Journal of Economic Perspectives—Volume 39, Number 3—Summer 2025—Pages 27–44

Curbing Rising Housing Costs: A Model-Based Policy Comparison

Boaz Abramson and Tim Landvoigt

Housing is becoming less affordable in the US economy, as housing prices have grown strongly relative to incomes. In 1980, the median house price was $47,500 and the median household income was $18,200. By 2020, the median house price rose to $250,000, reflecting an annual growth rate of approximately 4.2 percent. At the same time, median income grew on average by only 3.5 percent annually, reaching $73,000 in 2020. The divergence between house prices and income has been particularly pronounced since the Great Recession of 2007–2009. Between 2012 and 2022, while income has continued to grow at a steady annual pace of 3.5 percent, house prices appreciated at a much faster pace of 5.8 percent annually.

The ratio of house price to income, a common measure of (lack of) affordability, is higher for lower-income homeowners. This is illustrated in the left panel of Figure 1. In 2020, the price-to-income ratio was above ten for owners in the bottom income decile, compared to only three for the highest-income owners. In addition, the rise in the price-to-income ratio has been more dramatic for lower-income homeowners since 2000. A similar pattern is observed in the rental market, as illustrated by the right panel of Figure 1. Since 2000, rents have grown faster relative to renters' income, and particularly so for the lowest income renters.

■ *Boaz Abramson is Assistant Professor of Business, Columbia Business School, New York City, New York. Tim Landvoigt is Associate Professor of Finance, Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania. Their email addresses are ba2695@columbia.edu and timland@wharton.upenn.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.20241427.

Figure 1
Housing Affordability

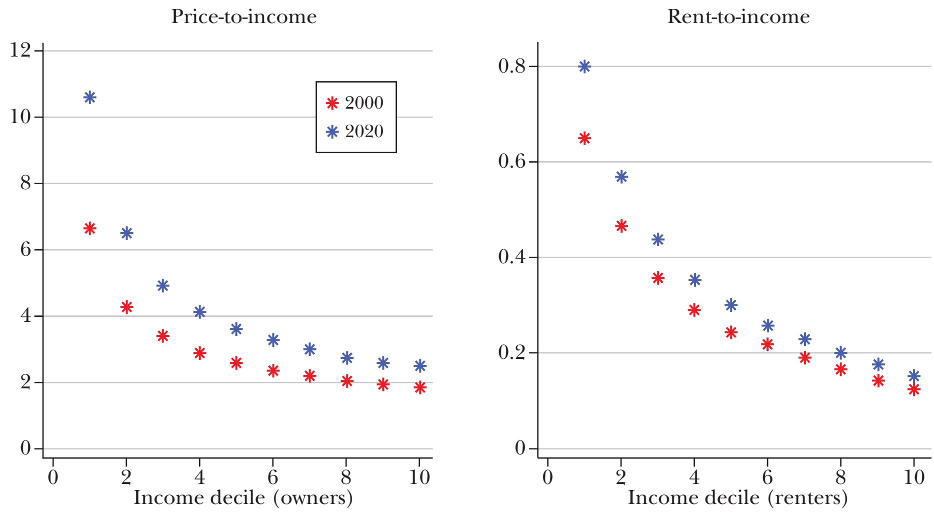

Source: Tabulations are based on American Community Survey data.
Note: The left (right) panel plots a bin-scatter of the price-to-income (rent-to-income) ratio as a function of homeowners (renters) household income in 2000 (in red) and in 2020 (in blue).

Numerous policies have been proposed to make housing more affordable. Most of these policies fall in one of two categories. For economists, at least, perhaps the most obvious remedy for high prices is to increase the supply of housing. Examples of supply-side policies include easing of housing supply regulations, tax credits for developers, and public housing. An alternative set of policies focus on making mortgages more available and less expensive. Examples include down-payment assistance for first-time buyers, mortgage interest deductions, subsidies for mortgage interest payments, or government guarantees of mortgage credit.

At a high level, both types of policies have intuitive appeal. Increasing supply should lower house prices, and cheaper mortgage credit should reduce housing costs for homeowners at a given level of house prices. However, as is often the case, the effects of these policies on housing affordability are less obvious. While there is widespread agreement that more housing construction is needed, it is not clear what type of housing should be constructed and in which locations. Cheap and easily available mortgage credit might cause upward pressure on prices through increased demand for housing and hence exacerbate the affordability problem instead of solving it. Should policymakers ease housing supply regulations, subsidize construction of small apartments, or subsidize credit for first-time homebuyers? How effective are these policies and how can they be compared?

To answer these questions, we develop a model of a city's housing market that allows us (1) to understand the sources of rising price-to-income ratios and (2) to evaluate different policies in a unified framework. Using the model, we can evaluate policies, even if these policies have yet to be implemented or were implemented in a way that does not lend itself to directly observing causal effects in the data.[1] In the model, households optimally choose consumption of housing and other goods given their budget. These households differ by their level of wealth. The supply of housing is differentiated by its level of quality. In equilibrium, the wealthiest households end up owning the highest quality houses and the poorest households end up renting the lowest quality houses. This is known as an "assignment equilibrium" in housing markets (Määttänen and Terviö 2014; Landvoigt, Piazzesi, and Schneider 2015).[2]

To apply the model to data, we calibrate the key parameters—most prominently the distributions of wealth and housing quality—so that the model matches the observed house prices and rents in San Francisco in 2013.[3] The calibrated model shows that rising wealth dispersion, together with stagnating housing supply, can explain the increase in price-to-income and rent-to-income ratios observed in Figure 1. Why did this confluence of trends likely contribute to the rise in prices and rents? Rising wealth inequality combined with stagnating supply mean that poor households end up competing with richer households for the same housing units.

