which a high level of government can rule on specific development applications. The appointed Ontario Land Tribunal and its predecessors have overruled many local planning decisions, even for proposed developments that violate municipal zoning codes. This has contributed to the city of Toronto's 4.9% annualized growth rate in dwellings in multifamily buildings over 2001-2021, far outpacing that for large US cities (ab Iorwerth et al., 2025).

One common upzoning policy is to require municipalities to allow accessory dwelling units (ADUs), additional small housing units added to existing properties. The state of California now requires all municipalities to allow the construction of ADUs with expedited permitting. As it is small in scale, this type of gentle densification is less likely to raise political opposition, as in the Massachusetts transit-oriented development example. Davidoff et al. (2022) examines the effects of the 2009 law that allowed ADUs (in the form of laneway homes) in 95% of single-family zoned areas in Vancouver, Canada. Interestingly, this study finds no “own-lot” effects of new laneway homes. In the Vancouver context, it seems that the construction cost of laneway homes is about the same as the present value of their net operating income. These units are so small that the fixed costs of development apparently make them poor investments. It is no wonder that there have not been many ADU constructions after this law change. These ADUs are found to impart negative externalities on neighboring properties, as measured through hedonic price responses. Negative capitalization is about 2.5% of value for the priciest quartile of neighboring homes but near zero for the lowest value quartile.

While the potential benefits from upzoning are arguably more significant for cities in developing economies, only a few studies examine such episodes. Gechter and Tsivanidis (2023) examines the redevelopment of the Mumbai Mills district, which covers 15% of the central area of Mumbai. These 60 textile mills were mostly replaced with high-rise residential and commercial developments, representing a large increase in Mumbai’s housing supply. The surrounding slum areas were not upzoned but could be redeveloped, with some compensation paid to incumbent residents. This study finds large positive externalities from the mills’ redevelopment, with property values increasing and slum redevelopment displacing low-income residents in adjacent neighborhoods. In the context of a quantitative spatial model, these responses can be rationalized with endogenous local amenities that positively depend on the skills mix of nearby residents. Overall, this redevelopment benefited high-skilled city residents who gained from increased housing supply and affordability at the expense of some (displaced) low-skilled incumbent residents near the redeveloped mills. This evidence of the gentrification cost of large-scale development in low-income neighborhoods perhaps helps to explain the opposition to neighborhood redevelopment seen in many contexts.

5.6 Property taxation and the political economy of land use regulation

We have seen how property owners have several motivations to restrict development and densification. Adverse external effects of density reduce local amenities, incentivizing property owners to resist and limit nearby development. Such negative spillovers can be greater in high-income neighborhoods. These are negative “external effects” in the language of Turner et al. (2014). However,

“own-lot” effects incentivize owners of vacant land to oppose land use restrictions provided that the countervailing “supply effects” are not too large.

Zoning rules are often set at the local level. In the US, zoning decisions are nominally made at the municipality level, although in practice even more local interests often dominate. For example, “aldermanic privilege” in Chicago means that almost all zoning proposals must receive the support of the local city alderman to be implemented. As the interests of local owner-occupants, renters, absentee landlords, and absentee owners of vacant land or properties ripe for redevelopment usually do not align, the interest group with the most political power often influences zoning rules and development approvals at the expense of some other groups. For this reason, areas of large cities that have many renters tend to be more development-friendly, at least for subsidized housing, than built-up suburbs that are overwhelmingly populated by owner-occupant households. Unbuilt areas at urban fringes with a lot of vacant land also tend to be development-friendly, given own-lot considerations.⁷⁷

Hilber and Robert-Nicoud (2013) is one attempt to conceptualize how two of these competing interests can lead to variation in development restrictions across locations. This study considers a game played by resident owners of developed properties and absentee owners of vacant land in a system of jurisdictions in which each jurisdiction makes its own decision about development restrictions. These two players have opposite incentives to advocate for land use regulation. The supply effect incentivizes developed property owners to support higher regulatory taxes (Glaeser et al., 2005b; Glaeser and Gyourko, 2018) while the own-lot effect incentivizes owners of undeveloped land to oppose them. These two players influence the planning board by “bribing” them proportionately to aggregate gains or losses in property value associated with enacted regulatory taxes. In equilibrium, the planning board chooses regulation to maximize total land rents plus regulatory taxes, which only involves a choice of nonzero taxes because jurisdictions compete for residents. Locations in high residential demand (because of good amenities, for example) get more built up in equilibrium, thereby leading to less undeveloped land and higher regulation. This pattern is also observed in the data across US metropolitan areas.

The observation that strategic interactions across municipalities inhibit housing development is also supported by quasi-experimental evidence. Looking at forced amalgamations of municipalities of less than 5,000 inhabitants in France after 2010, Tricaud (2025) finds, in a difference-in-differences empirical setup, that construction permits increase by 12.5% in forcibly amalgamated municipalities. However, no such increase occurs after voluntary consolidation. Moreover, more urban municipalities that were amalgamated and had the greatest construction response experienced no change and housing values and, if anything, improvements in public services. Tricaud (2025) takes this as evidence that

⁷⁷ Developers’ interests matter as well. Using a regression discontinuity empirical design for close elections, Yu (2022) finds that residential developers who donate to candidates elected mayor develop more housing.

resistance to amalgamation was primarily about residents' concerns about negative externalities from new housing development rather than about negative capitalization effects or reductions in the quality of local public goods. Mast (2024) demonstrates that US municipalities that move from at-large to ward-based representation in their city council reduce the number of approved permits for new housing units by 20%. Consistent with the idea that higher-income homeowners are most resistant to new development, responses are larger in high-income jurisdictions.

When it comes to decisions about local land use restrictions, property taxes also play an important role. So does the quality of local public goods provision, including schools, parks, and police. Decisions about zoning, property taxes, and local public goods are typically made jointly. A line of thinking going back to Tiebout (1956) concludes that property taxes are not distortionary if they are used to fund local public goods. Households with heterogeneous demand for local public goods "vote with their feet" to live in the jurisdiction that is right for them. This view does not explicitly consider zoning, housing supply, or within-jurisdiction resident heterogeneity. However, standard logic from public economics is that taxation of elastically supplied capital (or structures) is distortionary (Harberger, 1954) and leads to an inefficiently small capital share in housing provision. For this reason, the "Henry George Theorem" states that property taxes should only be on the land component of housing, as the fixed supply of land means that such taxes are not distortionary.⁷⁸

Building on this Georgian idea, Hamilton (1975, 1976) articulates a positive role for land use regulation. Imagine an environment, as is standard, in which the provision of the local public good to each household has the same cost to the local government. A result extending back to Oates (1972) is that identical head taxes efficiently fund the provision of public goods. However, head taxes are not an option in most jurisdictions, which must fund operations using proportional property taxes. If properties are heterogeneous in value, under proportional property taxation, low-value properties pay less in tax and are subsidized by the taxes levied on high-value properties. By stipulating that all properties in a community must be identical, zoning can eliminate this fiscal externality. That is, under a second-best environment in which head taxes are not possible, zoning can manipulate the built environment so that property taxes replicate head taxes. Brueckner (2023) formalizes these ideas, demonstrating efficiency with pre-zoned communities. However, inefficiencies arise if households can choose the amount of housing to consume/build for themselves.

These ideas are developed further in the computational multi-community model with household heterogeneity developed by Calabrese et al. (2007). This

⁷⁸ Given the common reality of fixed lot sizes, this logic clearly holds. With variable lot sizes, it only holds if there are no density externalities. Separate taxation of land and structure is common in Eastern Europe. In the US, Allegheny County (PA) used to assess the two separately. See also Behrens et al. (2015) for the genesis of the modern versions of the Henry George Theorem, key references, and some limitations of the mechanism.

model proceeds as a three-stage game. Households first buy land in one community. In the second stage, they choose the property tax rate and minimum housing requirement by majority vote. In the final stage, they relocate or adjust land holdings, build housing on their land, pay taxes, and consume. This timing structure skirts the problem of jointly determining an equilibrium community composition together with a zoning and taxation regime. The model analytically delivers endogenous equilibrium zoning in all but the lowest-income jurisdiction. Computationally, equilibrium zoning is very stringent but welfare-enhancing relative to an environment in which zoning is prohibited. However, zoning also induces stratification across communities, increasing inequality. Low-income households are zoned out of richer communities (“exclusionary zoning”) and are left with a low tax base and low-quality local public goods. Zoning eliminates the negative jurisdictional choice externality that the poor would impose if they were able to buy a small amount of housing in a richer community. However, it saddles them with lower-quality public goods.⁷⁹ As we discuss further in Section 6, zoning in one jurisdiction can raise housing costs in other jurisdictions through aggregate land constraints.

Fischel (2001) recasts this logic into a framework in which property owners vote to maximize the value of their own properties, the “homevoter hypothesis”. Given that areas are built up, this means enacting rules that limit negative externalities across properties. For these reasons, Fischel (2001)’s view is that “zoning is an essential ingredient of municipal formation and function”. Several empirical studies provide evidence that property owners vote to raise property taxes if it increases their property values in the context of vouchers for private schools and the construction of nearby stadiums (Brunner et al., 2001; Brunner and Sonstelie, 2003; Dehring et al., 2008). Hilber and Mayer (2009) document that school spending is higher in districts with less vacant land, consistent with the idea that municipalities with less elastic housing supply are more incentivized to tax themselves to support local public goods. These taxes only pay off in higher property values if zoning rules restrict free-riding and reduce externalities across properties.

Relatedly, Krimmel (2021) demonstrates how stricter zoning has been used to guard against the congestion of local public goods when municipalities lose control over property tax rates. After a reduction in California municipalities’ taxing power that came with school finance equalization in the 1970s, he shows that municipalities enacted stricter zoning regimes. These regimes required larger lots and more housing services per dwelling, thereby limiting the entry of new students into local school districts, with stronger such zoning responses in higher-spending localities.

⁷⁹ Barseghyan and Coate (2016) develop a dynamic model with household income heterogeneity, two communities, and two possible home sizes with similar analytical conclusions. We also note that zoning regimes have been found to target racial exclusion. For example, when Chicago established its first zoning code in 1923, the city assigned industrial zoning to minority residential areas much more than to White areas (Shertzer et al., 2016).

Still, in the spirit of the Henry George Theorem, property taxes can also fund local infrastructure like roads and transit. As infrastructure typically results in increases in property values through capitalization, a commonly proposed taxing model is “value capture”, in which capital gains that accrue because of new publicly provided infrastructure are taxed to fund it. For example, Gupta et al. (2022) find that New York City’s 2nd Avenue subway construction could have been fully funded with full value capture. Given current tax rates, however, additional associated property tax revenue implied by this extension of the subway will be well below its construction cost.

