Quality, Inequality and Understanding Housing Markets *

Schuyler Louie
UC Irvine

John Mondragon
FRBSF

Rami Najjar
Northwestern University

Johannes Wieland
UCSD, FRBSF, and NBER

August 2026

Abstract

Most housing research relies on a single-margin model where a single housing price relates to a single housing quantity. We generalize this approach by introducing income inequality and housing quality. Two important challenges to studying housing supply (and demand) emerge: First, there is no singular housing supply function, instead there is a continuum of unit-supply functions indexed by quality. This implies that supply elasticity estimates are local average treatment effects (LATEs), not structural elasticities. Second, spillovers common in spatial settings may preclude even the LATE interpretation of estimates, which are instead local general equilibrium objects. Critically, these issues cannot be bypassed by using the land share of value to measure supply constraints as it is not even a diagnostic for supply constraints in the standard model. We suggest that understanding heterogeneity in incomes and housing services quality is likely critical to understanding housing markets more generally.

*Authors' emails: Louie: schuyljl@uci.edu; Mondragon: john.mondragon@sf.frb.org; Najjar: rnajjar@u.northwestern.edu; Wieland: jfwieland@ucsd.edu. Thanks to Dominick Bartelme, Max Buchholz, Tom Davidoff, Cameron Murray, Jordan Rosenthal-Kay, Adam Shapiro, Stephen Terry, Dan Wilson and Anthony Yezer for comments. The views in this paper do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. All errors are our own.

1 Introduction

> "This surplus rent is the price which the inhabitant of the house pays for some real or supposed advantage of the situation. In country houses, at a distance from any great town, where there is plenty of ground to chuse upon, the ground-rent is scarce anything, or no more than what the ground which the house stands upon would pay, if employed in agriculture. Ground-rents are generally highest in the capital, and in those particular parts of it where there happens to be the greatest demand for houses, whatever be the reason of that demand, whether for trade and business, for pleasure and society, or for mere vanity and fashion." – Adam Smith (1776)

An enormous amount of research estimates housing supply elasticities to quantify regulatory constraints, address questions of housing affordability and understand how cities develop (see Molloy, 2020 and Saiz, 2023 for surveys). The standard housing framework in this literature has a single margin of adjustment: all changes to housing demand and supply operate through a single price and quantity of housing, with quantity typically (although not exclusively) measured as units or floor space.[1] We trace out the implications of introducing three realistic features of housing markets: first, some households are rich and some are poor; second, rich households purchase higher-quality housing than poor households; third, high-quality housing is harder to produce, and so more expensive, than low-quality housing. In this setting, there are many unit-supply functions, each indexed by unit "quality", where quality reflects both physical margins (size, finishes, and so on) and location-specific amenity quality. Thus, the standard single-margin supply and demand functions that relate changes in prices to changes in the number of units alone are misspecified as prices also encode information about the supply and demand of unit quality.

Incorporating the quality margin (which has a long and established history in much of the housing literature, see Smith, 1776, Rosen, 1974, Rosenthal, 2014, Landvoigt, Piazzesi, and Schneider, 2015, Epple, Quintero, and Sieg, 2020) and inequality, has important implications for understanding housing markets. First, because there is no single housing supply elasticity, standard empirical methods will at best recover a LATE of different housing supply functions (Imbens and Angrist, 1994). Different types of demand shocks will imply different supply elasticity estimates, even when the true supply functions are identical.[2] Second, with equilibrium spillovers common in spatial settings, supply elasticity estimates may not even have a LATE interpretation. Demand shocks to one household or margin can, through general equilibrium effects, alter demand for other households or margins, so that the weights on some housing supply functions are negative or stable unit treatment value assumptions are violated. Instead these estimates are local general equilibrium effects reflecting the complex interplay of quality demand, inequality, and spillovers. Finally, we point out that these issues are not bypassed by the popular approach exploiting the difference between price and measured construction cost – or equivalently "excessively" high land share of value (Glaeser and Gyourko, 2018) – as a quasi-sufficient statistic for housing supply constraints. We show that the land share of value does not

[1] This includes any framework that can be reduced to $\widehat{H} = \psi\widehat{P}$, where hats indicate percentage changes and ψ gives the elasticity of the supply function so that a single margin completely summarizes the equilibrium relationship between housing prices and quantities. A few examples where quantities are measured with units include Topel and Rosen (1988), Saiz (2010), Chodorow-Reich, Guren, and McQuade (2024), or Gorback and Keys (2020). Baum-Snow and Han (2024) emphasize floor space while also providing unit elasticities. Combes, Duranton, and Gobillon (2021) incorporate quality variation, but restrict their analysis to new single-family homes and assume housing is a homogeneous good within a specific segment, thus avoiding the aggregation question. Epple, Quintero, and Sieg (2020) provide a quantitative and dynamic model closely related to our framework and which features the same challenges we discuss here.

[2] This conclusion also affects work attempting to measure the productivity of housing production such as Goolsbee and Syverson, 2023, D'Amico, Glaeser, Gyourko, Kerr, and Ponzetto, 2024 or the popular blog "Construction Physics" (see here). Without accounting for differences in housing services quality, productivity estimates will be biased at best (Jorgenson and Griliches, 1967).

diagnose supply constraints even in the single-margin model, but instead varies with demand and production technology in any non-degenerate situation.