Having shown that the model performs well in matching the status quo observed in the data, we then use it to evaluate how supply-side and demand-side policies would affect price-to-income and rent-to-income ratios. Here, we want to emphasize a few conclusions from our analysis upfront. First, measures that target the demand side of the housing market are mostly ineffective: in our model, a down-payment subsidy for potential home buyers increases upwards pressure on house prices, especially in those segments of the market in which buyers benefit most from the subsidies. The model strongly suggests that any solution to the affordability crisis involves more supply. Indeed, we find that additional housing supply is always most beneficial for the poorest residents of the city. One perhaps surprising conclusion is that new construction in the high-end segments of the market decreases price-to-income ratios by more in *all segments* of the housing market compared to new construction in bottom-end segments. This is because new construction in high-end segments reduces competition for units in medium- and low-quality housing segments, improving affordability throughout the wealth distribution.

[1] A large literature in economics studies how housing supply or mortgage subsidies affect house prices. For example, several papers analyze how changes in construction activity for different types of housing at the city level have affected housing prices (Baum-Snow and Marion 2009; Glaeser and Ward 2009; Mense 2025; Pennington 2021; Asquith, Mast, and Reed 2023).

[2] An early model of segmented housing markets is Ortalo-Magné and Rady (2006). More recent examples include Nathanson (2023), Nikolakoudis (2024), Mense (2025), and Fonseca, Liu, and Mabille (2025).

[3] The San Francisco metro area serves as the main example in our analysis. In Supplemental Appendix D, we perform the identical analysis for Chicago and Dallas, and highlight differences and commonalities.

Construction in low-tier segments is similarly beneficial for the poorest residents, but it improves affordability much less for middle-income and wealthy households.[4]

A Model of Segmented Housing Markets

The main building blocks of our model are households, who are the residents of the city we analyze, and a housing market structure, which specifies the type of residences available for purchase or rent and how these properties are traded. We describe all model elements in greater detail, using graphical tools.[5]

Households

Households in the model maximize utility over two life stages, referred to as "today" and "tomorrow," but which can be considered to be roughly a decade apart. Households are identical except for initial wealth, which fully determines their choices—households with the same wealth behave identically. With only two periods, we do not distinguish between income and financial assets; the sum of both is referred to as "wealth."

Utility depends on two goods: housing services and a composite of all other consumption. Utility increases with either good, but at a decreasing rate. The household's objective is a weighted average of utility across both periods, with full weight on today and a smaller weight on tomorrow. Households simultaneously choose total consumption in each period, the share allocated to housing, and savings.

Their decisions are constrained by today's total wealth, which must cover today's consumption, housing costs (rent or purchase), and savings. Tomorrow, households receive savings with interest, sell their house if owned, and gain additional wealth. These sources fund tomorrow's consumption. Households can save a positive amount, interpreted as a savings account. If purchasing a home, they can have negative savings, which we interpret as borrowing via a mortgage, subject to a loan-to-value limit. Renters cannot borrow.

The solution to the household's problem depends on the utility function, model parameters (like the interest rate), and house prices and rents. Model parameters are ultimately chosen to match relevant aspects of the data in a process called calibration, which we describe below.

[4] Other studies have reached similar conclusions. For example, Nathanson (2023) argues, using a different type of model, that low-end construction may not be the best policy to improve affordability in the Boston-Cambridge-Newton, MA-NH, metropolitan area. Recent structural papers focusing on various aspects of housing affordability include Diamond and McQuade (2019), Hsieh and Moretti (2019), Anenberg and Kung (2020), Abramson (2021), Favilukis and Van Nieuwerburgh (2021), Favilukis, Mabille, and Van Nieuwerburgh (2023), İmrohoroğlu and Zhao (2022), Couture et al. (2023), Corbae, Glover, and Nattinger (2024), Abramson and Van Nieuwerburgh (2024), and Gupta, Hansman, and Mabille (2024).

[5] A full and formal mathematical description of the model is presented in Supplemental Appendix A.

Housing Markets

The other key component of the model is the residential housing market, where households either purchase or rent homes. This market is segmented by quality: lower-quality segments yield less housing services and are available for rent, while higher-quality segments offer more housing services and are only available for purchase. The market comprises N competitive submarkets, each representing a distinct quality level, with quality defined by the housing services entering a household's utility if the home is acquired or rented. Thus, each household's housing decision reduces to choosing a quality segment.

The model assumes market equilibrium, with prices determined by supply and demand in each segment. Supply is provided by real estate developers and construction firms outside the model. While it would be straightforward to include a construction sector explicitly, this analysis assumes the city government can direct supply distribution across segments. Demand in each segment results from the household optimization problem described above. Each household buys or rents exactly one property, so market clearing requires that the number of households in each segment matches the supply.

Equilibrium Structure

To illustrate how these market clearing conditions pin down the prices, consider a hypothetical example with three quality segments, illustrated in Figure 2. There is a total mass of households of 1, distributed uniformly between a lower bound of wealth W and an upper bound W. Supply of houses within each quality segment is given by H_0, H_1, and H_2, respectively. The two top housing segments correspond to the owner-occupied market, meaning that houses in these segments are offered for purchase, while the bottom segment corresponds to the rental market, meaning that houses in this segment are offered for rent. We normalize the price of living in the bottom segment so that it is less than the lower bound of the wealth distribution.

While solving for equilibrium outcomes in this setting may seem complex at first glance, it turns out that the ordering of households by wealth greatly simplifies the analysis. In particular, it is straightforward to show that any equilibrium in this model has a so-called "monotone assignment structure"; that is, households are assigned to houses by wealth, with the poorest households living in the bottom segment and the richest households in the top segment. After all, the willingness to pay of households is increasing in their wealth, and wealthier households are thus always willing to outbid poorer households competing for the same house.

To understand this equilibrium, focus on the households with threshold wealth levels W_1 and W_2. Households with wealth above W_2 strictly prefer buying a house in the top segment with quality h_2 and corresponding price p_2. Households in the interval $[W_1, W_2]$ strictly prefer buying a house with quality h_1 at price p_1. Households with wealth below W_1 are priced out of the owner market and forced to rent in segment 0. Households that are exactly at the threshold levels are exactly indifferent between the segments below or above.

