Online Appendix

June 2, 2025

A Model Description

There are two dates, $t = 0, 1$. The model consists of a continuum of households of total mass 1 that live in a city. Each resident household in the city occupies one housing unit. Households spend their income on consumption of housing services and other goods.

Housing Technology. In order to consume housing services, a household must buy or rent a house, which is a durable good that exists in N different quality levels. A house of quality segment $n \in \{1, \ldots, N\}$ yields housing services h_n, with ordering $h_n \leqslant h_{n+1}$. The N total segments consist of N_R rental segments and N_O owner-occupied segments, i.e. $N = N_R + N_O$. Segments are ordered such that the bottom segments are rentals and the top segments are owner-occupied. Houses in rental segments are only available for rent and houses in owner-occupied segments are only available for purchase. Each segment has a unique competitive market price p_n, which is the rent in R segments and the purchase price in O segments.

Household problem. Households are ordered by wealth $w_i \sim F$. Households receive this income w_i in period 0, and $\delta_w w_i$ in period 1, with δ_w parameterizing wealth growth. A house purchased in period 0 delivers services in both periods, 0 and 1. Purchased houses are sold in period 1 at price $\delta_p^n p_n$, where δ_p^n is a parameter governing house price growth in segment n. Rented houses require payment of the rent in each period, which we nest by setting $\delta_p^n = -1$ for rental segments. The value function of a household with wealth w is then

$$V(w) = \max_{c_0, c_1, h_n, b} \; u(c_0, h_n) + \beta u(c_1, h_n) \tag{1}$$

subject to

$$w \geqslant c_0 + p(h_n) - b, \tag{2}$$

$$c_1 \leqslant \delta_w w + \delta_p^n p(h_n) + Rb, \tag{3}$$

$$\theta^n p(h_n) \geqslant -b. \tag{4}$$

The period utility function in (1) satisfies the usual assumptions. The choice of housing comes down to picking a quality level $h_i \in \{h_n\}_{n=1}^N$. The price function $p(h_n)$ maps the quality segment choice into the vector of prices $\{p_n\}_{n=1}^N$. Households can borrow and save at interest rate R. However, borrowing is only possible in the form of mortgage debt and subject to an LTV constraint given in (4), with the parameter θ^n dictating maximum leverage in segment n ($\theta^n = 0$ in rental segments). The household problem gives rise to the optimal choices as function of wealth, $[c_0(w), c_1(w), h(w), b(w)]$.

Housing supply and market clearing. In each quality segment, real estate developers supply housing of mass H_n, with $\sum_{n=1}^{N} H_n \leqslant 1$. In addition to the $n = 1, \dots, N$ quality segments, households can choose to rent a residence of quality h_0 at price p_0 in both periods. Housing at the 0 quality level is perfectly elastically supplied at this price. Segment 0 stands in for homelessness, doubling up, or other types of marginal housing arrangements. Market clearing in the housing markets then requires that the mass of households choosing quality n is equal to the supply in this segment H_n.

$$H_n = \int \mathbb{1}_{[h(w)=h_n]} dF(w) \quad \text{for } n = 1, \dots, N. \tag{5}$$

If $\sum_{n=1}^{N} H_n < 1$, then the remaining households rent quality zero housing at price p_0.

Solution Approach. The housing decision of each household comes down to choosing a quality level n from a set of discrete options. Substituting in the budget constraints, the value function of a household choosing option $j \in \{0, \dots, N\}$ is

$$V_j(w) = \max_{b} \; u(w - p_j - b, h_j) + \beta u(\delta_w w + \delta_p^j p_j + Rb, h_j) \tag{6}$$

subject to $-b \leqslant \theta^j p_j$. Conditional on the housing choice, (6) clarifies that the household problem reduces to a constrained optimization problem in the net savings choice b. Solutions to this problem satisfy the usual first-order condition

$$u_c(c_0, h_j) \geqslant \beta R u_c(c_1, h_j)$$

jointly with the borrowing constraint. Since utility is increasing in both goods, and housing services are strictly increasing in quality, $h_n > h_{n-1}$, the model yields a well-known assignment equilibrium, such that there exist N thresholds in income $\{\hat{w}_n\}_{n=1}^{N}$, and

$$H_n = F(\hat{w}_{n+1}) - F(\hat{w}_n), \quad \text{for } n = 1, \dots, N-1, \tag{7}$$

$$H_N = 1 - F(\hat{w}_N), \tag{8}$$

$$H_0 = F(\hat{w}_1). \tag{9}$$

The housing choices in turn are implicitly defined through a set of indifference conditions. Households with wealth levels equal to the thresholds $\{\hat{w}_n\}_{n=1}^{N}$ are indifferent between choosing housing of quality n or $n - 1$, respectively:

$$V_n(\hat{w}_n) = V_{n-1}(\hat{w}_n), \quad \text{for } n = 1, \dots, N. \tag{10}$$