To summarize, under realistic assumptions a single housing supply (or demand) curve is not well defined and may be misleading.[3] We probably know much less about housing demand and supply than generally believed, but fore-fronting income and housing inequality should help advance our understanding of housing markets (Rodríguez-Pose and Storper, 2020, Buchholz, Kemeny, Randolph, and Storper, 2026, Louie, Mondragon, Najjar, and Wieland, 2026).

2 Two Margins and Infinite Supply Elasticities

To develop our arguments, we build a simple model of a "city" that incorporates three important and realistic elements: there is income inequality; housing quality varies; the supply of housing quality is relatively inelastic. These assumptions mean that utilitarian apartments in undesirable locations are easier to supply and so cheaper than large houses in prime locations, and that rich households will tend to live in the latter while poor households tend to live in the former.

Assume rich and poor households receive exogenous incomes y_r and y_p respectively. Households purchase non-housing consumption c (the numeraire) and a housing unit with quality q at price $p(q)$, where housing quality is a normal good and each household lives in one housing unit. Assume Cobb-Douglas utility over housing and non-housing consumption (Davis and Ortalo-Magné, 2011) with α giving the weight on housing.[4]

Assume that housing production is competitive. Units u are produced with capital that is perfectly elastically supplied at price r and quality is added at price r_q.[5] The city-wide supply of quality Q is determined by the supply function $Q = r_q^{\psi}$. Each housing supplier builds one unit of housing taking the price $p(q)$ as given and maximizes the profit function: $p(q)q - r - r_q q$. Zero profits implies that the hedonic pricing function is $p(q) = \frac{r}{q} + r_q = \frac{r}{q} + Q^{\frac{1}{\psi}}$.

Together these assumptions imply that housing quality demand for a household of type i is $q_i = \frac{\alpha(y_i - r)}{r_q}$. Let N_i be the number of households of type i, then in equilibrium total housing units are equal to total households $u = N_r + N_p$ and the quality market clears $N_r q_r + N_p q_p = Q$. Temporarily assume there is no endogenous migration. So incomes, population, and units are exogenously determined, leaving price and quality as endogenous outcomes.

This is a restrictive environment in which all changes in income and population are exogenous, so it would be straightforward to recover the elasticity of housing supply if such an object existed. Consider a situation where all households win the lottery. We have assumed that there is no migration, so all this additional

[3] Ultimately, this is a reiteration of classic results on the limits to aggregation (Felipe and Fisher, 2003; Fisher, 2005), which are particularly relevant to housing where the literature's emphasis on an aggregate housing supply curve is in sharp contrast to the vast differentiation across housing services observed in reality.

[4] Note that alternative assumptions about housing demand will affect the exact way shocks play out. We choose Cobb-Douglas for simplicity and to minimize our departures from the literature.

[5] That non-quality capital supply is perfectly elastic is just a simplification. In reality, it is likely that even the cheapest units face an upward-sloping supply curve if only because additional production requires additional labor, the supply of which is not perfectly elastic.

income will be spent on more housing quality and non-housing consumption in the city. Quality demanded will go up, which means the price $p(q)$ will increase. There is no change in the number of households or housing units by assumption. Figure 1 plots the change from this shock in the space of the log price index and log housing units as the vertical red line (see A.1). From the view of the single-margin model, these data incorrectly imply a perfectly inelastic supply curve, but in truth there has simply been no change in the demand for units at the same time that quality demand has increased.

Now imagine there is a large influx of poor migrants who can only afford very small amounts of housing quality. The supply of cheap housing increases to accommodate this inflow. The change in demand for housing quality (holding the price fixed) is given by $\Delta N_p q_p = \Delta N_p \frac{\alpha(y_i - r)}{r_q}$. If poor incomes are sufficiently low then the change in quality demand will be negligible and this expression will be small. Figure 1 plots this shock as the horizontal blue dashed line. From the single-margin view, these data incorrectly imply a very elastic supply curve, but there has merely been a large increase in the demand for low-quality, easily-supplied units.

Let us take stock. We have subjected our framework to two exogenous shocks: an increase in per capita incomes and an increase in population. Both of these experiments are empirically valid in that there are no confounding supply-side shocks. In the standard framework, either of these shocks would allow us to estimate the same elasticity of housing supply because all changes in demand operate along the same margin.[6] But Figure 1 shows that we will estimate different supply elasticities depending on which demand shock we use. In fact, any exogenous shock to incomes and population can be cast as a combination (LATE) of these two cases, allowing us to estimate a "supply elasticity" anywhere in the shaded region bounded by the two extreme cases.

Why is this happening? In our environment there is no such thing as the singular or aggregate housing supply elasticity. Instead there is a continuum of unit-supply elasticities indexed by the level of housing quality (see A.2). The empirical unit-price elasticity (at the market-weighted price index P) is the following complex function of changes in incomes (adjusted for quality-share of expenditure) $\omega_i \widehat{y}_i$, population $\widehat{N}_i$, population shares ϕ_i, shares of aggregate quality consumed θ_i, the rich market share of housing $s_r = N_r p(q_r)/(N_r p(q_r) + N_p p(q_p))$ and the share of house price due to quality $\theta(q)$ (see A.1):

$$\underbrace{\frac{\widehat{u}}{\widehat{P}}}_{\text{Empirical elasticity}} = \underbrace{\frac{\phi_r \widehat{N}_r + (1-\phi_r)\widehat{N}_p}{\theta_r(\widehat{N}_r + \omega_r \widehat{y}_r) + (1-\theta_r)(\widehat{N}_p + \omega_p \widehat{y}_p)}}_{\text{Shock dependence}} \underbrace{\left(\frac{\psi + 1}{s_r \theta(q_r) + (1-s_r)\theta(q_p)}\right)}_{\text{State dependence}}.$$