Figure 2
Assignment Equilibrium Structure

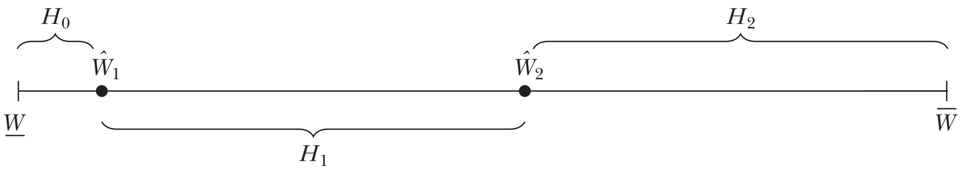

Note: This figure illustrates the assignment equilibrium with a uniform wealth distribution and three housing segments. The equilibrium is characterized by threshold levels of wealth W_1 and W_2.

How are the threshold levels determined? Because we know that each household buys exactly one house, and how many houses exist in each segment, the threshold levels can be directly calculated. The prices p_1 for housing level H_1 and p_2 for housing quality level H_2 are then determined such that the households at the thresholds are indifferent between neighboring segments. Solving the model comes down to computing the prices that solve the equilibrium conditions using numerical methods.

Applying the Model to Data

We now bring this model to the data. Because the model only has two periods, we focus our analysis on the long-run equilibrium in the housing market: one period of time in the model corresponds to ten years. Thus, our model-based projections apply to long-term effects. We believe this focus is justified, because residential real estate is long-lived and changes to the housing stock are very persistent. We first establish a baseline by matching the model to a specific city in a specific year: the San Francisco-Oakland-Hayward Metropolitan Statistical Area in 2013. This urban area has had one of the most severe affordable housing problems in the country; the San Francisco-Oakland-Hayward MSA ranked among the top ten in terms of median housing price-to-income ratio in every year between 2012 and 2022.[6] We choose 2013 as baseline year, because it marks the end of the recovery period from the 2008 housing bust and allows us to analyze the impact of housing policy choices in the subsequent decade. We measure realized changes in the distributions of (1) housing supply and (2) the wealth of residents over the 2013–2022 period. By feeding these measured changes as inputs into the model, we investigate whether realized shifts in wealth and housing supply can explain the decline in affordability

[6] Our calculations here are based on data from the American Community Survey. We perform the same analysis for Dallas and Chicago in Supplemental Appendix D, thus demonstrating the broader applicability of this approach.

ratios over the 2013–2022 period. The result is that we have a representation of San Francisco in 2022 in the calibrated model. In the next section, we study several experiments within this 2022 version of the model.

Matching the Model to Baseline Data

As a starting point to establish a baseline for analysis, we calibrate the model to the San Francisco-Oakland-Hayward Metropolitan Statistical Area in 2013. The goal of the calibration is to ensure that the model matches the economic environment that is being studied. Because we are studying housing markets, the key empirical moments to be matched are house prices, rents, and the allocation of households into housing units.[7]

In the model, prices and housing allocation are determined by the intersection of demand and supply. Demand is determined by households' wealth and preferences. We calibrate the wealth distribution of households to match the empirical wealth distribution of households in the data. Using the American Community Survey (ACS) and the Survey of Consumer Finances (SCF), we estimate the wealth distribution of households in San Francisco in 2013. The resulting wealth distribution is illustrated in blue in the top left panel of Figure 4. We calibrate household preferences using standard values from the literature on housing economics. Supply is governed by the distribution of the housing stock across quality segments. We assume that there are 13 quality segments in the city. Houses in the top ten segments are houses that households can own. They account for 54.7 percent of the total housing stock, which is the observed home-ownership rate in San Francisco in 2013. Each of the top ten segments is assumed to be of equal size.

Housing in the bottom three quality segments are residences that households can rent. The top two segments of the rental market correspond to the formal rental market and are assumed to be of equal size. The lowest segment corresponds to informal renting arrangements and accounts for 1.5 percent of the housing stock, which is the observed homelessness rate in the data (Abramson 2021; Abramson and Van Nieuwerburgh 2024). The distribution of the housing stock across quality segments is illustrated by the blue bars in the top right panel of Figure 4. The calibration of housing supply ensures that the model matches the homeownership rate, the rentership rate, and the share of households experiencing housing insecurity—all important statistical moments for studying housing affordability.

Finally, we ensure the model matches observed house prices and rents. To do so, we estimate the house quality in each of the housing segments so that the model implied prices are in line with the data. We estimate the median house price in each segment of the owner-occupied market using Corelogic, a private vendor that compiles data on the universe of US housing transactions and property tax records. Median rents are measured using American Community Survey data. The intuition for the estimation is straightforward. Given household preferences, wealth

[7] We discuss the calibration in detail in Supplemental Appendix B.

distribution, the distribution of the housing stock across segments, and the house quality in each segment, we can solve for the prices that equilibrate all housing markets. All else equal, different vectors of house qualities translate to different vectors of house prices.

How Good Is the Model?

At this point in the analysis, we confront the question of whether the model provides an adequate account of the economic environment that is being studied. We will answer this question in two steps.

The first step is evaluating the model's fit to targeted data moments. These are moments that the model was specifically calibrated to match—in our context, house prices and rents. As explained above, the vector of house qualities is estimated to target the median house price within each segment of the owner-occupied market as well as the median rent within each segment of the rental market. The left panel of Figure 3 compares the model-implied prices and rents (in blue circles) to the targeted prices and rents in the data (in red dots). Clearly, the model fits the data well. This result is not especially surprising, as the model was fitted to the 2013 San Francisco data. But of course, if the model did not fit the data well, it would indicate that it is missing something crucial.