The housing choices of the households between thresholds then follow the assignment structure

$$h_i = h_n \quad \text{if } \hat{w}_{n+1} > w_i \geqslant \hat{w}_n, \text{ for } n = 1, \dots, N-1,$$
$$h_i = h_N \quad \text{if } w_i \geqslant \hat{w}_N,$$
$$h_i = h_0 \quad \text{if } w_i < \hat{w}_1.$$

Assignment Equilibrium. The equilibrium can therefore be fully characterized by a set of thresholds $\{\hat{w}_n\}_{n=1}^N$, prices $\{p_n\}_{n=1}^N$, and savings choices $\{b_n^+, b_n^-\}_{n=1}^N$ that satisfy the market clearing conditions (7) – (9), the indifference conditions (10), and maximize the conditional value function in (6). Note that with each threshold, we have two associated savings choices, which are

$$b_n^+ = \underset{b}{\operatorname{argmax}} \; u(\hat{w}_n - p_{n+1} - b, h_{n+1}) + \beta u(\delta_w \hat{w}_n + \delta_p^{n+1} p_{n+1} + Rb, h_{n+1}) \;\; s.t. \; -b \leqslant \theta^{n+1} p_{n+1}, \tag{11}$$

$$b_n^- = \underset{b}{\operatorname{argmax}} \; u(\hat{w}_n - p_n - b, h_n) + \beta u(\delta_w \hat{w}_n + \delta_p^{n} p_n + Rb, h_n) \;\; s.t. \; -b \leqslant \theta^{n} p_n. \tag{12}$$

The model can be solved numerically as a system of $4N$ nonlinear equations in equally many unknowns.

B Data and Calibration

This appendix provides more detail on the model quantification and the data that disciplines it. We begin by discussing the quantification of the baseline model. Table 1 presents the calibrated parameters of this model.

Preferences. We calibrate household preferences based on standard values from the literature. For the utility function, we specify a CRRA utility over a Cobb-Douglas aggregator of housing services and other consumption (Landvoigt, Piazzesi and Schneider, 2015). The relative risk aversion γ is set to 2, which is within the range that is considered plausible by Mehra and Prescott (1985). We set weight on housing services to $\eta = 0.3$ as in Abramson (2021). We assume that the two model periods, today and tomorrow, span a 10-year period in the data. The discount factor is assumed to be 0.9. The interest rate is consistent with the discount rate (that is, $R = 1/\beta$).

Wealth. We calibrate the wealth distribution in the model, $F(w)$, to match the empirical wealth distribution of households in the data. To estimate the empirical wealth distribution of households in San Francisco, we combine two data sources. The first is the 2013 American Community Survey (ACS), which surveys a representative sample of households in the US. For each household, the ACS records the household's annual income, whether the household is an owner-occupier or renter, whether the household has a mortgage and if so what are the monthly payments on the mortgage, the age of the head of the household, and the MSA in which the household resides. A drawback of the ACS is that it does not record households' financial assets, which are part of households' wealth.

The second data source is the 2013 Survey of Consumer Finances (SCF). The advantage of the SCF, which is also nationally representative, is that it records households' financial assets. Financial assets are defined by the SCF as the balance between total assets and total debt, coded as "networth". The drawback relative to the ACS is that the SCF does not identify the MSA of the respondent's residence. However, since both the SCF and the ACS record households' income, home-ownership status, whether the household has a mortgage, the monthly payments on the mortgage, and the age of the head of the household, we can estimate the relationship between financial assets and these covariates in the SCF and use it to impute households' assets in the ACS.

For imputation, we use the predictive mean matching (PMM) imputation method (Rubin, 1986). PMM first estimates a linear regression of the outcome variable (in our case, financial assets) on a set of covariates (in our case, income, home-ownership status, age, whether the household has a mortgage, and the monthly payments on the mortgage) for the set of observations where the outcome variable is not missing (in our case, the SCF data). Second, it obtains linear predictions of the outcome variable for observations where the outcome variable is missing (in our case, the ACS data). For each observation where the outcome variable is missing, it then uses the linear prediction as a distance measure to form a set of nearest neighbors (observations where the outcome variable is non-missing, in our case the SCF data). Finally, it randomly draws an imputed value from this set. Using the PMM method, we impute financial assets for all households in San Francisco. Household wealth is then defined as the sum of income and (imputed) financial assets. The calibrated wealth distribution is illustrated in blue in the top left panel of Figure 4. We assume, for simplicity, that wealth does not grow (i.e. $\delta^w = 1$).