The empirical elasticity depends both on the kind of shock hitting the city (it is shock dependent) and the state of the housing market when the shock hits (it is state dependent). Intuitively, utilitarian units in undesirable locations have a very low quality-cost share $\theta(q)$ and the share of quality expenditures $(1-\theta_r)$ is small, so the empirical elasticity will appear high if most households are poor (s_r low) or if the shocks primarily affect poor households. Alternatively, large single-family homes in the best locations have a high quality-cost share and rich households consume a large share of quality, so the elasticity will appear much lower if most

[6] Note that it can even be the case that an aggregate supply function for housing services exists and that housing may have multiple margins with perfectly elastic substitution between the margins, but if housing demand is non-homothetic then different demand shocks may still recover different elasticity estimates. For example, if households prefer more units when they are poor and more quality when they are rich, then housing supply will look elastic in terms of units when poor and inelastic when rich. In this case the issue would be a lack of identification, not misspecification. Note that this framework, because it allows for multiple margins, is distinct from the standard approach in the literature tested in Louie, Mondragon, and Wieland (2025). Thanks to Jordan Rosenthal-Kay for emphasizing this point. See also Rosenthal-Kay (2024) and Louie, Mondragon, and Wieland (2025) Section A5.

of the market is rich or if the shock primarily affects rich households. In general, empirical estimates recover different LATEs, which may still be interesting of course, but which are not structural supply parameters.

It is unlikely that these issues are mere theoretical curiosities. We illustrate this point with three estimates of MSA-level unit-supply elasticities using data on house prices and units from 2000 to 2020 and three different "instruments" for house prices: per capita income growth, population growth and total income growth (these estimates are reported in Louie, Mondragon, and Wieland, 2025, Section A5). We plot the implied inverse elasticities in Figure 1. Per capita income growth suggests that unit-supply is quite inelastic, population growth suggests unit-supply is much more elastic and total income growth produces an estimate between the other two. These results are consistent with our LATE interpretation that different kinds of demand will manifest differently in the price/unit space depending on how that demand is distributed along the quality margin. These different kinds of shocks probably exist in reality. For example, the enormous income growth among the tech-affiliated population in the San Francisco Bay Area could correspond to our lottery example, while large flows of relatively low-income immigrants could match pandemic-era immigration effects in the US. Critically, in this more realistic framework, attempting to estimate a unit-supply elasticity is fundamentally an undefined exercise until one adequately isolates shocks to a specific level of quality (see A.2).[7]

3 Spillovers

Additional challenges arise when we incorporate endogenous spillovers common in spatial equilibrium frameworks. First, we introduce endogenous migration by assuming that the number of each type of household in the location is determined by real income, an amenity A, and an elasticity of migration η: $N_i = \left(A\frac{(y-r)}{r_q^\alpha}\right)^\eta$. We continue to take incomes as given, but population, units and prices are now endogenous.

Reconsider our lottery experiment, but now assume that only rich households are eligible to win. This prospect draws in rich households, driving up house prices. Higher prices reduce real incomes for poor households, who migrate out of the city. It can easily be the case that the outflow of poor households is larger than the inflow of rich households so that prices increase but population and units actually decline (see A.3). Figure 2 draws this case, which is often called "gentrification" (Guerrieri, Hartley, and Hurst, 2013) and could describe the post-pandemic experience in California, where real house prices increased while population stagnated or even fell.

Rich migration raises prices, which is a pecuniary spillover that pushes down demand by the poor, so that the correlation between prices and units is uninformative about the different supply curves faced by the rich and poor. In terms of Imbens and Angrist (1994), this housing demand shock puts negative weight on the low-quality housing supply curve(s), failing the monotonicity requirement and invalidating even a LATE interpretation of the estimate.

Alternatively, if rich households consume services provided by poor households so that poor incomes are increasing in rich incomes, then the gains from the lottery will be shared and can cause in-migration for both types of households (see A.4). This is shown in the upward-sloping blue line labeled "Positive Income Spillovers" in Figure 2. This violates the stable unit treatment value assumption, not the monotonicity requirement, so the resulting estimate is again not a LATE but rather a local general equilibrium effect.

[7] See Broxterman, Liu, and Yezer (2025) for additional reasons why it may be difficult to estimate housing supply elasticities and Broxterman, Larson, and Yezer (2026) for constructing a market-level price index using an appropriate measure of housing services (such as floor space) aggregated appropriately.

Now imagine that only poor households are eligible for the lottery, that rich households prefer living with other rich households and that poor households are indifferent. Higher incomes draw in poorer households, but this causes out-migration from richer households. If inequality is high and the rich really dislike living with the poor, then house prices may actually fall at the same time that there is an increase in housing units (see A.5). This is shown as the downward-sloping dashed line in Figure 2 and could be an illustration of "capital flight" or "white flight" in neighborhoods experiencing demographic transitions (Boustan, 2010). In this case, the negative migration spillovers induce negative weights and again void any LATE interpretation.

How can exogenous increases in "demand" for certain types of housing sometimes lead to increases in prices and declines in units and other times lead to increases in units and declines in prices? These seemingly perverse outcomes arise naturally from the facts that demand is not monolithic and that poor and rich households are connected to each other through local housing and labor markets. So an increase in one group's incomes may actually create a decline in housing units demanded for the other groups. Thus, research designs that exploit exogenous "demand" shocks must confront the problem that the resulting estimates are either LATEs or local general equilibrium responses, not structural parameters.