A second step to establish model credibility is to evaluate the model's fit to data that the calibration of the model did not specifically target, but that are important for the context under consideration. In our context, two important affordability measures are the price-to-income ratio and the rent-to-income ratio. The middle panel of Figure 3 plots a bin-scatter of the price-to-income ratio as a function of homeowners' income in the model (in blue) and in the data (in red, computed based on the 2013 American Community Survey).[8] The right panel of Figure 3 similarly compares the rent-to-income ratio as a function of renters' income in the model and in the data. Both panels illustrate that the model closely aligns with the data. In addition, the downward-sloping pattern illustrates that concerns about housing affordability are especially prevalent for the bottom of the income distribution of owners and renters.[9]

[8] In the model, income is a component of household wealth and does not play an independent role. When solving the model, we draw household wealth from the calibrated (empirical) wealth distribution but do not explicitly track how much of each household's wealth is due to income. To compute the price-to-income ratio in the model, we do explicitly track household incomes. We do so by jointly drawing household wealth and income from the empirical joint distribution of (imputed) wealth and income calculated from the 2013 ACS.

[9] The rent-to-income in the seventh decile of the renter distribution of the model is higher than in the sixth decile because the renter market is separated into two discrete segments, while the income distribution is continuous. Because renter households sort into the two segments of the rental market based on their income, there is an income cutoff such that renter households with an income above the threshold rent in the upper segment, while renter households with an income below the threshold rent in the bottom segment. The rent-to-income ratio of the renter household just above the cutoff is therefore higher than that of the renter just below the cutoff.

Figure 3
Quantitative Model

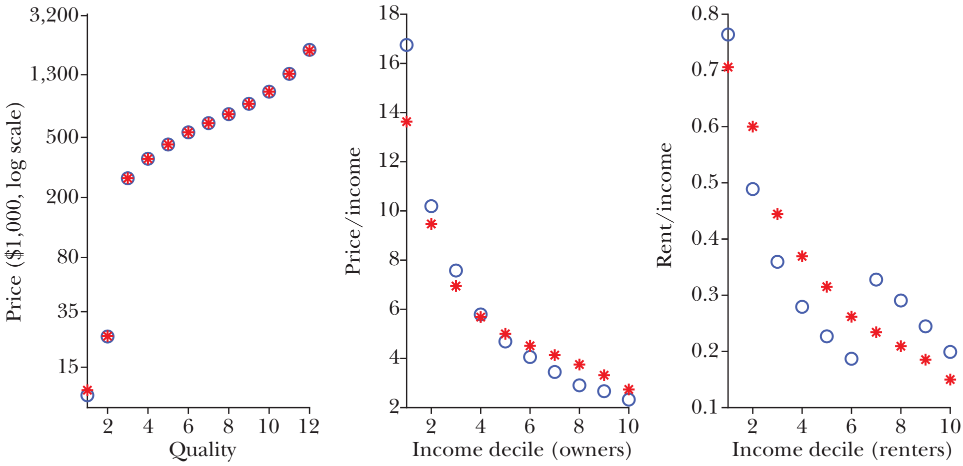

Source: Data moments are computed based on the American Community Survey and Corelogic data.
Note: This figure illustrates equilibrium in the quantified model. The left panel shows equilibrium prices for each segment in the model (blue) and in the data (red). The mid (right) panel plots a bin-scatter of the price-to-income (rent-to-income) ratio as a function of homeowners (renters) household income in the model (in blue) and in the data (in red). The estimation is described in more detail in Supplemental Appendix B.

Drivers of the Affordability Crisis

We have calibrated the model to San Francisco in 2013. Next, we use the calibrated model to quantify how changes in the wealth and housing supply distributions between 2013 and 2022 affect affordability.

In 2022, the wealth distribution shifts upward: existing residents become wealthier, along with in-migration of high-wealth households. This change is illustrated in the top-left panel of Figure 4 by the shift in the empirical wealth distribution from 2013 (in blue) to 2022 (in red). To give a sense of magnitude, the median real wealth in San Francisco more than doubled from $270,000 in 2013 to $644,000 in 2022. The wealth distribution in 2022 is estimated by combining data from the American Community Survey and the Survey of Consumer Finances, following the same steps described for the baseline calibration.

The distribution of houses across the 13 housing segments in San Francisco also changed. First, the share of houses supplied to the owner-occupied market (the top ten housing quality segments) increased from 54.7 percent (the ownership rate in 2013) to 57.1 percent (the ownership rate in 2022). For the two rental market segments, we maintain the baseline assumption that they are of equal size. Further, we estimate that the size of the lowest segment has increased from 1.5 percent to 2.4 percent. The increase in the size of the bottom segment is estimated so that the increase in rent burden in the bottom income decile between 2013 and 2022 implied by the model matches the increase observed in the American Community Survey data, which is 11 percent. The increase in the size of the first segment implies that between 2013 and 2022 the share of the San Francisco population who are either homeless or live in informal renting arrangements, such as doubling-up, has increased by 0.9 percentage points. In other words, on net, housing supply has not kept up with population growth.

Finally, within the top ten owner-occupied segments, we estimate that supply has shifted from the middle segments to the bottom and top segments, as illustrated by the top-right panel of Figure 4. Our estimate of this change in the quality distribution within the owner market requires a method to separate supply and demand factors in house price changes from 2013 to 2022. This task is notoriously difficult, and our approach establishes an upper bound on the shifts in relative supply by effectively attributing all relative price changes (that are not due to the aggregate house price trend in San Francisco) to supply changes.[10] Even though we compute an upper bound, Figure 4 shows that the changes in relative supply are small in an absolute sense, with the share of supply coming from the lowest and highest segments gaining a few percentage points from the middle.