As property taxation funds most local services in the US, interesting political economy dynamics across levels of government ensue, with implications for housing supply. California’s Proposition 13, enacted in 1978, limits property taxes to 1% of assessed value and allows municipalities to raise assessed values by at most 2% per year on incumbent property owners. However, properties with new owners or that are new construction can be re-assessed to market values. Proposition 13 is thus a form of enforced centralization, in which some fiscal responsibility for local service provision is uploaded from municipalities to the state. Increases in property prices thus imply increases in the implicit cost of moving, which should lower mobility rates.⁸⁰ With higher property prices, there is also a related increase in the implicit cost of buying new homes, given that they come with higher property taxes and a disincentive for property redevelopment. The literature has not systematically examined the housing supply response to Proposition 13 to date.

5.7 Taxation and housing policy

In addition to property taxes, common housing tax instruments include transfer taxes (stamp duties), which are usually levied on buyers and set to be progressive in the sales price, and capital gains taxes, which are also levied after transactions but are paid by sellers. Some countries also allow some tax filers to deduct mortgage interest from their tax liability and/or offer various tax incentives for home buying. Finally, there are tax subsidies for developers to build or renovate subsidized housing. In the US, the main program in this category is the Low Income Housing Tax Credit (LIHTC). Newly constructed or renovated housing units subsidized through the LIHTC program are subject to capped rents for at least 15 years and can only be rented to households with incomes below thresholds that depend on area median incomes.

Transfer taxes (stamp duties) are universal in the United Kingdom and much of the Commonwealth, including Canada and Australia. They also exist in many European countries, and most US states. By disincentivizing housing transactions, these taxes can generate misallocation of households to dwelling units.

⁸⁰ In practice, there is mixed evidence on the mobility response to this law (Ferreira, 2010; İmrohoroglu et al., 2018).

The associated additional cost of moving may be expected to keep small families in large homes for longer, disincentivize moves across labor markets to look for and accept better employment opportunities and reduce the supply of potential movers into newly constructed dwellings. The reason is that some housing transactions that would be mutually agreeable to buyers and sellers absent the tax do not generate sufficient surplus to cover the cost of the tax. The result is reduced liquidity in the housing market.

The existing literature on land transfer taxes finds sizable effects of such taxes on the number and pricing of transactions. Using variation in UK stamp duty schedules from changes that occurred over the 2003-2013 period, Best and Kleven (2018) uncover various pieces of evidence to this effect. Sharp increases in tax rates at each of five different price thresholds, meaning that buyers pay thousands of pounds more in tax for a one pound higher sales price, also facilitate identification of impacts of this tax.⁸¹ These notches generate missing masses in the sales price distribution and bunching at just under the price threshold associated with each notch in the tax schedule. Bunching estimates in Best and Kleven (2018) indicate that these tax notches reduce sales prices of the most responsive transactions by two to five times the marginal tax paid at the notch, which ranges from 1,250 to 40,000 pounds. In addition, a temporary one percentage point reduction in stamp duties on homes sold for 125,000 to 175,000 pounds in 2008-2009 induced 20% more transactions, primarily because of time-shifting, but also from trades that were newly mutually beneficial between buyers and sellers (the “extensive margin” effect). As a result, this tax holiday was a successful fiscal stimulus.⁸²

In related work, Kopczuk and Munroe (2015) use the fact that New York and New Jersey impose a transaction tax of 1% on homes that sell for at least one million dollars and several other features of the nonlinear tax schedule in these states to study how transaction taxes influence housing market liquidity. Again using a bunching estimator, this study also finds that fewer total transactions take place as a result of this tax treatment. Unraveling can occur because potential matches with low surplus within the region of the notch can be abandoned by sellers waiting longer to receive higher prices or by buyers waiting longer for lower prices. Overall, Kopczuk and Munroe (2015) find that 0.7% of transactions do not occur because of this tax notch. Through the lenses of models of buyer-selling bargaining, both Besley et al. (2014) and Kopczuk and Munroe (2015) find that sellers have more bargaining power and face greater tax incidence than buyers as a result.

Finally, Dachis et al. (2012) examine the consequences of the imposition of a 1.1% land transfer tax in Toronto, Canada, in 2008. Using a difference in discontinuity empirical strategy, in which distance to the Toronto municipality border is the running variable, they find that these taxes led to a 14% decline in

⁸¹ For example, the two percentage point increase in tax at 250,000 pounds increases the liability by 5,000 pounds for a one pound increase in price.

⁸² Besley et al.’s (2014) earlier look at the same episode generally confirms these results.

transactions each year and were capitalized into lower sales prices about one for one. As a result, land transfer tax revenue crowds out subsequent property tax revenue. In the context of a simple location choice model, an increased Toronto property tax rate would have been more efficient at raising the same amount of additional revenue. In particular, model calibrations indicate a 19.5 million dollar aggregate loss each year because of transactions that do not happen. This 12.5% of estimated land transfer tax revenue is a social loss that reflects distorted mobility decisions. Property taxes could have been set to raise the same amount of revenue without this distortion. This evidence is independently corroborated by Han et al. (2022), who additionally argue through the structure of a search model that the tax generates an inefficient shift from owning to renting with associated deadweight losses of 32% of revenue.

The other common transfer tax on real estate is the capital gains tax on primary residence. While most OECD countries fully exempt these capital gains from taxation, nine do levy this tax, including the US (Andrews, 2011). Using the shift in 1997 from a broader applicability of the tax to the current less stringent regime for identification, Cunningham and Engelhardt (2008) and Shan (2011) find large mobility and housing market liquidity impacts for affected properties.

The connection between mobility frictions induced by taxation and housing supply is indirect but straightforward. As an example, transaction taxes lock in high-income retirees in San Francisco when they would otherwise move to Florida, where new housing can be built easily. This phenomenon will be strongest in locations with the most inelastic supply of available properties, as these are the locations with the highest prices and greatest capital gains. Hence, besides misallocating housing, transaction taxes reduce supply through mobility frictions.

Many governments also provide various types of incentives to support housing demand. For example, half of OECD countries allow deductions of mortgage interest from taxable income in some situations (Andrews, 2011). In the US before 2017, when eligibility was tightened, the mortgage interest deduction amounted to a 90 billion dollar annual tax expenditure (Sommer and Sullivan, 2018). Other common policies are grants or tax credits for first-time home buyers, direct provision of home loans at subsidized interest rates, taxpayer-backed mortgage securitization enterprises, including Fannie Mae and Freddie Mac, and subsidized mortgage insurance.⁸³

These policies that directly support housing demand will raise prices and reduce affordability, especially in supply-constrained markets (Favara and Imbs, 2015; Justiniano et al., 2019). They can also influence the types of new construction. As many of these policies work on the intensive margin, they subsidize the purchase of larger, more expensive homes. This demand is then fulfilled by

⁸³ On the other side, some countries, including Canada, also limit demand for owner-occupied units by banning or taxing foreign buyers, requiring qualifying borrowers to pass mortgage stress tests, and limiting entry into mortgage credit markets.

builders with more such new constructions, which could contribute to affordability challenges. Moreover, they will tend to shift demand toward owner-occupied and away from rental properties, which can also influence the types of dwelling units that developers build.

Evidence in the literature supports all of these claims. Hilber and Turner (2014) find that the US mortgage interest deduction boosts homeownership of higher-income households in elastically supplied markets yet reduces homeownership in inelastic markets. Sommer and Sullivan (2018)'s calibrations of a dynamic structural model of owning and renting for the US indicate that removing the mortgage interest deduction would result in substantially lower house prices, allowing for more lower-income households to become homeowners.⁸⁴ Using a reduction in the generosity in the mortgage interest deduction for medium- and high-income households in Denmark, Gruber et al. (2021) finds no evidence that this policy influences homeownership rates for these groups but compelling evidence that it increases house prices and incentivizes the purchase of larger homes.

The tax code is also commonly used to incentivize the construction of rental housing, including in the LIHTC program mentioned above. As these housing developments are more likely to be in low-income neighborhoods with market rents below levels needed to justify construction costs, it is logical that these subsidies may not crowd out much private construction. However, it is revealing that developers have a large amount of excess demand for these tax credits (Baum-Snow and Marion, 2009). This study finds up to 60% crowd out of market rental construction in gentrifying neighborhoods, though little crowd out in other settings at the census block group level. Eriksen and Rosenthal (2010) finds substantially greater crowd-out when looking at a 10-mile radius spatial scale. Using variation in developer application rates as a function of the size and availability of tax credits and detailed parcel-level information about development, Soltas (2024a) estimates about 80% crowd-out at the parcel level and that about 50% of the tax credits ending up as developer profits. Across all subsidized housing programs, Sinai and Waldfogel (2005) find about 30%-50% crowd-out at the metropolitan area level, arguing that aggregate supply effects of housing policies are more appropriately measured at the market level. Conceptually, we expect more crowd-out, the less elastic is rental housing demand.

The LIHTC program has replaced government-built public housing as the main source of project-based subsidized housing in the US. In recent decades, it has accounted for about one-fifth of all multifamily rental construction nationwide (Soltas, 2024a). Overall, the US subsidized housing system covers about 8.5 million dwelling units, almost 20% of the overall rental stock. Of these, 3.6 million are LIHTC units, with about 300,000 added each year, 0.9 million are public housing, 2.8 million are Section 8 person-based vouchers, and 1.3 million are Section 8 project-based vouchers (Department of Housing and Urban Development, 2023). These latter two categories are rental housing vouchers

⁸⁴ A central parameter in this study is the housing supply elasticity, for which they use 0.9.

that cover the difference between apartment rent and 30% of household income, provided the former is at or below a designated fair market rent. Person-based vouchers are portable across apartments, whereas project-based vouchers are tied to dwelling units.⁸⁵

An alternative policy instrument to housing subsidies for low-income families is rent control or rent stabilization, which constrains rents to be below market levels. This policy is common around the world but has been eliminated in several US jurisdictions in recent decades, including three in the Boston area and in San Francisco. Because rent control depresses prices, it disincentivizes rental housing development. For this reason, New York and Toronto only limit rent growth in most dwellings constructed before 1974 and 2018, respectively. In addition, rent controls generate shortages and deadweight losses because some positive surplus rental market transactions are prohibited from taking place. Rent control is also likely to generate misallocation of renters to apartments. Some renters stay in apartments for which they are only willing to pay slightly more than controlled rents. Other potential renters have a much higher willingness to pay for these apartments. Glaeser and Luttmer (2003) find substantial evidence of such misallocation in New York City.

One potential justification for rent control is the insurance value of predictable rents (Diamond et al., 2019). Favilukis et al. (2023) provides quantitative evidence that the neediest households enjoy substantial associated net welfare gains from rent control. Hilber and Schöni (2022) provides additional context and perspective about housing policy, including rent control, in the US and around the world.

6 Top-down approaches to housing supply

When housing is unaffordable, households either consume less of it or elect to live in alternative locations where housing is cheaper, even though these locations may be less productive or provide lower consumer amenities. In extreme cases, households become homeless, a topic we do not address in this chapter, or households fail to form in the first place. To capture the full social cost of the unaffordability of housing, a general equilibrium model is needed. Only a general equilibrium model can track the welfare consequences of migration responses to changes in land use regulations, as these responses will work through changes in housing prices, amenities, and productivity in all locations.