4 Land Share of Value is not a Sufficient Statistic for Supply Constraints

The discussion above shows that a common way of thinking about housing supply faces serious difficulties. But a distinct method suggests it might be possible to sidestep these concerns by using the gap between the price of housing and the measured cost of building a housing unit in a "competitive" or "unconstrained" environment (Glaeser and Gyourko, 2003; Glaeser and Gyourko, 2018). Thus, the "excess" land share of value (the residual after removing "competitive" land and construction costs from observed values) is interpreted as a quasi-sufficient statistic for housing supply constraints. However, we show that this conclusion does not follow even within the standard single-margin framework, much less a setting with multiple margins of adjustment.[8]

Assume a standard neoclassical production function for housing in a city where housing production uses land and capital (see A.6). Let p_l be the price of land and p be the price of a housing unit. Following Glaeser and Gyourko (2003), assume that the supplier pays a regulatory cost τ for every unit of housing produced so that they only receive $p - \tau$. Let ϵ be the land-capital elasticity of substitution, $\sigma_l = p_l l / pH$ be the gross cost share of land in production, $\sigma_{l,\tau} = p_l l/(p-\tau)H$ be the net land share, $\sigma_{k,\tau}$ the net capital share, ψ the elasticity of land supply, and η the migration elasticity into the city (assume one kind of household).[9] Then log differences in the gross land share of value are approximated as the following function of log differences in incomes, amenities (A) and regulatory taxes (as a share of the house price)

$$\widehat{\sigma}_l = \left(\frac{1-\sigma_l-\sigma_{k,\tau}\epsilon}{\psi+\alpha\eta\sigma_l+\sigma_{k,\tau}\epsilon}\right)(\eta(\widehat{A}+\widehat{y})-\alpha\eta\widehat{\tau})-\widehat{\tau}.$$

The land share ratio varies across time and space as a complex function of structural elasticities, incomes,

8 See Murray (2021), Somerville (2005) and O'Flaherty (2003), which is a discussion (available here) of Glaeser and Gyourko (2003) for related and more comprehensive critiques of this approach. We do not here discuss the important issue of relating marginal to average land values, which is addressed in those critiques.

9 Our setup here differs from that in Glaeser and Gyourko (2003) in that we do not include an intensive margin of land per unit of housing, which simplifies the problem. Including this margin would not change our conclusion.

amenities, regulatory taxes and cost shares.[10] If we assume that there are no regulatory taxes, then gross cost shares are equal to net cost shares and the land share of value will still vary with demand except in the case of Cobb-Douglas production ($\epsilon = 1$) or infinite land supply. However, if regulatory taxes do exist then land value shares will vary with fundamentals even if production is Cobb-Douglas.

Regulations or geography can affect the land share of value through the land supply elasticity, the ease of substitution between land and capital or tax-like penalties. But because the land share of value varies with, *inter alia*, demand and the production technology, it is not a sufficient statistic for supply constraints. Introducing additional margins of adjustment only complicates this picture further (for example, if house size is a normal good then wages and amenities will have distinct effects on the land share of value). This discussion also ignores the possibility that development taxes, zoning and urban planning more generally may actually improve amenity values, further complicating the empirical relationship (Lin, 2024, Yezer, 2026).

5 Conclusion

We raise a number of challenges to some of the standard approaches to studying housing supply employed for the last few decades. These challenges are serious, but almost certainly they are not insurmountable. A number of avenues seem important going forward. First, identifying and quantifying the various dimensions of housing quality (size, commute length, finish quality, neighborhood amenities and so on) is a necessary step to understanding the supply of housing quality. Some forms of housing quality, like high-end appliances, are likely to be quite elastic while others, like living in a large house in the "best" neighborhood, are inherently scarce. Second, understanding the demand-drivers for all these margins will help identify the kind of empirical variation needed to learn about quality-specific supply functions. Finally, a greater understanding of how households are connected in local spatial equilibrium and how these connections mediate the effects of various shocks is likely to be first order.

[10] Notice that higher regulatory taxes reduce land share of value as constraints on development effectively make land too cheap. Glaeser and Gyourko (2003) and Glaeser and Gyourko (2018) primarily (but not exclusively) examine the residual of prices net of construction costs, which combines both the pure land share and the regulatory taxes share. Deviations in this ratio are still complex functions of demand and production which vary for many reasons apart from regulatory taxes: $\widehat{\sigma_l + \frac{\tau}{p}} = \sigma_l \widehat{\sigma}_l + \widehat{\tau} - \sigma_\tau \widehat{p}$ where $\sigma_\tau = \frac{\tau}{p}$. We focus on the true land share of value to emphasize that it will not generally be constant.

References

Baum-Snow, Nathaniel and Lu Han (2024). "The microgeography of housing supply". *Journal of Political Economy* 132.6, pp. 1897–1946.

Boustan, Leah Platt (2010). "Was postwar suburbanization "white flight"? Evidence from the black migration". *The Quarterly Journal of Economics* 125.1, pp. 417–443.

Broxterman, Daniel, William Larson, and Anthony Yezer (2026). "Characteristics of a Sufficient Statistic to Measure City Housing Prices". Ed. by Measurement of Housing and the Housing Sector.

Broxterman, Daniel, Yishen Liu, and Anthony Yezer (2025). "Why we still don't know much about housing supply elasticity". *Real Estate Economics*.

Buchholz, Maximilian, Tom Kemeny, Gregory F Randolph, and Michael Storper (2026). "Inequality, not regulation, drives America's housing affordability crisis".