The bottom panels of Figure 4 illustrate how the observed changes in the supply of housing and in the wealth of households affected housing prices, price-to-income, and rent-to-income ratios in San Francisco. Relative to 2013 (in blue), house prices and rents increase in 2022 (in red), and housing becomes less affordable for both owners and renters, especially among the lowest-income households. These patterns mimic the national patterns presented in Figure 1. We have further verified that the model's predictions regarding prices, price-to-income and rent-to-income ratios in San Francisco in 2022 are closely in line with the data, providing further model validation.[11]

Through the lens of the model of utility-maximizing households, a housing market divided into quality segments, and an assignment equilibrium, we can now explain this deterioration in price-to-income and rent-to-income measures. All else equal, the increase in households' wealth drives up house prices. Intuitively, when households have more funds at their disposal, but the distribution of house qualities is largely unchanged, prices must increase across the quality distribution. The rise in wealth can in theory be mitigated by supplying more housing, but in fact, the overall supply of housing decreased (as illustrated in Figure 4 by the increase in the size of the bottom segment), further deepening the affordability problem in the rental market. The modest increase in supply in the top segments does little to improve affordability because it comes at the expense of supply in other segments. As illustrated in the next section, a more substantial increase in supply in the top segments that increases overall housing supply can alleviate the affordability crisis.

[10] We provide details on this procedure in Supplemental Appendix B.1.
[11] This validation is spelled out in Supplemental Appendix C.1.

Figure 4
Drivers of the Affordability Crisis

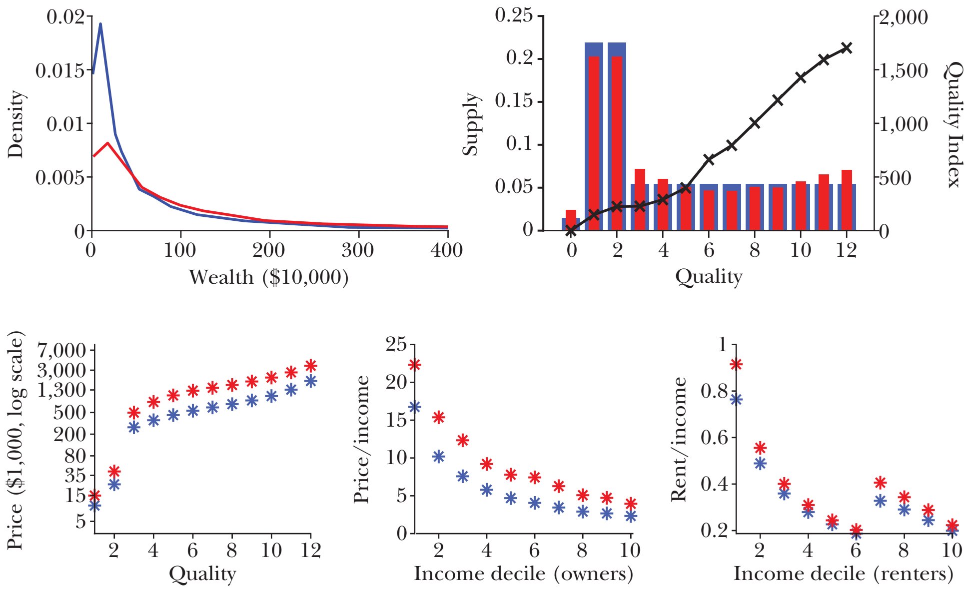

Note: This figure illustrates the effect of the observed change in housing supply and in the wealth distribution between 2013 and 2022. The top-left graph plots the density of the calibrated wealth distributions. The top-right panel shows the distributions of housing supply across segments (bars, left axis) and the quality index associated with housing segments (line, right axis). The bottom-left panel shows equilibrium prices for each segment. The bottom-middle (right) panel plots a bin-scatter of the price-to-income (rent-to-income) ratio as a function of homeowners' (renters') household income. Blue corresponds to the baseline 2013 economy and red corresponds to the 2022 economy.

Evaluating Proposed Solutions to High Housing Prices

How can policymakers lower the cost of housing? In this section, we conduct a sequence of numerical experiments within our quantitative model to evaluate the implications of three commonly proposed policies. More specifically, we identify different parameters of the model that represent policy choices. We vary these parameters, one at a time, while keeping all other parameters fixed at their 2022 values. We then analyze the effects of these changes on house prices, rents, and the allocation of housing.

Adding Housing Supply at the Bottom

Increasing the supply of housing should lead to lower house prices. But what type of housing that should be constructed? To answer this question, we compare the effects of increasing housing supply in the bottom segment to the effects of increasing housing supply in the top segment. In both experiments, we

Figure 5
Results of Three Experiments

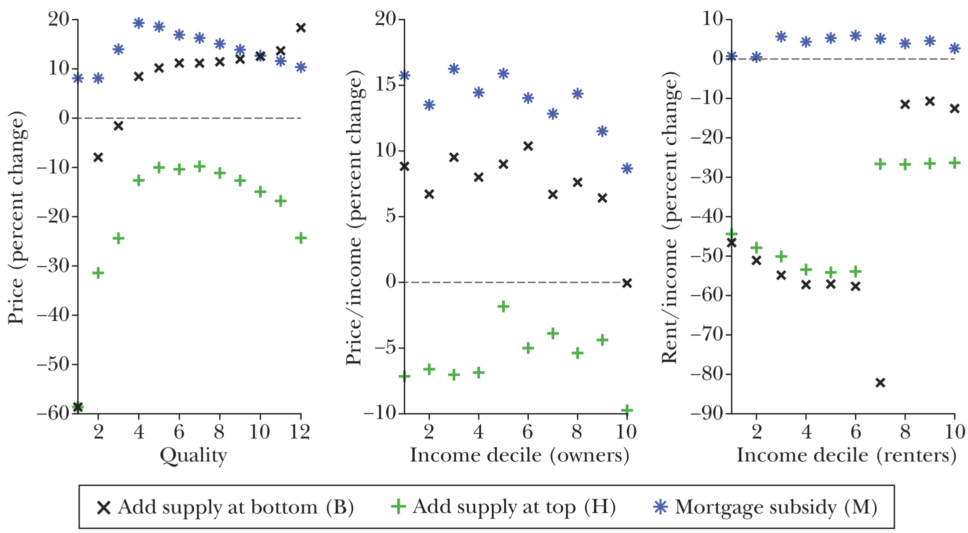

Note: This figure illustrates the effect of three experiments: the black crosses show the effect of an increase in supply at the bottom of the rental market; the green plus signs show the effect of an increase in supply at the top quality segment of the housing market; and the blue asterisks show the effect of a mortgage subsidy. The left panel shows the percent change in equilibrium prices. The middle (right) panel plots the percent change in the price-to-income (rent-to-income) ratio as a function of homeowners' (or renters') income.