We know of no assessment focusing on the aggregate effects of households consuming too little housing because of our inability to build housing

⁸⁵ Using identifying variation from updates in the calculation of metropolitan area-wide fair market rents, Collinson and Ganong (2018) provide evidence that landlords capture a large amount of the surplus from housing vouchers and that the supply of affordable apartments is thus generally quite inelastic, even in markets where new construction is elastically supplied.

cheaply and efficiently.⁸⁶ An extremely simple back-of-the-envelope calculation nonetheless suggests large effects. To be roughly consistent with the data and results established above, assume Cobb-Douglas preferences for housing with an expenditure share of 0.25 and a Cobb-Douglas production function for housing with a share of capital of 0.65. If construction was 20% more productive – a higher B using our notations and a conservative figure given our discussion in Section 3.6 – Eq. (17) then implies that builders would provide 68% more housing per unit of land. Following what is reported in Fig. 3, we also conservatively assume a growth in the number of net units of 1.5% annually or 16% after ten years. Under these assumptions, a 20% higher productivity in construction would amount to an extra 11% added to the stock of housing after ten years. In turn, this higher productivity in construction would lead to a 2.6% gain in welfare, only from being able to consume more housing. Obviously, this is only a partial equilibrium calculation, and a full general equilibrium setting would need to be developed to explore this question more fully.

Recent work has focused on the spatial misallocation effects of land use regulations. This strand of research has attracted a lot of attention in recent years following Hsieh and Moretti's (2019) initial impulse and subsequent research by Duranton and Puga (2023), Parkhomenko (2023), and other references cited below.⁸⁷ In this line of work, stringent land use regulations in prosperous cities act as a barrier to migration from less prosperous cities and create large aggregate costs. This feature is consistent with the results in Howard and Liebersohn (2021) who show that high housing prices in superstar cities are caused by increasingly stringent land use regulations aimed at fending off the arrival of newcomers.

In the model proposed by Duranton and Puga (2023), cities face a trade-off between agglomeration economies and urban costs. Both the benefits and costs of cities increase with their populations. Per-capita consumption is hump-shaped in the number of residents. It initially increases with city population, as agglomeration economies dominate, before it starts to decline as urban costs eventually dominate. This trade-off is resolved by incumbent residents who decide how large their city should be in each period. They unanimously choose the population of their city to maximize consumption per capita at the top of the hump, as in Albouy et al. (2019).⁸⁸

⁸⁶ There are several reasons behind this lack of interest. The widespread realization that the productivity performance of the construction sector has been dismal, as discussed in Section 3.6, is fairly recent. Most of the literature has instead focused on situations where some locations are heavily regulated and others less, which is more in line with the context surrounding land development reviewed in Section 5.

⁸⁷ There are several controversies around the results of Hsieh and Moretti (2019), with, in particular, a comment by Greaney (2019) making strong counter-claims. This debate has not been settled at the time of the writing of this chapter.

⁸⁸ In the absence of a mechanism to regulate city creation and city population, multiple equilibria abound due to the coordination failure caused by agglomeration economies. The most productive locations may not develop as a city. In addition, the benefits of the most productive locations may get crowded out by newcomers moving in until consumption is equalized between that city and the location offering the lowest level of consumption, leading to possibly immense and wasteful overcrowding. See Behrens and Robert-Nicoud (2015) for further elaborations.

Residents then impose their preferred population size by making newcomers pay a wasteful permitting cost. The empirical counterpart of this cost is taken to be the wedge between the market price of an existing house and its replacement cost on unpermitted land. This political-economy mechanism is akin to the homevoter hypothesis of Fischel (2001) discussed above. It resolves the trade-off between the agglomeration benefits of cities and their urban costs in a privately optimal way for each city. The spatial equilibrium between cities is simple to construct. The location with the highest productivity fundamentals will develop the largest city, followed by the second highest, etc. This process of city creation continues until the remaining households prefer to live in a rural location operating under decreasing returns to scale rather than establish a new city on the most productive of the remaining potential locations for a city.

By operating at their privately optimal scales enforced by permitting costs, cities – especially the most productive ones – generate an inefficiency, which, in turn, implies large social costs. To see why, consider the relocation of a small group of households from the least productive city or the residual rural area to the largest city. Given the gap in consumption between the largest city and rural area, this relocation generates a large discrete gain in consumption for relocating households and a small marginal loss for incumbent residents from whom urban costs grow marginally larger than agglomeration benefits. In equilibrium, the largest cities are too small, there are too many small cities, and the rural population is too large.

The main counterfactual in Duranton and Puga (2023) considers what would have happened if the seven superstar cities defined in Section 2 had built new housing between 1980 and 2010 at the same rate as the metropolitan area at the 75th percentile for building permits. That percentile corresponds to about 0.81 permits per dwelling since 1980 relative to, for instance, 0.22 in New York. Using estimated parameter values summarizing agglomeration forces and urban costs along with calibrated values from the literature capturing the magnitude of decreasing returns in rural areas, this counterfactual implies that an additional 17 million residents would live in these seven superstar cities. The largest increase in population occurs in New York, which would have received an additional 7.5 million residents between 1980 and 2010.

This (counterfactual) increase in population in the largest and most productive cities leads to an aggregate gain in total output close to 8%. This gain arises from (i) the relocation of 17 million people to more productive places, (ii) stronger agglomeration effects in superstar cities after the arrival of these newcomers, and (iii) a higher output per worker in (left-behind) rural areas operating under decreasing returns to scale. However, a large part of these production gains are dissipated into increasing urban costs. The aggregate gain in consumption (welfare) is 2.1%.⁸⁹ Newcomers and rural dwellers experience consumption

⁸⁹ The standard distinction between output (GDP) and welfare (consumption) is particularly salient here given the importance of unavoidable urban costs when cities grow in population.

gains of 6.5%. However, incumbent residents of superstar cities experience consumption losses from increased urban costs, which more than offset higher agglomeration benefits. These losses are nonetheless minimal: 0.05% at most in New York. The reason behind these small losses is that the difference between the agglomeration elasticity and the urban cost elasticity is small, 0.04, with the parameter values used in Duranton and Puga (2023). Put differently, the hump shape of city consumption as a function of population size is fairly flat after its maximum. Cities can thus grow significantly larger than their privately optimal scale at a low net cost for their incumbent residents.

Two further results are of interest. First, aggregate income and consumption inequalities decline as the income and consumption of (poorer) rural residents increase. Second, the relaxation of land uses regulations in superstar cities to allow for more dwellings to be built also leads to a relaxation of land use regulations in other cities, as the migration pressure toward these cities is also being eased.

The conclusion that relaxing restrictions on land development and new constructions could lead to large aggregate gains is pervasive in the literature. Magnitudes comparable to those of Duranton and Puga (2023), or larger, also appear in Ganong and Shoag (2017), Herkenhoff et al. (2018), Hsieh and Moretti (2019), and Parkhomenko (2023).⁹⁰ This conclusion is, however, disputed by Glaeser and Gyourko (2018), who argue that the aggregate social cost of restrictive land use regulations is more like 2% instead of 10% or more, as implied by the rest of the literature.

The key difference regards how local economies respond to a population shock. Glaeser and Gyourko (2018) argue that labor demand at the city level is fairly inelastic so that the large wage differences we observe between cities quickly disappear as we allow households to move from low-wage locations to high-wage but tightly regulated cities. In turn, the gains from such a reduction in spatial misallocation are limited and even vanish as labor demand becomes totally inelastic. It is true that the gains from weakening land use regulations where they bind the most crucially depend on how local firms respond to the arrival of new workers.⁹¹ The labor literature cited in Glaeser and Gyourko (2018) is supportive of a downward-sloping labor demand for existing local firms. However,

⁹⁰ Fafilukis et al. (2023) also find a large supply response from relaxing land use regulations in New York while taking into account that regulations detrimental to supply are also motivated by protecting incumbent renters from the price risks associated with renting as discussed above. Babalievsky et al. (2023) explore the aggregate costs from the regulations of commercial real estate. They claim welfare gains of 3 to 6% from deregulating an asset class that represents about 20% to aggregate capital.

⁹¹ While general equilibrium models of multiple cities often assume that labor is freely mobile across cities, a large literature on migrations reports large migration costs (see Jia et al. 2023, for a review). In a world where labor supply (population) is highly inelastic across cities, land use regulations might not be binding. This is unlikely given the evidence of a strong demand to live in large cities, in particular superstar cities (Howard and Liebersohn, 2021). Diamond's (2016) preferred estimates for the elasticity of city population to city wages are 3.3 for non-college workers and 5.0 for college workers. Analogous estimates for rents are -2.9 for non-college workers and

these estimates are probably best viewed as “all else equal”: how many workers will firms hire for a given wage in a partial equilibrium setting? Instead, the relevant “labor demand curve” at the city level should include equilibrium effects since newcomers will bring their demand for local goods, some of them will create new firms, and they may induce some firms to relocate. In line with these effects, the data are supportive of a modestly upward-sloping city wage with city population, with at least some of this relationship being causal (Combes and Gobillon, 2015). This is what we should expect with agglomeration effects, at least in the long run.

However, labor is not homogeneous and is not located proportionately by skill level across cities. The most productive cities are also where housing is the most expensive. As a result, these cities disproportionately attract highly skilled workers who benefit the most from their agglomeration effects and are able to pay the cost of living there (Behrens et al., 2014). This sorting of workers across cities also implies that a large part of the urban wage premium is due to sorting. Urban amenities also play an important role in household location choices and are most likely affected by the skill composition of cities. Diamond (2016) proposes a model with endogenous city amenities that are increasing in the local fraction of college-educated households. After a skill-biased labor demand shock, endogenous amenities magnify changes in the skill composition of cities, making more skilled cities all the more attractive (and expensive) to other skilled workers.

In turn, these features have implications for the analysis of land use regulations in general equilibrium. Macek (2024) incorporates endogenous local amenities in his quantitative evaluation of the consequences of relaxing minimum lot size (MLS) zoning restrictions. He develops a static model with neighborhood and worker heterogeneity that is quantified to census block groups across the United States. The location choices of workers depend locally on housing costs, wages, and amenities, which are endogenous to neighborhood composition. MLS restrictions negatively impact productivity and wages by limiting the population of the most productive locations. In addition, these restrictions induce the sorting of less skilled workers into less restrictively zoned neighborhoods.⁹² By keeping less-skilled workers out of more productive locations, MLS regulations cause more high-skilled workers to prefer these locations. Hence, MLS zoning tends to concentrate more skilled workers where they are the

–2.2 for college workers. Given differences in housing prices across cities, these estimates are large enough to induce very significant population increases if housing was made more affordable in the most prosperous cities. See also Beaudry et al. (2014) for additional estimates. In short, even though it ultimately limits how many households want to move to a city, the slope of the labor supply across cities is unlikely to be an issue in the type of counterfactual considered in Duranton and Puga (2023) and discussed above.