Chodorow-Reich, Gabriel, Adam M Guren, and Timothy J McQuade (2024). "The 2000s housing cycle with 2020 hindsight: A neo-kindlebergerian view". *Review of Economic Studies* 91.2, pp. 785–816.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2021). "The production function for housing: Evidence from France". *Journal of Political Economy* 129.10, pp. 2766–2816.

D'Amico, Leonardo, Edward Glaeser, Joseph Gyourko, William R Kerr, and Giacomo AM Ponzetto (2024). *Why has construction productivity stagnated? The role of land-use regulation*. Tech. rep. National Bureau of Economic Research.

Davis, Morris A and François Ortalo-Magné (2011). "Household expenditures, wages, rents". *Review of Economic Dynamics* 14.2, pp. 248–261.

Epple, Dennis, Luis Quintero, and Holger Sieg (2020). "A new approach to estimating equilibrium models for metropolitan housing markets". *Journal of Political Economy* 128.3, pp. 948–983.

Felipe, Jesus and Franklin M Fisher (2003). "Aggregation in production functions: what applied economists should know". *Metroeconomica* 54.2-3, pp. 208–262.

Fisher, Franklin M (2005). "Aggregate production functions-a pervasive, but unpersuasive, fairytale". *Eastern Economic Journal* 31.3, pp. 489–491.

Glaeser, Edward and Joseph Gyourko (2003). "The Impact of Building Restrictions on Housing Affordability". *Economic Policy Review* 9.2.

– (2018). "The economic implications of housing supply". *Journal of economic perspectives* 32.1, pp. 3–30.

Goolsbee, Austan and Chad Syverson (2023). *The strange and awful path of productivity in the US construction sector*. Tech. rep. National Bureau of Economic Research.

Gorback, Caitlin S and Benjamin J Keys (2020). *Global capital and local assets: House prices, quantities, and elasticities*. Tech. rep. National Bureau of Economic Research.

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

Imbens, Guido W and Joshua D Angrist (1994). "Identification and Estimation of Local Average Treatment Effects". *Econometrica* 62.2, pp. 467–475.

Jorgenson, Dale Weldeau and Zvi Griliches (1967). "The explanation of productivity change". *The review of economic studies* 34.3, pp. 249–283.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider (2015). "The housing market (s) of San Diego". *American Economic Review* 105.4, pp. 1371–1407.

Lin, Chuanhao (2024). "Do households value lower density: Theory, evidence, and implications from Washington, DC". *Regional Science and Urban Economics* 108, p. 104023.

Louie, Schuyler, John Mondragon, Rami Najjar, and Johannes Wieland (2026). "Housing Affordability and Housing Demand". *FRBSF Economic Letter* 2026.03, pp. 1–5.

Louie, Schuyler, John Mondragon, and Johannes Wieland (2025). *Supply constraints do not explain house price and quantity growth across us cities*. Tech. rep. National Bureau of Economic Research.

Molloy, Raven (2020). “The effect of housing supply regulation on housing affordability: A review”. *Regional science and urban economics* 80.C, pp. 1–5.

Murray, Cameron K (2021). “Marginal and average prices of land lots should not be equal: A critique of Glaeser and Gyourko’s method for identifying residential price effects of town planning regulations”. *Environment and Planning A: Economy and Space* 53.1, pp. 191–209.

O’Flaherty, Brendan (2003). “Commentary (on Glaeser and Gyourko)”. *Federal Reserve Bank of New York Policy Review* 9.2, pp. 41–43.

Rodríguez-Pose, Andrés and Michael Storper (2020). “Housing, urban growth and inequalities: The limits to deregulation and upzoning in reducing economic and spatial inequality”. *Urban Studies* 57.2, pp. 223–248.

Rosen, Sherwin (1974). “Hedonic prices and implicit markets: product differentiation in pure competition”. *Journal of political economy* 82.1, pp. 34–55.

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, pp. 687–706.

Rosenthal-Kay, Jordan (2024). *Urban costs around the world*. Tech. rep. Working paper.

Saiz, Albert (2010). “The geographic determinants of housing supply”. *The Quarterly Journal of Economics* 125.3, pp. 1253–1296.

– (2023). “The global housing affordability crisis: Policy options and strategies”. *MIT Center for Real Estate Research Paper* 23/01.

Smith, Adam (1776). *The wealth of nations*.

Somerville, C Tsuriel (2005). *Zoning and Affordable Housing: A Critical Review of Glaeser and Gyourko’s Paper*. Canada Mortgage and Housing Corporation.

Topel, Robert and Sherwin Rosen (1988). “Housing investment in the United States”. *Journal of political Economy* 96.4, pp. 718–740.

Yezer, Anthony M (2026). “Can Land Use Regulation Lower the Price of Housing in Cities?” *Journal of Regional Science*.