Housing supply. We consider a city with 13 housing segments. Houses in the top 10 segments are houses that households can own and houses in the bottom three quality segments are houses that households can rent. Mathematically, this is captured by the assumption that $\delta_p^n = 0$ for $n = 0, 1, 2$ and $\delta_p^n = 1$ for $n = 3, ..., 10$. In terms of the distribution of houses across segments, we assume the top 10 segments account for 54.7% of the total housing stock, which is the observed home-ownership rate in San Francisco in the 2013 ACS data. Each of the top 10 segments is assumed to be of equal size. That is, $H_3 = H_4 = ... = H_{12} = 0.0547$. The second and third segments each account for 21.9% of the housing stock ($H_1 = H_2 =$

0.219), and the lowest segment accounts for the remainder 1.5% ($H_0 = 0.015$). We conceptualize the bottom segment as corresponding to informal renting arrangements such as doubling-up and homeless shelters. The distribution of the housing stock across quality segments is illustrated by the blue bars in the top right panel of Figure 4.

Borrowing. We set the maximum loan-to-value in each segment of the market, θ^n, as follows. For the first three segments, which correspond to the rental market, $\theta^n = 0$. For the first segment of the owner market (which accounts for 10% of the owner-occupied market) we set the maximum loan-to-value to the average loan-to-value ratio of homeowners in the bottom decile of the national wealth distribution, as computed from the 2013 SCF data. For the second segment, we set the maximum loan-to-value to the average loan-to-value of homeowners in the second decile of the wealth distribution, and so on. The values of θ^n are reported in Table 1.[1]

Housing Quality. The quality of housing in each segment, $\{h_0, ..., h_{12}\}$, is estimated endogenously. In particular, our strategy is to estimate house qualities so that the model implied prices match the observed prices in the data. The intuition is simple. Given household preferences, wealth distribution, the distribution of the housing stock across segments and the quality of houses 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. We numerically solve for the house qualities that align the model prices perfectly with the data.

We measure house prices in the owner-occupied market (i.e. in the top 10 segments of the market) using Corelogic. Corelogic is a private vendor that compiles the universe of housing transactions as well as property tax records from across the US. We begin with the transaction data. For each transaction, we observe the transaction price, the address of the unit transacted, a rich set of physical characteristics of the unit and whether the unit is owner-occupied. A unit is defined as owner-occupied if and only if the mailing address of the registered owner is the same as the property address. Focusing on the universe of transactions of owner-occupied houses in San Francisco in 2013, we estimate a hedonic house price model that links the physical characteristics and location of transacted houses to their house price. In particular, we estimate the following specification:

$$p_i = \beta X_i + u_i, \tag{13}$$

where p_i is the (log) price of unit i and X_i is a vector of controls that includes zip-code fixed effects,

[1] These average loan-to-value ratios are of course much lower than maximum loan-to-value ratio available to new home buyers.

indicators for the the property type (e.g. apartment, single-family, condo, etc.), indicators for the number of beds and baths, for the age of the building, for the unit's size, as well as indicators for whether the unit has various amenities such as air-conditioning, central heating, parking, garage. We also include all possible triple interactions between the zip-code fixed effect, the property type, and each of the other indicators. The R^2 from this hedonic regression is 0.75. We denote the estimated coefficients from this regression by $\hat{\beta}$.

We then turn to the Corelogic tax data, which records the address and physical characteristics of *all* housing units in San Francisco in 2013 (not only those sold) and their occupancy status, but does not record their price. To impute the price of all owner-occupied housing units in the tax data, we use the hedonic model that we estimated on the transaction data. That is, for each observation j in the tax data, we impute $\hat{p}_j = \hat{\beta} X_i$. Finally, we divide the stock of owner-occupied houses to 10 deciles and compute the median house price within each segment. Denote the 9 price cutoffs that define the 10 segments of the owner market by $\{\overline{p}_1, ..., \overline{p}_9\}$.

To measure rents in the top two rental segments, we use the 2013 ACS data, which records rent payments for renter households. Namely, we sort renters in San Francisco in 2013 based on their rent and compute the median annual rent in each of the two halves of the renter distribution. Finally, we assume rent in the bottom segment (which corresponds to informal rental arrangements and is unobserved) is $\$6,000$ annually and normalize house quality in this segment to one. Given the targeted prices in the data, we estimate all remaining house qualities so that the model implied prices are in line with the data. The house qualities that are estimated so that the model implied prices match the observed prices in the data are illustrated by the solid line in the top right panel of Figure 4.

B.1 Model Calibration to 2022

Having calibrated the model to San Francisco in 2013, we then measure **realized changes** in the distributions of (i) housing supply and (ii) the wealth of residents over the 2013-2022 period. By feeding these measured changes as inputs into the model, we identify the causes of the decline in affordability ratios over the 2013-2022 period. The result is that we have a representation of San Francisco in 2022 in the calibrated model. This section describes the measurement of the wealth and housing supply distribution in 2022.

Wealth distribution in 2022. We calibrate the wealth distribution to match the empirical wealth distribution of households in San Francisco in 2022. We estimate the empirical wealth distribution using ACS and SCF data from 2022, following the same steps described above for the baseline calibration. The calibrated wealth distribution for the 2022 model is illustrated in red in the top left panel of Figure 4.