increase the total supply of housing in the economy by 2 percent, which in the San Francisco-Oakland-Hayward area accounts for approximately 33,000 additional housing units. We envisage this change happening exogenously—say, via local government decisions on zoning, permitting, and tax credits that facilitate construction of housing.

In our first experiment, the newly built housing units are in the bottom of the two segments of the rental market. In this experiment, we further increase supply in the bottom segment by converting 10 percent of the housing units in the top housing segment to units in the bottom rental segment. Overall, the changes resulting from this policy, which we refer to as policy B, are illustrated by the black crosses in the three panels of Figure 5. The left panel shows that the policy is highly effective in reducing rents in the bottom rental segment. Indeed, the expanded supply of these more affordable units allows a substantial share of households previously forced to reside in the "homelessness" segment to move to higher-quality rental housing. There are moderate spillovers of the supply expansion to the top rental segment and the lowest owner-occupied segment, lowering prices in those segments slightly. However, by shrinking the supply of high-quality owner-occupied units, this policy causes prices in the upper quality segments to rise. This effect may seem counterintuitive at first, but it is a robust outcome of the economic logic embedded in the model. When we shrink housing supply at the very top end of the housing market, the marginal household that is indifferent between the topmost and the second-best segment is now wealthier. By reducing the quantity of luxury housing, wealthier households are forced into the neighboring lower-quality segments, where they drive up prices. The black dots in the other two panels of Figure 5 show how these price changes reduce the rent/income ratio of renters, but raise price/income ratios of homeowners.

Experiment 2: Adding Housing Supply at the Top

In our second experiment, the newly built housing units are in the top segment of the owner-occupied market. In this experiment, we further increase supply in the top segment by converting 10 percent of the rental units in the bottom segment of the rental market to units in the top housing segment. At an intuitive level, this conversion can be thought of as investing in the quality of the existing housing stock. The changes caused by this policy, which we refer to as policy H, are illustrated by the green plus signs in Figure 5.

The left panel shows, perhaps surprisingly, that rents and prices drop in all segments due to this policy. In particular, expanding high-quality units is as effective at making rental units more affordable as expanding the supply of low-cost rental units. However, this policy also leads to large reductions in house prices in the topmost segments, and moderate declines in middle segments. By raising supply in the highest segment, we prevent the wealthiest households from competing for units in lower segments, which in turn reduces demand in all lower segments, causing a "trickle-down" effect. The middle and right panel show that price/income ratios and rent/income ratios decline substantially. This result contrasts with the previous experiment, where price/income ratios increased. In short, increasing the supply of housing in the top segment is more effective at reducing house-price-to-income ratios than adding supply in the bottom rental segment.[12]

Experiment 3: Mortgage Availability

Another policy that is often proposed by policymakers is subsidizing homeownership, for example by providing down-payment assistance for first-time buyers or by lowering property taxes. In this section, we evaluate the effects of these policies by simulating an economy where owners are granted $100,000 towards the purchase of their house. The main result is that subsidizing ownership without building more houses unintentionally worsens housing affordability in the city. The

[12] Increasing housing supply might lead to an influx of migrants from other cities (French and Gilbert 2023). In the model, because the mass of residents is normalized to 1, population growth corresponds to a drop in housing supply. Such in-migration would therefore dampen the effect of new construction on prices and incumbents' welfare.

changes caused by this policy, which we refer to as policy M, are illustrated by the blue asterisks in Figure 5.

The left panel plots the price increases resulting from this pure demand-side policy. The economic intuition is straightforward: transferring wealth to homebuyers increases demand without increasing supply. The consequence are higher prices across all owner-occupied segments. Perhaps surprisingly, subsidizing homeownership through down-payment assistance also makes rental housing more expensive. The intuition is that increasing demand in upper segments of the housing market trickles down to lower segments. By effectively making the marginal home buyer wealthier, the down-payment subsidy also affects the wealth of the marginal household who is indifferent between renting in the top rental segment and buying in the lowest owner segment, which translates to higher rents.

As we can see in the middle and right panels of Figure 5, price-to-income and rent-to-income ratios rise unanimously as result of a subsidized down-payment policy. The lesson is that only policies that raise supply (or decrease demand) will make housing more affordable. Boosting demand for owner-occupied housing through subsidies will boost demand and raise prices, unless the policy also increases supply at the same time.

Computing Welfare

The model allows for a direct welfare comparison of the three potential policies. As illustrated in Figure 5, each of the proposed policies affects equilibrium house prices and rents. This means that household utility, which increases with the quality of the house a household lives in as well as with the amount of other consumption goods, is also affected by these policies. For example, all else equal, a policy that drives households to worse quality housing lowers utility and dampens welfare. To compare the welfare effects of the different policies, we compute a measure called "compensating variation," which is the amount of wealth a household would need to be paid under each policy so that it has equal utility relative to the baseline economy. We display the measures of compensating variation implied by the proposed policies in Figure 6, with compensating variation in absolute dollars on the left and as a share of wealth on the right.