6 Figures

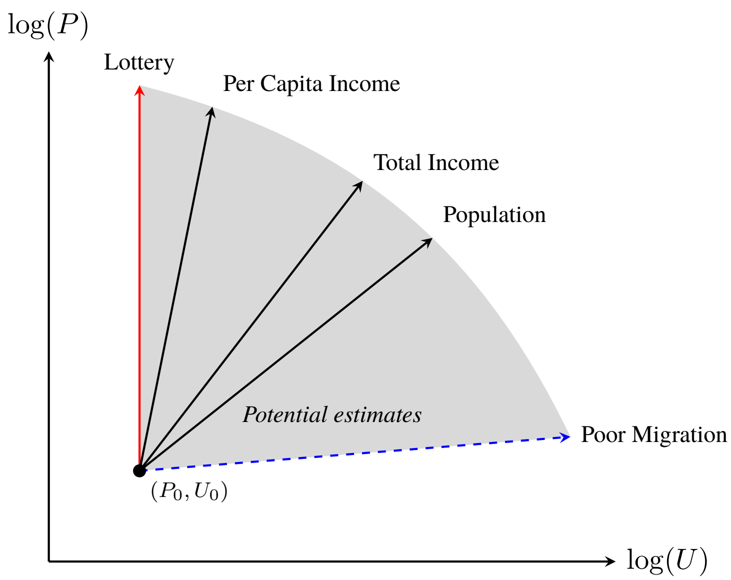

Figure 1: Housing Supply Estimates with Two Margins and Exogenous Demand Shocks

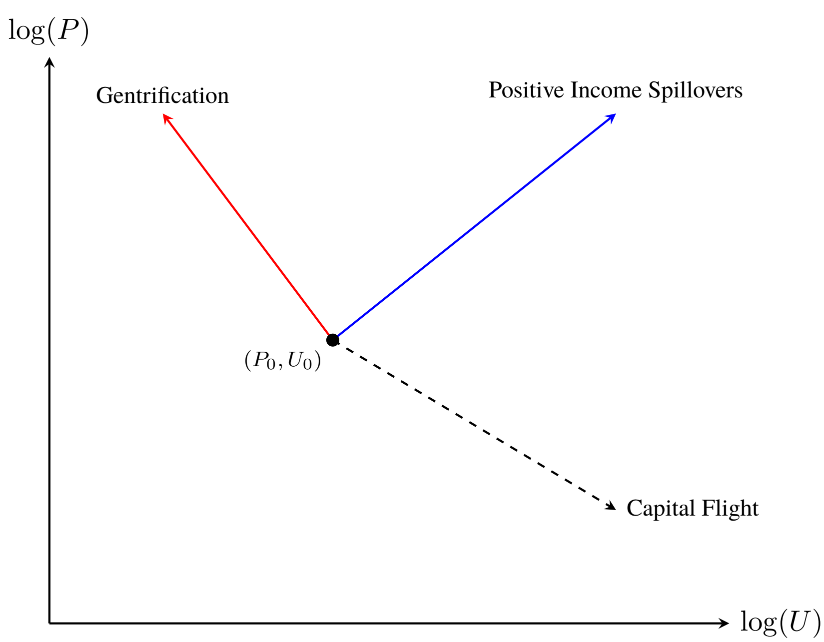

Figure 2: Housing Supply Estimates with Exogenous Demand Shocks and Spillovers

A Appendix

A.1 Baseline Model Details

Here we present the general expression for the change empirical unit-price elasticity as a function of changes in incomes and population. Note that the linearized equilibrium relationships are given by the following where $\phi_r = N_r/(N_r + N_p)$ is the rich share of population and $\theta_r = N_r q_r/(N_r q_r + N_p q_p)$ is the rich share of quality expenditures

$$\begin{aligned}
\widehat{u} &= \phi_r \widehat{N}_r + (1-\phi_r)\widehat{N}_p \\
\widehat{q_i} &= \omega_i \widehat{y_i} - \widehat{r}_q, \\
\omega_i &= \frac{y_i}{y_i - r}, \\
\widehat{Q} &= \theta_r(\widehat{N}_r + \widehat{q}_r) + (1-\theta_r)(\widehat{N}_p + \widehat{q}_p) \\
\widehat{Q} &= \psi \widehat{r}_q, \\
\widehat{p}(q) &= \theta(q)\widehat{r}_q, \\
\theta(q) &= \frac{r_q q}{r + r_q q}.
\end{aligned}$$

We can use these expressions to solve for the change in the equilibrium price of quality

$$\widehat{r}_q = \frac{\theta_r(\widehat{N}_r + \omega_r \widehat{y}_r) + (1-\theta_r)(\widehat{N}_p + \omega_p \widehat{y}_p)}{\psi + 1}.$$

Then house prices at quality q are simply:

$$\widehat{p}(q) = \theta(q)\frac{\theta_r(\widehat{N}_r + \omega_r \widehat{y}_r) + (1-\theta_r)(\widehat{N}_p + \omega_p \widehat{y}_p)}{\psi + 1}.$$

We can evaluate the unit-price elasticity at the market-level price index, defined as $P = \phi_r p(q_r) + (1 - \phi_r)p(q_p)$ so that $\widehat{P} = s_r \widehat{p}(q_r) + (1-s_r)\widehat{p}(q_p)$ where $s_r = \phi_r p(q_r)/P$:

$$\frac{\widehat{u}}{\widehat{P}} = \frac{\phi_r \widehat{N}_r + (1-\phi_r)\widehat{N}_p}{\theta_r(\widehat{N}_r + \omega_r \widehat{y}_r) + (1-\theta_r)(\widehat{N}_p + \omega_p \widehat{y}_p)} \left(\frac{\psi + 1}{s_r \theta(q_r) + (1-s_r)\theta(q_p)} \right).$$

This expression is a LATE where the quality and substitution elasticities are weighted by the relative importance of shocks to rich and poor demands for quality and units (the substitution elasticity is assumed to be one in Cobb-Douglas demand).