Table 1: Model Parameters

Parameter	Value
	Preferences
Utility function $u(c,h)$	$\frac{(c^{1-\eta}h^{\eta})^{1-\gamma}}{1-\gamma}$
Risk aversion γ	2
Weight on housing η	0.3
Discount factor β	0.9
	Technology
House price growth $\{\delta_p^0, ..., \delta_p^{12}\}$	{0,0,0,1,1,1,1,1,1,1,1,1,1}
Wealth growth δ^w	1
Interest rate R	$1/\beta$
Maximum LTV $\{\theta^0, ..., \theta^{12}\}$	{0 0 0 0.75 0.65 0.56 0.44 0.43 0.39 0.35 0.30 0.29 0.25}
Housing supply $\{H_0, ..., H_{12}\}$	Figure 4 (top right panel, in blue)
Housing quality $\{h_0, ..., h_{12}\}$	Figure 4 (top right panel, in black)
Wealth distribution $F(w)$	Figure 4 (top left panel, in blue)

Table 1 presents the parameters of the baseline model.

Housing supply in 2022. We calibrate the 2022 distribution of housing supply across the 13 segments, $H_0, H_1, ..., H_{12}$, as follows. We begin by setting the total supply in the top 10 segments of the market, $\sum_{n=3}^{12} H_n$, to 0.571. This is consistent with the observed home-ownership rate in San Francisco in 2022, which is 57.1%.

Next, we calibrate the distribution of housing supply within the owner market, $H_3, ..., H_{12}$. To do so, we assign each owner-occupied housing unit in 2022 to one of the 10 segments of the owner-occupied market and measure the share of units assigned to each segment. The basic idea is to impute the price that each housing unit in 2022 would have sold for in 2013, and then use the price cutoffs $\{\overline{p}_1, ..., \overline{p}_9\}$ that define the segments of the owner market (and that were computed as part of the baseline calibration in terms of 2013 dollars) for assignment. We implement this as follows. First, we impute the price of all owner-occupied housing units in San Francisco in 2022 using the Corelogic transaction and tax data, following the same steps described for the baseline calibration. Second, we deflate these prices to 2013 dollars by dividing them by the growth rate of house prices in San Francisco between 2013 and 2022 according to the Case-Shiller house price index (https://fred.stlouisfed.org/series/SFXRSA). Third, we assign each owner-occupied housing unit in 2022 to the quality segment it belongs to based on its deflated 2013 price. For example, if the imputed 2013 price of unit i is below the cutoff $\overline{p}_1$, it is assigned to the first segment of the owner market.

Finally, we calibrate the distribution of housing supply within the renter market, H_0, H_1, H_2. We maintain the baseline assumption that the two rental market segments are of equal size. We assume that the size of the first segment, H_0 has increased to 0.024. The increase in the size of the bottom segment is calibrated 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 ACS data, which is 11 percentage points. This implies that the sizes of the second and third segments in the 2022 calibration is 0.20. Overall, the calibrated housing supply distribution in 2022 is illustrated by the red bars in the top right of Figure 4.