⁹² This sorting occurs because MLS restrictions in desirable neighborhoods reduce housing affordability for less skilled workers, either by constraining them to purchase more housing services than they would otherwise choose or by reducing aggregate housing supply, thereby raising costs.

most productive, thereby enhancing aggregate output.⁹³ Unlike in the models previously discussed with homogeneous workers, eliminating MLS zoning only marginally changes aggregate output because it induces a re-sorting of workers that counteracts agglomeration effects. Strikingly, eliminating MLS zoning in a productive city like San Francisco induces less skilled workers to flow into the city, activating negative endogenous amenities, which encourage more skilled workers to leave. In turn, this exit of skilled workers depresses property values and leaves many worse off. Nonetheless, housing affordability and welfare improve for low- and middle-income renters.

Because they limit the development of the most productive cities, stringent land use regulations may also have dynamic implications through their effects on long-term growth. Duranton and Puga (2023) in an overlapping-generations model and Crews (2024) with a more developed dynamic framework can explore the consequences of land use regulations for human capital accumulation. Larger cities offer better opportunities for learning and accumulating human capital. Large cities may attract more skilled workers, as just argued, but they will also produce them.

Duranton and Puga (2023) first propose a dynamic accounting exercise. Agglomeration economies play two roles in the growth process. First, through faster learning when working in cities, they directly foster the accumulation of human capital. Second, the same agglomeration economies will also magnify the size of cities, which in turn promotes human capital acquisition and accumulation. Overall, Duranton and Puga (2023) find that agglomeration effects account for 14 basis points of annual growth or 7% of the annual growth rate between 1950 and 2010. While this is not directly informative of the dynamic costs of land use regulations in general equilibrium, it is indicative that there are long-term costs to limiting the growth of the most productive cities because they are key breeding grounds for learning and accumulating human capital. In another counterfactual exercise, Crews (2024) finds that lifting restrictions to moving to the most skilled cities would raise growth by 13 basis points annually. Although speculative, these results nonetheless point to costs to land use restrictions that potentially go well beyond the static costs explored above.

7 Conclusions and future research directions

Recent declines in housing affordability have prompted a renewed interest in research on housing supply. In the US, the rate of new construction has slowed considerably since 2010, and estimated housing supply elasticities have been declining, particularly in the most productive locations. Our survey of the literature reveals a number of potential reasons for these changes. Productivity in the construction sector has declined, land use regulation has grown more stringent, and there is less land available for development in the most productive

⁹³ See Rossi-Hansberg et al. (2019) for a related model with strong complementarities between skills and local productivity.

and highest amenity locations. While incumbent property owners in the best locations are incentivized to enact land use regulations to protect their property values, countless others would benefit from reduced regulation through lower housing costs and improved access.

Our overview of facts on quantities, prices, and housing densities in the US confirms casual observations about declining affordability. As housing prices and rents have risen more rapidly than household incomes, rates of new construction have fallen precipitously since 2010. To compensate, dwellings are being removed from the housing stock at lower rates, especially in the costliest housing markets. Relative to a baseline rate of dwelling unit teardowns and abandonment of 0.7 percent per year in small and rural US counties, the corresponding depreciation rate is 0.2 percent per year in the “superstar cities” where housing is most expensive. This 0.5 percentage point difference contributes about 50% of “new” housing supply in superstar cities in comparison to small and rural areas. Along with lower construction rates, there has been a rapid decline after 2000 in the rate of housing densification in suburban areas. In most locations, redevelopment into denser land use has become very rare, an observation also reflected in very low estimated housing supply elasticities for redevelopment.

Our overview of the literature considers both the construction and land development components of housing supply responses to rising prices. We lay out a standard model of housing production in which competitive developers install capital on heterogeneous lots to build housing. Consensus estimates from this literature reveal approximately Cobb-Douglas production with a capital share of 0.65. This implies an intensive margin supply elasticity with respect to the price of housing services of about 2. An independent literature finds sharply declining productivity in the construction sector, which can help to explain reduced construction rates despite rising prices. More research is needed to better understand the reasons for this declining productivity and potential policy remedies.

One challenge faced by this neoclassical approach is that it has nothing to say about the timing of construction. While there is a rich theoretical literature about real development options and qualitative forces influencing the timing of construction, there is less empirical and quantitative work on this topic. The small amount of empirical work that exists is not well-suited for aggregation or incorporation into a general equilibrium framework. As a result, unlike their static counterparts, dynamic market-level supply elasticities have not been systematically characterized. Relatedly, the literature has very little of a handle on the mechanisms driving very low estimated housing redevelopment elasticities. With an emerging consensus that reductions in land use regulation are needed to promote more new construction, more research is needed on the redevelopment process, including its dynamics and aggregation to market supply elasticities.

A different challenge for the existing literature is how to handle the fact that dwellings are discrete and of heterogeneous sizes and qualities. The use of assignment models as a vehicle to match heterogeneous households to heterogeneous dwellings is a start and has yielded some useful insights. However,

existing research in this area still does not have a good command of the housing development decision. Little is known about the extent to which new constructions of different sizes and qualities can be rationalized with the same approach as has been developed for the construction of units of housing services. Additional research in this area would be particularly fruitful, as it would speak to price responsiveness to demand shocks across the distribution of housing quality. We also hope that future research develops tractable and empirically quantifiable frameworks for analyzing the efficacy of subsidized housing policies targeting groups for which developers do not find it profitable to build market-rate units. Complicating such inquiries are the existence of externalities across homes and households that operate within neighborhoods. Incorporating such spatial relationships and location choice into assignment models would further facilitate their broader application.

Land development makes up the extensive margin of housing supply. A flood of recent empirical research analyzes the consequences of supply constraints through land use regulations. The literature has successfully confronted both systematic measurement and the credible estimation of the consequences of various types of land use regulation in targeted and well-identified settings. A convincing articulation of the idea that land use regulations can raise property values through an aggregate supply effect is critical for understanding the political economy of land use regulation. This motivation comes in addition to property owners' incentives to restrict development in order to limit negative fiscal externalities from lower taxes levied on low-value properties in the same jurisdiction along with associated potential negative social and density externalities. There is a strong consensus in the empirical literature that stricter land use regulation impedes development. Conditional on development, regulations and the tax system push toward the construction of fewer and larger dwellings.

Well-identified reduced-form evidence on the consequences of land use regulations for property values and construction feeds into an active literature that uses general equilibrium models to assess the welfare consequences of reducing regulations. This literature finds benefits, though the magnitude and composition of these benefits depend crucially on the system of migration elasticities across locations and the data-generating processes for local productivities and amenities. Active research in progress extends these models with richer treatments of dynamics and richer agent heterogeneity. In future research, these classes of models could be successfully repurposed to evaluate the distribution of returns to investing in improvements in construction technology. In addition, research that applies similar models to evaluate possibilities for feasibly compensating incumbent property owners for losses from land use deregulation would be fruitful.

References

Aastveit, Knut Are, Albuquerque, Bruno, Anundsen, André K., 2023. Changing supply elasticities and regional housing booms. *Journal of Money, Credit, and Banking* 55 (7), 1749–1783.

ab Iorwerth, Aled, Baum-Snow, Nathaniel, Macek, James, 2025. Housing supply and affordability in Canada. *Canadian Journal of Economics* 58.

Ackerberg, Daniel, Benkard, C. Lanier, Berry, Steven, Pakes, Ariel, 2007. Econometric tools for analyzing market outcomes. In: Heckman, James J., Leamer, Edward E. (Eds.), *Handbook of Econometrics*, Volume 6A. North-Holland, Amsterdam, pp. 4271–4276.

Ahlfeldt, Gabriel M., Baum-Snow, Nathaniel, Jedwab, Rémi, 2023. The skyscraper revolution: Global economic development and land savings. CESifo Working Paper.

Ahlfeldt, Gabriel M., McMillen, Daniel P., 2013. New estimates of the elasticity of substitution of land for capital. Processed. London School of Economics.

Ahlfeldt, Gabriel M., McMillen, Daniel P., 2018. Tall buildings and land values: height and construction cost elasticities in Chicago, 1870–2010. *Review of Economics and Statistics* 100 (5), 861–875.

Albouy, David, 2009. The unequal geographic burden of federal taxation. *Journal of Political Economy* 117 (4), 635–667.

Albouy, David, Behrens, Kristian, Robert-Nicoud, Frédéric, Seegert, Nathan, 2019. The optimal distribution of population across cities. *Journal of Urban Economics* 110, 102–113.

Albouy, David, Ehrlich, Gabriel, 2018. Housing productivity and the social cost of land-use restrictions. *Journal of Urban Economics* 107, 101–120.

Albouy, David, Ehrlich, Gabriel, Shin, Minchul, 2018. Metropolitan land values. *Review of Economics and Statistics* 100 (3), 454–466.

Allen, Steven G., 1985. Why construction industry productivity is declining. *Review of Economics and Statistics* 67, 601–669.

Almagro, Milena, Chyn, Eric, Stuart, Bryan A., 2023. Urban renewal and inequality: Evidence from Chicago's public housing demolitions. Technical report. National Bureau of Economic Research.

Almagro, Milena, Domínguez-Iino, Tomás, 2025. Location sorting and endogenous amenities: Evidence from Amsterdam. Technical report.

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

Altshuler, Alan A., Gómez-Ibáñez, José A., 2000. *Regulation for Revenue: The Political Economy of Land Use Exactions*. Brookings Institution Press, Washington DC.

Ambrose, Brent W., Coulson, N. Edward, Yoshida, Jiro, 2015. The repeat rent index. *Review of Economics and Statistics* 97 (5), 939–950.

Anagol, Santosh, Ferreira, Fernando, Rexer, Jonah, 2025. Estimating the economic value of zoning reform. *American Economic Journal: Economic Policy*. In press.

Anas, Alex, Pines, David, 2008. Anti-sprawl policies in a system of congested cities. *Regional Science and Urban Economics* 38 (5), 408–423.

Anas, Alex, Rhee, Hyok-Joo, 2006. Curbing excess sprawl with congestion tolls and urban boundaries. *Regional Science and Urban Economics* 36 (4), 510–541.

Andrews, Dan, 2011. Housing markets and structural policies in oecd countries.

Arcidiacono, Peter, Miller, Robert A., 2011. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79 (6), 1823–1867.

Arnott, Richard, 1987. Economic theory and housing. In: Mills, Edwin S. (Ed.), *Handbook of Regional and Urban Economics*, vol. 2. Elsevier, North-Holland, pp. 959–988.

Arnott, Richard, 1995. Time for revisionism on rent control? *The Journal of Economic Perspectives* 9 (1), 99–120.

Arnott, Richard J., Braid, Ralph M., 1997. A filtering model with steady-state housing. *Regional Science and Urban Economics* 27 (4–5), 515–546.