If population changes are zero (as in the lottery example) but income changes are positive, then this elasticity will be zero (and the inverse elasticity will approach infinity). If only the population of poor households changes, then

$$\frac{\widehat{u}}{\widehat{P}} = \frac{1-\phi_r}{1-\theta_r} \left(\frac{\psi + 1}{s_r \theta(q_r) + (1-s_r)\theta(q_p)} \right).$$

In this case the resulting estimate will be a weight on the quality elasticity (indexed by the price level) and the substitution elasticity where the weight is the ratio of the poor share of population over the poor share of quality expenditures. If the poor are very poor, then the elasticity will be large and the inverse elasticity will be very small.

A.2 Quality-Constant Supply Curves

We can also use this framework to clarify what quality-constant supply curves look like. Hold the level of quality fixed so that $q = q^*$ and then log-linearize the hedonic price index

$$\widehat{p}(q^*) = \theta(q^*)\widehat{r}_q = \frac{\theta(q^*)}{\psi}\widehat{Q},$$

where $\theta(q^*) = r_q q^*/(r + r_q q^*)$ is the quality share of cost for that type of housing and the second equation uses the aggregate quality supply function. Then assume a discretization of the initial quality distribution so that the log-linearized change in aggregate quality can be represented as the following (remembering that we assume no change in quality at the point of interest)

$$\widehat{Q} = s(q^*)\widehat{u}(q^*) + \sum_{q \neq q^*} s(q)(\widehat{u}(q) + \widehat{q}),$$

where $s(q^*) = u(q^*)q^*/Q$. This gives us

$$\widehat{p}(q^*) = \frac{\theta(q^*)}{\psi}\left(s(q^*)\widehat{u}(q^*) + \sum_{q \neq q^*} s(q)(\widehat{u}(q) + \widehat{q})\right).$$

If we are willing to apply a ceteris paribus assumption that no other quality segment is affected, then the summation terms drop out and we can recover a meaningful unit-supply elasticity function

$$\frac{\widehat{u}(q^*)}{\widehat{p}(q^*)} = \frac{\psi}{s(q^*)\theta(q^*)}.$$

If this segment's share of aggregate quality is small or if the quality-share of cost is small, then this ratio will be very large and the unit supply will be very elastic. Alternatively, if the share of quality consumed in this margin is high and if quality is a large part of the cost-share, then the elasticity will approach the supply elasticity of aggregate quality. This is analogous to the supply-functions in Epple, Quintero, and Sieg (2020) (equation 28), except they assume a constant elasticity across all levels of quality, which simplifies their problem but does not avoid the aggregation issues we discuss.

A.3 Rich Lottery with Migration

Given the migration equation $N_i = \left(\frac{y_i - r}{r_q^\alpha}\right)^\eta$ (we drop the amenity term for this exercise), linearized labor supply is

$$\widehat{N}_i = \eta(\omega_i \widehat{y}_i - \alpha \widehat{r}_q).$$

Assume that all rich households are eligible for the lottery so that $\widehat{y}_r > 0$, but there are no positive spillovers to poor income so that $\widehat{y}_p = 0$. The change in the demand for quality is

$$\widehat{q}_i = \omega_i \widehat{y}_i - \widehat{r}_q.$$

The equilibrium change in house price index (holding quality constant) is then

$$\widehat{p}_q = \frac{\theta(q)}{\psi}\widehat{Q}.$$

Finally, the change in aggregate quality is the following where $\theta_r = N_r q_r/(N_r q_r + N_p q_p)$ is the rich share of quality consumption

$$\widehat{Q} = \theta_r(\widehat{N}_r + \widehat{q}_r) + (1-\theta_r)(\widehat{N}_p + \widehat{q}_p).$$

Using these equations, we can solve for the equilibrium change in prices and total population. First, substituting population changes and quality demand for each group into the market clearing condition for quality and re-arranging gives

$$\widehat{Q} = \theta_r\omega_r(\eta+1)\widehat{y}_r - (\alpha\eta+1)\widehat{r}_q.$$

Combining this with the change in the quality-constant hedonic index we get

$$\widehat{r}_q = \frac{\theta_r\omega_r(\eta+1)}{\psi+\alpha\eta+1}\widehat{y}_r. \Rightarrow \widehat{p} = \theta(q)\frac{\theta_r\omega_r(\eta+1)}{\psi+\alpha\eta+1}\widehat{y}_r.$$

Let $\phi_r = N_r/(N_r+N_p)$ be the population share of the rich. Then using the migration equations we can see that equilibrium change in units is simply:

$$\widehat{u} = \phi_r\eta\omega_r\widehat{y}_r - \alpha\eta\widehat{r}_q = \left(\phi_r - \alpha\frac{\theta_r(\eta+1)}{\psi+\alpha\eta+1}\right)\eta\omega_r\widehat{y}_r.$$

There are many scenarios in which this term will be negative. If the rich are very rich and so consume lots of quality, if they are relatively small but still sufficiently rich, if the elasticity of quality is relatively low, and so on. Of course, there may be scenarios in which this term will be positive. Either way, notice that the ratio of the unit effect to the price (at average quality) effect does not simply reduce to a structural unit-supply elasticity:

$$\frac{\widehat{u}}{\widehat{p}} = \frac{\left(\phi_r - \alpha\frac{\theta_r(\eta+1)}{\psi+\alpha\eta+1}\right)}{\frac{\theta_r(\eta+1)}{\psi+\alpha\eta+1}}\frac{\eta}{\theta(q)}.$$