C Additional Figures

Figure C.1: Quantitative Model - 2022

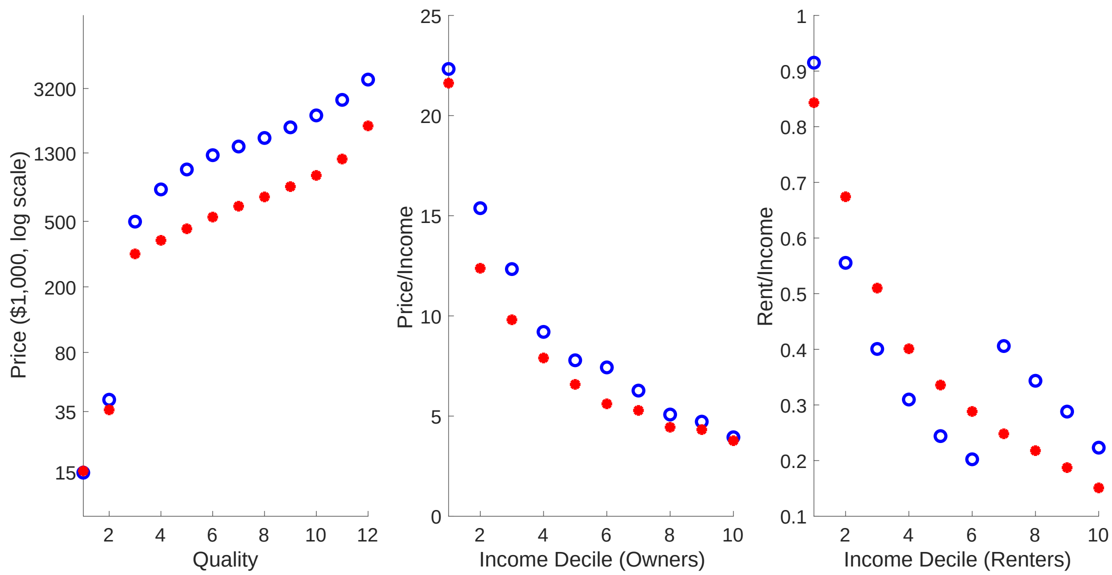

Notes: This figure illustrates the model equilibrium in 2022. All model parameters are set at their baseline values, except for the wealth and housing supply distributions which are estimated from the 2022 data, as explained in the main text. The left panel shows equilibrium prices for each segment in 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). Data moments are computed based on the 2022 ACS.

D Additional Cities

This section replicates our analysis for two additional cities - Chicago and Dallas. While these cities are notably different from San Francisco in terms of their population and housing markets, the analysis suggests that the drivers of the affordability crisis across the three cities are common. In all three cities, a rise of wealth inequality, accompanied by little new construction, raises prices and rents and leads to a deterioration in affordability.

For each of the additional cities, we begin by matching the model to the data in 2013. We calibrate the wealth distribution in the model to match the empirical wealth distribution of households in the city, following the procedure discussed in Appendix B. Household preference parameters are fixed across cities and are listed in Table 1. We assume that there are 13 housing segments in a city, where houses in the top 10 segments are houses that households can own and houses in the bottom three segments are houses for rent. In terms of the distribution of houses across segments, we assume the top 10 segments account for a percentage of the total housing stock that is equal to the observed home-ownership rate in the city in the 2013 ACS data, that each of the top 10 segments is of equal size, that the lowest segment accounts for 1.5% of the housing stock, and that segments 2 and 3 are of equal size. Following the estimation procedure described in Appendix B, we estimate the house qualities to match prices and rents in the data. Figure D.1 (Figure D.2) illustrates the equilibrium in the quantitative model for Chicago (Dallas). The model matches well both targeted and non-targeted moments.

Next, for each city, we quantify how changes in the wealth and housing supply distributions between 2013 and 2022 impact affordability through the lens of model. We measure changes in the wealth and housing supply distributions following the procedure described in detail in Appendix B.1. The calibrated wealth distribution in Chicago (Dallas) in 2022 is illustrated in red in the top left panel of Figure D.3 (D.4). The calibrated housing supply distribution in Chicago (Dallas) in 2022 is illustrated by the red bars in the top right of Figure D.3 (D.4). Figures D.5 and D.6 show that the model's predictions regarding prices, price-to-income and rent-to-income ratios in Chicago and Dallas in 2022 are closely in line with the data, providing further model validation.

As with the case of San Francisco, Dallas and Chicago have also experienced an increase in wealth inequality between 2013 and 2022 alongside a decline in housing affordability. Our model provides a rational for this co-movement. All else equal, the increase in households' wealth drives up housing prices. The rise in wealth can be mitigated by supplying more high-quality houses, but in practice Chicago and Dallas have not increased the supply in the top segments sufficiently. In fact, the overall supply of housing relative to population has decreased (as illustrated by the increase in the size of the bottom segment). Together, these have led to a deepening affordability crisis in both the owner-occupied and rental market, and particularly so for the poorest households (as illustrated by the bottom panels of Figures D.3 and D.4).