The down-payment subsidy performs worst, because it makes housing uniformly more expensive and therefore lowers households' utility. This translates to a negative compensating variation for all households (blue asterisks). Adding supply to the bottom segment (black crosses) leads to positive welfare gains for relatively poor households but to welfare losses for relatively rich households. For example, the compensating variation is 35 percent of wealth for households in the second percentile of wealth but is negative 3 percent for the richest households. In contrast, adding supply to the top (green plus signs) leads to similarly large welfare gains for poor households, but also improves welfare of rich households. Thus, within our model, this policy is a rare case of a so-called "Pareto superiority," implying that one policy change makes everyone involved better off relative to an alternative policy.

Figure 6
Measuring Welfare

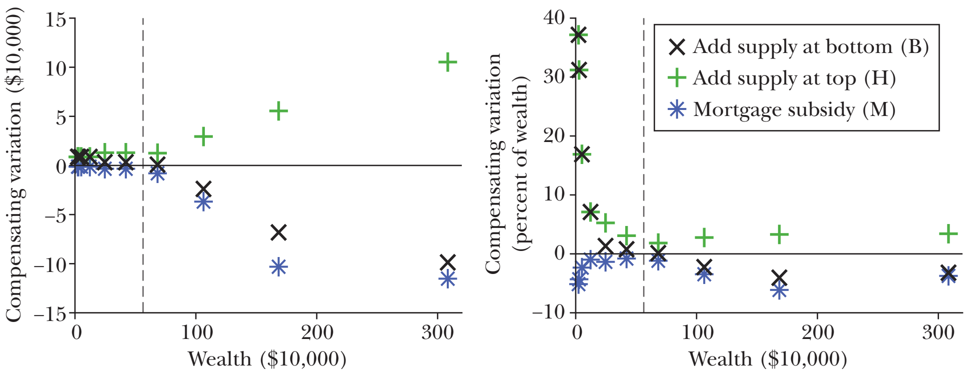

Source: Authors' creation.
Note: This figure shows the average compensating variation that households would need to receive in the baseline 2022 model in order to be as well off as under each of the proposed policies. A positive compensating variation means that the household is better off under the policy relative to the baseline. The estimates are binned by wealth decile. The left panel shows the compensating variation in 2013 dollars, and the right panel as a percent of household wealth.

Underlying Assumptions

Our calibrated model is complex enough to capture key heterogeneity in household wealth and in housing qualities, yet simple enough to yield clear insights. It reflects how interventions in one segment of the housing market affect the entire market. However, as is the case with most models, the model has limitations due to simplifying assumptions.

Notably, it lacks realistic transition dynamics, operating over just two periods—today and tomorrow. This makes it best suited for analyzing the long-run effects of persistent housing policies, which are slow to implement but have lasting impacts. The model assumes instantaneous housing reallocation as higher-wealth households move into newly supplied higher-quality segments, freeing up units for others. In reality, this process may unfold over years, especially until it reaches the lower end of the market, which is crucial for improving rental affordability. A mitigating factor is that prices typically adjust quickly in anticipation of future supply shifts, potentially reducing housing costs before the long-run equilibrium is reached.

Second, the model omits capital gains from home sales. It assumes lower house prices are always beneficial, as they reduce the cost of housing services. In reality, major shifts in household-unit assignments involve widespread buying and selling. For existing homeowners, falling prices due to increased supply can mean financial losses, helping explain resistance to new construction. Incorporating dynamics into the model would allow us to capture these capital gains and losses. Some studies already explore how policy reforms affect capital gains for heterogeneous households: useful starting points are Floetotto, Kirker, and Stroebel (2016) and İmrohoroğlu, Matoba, and Tüzel (2018). Taking these effects into account could change the compensating variation number presented in Figure 6.

Third, our model does not include a construction sector. In reality, increasing housing supply requires developers and builders to be compensated at market rates. By abstracting from construction costs, the model assumes uniform costs across housing segments. However, evidence from affordable housing mandates and rent control shows that developers tend to favor high-end construction. Without regulation, new building would likely occur in the middle- or upper-market segments—consistent with our policy simulations. In contrast, mandates to build lower-quality housing at below-market prices may distort incentives and reduce overall supply over time (Diamond, McQuade, and Qian 2019; Krimmel and Wang 2023).

A further implication of excluding the construction sector is our assumption that supply remains fixed when analyzing mortgage subsidies. In practice, rising prices often prompt developers to build more, but this response varies by city due to natural and regulatory constraints. Supply elasticity captures this variation, with places like San Francisco exhibiting very low responsiveness. Saiz (2010) estimates a medium- to long-run elasticity of 0.66 for San Francisco, while Baum-Snow and Han (2024) estimate a lower long run elasticity of 0.27,[13] supporting our fixed-supply assumption. Additionally, our model does not consider how mortgage subsidies are funded. Because they likely draw from limited tax revenues, accounting for funding costs would further reduce the policy's net welfare benefits.

A final caveat is that our model analyzes the housing market of a city in isolation. Our approach accounts for realized changes in population size over time, but not endogenous in- or out-migration in response to housing becoming cheaper or more expensive. Any policy that lowers housing costs in San Francisco will attract new residents, which will increase housing pressures again.

Conclusion

We use an equilibrium model of a local housing market to examine the causes of and potential solutions to housing unaffordability. Our analysis indicates that effective policies must increase housing supply. While supply interventions may target specific market segments, their effects extend across the entire housing market. Notably, expanding high-end housing supply reduces competition for mid- and lower-tier units, improving affordability more broadly. A straightforward long-term strategy is for local governments to permit and even support developers in building the housing they find most profitable—often at the higher end of the market.

[13] Baum-Snow and Han (2024) estimate supply elasticities at the level of census tracts. We obtain the aggregate elasticity by averaging across their tract-level estimates. The elasticities measure the percentage change in supply for a 1 percent increase in housing prices.

As for any structural analysis, our conclusions depend on the validity of our modeling assumptions. In reality, household decisions are shaped by more than wealth, and housing quality is not one-dimensional. A more complex model could incorporate these nuances and address some of the omissions we discussed above, potentially altering some findings. Still, our framework likely captures the key economic mechanisms that would persist even in a richer model.