A.4 Rich Lottery with Income Spillovers

Now assume that when the rich get richer some of their income trickles down to the poor so that $\widehat{y}_p = \gamma\widehat{y}_r$ where $\gamma > 0$. Following the same steps but now substituting for poor incomes with the multiplier on changes in the rich incomes, we get

$$\widehat{p} = \frac{(\eta+1)(\theta_r\omega_r + \gamma(1-\theta_r)\omega_p)}{\psi+\alpha\eta+1}\theta(q)\widehat{y}_r,$$

and

$$\widehat{u} = \eta\left(\phi_r\omega_r + \gamma(1-\phi_r)\omega_p - \frac{\alpha(\eta+1)(\theta_r\omega_r + \gamma(1-\theta_r)\omega_p)}{\psi+\alpha\eta+1}\right)\widehat{y}_r.$$

If the income multiplier is sufficiently large then units and population will both increase, although the same gentrification dynamic can still moderate the effect on unit growth. Again, the IV estimate would not recover an object interpretable as a unit-supply elasticity and the negative weights invalid any potential interpretation as a LATE:

$$\frac{\widehat{u}}{\widehat{p}} = \frac{\left(\phi_r \omega_r + \gamma(1-\phi_r)\omega_p - \frac{\alpha(\eta+1)(\theta_r\omega_r + \gamma(1-\theta_r)\omega_p)}{\psi+\alpha\eta+1}\right)}{\frac{(\eta+1)(\theta_r\omega_r+\gamma(1-\theta_r)\omega_p)}{\psi+\alpha\eta+1}} \frac{\eta}{\theta(q)}.$$

A.5 Poor Lottery with Capital Flight

Now introduce a type-specific amenity into the migration equation so that

$$\widehat{N}_i = \eta(\widehat{A}_i + \omega_i \widehat{y}_i - \alpha \widehat{r}_q),$$

and A_r is assumed to be increasing in the share of the population that is rich. We parameterize this as

$$\widehat{A}_r = \lambda(\widehat{N}_r - \widehat{u})$$

where $\lambda > 0$ captures the preference for living with more rich people and we assume $\widehat{A}_p = 0$ as poor people are indifferent.

The equilibrium change in the price of quality is now

$$\widehat{r}_q = \frac{(1-\theta_r)(\eta+1) - \frac{\theta_r \eta^2 \lambda(1-\phi_r)}{1-\eta\lambda(1-\phi_r)}}{\psi + \eta\alpha + 1} \omega_p \widehat{y}_p,$$

and the change in units is

$$\widehat{u} = \left(\frac{\eta(1-\phi_r)(1-\eta\lambda)}{1-\eta\lambda(1-\phi_r)} - \eta\alpha \frac{(1-\theta_r)(\eta+1) - \frac{\theta_r\eta^2\lambda(1-\phi_r)}{1-\eta\lambda(1-\phi_r)}}{\psi+\eta\alpha+1}\right)\omega_p \widehat{y}_p.$$

Notice that if the disamenity from the poor is strong enough then prices will actually fall, despite the rise in incomes and housing demand from the poor. This decline in prices will actually induce additional increase in unit demand beyond that induced by the initial increase in incomes. Intuitively, the capital flight reduces pressure on housing costs, which induces in-migration.

A.6 Land Share of Value

Let the demand function for housing in a city be given as $N = \left(A \frac{y}{p^\alpha}\right)^\eta$ where households in the city solve the same maximization problem across housing and non-housing consumption as above.

The production function for housing units is given as

$$H^S = F(l, k),$$

where F satisfies the standard neoclassical assumptions. We introduce a "regulatory" wedge ala Glaeser and Gyourko (2003) so that the supplier of housing faces a tax τ for every unit of housing supplied

$$\max_{k,l} (p - \tau) F(k, l) - rk - p_l l.$$

Let $\sigma_{l,\tau} = p_l l/(p-\tau)H$ be the net land share of cost and $\sigma_{k,\tau}$ be defined similarly and σ_l be the gross cost share $\sigma_l = p_l l/pH$.

We know that in a zero-profit equilibrium the price must be given by

$$p = \frac{p_l l}{H} + \frac{rk}{H} + \tau.$$

Letting $\widehat{\tau} \equiv \mathrm{d}\tau/p$ be the change in tax as a share of the price, we get that the log deviation of the price is

$$\widehat{p} = \sigma_l \widehat{p_l} + \widehat{\tau}.$$

The linearized household demand, land supply and housing production function are given as

$$\begin{aligned}
\widehat{H} &= \eta(\widehat{A} + \widehat{y} - \alpha\widehat{p}) \\
\widehat{l} &= \psi\widehat{p_l} \\
\widehat{H} &= \widehat{l} + \sigma_{k,\tau}\widehat{(k/l)} = \widehat{l} + \sigma_{k,\tau}\epsilon\widehat{p_l}.
\end{aligned}$$

We can use these and the hedonic price to derive the log deviation of the price of land is given as

$$\widehat{p_l} = \frac{\eta(\widehat{A} + \widehat{y}) - \alpha\eta\widehat{\tau}}{\psi + \alpha\eta\sigma_l + \sigma_{k,\tau}\epsilon}.$$

We can now solve for deviations in the average land share of value $\sigma_l = p_l l/pH$:

$$\widehat{\sigma_l} = \left(\frac{1 - \sigma_l - \sigma_{k,\tau}\epsilon}{\psi + \alpha\eta\sigma_l + \sigma_{k,\tau}\epsilon}\right)(\eta(\widehat{A} + \widehat{y}) - \alpha\eta\widehat{\tau}) - \widehat{\tau}.$$