Figure D.1: Quantitative Model - Chicago

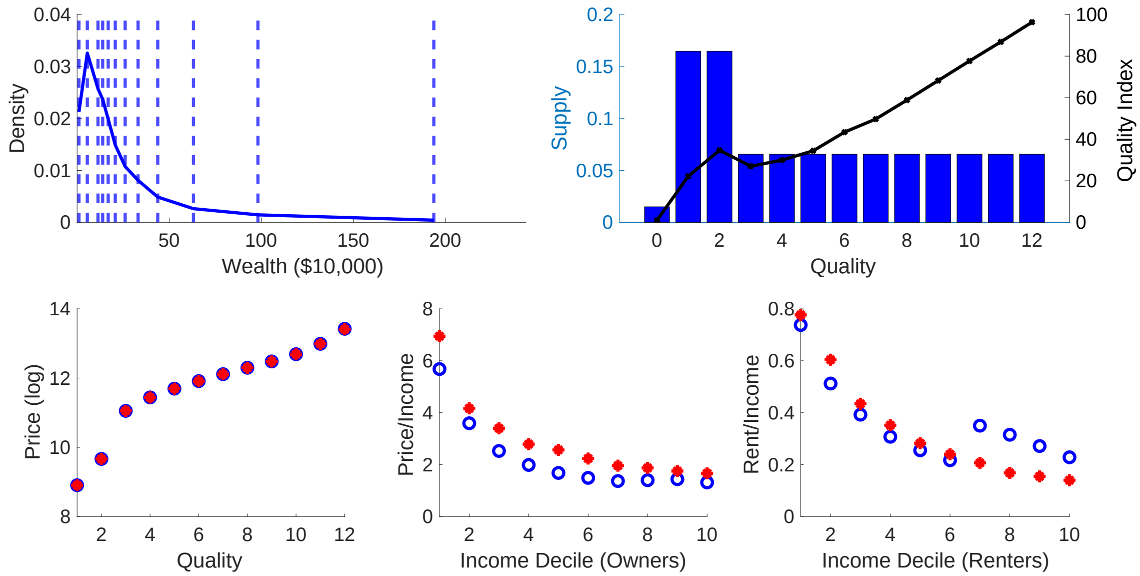

Notes: This figure illustrates the equilibrium in the quantified model for Chicago. The top left graph plots the density of the wealth distribution, estimated based on ACS and SCF data. The top right panel shows the distribution of housing supply across segments (bars, left axis), and the quality index associated with housing segments (line, right axis). The bottom left panel shows equilibrium prices for each segment in the model (blue) and in the data (red). The bottom 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). Data moments are computed based on the 2013 ACS.

Figure D.2: Quantitative Model - Dallas

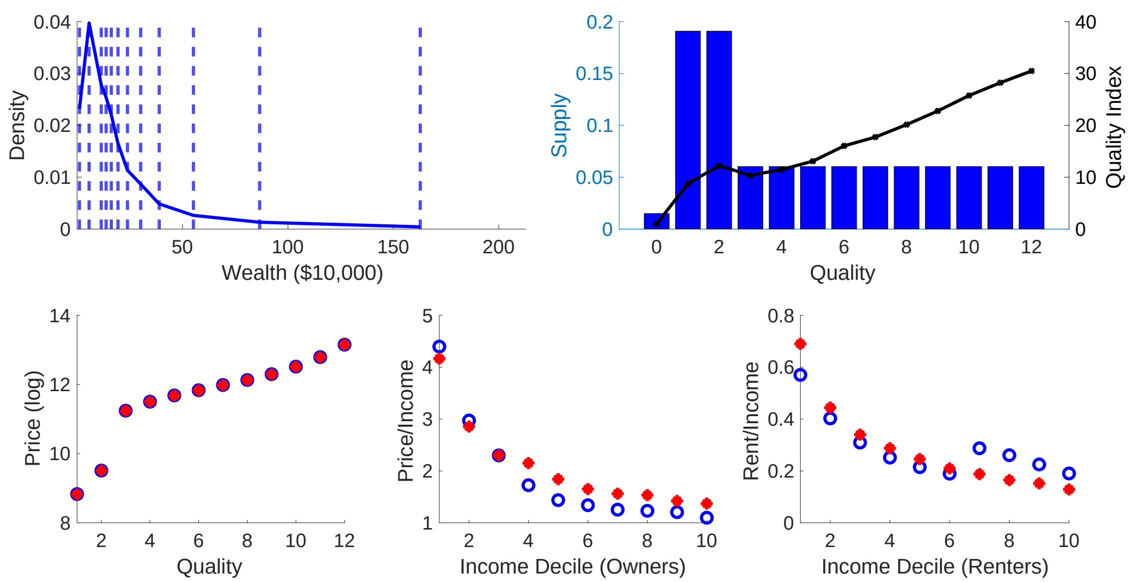

Notes: This figure illustrates the equilibrium in the quantified model for Dallas. The top left graph plots the density of the wealth distribution, estimated based on ACS and SCF data. The top right panel shows the distribution of housing supply across segments (bars, left axis), and the quality index associated with housing segments (line, right axis). The bottom left panel shows equilibrium prices for each segment in the model (blue) and in the data (red). The bottom 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). Data moments are computed based on the 2013 ACS.