Artmann, Martina, Inostroza, Luis, Fan, Peilei, 2019. Urban sprawl, compact urban development and green cities. How much do we know, how much do we agree?.

Asker, John, Collard-Wexler, Allan, De Loecker, Jan, 2014. Dynamic inputs and resource (mis) allocation. *Journal of Political Economy* 122 (5), 1013–1063.

Asquith, Brian J., Mast, Evan, Reed, Davin, 2023. Local effects of large new apartment buildings in low-income areas. *Review of Economics and Statistics* 105 (2), 359–375.

Autor, David H., Palmer, Christopher J., Pathak, Parag A., 2014. Housing market spillovers: evidence from the end of rent control in Cambridge, Massachusetts. *Journal of Political Economy* 122 (3), 661–717.

Bababievsky, Fil, Herkenhoff, Kyle F., Ohanian, Lee E., Prescott, Edward C., 2023. The impact of commercial real estate regulations on US output. Technical report, National Bureau of Economic Research Working Paper 31985.

Bar-Ilan, Avner, Strange, William C., 1996. Investment lags. *American Economic Review* 86 (3), 610–622.

Barr, Jason M., 2016. *Building the Skyline: The Birth and Growth of Manhattan's Skyscrapers*. Oxford University Press, New York City (NY).

Barseghyan, Levon, Coate, Stephen, 2016. Property taxation, zoning, and efficiency in a dynamic Tiebout model. *American Economic Journal: Economic Policy* 8 (3), 1–38.

Bartik, Alexander, Gupta, Arpit, Milo, Daniel, 2024. The costs of housing regulation: Evidence from generative regulatory measurement. Technical report. Working Paper. New York University, Stern School of Business.

Baum-Snow, Nathaniel, 2023. Constraints on city and neighborhood growth: the central role of housing supply. *The Journal of Economic Perspectives* 37 (2), 53–74.

Baum-Snow, Nathaniel, Han, Lu, 2024. The microgeography of housing supply. *Journal of Political Economy* 132 (6), 1897–1946.

Baum-Snow, Nathaniel, Marion, Justin, 2009. The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics* 93 (5–6), 654–666.

Bayer, Patrick, Ferreira, Fernando, McMillan, Robert, 2007. A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy* 115 (4), 588–638.

Bayliss, Simon, Bergin, Rory, 2020. The modular world. In: Bayliss, Simon, Bergin, Rory (Eds.), *The Modular Housing Handbook*. RIBA Publishing, London (UK), pp. 116–139.

Beaudry, Paul, Green, David A., Sand, Benjamin M., 2014. Spatial equilibrium with unemployment and wage bargaining: theory and estimation. *Journal of Urban Economics* 79 (1), 2–19.

Been, Vicki, Ellen, Ingrid Gould, Gedal, Michael, Glaeser, Edward, McCabe, Brian J., 2016. Preserving history or restricting development? The heterogeneous effects of historic districts on local housing markets in New York City. *Journal of Urban Economics* 92, 16–30.

Behrens, Kristian, Duranton, Gilles, Robert-Nicoud, Frédéric, 2014. Productive cities: sorting, selection, and agglomeration. *Journal of Political Economy* 122 (3), 507–553.

Behrens, Kristian, Kanemoto, Yoshitsugu, Murata, Yasusada, 2015. The Henry George theorem in a second-best world. *Journal of Urban Economics* 85, 34–51.

Behrens, Kristian, Robert-Nicoud, Frédéric, 2015. Agglomeration theory with heterogeneous agents. In: Duranton, Gilles, Henderson, Vernon, Strange, William (Eds.), *Handbook of Regional and Urban Economics*, Volume 5A. Elsevier, Amsterdam, pp. 171–245.

Ben-Moshe, Dan, Genesove, David, 2024. Regulation and Frontier Housing Supply. Processed. The Hebrew University of Jerusalem.

Ben-Shahar, Danny, Deng, Yongheng, Solganik, Eyal, Somerville, Tsur, Hongjia, Zhu, 2022. The value of vertical status: Evidence from the real estate market. Processed. University of British Columbia.

Bergman, Peter, Chetty, Raj, DeLuca, Stefanie, Hendren, Nathaniel, Katz, Lawrence F., Palmer, Christopher, 2024. Creating moves to opportunity: experimental evidence on barriers to neighborhood choice. *American Economic Review* 114 (5), 1281–1337.

Bertaud, Alain, Brueckner, Jan K., 2005. Analyzing building-height restrictions: predicted impacts and welfare costs. *Regional Science and Urban Economics* 35 (2), 109–125.

Besley, Timothy, Meads, Neil, Surico, Paolo, 2014. The incidence of transaction taxes: evidence from a stamp duty holiday. *Journal of Public Economics* 119, 61–70.

Best, Michael Carlos, Kleven, Henrik Jacobsen, 2018. Housing market responses to transaction taxes: evidence from notches and stimulus in the uk. *The Review of Economic Studies* 85 (1), 157–193.

Braid, Ralph M., 1981. The short-run comparative statics of a rental housing market. *Journal of Urban Economics* 10 (3), 286–310.

Bratu, Cristina, Harjunen, Oskari, Saarimaa, Tuukka, 2023. City-wide effects of new housing supply: evidence from moving chains. *Journal of Urban Economics: Insights* 133, 103528.

Bronin, Sara C., Markley, Scott, Fader, Aline, Derickson, Evan, 2023. How to make a zoning atlas 2.0: the official methodology of the National Zoning Atlas. Technical report. Cornell University, Department of City and Regional Planning.

Brooks, Leah, Liscow, Zachary, 2023. Infrastructure costs. *American Economic Journal: Applied Economics* 15 (2), 1–30.

Brooks, Leah, Lutz, Byron, 2016. From today's city to tomorrow's city: an empirical investigation of urban land assembly. *American Economic Journal: Economic Policy* 8 (3), 69–105.

Brueckner, Jan, Fu, Shihe, Gu, Yizhen, Zhang, Junfu, 2017. Measuring the stringency of land use regulation: the case of China's building height limits. *Review of Economics and Statistics* 99 (4), 663–677.

Brueckner, Jan K., 2000. Urban sprawl: diagnosis and remedies. *International Regional Science Review* 23 (2), 160–171.

Brueckner, Jan K., 2007. Urban growth boundaries: an effective second-best remedy for unpriced traffic congestion? *Journal of Housing Economics* 16 (3–4), 263–273.

Brueckner, Jan K., 2023. Jue insight: zoning and property taxation revisited—was Hamilton right? *Journal of Urban Economics* 133, 103393.

Brueckner, Jan K., Leather, David, Zerecero, Miguel, 2024. Bunching in real-estate markets: regulated building heights in New York City. *Journal of Urban Economics* 143, 103683.

Brueckner, Jan K., Rosenthal, Stuart S., 2009. Gentrification and neighborhood housing cycles: will America's future downtowns be rich? *Review of Economics and Statistics* 91 (4), 725–743.

Brueckner, Jan K., Singh, Ruchi, 2020. Stringency of land-use regulation: building heights in US cities. *Journal of Urban Economics* 116, 103239.

Brunner, Eric, Sonstelie, Jon, 2003. Homeowners, property values, and the political economy of the school voucher. *Journal of Urban Economics* 54 (2), 239–257.

Brunner, Eric, Sonstelie, Jon, Thayer, Mark, 2001. Capitalization and the voucher: an analysis of precinct returns from California's Proposition 174. *Journal of Urban Economics* 50 (3), 517–536.

Büchler, Simon, Ehrlich, Maximilian v., Schöni, Olivier, 2021. The amplifying effect of capitalization rates on housing supply. *Journal of Urban Economics* 126, 103370.

Burchfield, Marcy, Overman, Henry G., Puga, Diego, Turner, Matthew A., 2006. Causes of sprawl: a portrait from space. *The Quarterly Journal of Economics* 121 (2), 587–634.

Cai, Hongbin, Wang, Zhi, Zhang, Qinghua, 2017. To build above the limit? Implementation of land use regulations in urban China. *Journal of Urban Economics* 98, 223–233.

Calabrese, Stephen, Epple, Dennis, Romano, Richard, 2007. On the political economy of zoning. *Journal of Public Economics* 91 (1–2), 25–49.

Calder-Wang, Sophie, Kim, Gi Heung, 2024. Algorithmic pricing in multifamily rentals: Efficiency gains or price coordination? Working Paper. Wharton School, University of Pennsylvania.

Capozza, Dennis R., Helsley, Robert W., 1990. The stochastic city. *Journal of Urban Economics* 28 (2), 187–203.

Cheshire, Paul C., Dericks, Gerard H., 2020. 'Trophy architects' and design as rent-seeking: quantifying deadweight losses in a tightly regulated office market. *Economica* 87 (348), 1078–1104.

Cheshire, Paul C., Hilber, Christian A.L., 2008. Office space supply restrictions in Britain: the political economy of market revenge. *The Economic Journal* 118 (529), F185–F221.

Chetty, Raj, Friedman, John N., Hendren, Nathaniel, Jones, Maggie R., Porter, Sonya R., 2018. The opportunity atlas: Mapping the childhood roots of social mobility. Technical report. National Bureau of Economic Research.

Clapp, John M., 1979. The substitution of urban land for other inputs. *Journal of Urban Economics* 6 (1), 122–134.

Clarke, Wyatt, Freedman, Matthew, 2019. The rise and effects of homeowners associations. *Journal of Urban Economics* 112, 1–15.

Collinson, Robert, Ganong, Peter, 2018. How do changes in housing voucher design affect rent and neighborhood quality? *American Economic Journal: Economic Policy* 10 (2), 62–89.

Combes, Pierre-Philippe, Duranton, Gilles, Gobillon, Laurent, 2019. The costs of agglomeration: house and land prices in French cities. *The Review of Economic Studies* 86 (4), 1556–1589.

Combes, Pierre-Philippe, Duranton, Gilles, Gobillon, Laurent, 2021. The production function for housing: evidence from France. *Journal of Political Economy* 129 (10), 2766–2816.

Combes, Pierre-Philippe, Duranton, Gilles, Gobillon, Laurent, Gorin, Clément, 2024. Measuring land use changes by (machine) learning from historical maps. Mimeo graphed. University of Pennsylvania.

Combes, Pierre-Philippe, Duranton, Gilles, Gobillon, Laurent, Roux, Sébastien, 2010. Estimating agglomeration economies with history, geology, and worker effects. In: Glaeser, Edward L. (Ed.), *The Economics of Agglomeration*. National Bureau of Economic Research, Cambridge, MA, pp. 15–65.

Combes, Pierre-Philippe, Gobillon, Laurent, 2015. The empirics of agglomeration economies. In: Duranton, Gilles, Henderson, Vernon, Strange, William (Eds.), *Handbook of Regional and Urban Economics*, Volume 5A. Elsevier, Amsterdam, pp. 247–348.

Cosman, Jacob, Davidoff, Tom, Williams, Joe, 2018. Housing appreciation and marginal land supply in monocentric cities with topography. Working paper. University of British Columbia.