■ *We thank Nate Baum-Snow, Jonathan Parker, Monika Piazzesi, Martin Schneider, and Timothy Taylor for helpful comments and Adithya Raajkumar for excellent research assistance.*

References

Abramson, Boaz. 2021. "The Equilibrium Effects of Eviction Policies." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4112426.

Abramson, Boaz, and Stijn Van Nieuwerburgh. 2024. "Rent Guarantee Insurance." NBER Working Paper 32582.

Abramson, and Tim Landvoigt. 2025. *Data and Code for: "Curbing Rising Housing Costs: A Model-Based Policy Comparison."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E2334017V1.

Anenberg, Elliot, and Edward Kung. 2020. "Can More Housing Supply Solve the Affordability Crisis? Evidence from a Neighborhood Choice Model." *Regional Science and Urban Economics* 80: 103363.

Asquith, Brian J., Evan Mast, and Davin Reed. 2023. "Local Effects of Large New Apartment Buildings in Low-Income Areas." *Review of Economics and Statistics* 105 (2): 359–75.

Baum-Snow, Nathaniel, and Lu Han. 2024. "The Microgeography of Housing Supply." *Journal of Political Economy* 132 (6): 1897–946.

Baum-Snow, Nathaniel, and Justin Marion. 2009. "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods." *Journal of Public Economics* 93 (5–6): 654–66.

Board of Governors of the Federal Reserve System. 2013-2022. *Survey of Consumer Finances (SCF).* https://www.federalreserve.gov/econres/scfindex.htm (accessed June 18, 2025).

Corbae, Dean, Andrew Glover, and Michael Nattinger. 2024. "Equilibrium Evictions." NBER Working Paper 32898.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst. 2023. "Income Growth and the Distributional Effects of Urban Spatial Sorting." *Review of Economic Studies* 91 (2): 858–98.

Diamond, Rebecca, and Tim McQuade. 2019. "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development." *Journal of Political Economy* 127 (3): 1063–117.

Diamond, Rebecca, Tim McQuade, and Franklin Qian. 2019. "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco." *American Economic Review* 109 (9): 3365–94.

Díaz, Antonia, and Belén Jerez. 2013. "House Prices, Sales, and Time on the Market: A Search-Theoretic Framework." *International Economic Review* 54 (3): 837–72.

Favilukis, Jack, Pierre Mabille, and Stijn Van Nieuwerburgh. 2023. "Affordable Housing and City Welfare." *Review of Economic Studies* 90 (1): 293–330.

Favilukis, Jack, and Stijn Van Nieuwerburgh. 2021. "Out-of-Town Home Buyers and City Welfare." *Journal of Finance* 76 (5): 2577–638.

Floetotto, Max, Michael Kirker, and Johannes Stroebel. 2016. "Government Intervention in the Housing Market: Who Wins, Who Loses?" *Journal of Monetary Economics* 80: 106–23.

Fonseca, Julia, Lu Liu, and Pierre Mabille. 2025. "Unlocking Mortgage Lock-In: Evidence from a Spatial Housing Ladder Model." Working Paper.

French, Robert, and Valentine Gilbert. 2023. "Suburban Housing and Urban Affordability: Evidence from Residential Vacancy Chains." Unpublished.

Garriga, Carlos, and Aaron Hedlund. 2020. "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession." *American Economic Review* 110 (6): 1603–34.

Glaeser, Edward L., and Bryce A. Ward. 2009. "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston." *Journal of Urban Economics* 65 (3): 265–78.

Gupta, Arpit, Christopher Hansman, and Pierre Mabille. 2024. "Financial Constraints and the Racial Housing Gap." INSEAD Working Paper 2022/58/FIN.

Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics* 11 (2): 1–39.

İmrohoroğlu, Ayşe, Kyle Matoba, and Şelale Tüzel. 2018. "Proposition 13: An Equilibrium Analysis." *American Economic Journal: Macroeconomics* 10 (2): 24–51.

İmrohoroğlu, Ayşe, and Kai Zhao. 2022. "Homelessness." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4308222.

Krimmel, Jacob, and Betty Wang. 2023. "Upzoning with Strings Attached: Evidence from Seattle's Affordable Housing Mandate." *Cityscape* 25 (2): 257–78.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. "The Housing Market(s) of San Diego." *American Economic Review* 105 (4): 1371–407.

Määttänen, Niku, and Marko Terviö. 2014. "Income Distribution and Housing Prices: An Assignment Model Approach." *Journal of Economic Theory* 151: 381–410.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15 (2): 145–61.

Mense, Andreas. 2025. "The Impact of New Housing Supply on the Distribution of Rents." *Journal of Political Economy Macroeconomics* 3 (1): 1–42.

Nathanson, Charles G. 2023. "Trickle-Down Housing Economics." Unpublished.

Nikolakoudis, George. 2024. "The Economics of Segmented Housing Markets." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4251925.

Ortalo-Magné, François, and Sven Rady. 2006. "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints." *Review of Economic Studies* 73 (2): 459–85.

Pennington, Kate. 2021. "Does Building New Housing Cause Displacement? The Supply and Demand Effects of Construction in San Francisco." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.3867764.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel. 2020. "Segmented Housing Search." *American Economic Review* 110 (3): 720–59.

Rubin, Donald B. 1986. "Statistical Matching Using File Concatenation with Adjusted Weights and Multiple Imputations." *Journal of Business and Economic Statistics* 4 (1): 87–94.

S&P Dow Jones Indices LLC. 2013-2022. *S&P CoreLogic Case-Shiller CA-San Francisco Home Price Index [SFXRSA]*. https://fred.stlouisfed.org/series/SFXRSA (accessed June 18, 2025).

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–96.

US Census Bureau. 1980–2022. *American Community Survey (ACS)*. https://usa.ipums.org/usa-action/variables/group (accessed June 18, 2025).