Figure D.3: Drivers of the Affordability Crisis - Chicago

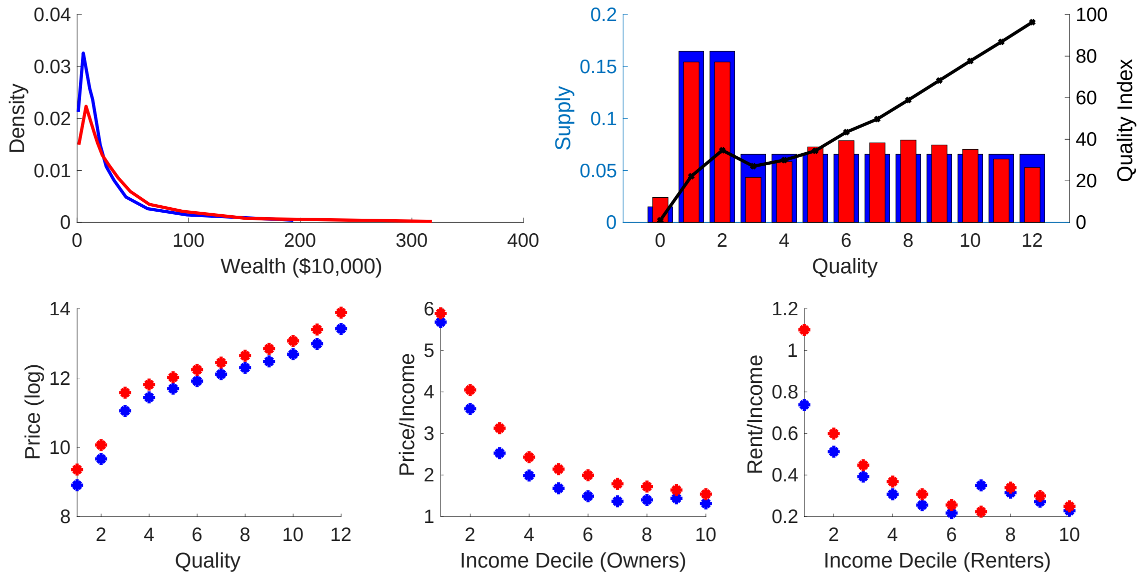

Notes: This figure illustrates the effect of the observed change in housing supply and in the wealth distribution between 2013 and 2022 in Chicago. The top left graph plots the density of the calibrated wealth distribution. The top right panel shows the distribution of housing supply across segments (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 mid (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.

Figure D.4: Drivers of the Affordability Crisis - Dallas

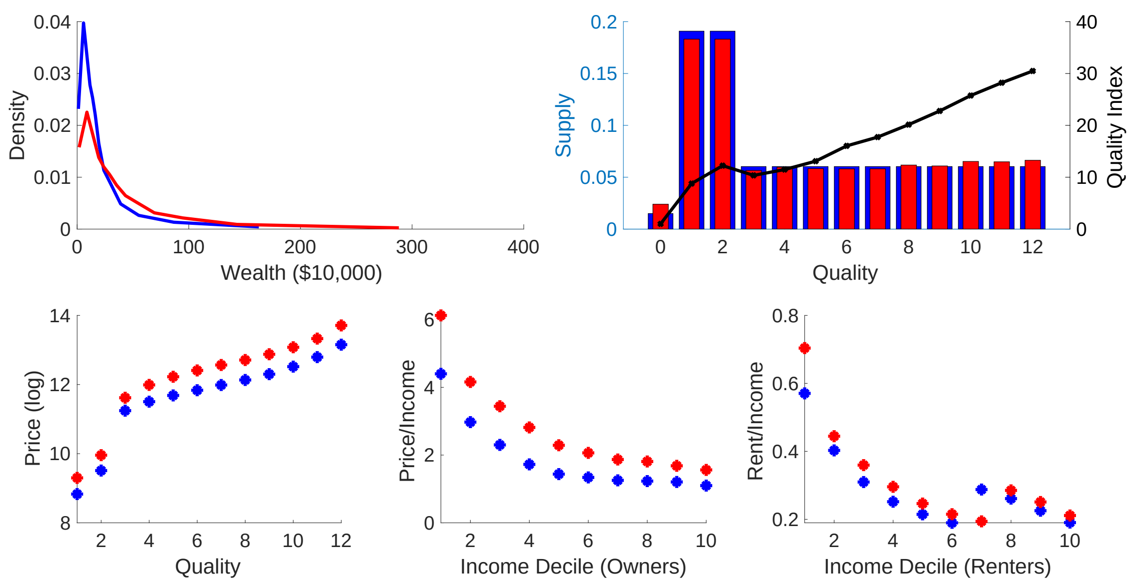

Notes: This figure illustrates the effect of the observed change in housing supply and in the wealth distribution between 2013 and 2022 in Dallas. The top left graph plots the density of the calibrated wealth distribution. The top right panel shows the distribution of housing supply across segments (bars, 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 mid (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.