Couture, Victor, Gaubert, Cecile, Handbury, Jessie, Hurst, Erik, 2024. Income growth and the distributional effects of urban spatial sorting. *The Review of Economic Studies* 91 (2), 858–898.

Crews, Levi, 2024. A dynamic spatial knowledge economy. Working Paper. University of California Los Angeles.

Cui, Tianfang, 2024. Did race fence off the American city? The great migration and the evolution of exclusionary zoning. Working paper. New York University, Robert F. Wagner School of Public Service.

Cunningham, Christopher R., 2007. Growth controls, real options, and land development. *Review of Economics and Statistics* 89 (2), 343–358.

Cunningham, Christopher R., Engelhardt, Gary V., 2008. Housing capital-gains taxation and homeowner mobility: evidence from the taxpayer relief act of 1997. *Journal of Urban Economics* 63 (3), 803–815.

Dachis, Ben, Duranton, Gilles, Turner, Matthew A., 2012. The effects of land transfer taxes on real estate markets: evidence from a natural experiment in Toronto. *Journal of Economic Geography* 12 (2), 327–354.

D'Amico, Leonardo, Glaeser, Edward L., Gyourko, Joseph, Kerr, William, Ponzetto, Giacomo A.M., 2024. Why has construction productivity stagnated? The role of land-use regulation. Working Paper. Harvard University.

Danton, Jayson, Himbert, Alexander, 2018. Residential vertical rent curves. *Journal of Urban Economics* 107, 89–100.

Davidoff, Thomas, 2013. Supply elasticity and the housing cycle of the 2000s. *Real Estate Economics* 41 (4), 793–813.

Davidoff, Thomas, 2015. Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors. Working paper. University of British Columbia, Sauder School of Business.

Davidoff, Thomas, Pavlov, Andrey, Somerville, Tsur, 2022. Not in my neighbour's back yard? Laneway homes and neighbours' property values. *Journal of Urban Economics* 128, 103405.

Davis, Donald R., Dingel, Jonathan I., 2020. The comparative advantage of cities. *Journal of International Economics* 123, 103291.

Davis, Morris A., Ghent, Andra C., Gregory, Jesse, 2024. The work-from-home technology boon and its consequences. *The Review of Economic Studies* 91, 3362–3401.

Davis, Morris A., Larson, William D., Oliner, Stephen D., Shui, Jessica, 2021. The price of residential land for counties, ZIP Codes, and census tracts in the United States. *Journal of Monetary Economics* 118, 413–431.

Davis, Morris A., Palumbo, Michael G., 2008. The price of residential land in large US cities. *Journal of Urban Economics* 63 (1), 352–384.

Dehring, Carolyn A., Depken II, Craig A., Ward, Michael R., 2008. A direct test of the homevoter hypothesis. *Journal of Urban Economics* 64 (1), 155–170.

Department of Housing, Urban Development, 2023. A picture of subsidized households in 2023. US Department of Housing & Urban Development, Office of Policy Development.

Desmet, Klaus, Parro, Fernando, 2025. Spatial dynamics. In: Donaldson, Dave, Redding, Stephen J. (Eds.), *Handbook of Regional and Urban Economics*, vol. 6. Elsevier North Holland, Amsterdam.

Diamond, Rebecca, 2016. The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *American Economic Review* 106 (3), 479–524.

Diamond, Rebecca, McQuade, Tim, 2019. Who wants affordable housing in their backyard? An equilibrium analysis of low-income property development. *Journal of Political Economy* 127 (3), 1063–1117.

Diamond, Rebecca, McQuade, Tim, Qian, Franklin, 2019. The effects of rent control expansion on tenants, landlords, and inequality: evidence from San Francisco. *American Economic Review* 109 (9), 3365–3394.

DiPasquale, Denise, 1999. Why don't we know more about housing supply? *Journal of Real Estate Finance and Economics* 18 (1), 9–23.

Dumm, Randy E., Sirmans, G. Stacy, Smersh, Greg, 2011. The capitalization of building codes in house prices. *Journal of Real Estate Finance and Economics* 42, 30–50.

Duranton, Gilles, Puga, Diego, 2014. The growth of cities. In: Aghion, Philippe, Durlauf, Steven (Eds.), *Handbook of Economic Growth*, Volume 2. North-Holland, Amsterdam, pp. 781–853.

Duranton, Gilles, Puga, Diego, 2015. Urban land use. In: Duranton, Gilles, Henderson, J. Vernon, Strange, William C. (Eds.), *Handbook of Regional and Urban Economics*, Volume 5A. North-Holland, Amsterdam, pp. 467–560.

Duranton, Gilles, Puga, Diego, 2023. Urban growth and its aggregate implications. *Econometrica* 91 (6), 2219–2259.

Duranton, Gilles, Turner, Matthew A., 2018. Urban form and driving: evidence from US cities. *Journal of Urban Economics* 108, 171–191.

Eckert, Fabian, Kleineberg, Tatjana, 2021. The geography of opportunity: Education, work, and intergenerational mobility across US counties. Working paper. University of California San Diego.

Ellickson, Robert C., 2020. The zoning straitjacket: the freezing of American neighborhoods of single-family houses. *Indiana Law Journal* 96, 395–427.

Apple, Dennis, Gordon, Brett, Sieg, Holger, 2010. A new approach to estimating the production function for housing. *American Economic Review* 100 (3), 905–925.

Apple, Dennis, Quintero, Luis, Sieg, Holger, 2020. A new approach to estimating equilibrium models for metropolitan housing markets. *Journal of Political Economy* 128 (3), 948–983.

Eriksen, Michael D., Orlando, Anthony W., 2022. Returns to scale in residential construction: the marginal impact of building height. *Real Estate Economics* 50 (2), 534–564.

Eriksen, Michael D., Rosenthal, Stuart S., 2010. Crowd out effects of place-based subsidized rental housing: new evidence from the lihtc program. *Journal of Public Economics* 94 (11–12), 953–966.

Esch, Thomas, Deininger, Klaus W., Jedwab, Rémi, Palacios-Lopez, Daniela, 2024. Outward and upward construction: a 3D analysis of the global building stock. Working paper. World Bank Group.

Ewing, Reid, Cervero, Robert, 2010. Travel and the built environment: a meta-analysis. *Journal of the American Planning Association* 76 (3), 265–294.

Favara, Giovanni, Imbs, Jean, 2015. Credit supply and the price of housing. *American Economic Review* 105 (3), 958–992.

Favilukis, Jack, Mabille, Pierre, Van Nieuwerburgh, Stijn, 2023. Affordable housing and city welfare. *The Review of Economic Studies* 90 (1), 293–330.

Favilukis, Jack, Van Nieuwerburgh, Stijn, 2021. Out-of-town home buyers and city welfare. *The Journal of Finance* 76 (5), 2577–2638.

Fernandes Rocha, Patrícia, Ferreira, Nuno Oliveira, Pimenta, Fernando, Pereira, Nelson Bento, 2022. Impacts of prefabrication in the building construction industry. *Encyclopedia* 3 (1), 28–45.

Ferreira, Fernando, 2010. You can take it with you: Proposition 13 tax benefits, residential mobility, and willingness to pay for housing amenities. *Journal of Public Economics* 94 (9–10), 661–673.

Fischel, William A., 2001. *The Homevoter Hypothesis*. Harvard University Press, Cambridge (MA).

Foerster, Andrew T., Hornstein, Andreas, Sarte, Pierre-Daniel G., Watson, Mark W., 2022. Aggregate implications of changing sectoral trends. *Journal of Political Economy* 130 (12), 3286–3333.

Fogli, Alessandra, Guerrieri, Veronica, Ponder, Mark, Prato, Marta, 2024. The end of the American dream? Inequality and segregation in US cities. Working Paper. Federal Reserve Bank of Minneapolis.

Francke, Marc K., van de Minne, Alex M., 2017. Land, structure and depreciation. *Real Estate Economics* 45 (2), 415–451.

French, Robert, Gilbert, Valentine, 2024. Suburban housing and urban affordability: Evidence from residential vacancy chains. Working paper. Harvard Kennedy School.

Ganong, Peter, Shoag, Daniel, 2017. Why has regional income convergence in the US declined? *Journal of Urban Economics* 102, 76–90.

Garcia, Daniel, Molloy, Raven, 2023. Can Measurement Error Explain Slow Productivity Growth in Construction? *Finance and Economics Discussion Series* 2023-052. Federal Reserve Board.

Gechter, Michael, Tsivanidis, Nick, 2023. Spatial spillovers from high-rise developments: Evidence from the Mumbai Mills. Working paper. University of California Berkeley, Haas School of Business.

Gerardi, Kristopher, Rosenblatt, Eric, Willen, Paul S., Yao, Vincent, 2015. Foreclosure externalities: new evidence. *Journal of Urban Economics* 87, 42–56.

Glaeser, Edward, Gyourko, Joseph, 2018. The economic implications of housing supply. *The Journal of Economic Perspectives* 32 (1), 3–30.

Glaeser, Edward L., Gyourko, Joseph, 2005. Urban decline and durable housing. *Journal of Political Economy* 113 (2), 345–375.

Glaeser, Edward L., Gyourko, Joseph, Saks, Raven, 2005a. Why is Manhattan so expensive? Regulation and the rise in housing prices. *The Journal of Law & Economics* 48 (2), 331–369.

Glaeser, Edward L., Gyourko, Joseph, Saks, Raven E., 2005b. Why have housing prices gone up? *American Economic Review* 95 (2), 329–333.

Glaeser, Edward L., Luttmer, Erzo F.P., 2003. The misallocation of housing under rent control. *American Economic Review* 93 (4), 1027–1046.

Glaeser, Edward L., Ward, Bryce A., 2009. The causes and consequences of land use regulation: evidence from Greater Boston. *Journal of Urban Economics* 65 (3), 265–278.

González-Pampillón, Nicolás, 2022. Spillover effects from new housing supply. *Regional Science and Urban Economics* 92, 103759.

Goolsbee, Austan, Syverson, Chad, 2023. The strange and awful path of productivity in the US construction sector. National Bureau of Economic Research Working Paper 30845.

Gorback, Caitlin S., Keys, Benjamin J., 2023. Global capital and local assets: House prices, quantities, and elasticities. Technical report. Working paper. University of Pennsylvania, Wharton School of Business.

Greaney, Brian, 2019. Housing constraints and spatial misallocation: Comment. Working paper. University of Washington, Department of Economics.

Greenaway-McGrevy, Ryan, Pacheco, Gail, Sorensen, Kade, 2021. The effect of upzoning on house prices and redevelopment premiums in Auckland, New Zealand. *Urban Studies* 58 (5), 959–976.

Greenaway-McGrevy, Ryan, Phillips, Peter C.B., 2023. The impact of upzoning on housing construction in Auckland. *Journal of Urban Economics* 136, 103555.

Greenwald, Daniel L., Guren, Adam, 2021. Do credit conditions move house prices? Technical report. National Bureau of Economic Research.

Gruber, Jonathan, Jensen, Amalie, Kleven, Henrik, 2021. Do people respond to the mortgage interest deduction? Quasi-experimental evidence from Denmark. *American Economic Journal: Economic Policy* 13 (2), 273–303.

Guerrieri, Veronica, Hartley, Daniel, Hurst, Erik, 2013. Endogenous gentrification and housing price dynamics. *Journal of Public Economics* 100, 45–60.

Gupta, Arpit, Van Nieuwerburgh, Stijn, Kontokosta, Constantine, 2022. Take the Q train: value capture of public infrastructure projects. *Journal of Urban Economics* 129, 103422.

Gyourko, Joe, Krimmel, Jacob, 2021. The impact of local residential land use restrictions on land values across and within single family housing markets. *Journal of Urban Economics* 126 (103374), 1–14.

Gyourko, Joseph, Hartley, Jonathan S., Krimmel, Jacob, 2021. The local residential land use regulatory environment across US housing markets: evidence from a new Wharton index. *Journal of Urban Economics* 124, 103337.

Gyourko, Joseph, Mayer, Christopher, Sinai, Todd, 2013. Superstar cities. *American Economic Journal: Economic Policy* 5 (4), 167–199.

Gyourko, Joseph, McCulloch, Sean E., 2024. The distaste for housing density. Technical report. National Bureau of Economic Research.

Gyourko, Joseph, Molloy, Raven, 2015. Regulation and housing supply. In: Duranton, Gilles, Henderson, Vernon, Strange, William (Eds.), *Handbook of Regional and Urban Economics*, Volume 5B. Elsevier, Amsterdam, pp. 1289–1337.

Gyourko, Joseph, Saiz, Albert, 2004. Reinvestment in the housing stock: the role of construction costs and the supply side. *Journal of Urban Economics* 55 (2), 238–256.

Gyourko, Joseph, Saiz, Albert, 2006. Construction costs and the supply of housing structure. *Journal of Regional Science* 46 (4), 661–680.

Gyourko, Joseph, Saiz, Albert, Summers, Anita A., 2008. A new measure of the local regulatory environment for housing markets: the Wharton residential land use regulatory index. *Urban Studies* 45 (3), 693–729.

Hamilton, Bruce W., 1975. Zoning and property taxation in a system of local governments. *Urban Studies* 12 (2), 205–211.

Hamilton, Bruce W., 1976. Capitalization of intrajurisdictional differences in local tax prices. *American Economic Review* 66 (5), 743–753.

Han, Lu, Ngai, Liwa Rachel, Sheedy, Kevin D., 2022. To Own or to Rent?: the Effects of Transaction Taxes on Housing Markets. Centre for Economic Policy Research.

Handbury, Jessie, Huges, Samuel K., Keys, Benjamin J., 2024. Segmentations and returns on rental housing. Working paper. University of Pennsylvania, Wharton School.

Harberger, Arnold C., 1954. Monopoly and resource allocation. *American Economic Review* 44 (2), 77–87.

Haughwout, Andrew, Orr, James, Bedoll, David, 2008. The price of land in the New York metropolitan area. *Current Issues in Economics and Finance* 14 (3).

Henderson, J. Vernon, Regan, Tanner, Venables, Anthony J., 2021. Building the city: from slums to a modern Metropolis. *The Review of Economic Studies* 88 (3), 1157–1192.

Henderson, J. Vernon, Venables, Anthony J., 2009. The dynamics of city formation. *Review of Economic Dynamics* 39 (2), 233–254.

Herkenhoff, Kyle F., Ohanian, Lee E., Prescott, Edward C., 2018. Tarnishing the golden and empire states: land-use restrictions and the us economic slowdown. *Journal of Monetary Economics* 93, 89–109.

Hilber, Christian, Robert-Nicoud, Frédéric, 2013. On the origins of land use regulations: theory and evidence from US metro areas. *Journal of Urban Economics* 75 (1), 29–43.

Hilber, Christian, Schöni, Olivier, 2022. Housing Policy and Affordable Housing. Centre for Economic Performance, London School of Economics and Political

Hilber, Christian A.L., Mayer, Christopher, 2009. Why do households without children support local public schools? Linking house price capitalization to school spending. *Journal of Urban Economics* 65 (1), 74–90.

Hilber, Christian A.L., Turner, Tracy M., 2014. The mortgage interest deduction and its impact on homeownership decisions. *Review of Economics and Statistics* 96 (4), 618–637.

Hilber, Christian A.L., Vermeulen, Wouter, 2016. The impact of supply constraints on house prices in England. *The Economic Journal* 126 (591), 358–405.

Hornbeck, Richard, Keniston, Daniel, 2017. Creative destruction: barriers to urban growth and the great Boston fire of 1872. *American Economic Review* 107 (6), 1365–1398.

Howard, Greg, Liebersohn, Jack, 2021. Why is the rent so darn high? The role of growing demand to live in housing-supply-inelastic cities. *Journal of Urban Economics* 124, 103369.

Howard, Greg, Liebersohn, Jack, Ozimek, Adam, 2023. The short- and long-run effects of remote work on US housing markets. *Journal of Financial Economics* 150 (1), 166–184.

Hsiao, Allan, 2023. Sea level rise and urban adaptation in Jakarta. Working paper. Princeton University.

Hsieh, Chang-Tai, Moretti, Enrico, 2019. Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics* 11 (2), 1–39.

Ihlanfeldt, Keith R., 2007. The effect of land use regulation on housing and land prices. *Journal of Urban Economics* 61 (3), 420–435.

İmrohoroglu, Ayşe, Matoba, Kyle, Tüzel, Şelale, 2018. Proposition 13: an equilibrium analysis. *American Economic Journal: Macroeconomics* 10 (2), 24–51.

Jackson, Kristoffer, 2016. Do land use regulations stifle residential development? Evidence from California cities. *Journal of Urban Economics* 91, 45–56.

Jackson, Kristoffer Kip, 2018. Regulation, land constraints, and California's boom and bust. *Regional Science and Urban Economics* 68, 130–147.

Jia, Ning, Molloy, Raven, Smith, Christopher, Wozniak, Abigail, 2023. The economics of internal migration: advances and policy questions. *Journal of Economic Literature* 61 (1), 144–180.

Justiniano, Alejandro, Primiceri, Giorgio E., Tambalotti, Andrea, 2019. Credit supply and the housing boom. *Journal of Political Economy* 127 (3), 1317–1350.

Kahn, Matthew E., Vaughn, Ryan, Zasloff, Jonathan, 2010. The housing market effects of discrete land use regulations: evidence from the California coastal boundary zone. *Journal of Housing Economics* 19 (4), 269–279.

Kalouptsidi, Myrto, 2014. Time to build and fluctuations in bulk shipping. *American Economic Review* 104 (2), 564–608.

Kok, Nils, Monkkonen, Paavo, Quigley, John M., 2014. Land use regulations and the value of land and housing: an intra-metropolitan analysis. *Journal of Urban Economics* 81, 136–148.

Kopczuk, Wojciech, Munroe, David, 2015. Mansion tax: the effect of transfer taxes on the residential real estate market. *American Economic Journal: Economic Policy* 7 (2), 214–257.

Koster, Hans R.A., 2024. The welfare effects of greenbelt policy: evidence from England. *The Economic Journal* 134 (657), 363–401.

Koster, Hans R.A., van Ommeren, Jos, Rietveld, Piet, 2013. Is the sky the limit? High-rise buildings and office rents. *Journal of Economic Geography* 14 (1), 125–153.

Krimmel, Jacob, 2021. Reclaiming local control: School finance reforms and housing supply restrictions. FRB Working Paper.

Krimmel, Jake, Wang, Betty X., 2024. Upzoning with strings attached: Evidence from Seattle's Affordable Housing Mandate. Working paper. Federal Reserve Board.

Kuhlmann, Daniel, 2021. Upzoning and single-family housing prices: a (very) early analysis of the Minneapolis 2040 plan. *Journal of the American Planning Association* 87 (3), 383–395.

Kulka, Amrita, 2019. Sorting into neighborhoods: the role of minimum lot sizes. In: Conference Paper, APPAM 41st Annual Fall Research Conference. Denver, CO.

Kulka, Amrita, Sood, Aradhya, Chiumenti, Nicholas, 2022. How to increase housing affordability: Understanding local deterrents to building multifamily housing. Technical report, Working Papers. Warwick University, Department of Economics.

Landvoigt, Tim, Piazzesi, Monika, Schneider, Martin, 2015. The housing market (s) of San Diego. *American Economic Review* 105 (4), 1371–1407.

Lebret, Daniel, Liu, Crocker H., Valentin, Maxence, 2025. Carrot and stick zoning. Working paper. Cornell University, SC Johnson College of Business.

Li, Xiaodi, 2021. Do new housing units in your backyard raise your rents? *Journal of Economic Geography* 22 (6), 1309–1352.

Liang, Linlin, Staveski, Adam, Horowitz, Alex, 2024. Minneapolis land use reforms offer a blueprint for housing affordability. Technical report. The Pew Charitable Trusts.

Listokin, David, Hattis, David B., 2005. Building codes and housing. *Cityscape* 8 (1), 21–67.

Liu, Crocker, Rosenthal, Stuart S., Strange, William C., 2018. The vertical city: rent gradients and spatial structure. *Journal of Urban Economics* 106, 101–122.

Liu, Liyi, McManus, Doug, Yannopoulos, Elias, 2022. Geographic and temporal variation in housing filtering rates. *Regional Science and Urban Economics* 93, 103758.

Lutz, Chandler, Sand, Ben, 2023. Highly disaggregated land unavailability. Working paper. York University, Department of Economics.

Määttänen, Niku, Terviö, Marko, 2014. Income distribution and housing prices: an assignment model approach. *Journal of Economic Theory* 151, 381–410.

Macek, James, 2024. Housing regulation and neighborhood sorting across the United States. Working Paper. University of Toronto, Rotman School of Business.

Mast, Evan, 2023. JUE insight: the effect of new market-rate housing construction on the low-income housing market. *Journal of Urban Economics* 133, 103383.

Mast, Evan, 2024. Warding off development: local control, housing supply, and NIMBYs. *Review of Economics and Statistics* 106 (3), 671–680.

Mayer, Christopher J., Somerville, C. Tsuriel, 2000. Land use regulation and new construction. *Regional Science and Urban Economics* 30 (6), 639–662.

McDonald, John F., 1979. *Economic Analysis of an Urban Housing Market*. Academic Press, New York.

McDonald, John F., 1981. Capital-land substitution in urban housing: a survey of empirical estimates. *Journal of Urban Economics* 9 (2), 190–211.

McMillen, Daniel P., McDonald, John F., 2002. Land values in a newly zoned city. *Review of Economics and Statistics* 84 (1), 62–72.