Figure D.5: Quantitative Model Chicago - 2022

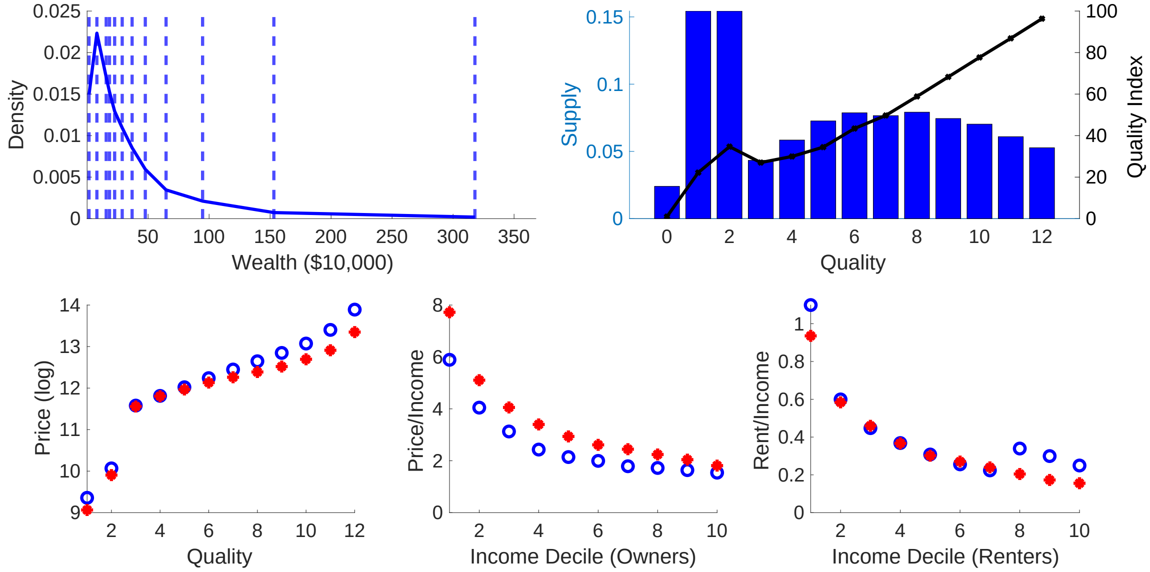

Notes: This figure illustrates the model equilibrium in 2022 for Chicago. The top left graph plots the density of the wealth distribution. The top right panel shows the distribution of housing supply across segments (bars, left axis), and the quality index associated with housing segments (line, right axis). The bottom left panel shows equilibrium prices for each segment in the model (blue) and in the data (red). The bottom 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). Data moments are computed based on the 2022 ACS.

Figure D.6: Quantitative Model Dallas - 2022

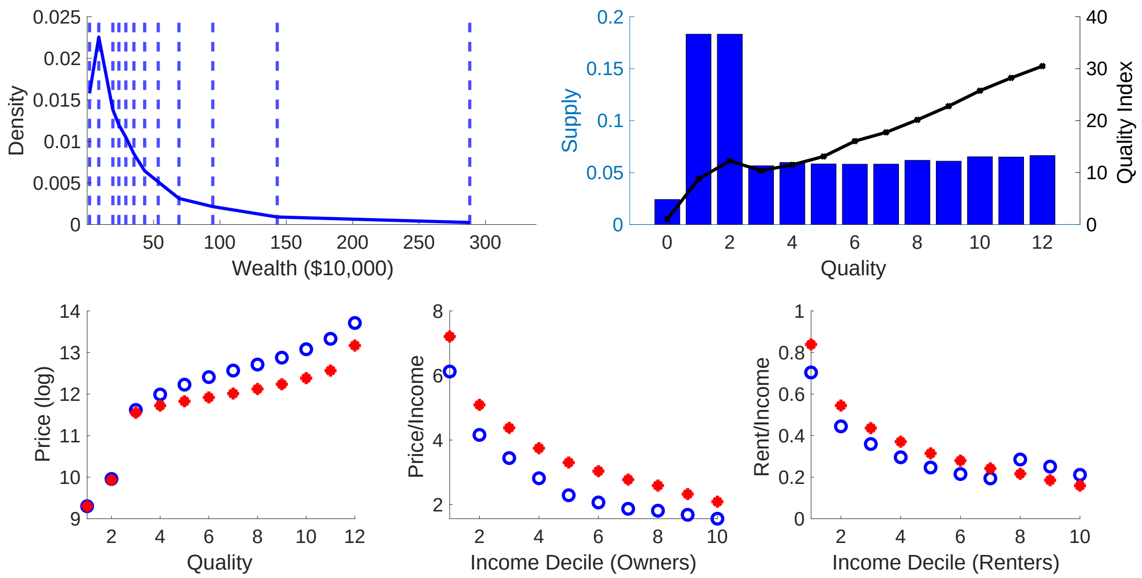

Notes: This figure illustrates the model equilibrium in 2022 for Dallas. The top left graph plots the density of the wealth distribution. The top right panel shows the distribution of housing supply across segments (bars, left axis), and the quality index associated with housing segments (line, right axis). The bottom left panel shows equilibrium prices for each segment in the model (blue) and in the data (red). The bottom 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). Data moments are computed based on the 2022 ACS.

References

Abramson, Boaz. 2021. "The Equilibrium Effects of Eviction Policies." *Available at SSRN 4112426*.

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

Mehra, Rajnish, and Edward C Prescott. 1985. "The equity premium: A puzzle." *Journal of monetary Economics*, 15(2): 145–161.

Rubin, Donald B. 1986. "Statistical matching using file concatenation with adjusted weights and multiple imputations." *Journal of Business & Economic Statistics*, 4(1): 87–94.