Mehrotra, Neil, Turner, Matthew A., Uribe, Juan Pablo, 2024. Does the US have an infrastructure cost problem? Evidence from the interstate highway system. *Journal of Urban Economics* 143, 103681.

Mense, Andreas, 2025. The impact of new housing supply on the distribution of rents. *Journal of Political Economy Macroeconomics* 3 (1), 1–42.

Mills, Edwin S., 1969. The value of urban land. In: Perloff, Harvey S. (Ed.), *The Quality of the Urban Environment. Resources for the Future*, Baltimore MD, pp. 231–253.

Mleczko, Matthew, Desmond, Matthew, 2023. Using natural language processing to construct a national zoning and land use database. *Urban Studies* 60 (13), 2564–2584.

Molloy, Raven, 2020. The effect of housing supply regulation on housing affordability: a review. *Regional Science and Urban Economics* 80 (C), 1–5.

Molloy, Raven, Nathanson, Charles G., Paciorek, Andrew, 2022. Housing supply and affordability: evidence from rents, housing consumption and household location. *Journal of Urban Economics* 129.

Murphy, Alvin, 2018. A dynamic model of housing supply. *American Economic Journal: Economic Policy* 10 (4), 243–267.

Muth, Richard F., 1964. The derived demand curve for a productive factor and the industry supply curve. *Oxford Economic Papers* 16 (2), 221–234.

Muth, Richard F., 1969. *Cities and Housing*. University of Chicago Press, Chicago.

Muth, Richard F., 1975. Numerical solution of urban residential land-use models. *Journal of Urban Economics* 4 (2), 307–332.

Nathanson, Charles G., 2023. Trickle-down housing economics. Working paper. Northwestern University, Kellogg School of Management.

Nikolakoudis, George, 2024. *The Economics of Segmented Housing Markets*. Unpublished. Princeton University.

Oates, Wallace E., 1972. Fiscal federalism. HarcourtBrace Jovanovich Inc.

Olsen, Edgar O., 1987. The demand and supply of housing service: a critical survey of the empirical literature. In: Mills, Edwin S. (Ed.), *Handbook of Regional and Urban Economics*, Volume 2. Elsevier, North-Holland, pp. 989–1022.

Ortalo-Magné, Francois, Rady, Sven, 2006. Housing market dynamics: on the contribution of income shocks and credit constraints. *The Review of Economic Studies* 73 (2), 459–485.

Oster, Sharon M., Quigley, John M., 1977. Regulatory barriers to the diffusion of innovation: some evidence from building codes. *Bell Journal of Economics* 8 (2), 361–377.

Overman, Henry G., Puga, Diego, Turner, Matthew A., 2008. Decomposing the growth in residential land in the United States. *Regional Science and Urban Economics* 38 (5), 487–497.

Owens III, Raymond, Rossi-Hansberg, Esteban, Sarte, Pierre-Daniel, 2020. Rethinking Detroit. *American Economic Journal: Economic Policy* 12 (2), 258–305.

Paciorek, Andrew, 2013. Supply constraints and housing market dynamics. *Journal of Urban Economics* 77, 11–26.

Parkhomenko, Andrii, 2023. Local causes and aggregate implications of land use regulation. *Journal of Urban Economics* 138, 103605.

Peng, Xuequan Elsie, 2023. The dynamics of urban development: Evidence from land use reform in New York. Working Paper. Department of Economics, University of Pennsylvania.

Pennington, Kate, 2021. Does building new housing cause displacement? The supply and demand effects of construction in San Francisco. Working paper. University of California, Berkeley.

Pines, David, Sadka, Efraim, 1985. Zoning, first-best, second-best, and third-best criteria for allocating land for roads. *Journal of Urban Economics* 17 (2), 167–183.

Quigley, John M., Raphael, Steven, 2005. Regulation and the high cost of housing in California. *American Economic Review* 95 (2), 323–328.

Quintero, Luis, 2023. Fewer players, fewer homes: Concentration and the new dynamics of housing supply. Working Paper 18-18. Johns Hopkins University, Carey Business School.

Redding, Stephen J., 2025. Quantitative urban economics. In: Donaldson, Dave, Redding, Stephen J. (Eds.), *Handbook of Regional and Urban Economics*, vol. 6. Elsevier North Holland, Amsterdam.

Roback, Jennifer, 1982. Wages, rents and the quality of life. *Journal of Political Economy* 90 (6), 1257–1278.

Rosenthal, Stuart S., 2014. Are private markets and filtering a viable source of low-income housing? Estimates from a “repeat income” model. *American Economic Review* 104 (2), 687–706.

Rosenthal, Stuart S., Strange, William C., 2008. The attenuation of human capital spillovers. *Journal of Urban Economics* 64 (2), 373–389.

Rossi-Hansberg, Esteban, Sarte, Pierre-Daniel, Owens III, Raymond, 2010. Housing externalities. *Journal of Political Economy* 118 (3), 485–535.

Rossi-Hansberg, Esteban, Sarte, Pierre-Daniel, Schwartzman, Felipe, 2019. Cognitive hubs and spatial redistribution. Technical report. National Bureau of Economic Research. #26267.

Ruggles, Steven, Flood, Sarah, Sobek, Matthew, Backman, Daniel, Chen, Annie, Cooper, Grace, Richards, Stephanie, Rodgers, Renae, Schouweiler, Megan, 2024. *Ipums usa: Version 15.0* [dataset]. Technical report. IPUMS, Minneapolis, MN.

Rust, John, 1987. Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher. *Econometrica*, 999–1033.

Sahn, Alexander, 2021. Racial diversity and exclusionary zoning: evidence from the great migration. *The Journal of Politics*. Forthcoming.

Saiz, Albert, 2010. The geographic determinants of housing supply. *The Quarterly Journal of Economics* 125 (3), 1253–1296.

Satttinger, Michael, 1993. Assignment models of the distribution of earnings. *Journal of Economic Literature* 19 (1), 831–880.

Schmitz, James A., 2020. Solving the housing crisis will require fighting monopolies in construction. *Working Papers* 773. Federal Reserve Bank of Minneapolis.

Schuetz, Jenny, Meltzer, Rachel, Been, Vicki, 2011. Silver bullet or Trojan horse? The effects of inclusionary zoning on local housing markets in the United States. *Urban Studies* 48 (2), 297–329.

Severen, Christopher, Plantinga, Andrew J., 2018. Land-use regulations, property values, and rents: decomposing the effects of the California coastal act. *Journal of Urban Economics* 107, 65–78.

Shan, Hui, 2011. The effect of capital gains taxation on home sales: evidence from the taxpayer relief act of 1997. *Journal of Public Economics* 95 (1–2), 177–188.

Shanks, Brendan, 2021. Land use regulations and housing development: Evidence from tax parcels and zoning bylaws in Massachusetts. *Working paper*. Ludwig Maximilien University Munich, Department of Economics.

Shertzer, Allison, Twinam, Tate, Walsh, Randall P., 2016. Race, ethnicity, and discriminatory zoning. *American Economic Journal: Applied Economics* 8 (3), 217–246.

Shertzer, Allison, Twinam, Tate, Walsh, Randall P., 2018. Zoning and the economic geography of cities. *Journal of Urban Economics* 105, 20–39.

Sinai, Todd, Waldfogel, Joel, 2005. Do low-income housing subsidies increase the occupied housing stock? *Journal of Public Economics* 89 (11–12), 2137–2164.

Soltas, Evan, 2024a. Tax incentives and the supply of low-income housing. *Working paper*. MIT.

Soltas, Evan J., 2024b. The price of inclusion: evidence from housing developer behavior. *Review of Economics and Statistics* 106 (6), 1588–1606.

Sommer, Kamila, Sullivan, Paul, 2018. Implications of us tax policy for house prices, rents, and homeownership. *American Economic Review* 108 (2), 241–274.

Song, Jaehee, 2024. The effects of residential zoning in US housing markets. *Working paper*. University of Colorado-Boulder.

Spader, Jonathan, 2024. Has housing filtering stalled? Heterogeneous outcomes in the American housing survey, 1985–2021. In: *Housing Policy Debate*, pp. 1–23.

Stacy, Christina, Hodge, Timothy, Komarek, Timothy, Davis, Chris, Noble, Eleanor, Morales-Burnett, Jorge, Rogin, Amy, 2023. The impact of inclusionary zoning and rent control on housing affordability and access to opportunity. In: 2023 APPAM Fall Research Conference. APPAM.

Stokes, H. Kemble, 1981. An examination of the productivity decline in the construction industry. *Review of Economics and Statistics* 64, 495–502.

Sveikauskas, Leo, Rowe, Samuel, Mildenberger, James, Price, Jennifer, Young, Arthur, 2016. Productivity growth in construction. *Journal of Construction Engineering and Management* 142 (10), 1–8.

Sweeney, James L., 1974. Quality, commodity hierarchies, and housing markets. *Econometrica* 42 (1), 147–167.

Syverson, Chad, 2011. What determines productivity? *Journal of Economic Literature* 49 (2), 326–365.

Tan, Ya, Wang, Zhi, Zhang, Qinghua, 2020. Land-use regulation and the intensive margin of housing supply. *Journal of Urban Economics* 115, 103199.

Thorsnes, Paul, 1997. Consistent estimates of the elasticity of substitution between land and non-land inputs in the production of housing. *Journal of Urban Economics* 42 (1), 98–108.

Tiebout, Charles M., 1956. A pure theory of local expenditures. *Journal of Political Economy* 64 (5), 416–424.

Titman, Sheridan, 1985. Urban land prices under uncertainty. *American Economic Review* 75 (3), 505–514.

Tricaud, Clemence, 2025. Better alone? Evidence on the costs of intermunicipal cooperation. *American Economic Journal: Applied Economics* 17 (1), 160–207.

Turner, Matthew A., 2005. Landscape preferences and patterns of residential development. *Journal of Urban Economics* 57 (1), 19–54.

Turner, Matthew A., Haughwout, Andrew, Van Der Klaauw, Wilbert, 2014. Land use regulation and welfare. *Econometrica* 82 (4), 1341–1403.

Wang, Xiao Betty, 2022. Housing market segmentation. Preprint. New York University.

Wheeler, Harrison, 2022. Locally optimal place-based policies: Evidence from Opportunity Zones. Working paper. University of California Berkeley, Department of Economics.

Yoshida, Jiro, 2016. Structure depreciation and the production of real estate services. Processed. Smeal College of Business, Penn State University.

Yu, Rui, 2022. Returns to political contributions in local housing markets. Technical report, Working Paper. University of Pennsylvania.

Yu, Yue, 2024. The local and aggregate effects of land-use regulation on farmland protection. Working paper. University of Toronto.

Zabel, Jeffrey, Dalton, Maurice, 2011. The impact of minimum lot size regulations on house prices in eastern Massachusetts. *Regional Science and Urban Economics* 41 (6), 571–583.

Zhang, Howard, 2022. Consumer cities: the role of housing variety. Working paper. Columbia University.