

Volatility in Home Sales and Prices: Supply or Demand? *

Elliot Anenberg[†] Daniel Ringo[‡]

December 16, 2022

Abstract

We use a housing search model and data on individual home listings to decompose fluctuations in home sales and price growth into supply or demand factors. Simulations of the estimated model show that housing demand drives short-run fluctuations in home sales and prices, while variation in supply plays only a limited role. We consider two implications of these results. First, we show that reduction of supply was a minor factor relative to increased demand in the tightening of housing markets during COVID-19. New for-sale listings would have had to expand 30 percent to keep the rate of price growth at pre-pandemic levels given the pandemic-era surge in demand. Second, we estimate that housing demand is very sensitive to changes in mortgage rates, even more so than comparable estimates for home sales. This suggests that policies that affect housing demand through mortgage rates can influence housing market dynamics.

*The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

[†]Board of Governors of the Federal Reserve System

[‡]Board of Governors of the Federal Reserve System

1 Introduction

Both home sales and home price growth are volatile and cyclical, generally rising together during booms and falling during busts. These fluctuations have important implications for economic activity, financial stability, and access to homeownership.¹ They are also frequent targets of policy making, as governments have a variety of policy options that can affect the demand for homes (e.g. first-time home buyer tax credits, mortgage subsidies, monetary policy) or the supply of homes available for sale (e.g. zoning reform, tax assessment restrictions, transfer taxes). To predict the outcomes of these policy choices, it is therefore important to understand the determinants of housing market volatility. The main objective of this paper is to estimate the extent to which short-run fluctuations in sales and price growth are driven by the demand for homes, or by the supply of homes for sale.

To decompose these fluctuations into supply or demand factors, we use a simple model in which the stock of active buyers and sellers produce market dynamics through a frictional housing search process. Our housing search model is motivated by Figures 1 and 2, which show that changes in home sales and prices are accompanied by changes in housing liquidity. Figure 1 shows that homes take longer to sell (i.e. a low rate of sale hazard) when the the stock of active sellers is high, and Figure 2 shows that the average time on market of for-sale homes is strongly negatively associated with home price growth. These correlations suggest an important role for market tightness, or the ratio of buyers to sellers, in explaining short-run housing dynamics. We therefore take the number of newly active sellers and buyers entering the market as the fundamental measures of demand and supply we will be investigating.² These inflows, balanced with the outflows due to successful matches and discouraged searchers, determine the market tightness in equilibrium.

An important empirical challenge is that we do not have data on the number of potential buyers actively searching for a home. Sellers advertise for-sale homes to buyers through platforms such as the multiple listings service (MLS), but active buyers generally do not record their presence or search activity. Consequently, we

¹For home prices and economic activity, see Aladangady (2017); Berger et al. (2017); Guren et al. (2021); Mian and Sufi (2011); Mian, Rao and Sufi (2013). For home sales, see Benmelech, Guren and Melzer (forthcoming); Karahan and Rhee (2019); Ortalo-Magne and Rady (2006).

²Newly active sellers include both builders of new construction and sellers of existing homes, with the latter typically accounting for a large majority of newly active sellers.

obtain data on the inflow of new listings for sale, the stock of active sellers, and the sale hazard rate from the United States MLS, and we use our model structure combined with these MLS data to estimate housing demand.

Using the model and estimates of housing demand, we consider simulations where we hold fixed either supply or demand, and allow the other to vary as in the data. We find that fluctuations in housing demand explain much more of the variation in home sales and price growth than do fluctuations in housing supply. In our preferred parameterization, fluctuations in demand explain essentially all of the variation in home sales, and 80% of the variation in prices, between 2002-2021. In other parameterizations, demand can be forced to play a somewhat less important role, but its strong contribution relative to supply is a robust result.

We then consider two implications of the results from our model. First, we show that the COVID-19 housing boom in the U.S. was driven by an increase in demand. Even though the supply of new for-sale listings fell sharply at the beginning of the pandemic, we show that reduction of supply was a minor factor relative to increased demand in explaining the tightening of housing markets over the first year of the pandemic. A policy concern during the pandemic has been that the sharp rise in house prices has exacerbated affordability pressures and increased financial stability risks. We use our model to estimate how much additional supply would be needed to offset the observed increase in demand so that house prices continued along their pre-pandemic trend, instead of accelerating. We find that a 30% increase in the monthly number of homes coming on to the market would have been necessary to keep up with the pandemic-era surge in demand. Since new construction typically accounts for about 15% of supply, our estimates imply that new construction would have had to increase by roughly 300% to absorb the pandemic-era surge in demand. This is a very large, unrealistic impulse to housing supply in the short-run, suggesting that policies aimed at reducing bottlenecks to new construction would have done little to cool the housing market during COVID-19.³

Second, we show that our estimate of housing demand is very mortgage rate elastic. We estimate that a one percentage point increase in the mortgage rate lowers housing demand by 10.4 percent. This is a larger demand sensitivity to rates than evidence using purely observable housing market variables suggests. In particular,

³In the long run, increasing new construction may be a more effective policy response. Longer run fluctuations in the housing market are beyond the scope of this paper.

higher mortgage rates also decrease home sales, but the semi-elasticity is 6, or about one-half the semi-elasticity for housing demand. Because search frictions effectively smooth the response of home sales to demand shocks over time, estimates of the short-term elasticity of home sales obscure some of the mortgage rate sensitivity of demand. A high mortgage rate sensitivity of demand combined with our main result showing that short-run housing market fluctuations are largely explained by demand suggest that policies that target mortgage rates are an effective way to influence short-run fluctuations in the housing market.⁴

Our paper is related to a large literature that takes a search-theoretic approach to modeling the dynamics of the housing market. Han and Strange (2015) provide a summary of this literature. Our model of random housing search is similar to a number of models in the literature – see, for example, Diaz and Jerez (2013), Guren and McQuade (2020), Krainer (2001), Novy-Marx (2009), Piazzesi and Schneider (2009), and Gabrovski and Ortego-Martí (2019). Our contribution, building on Anenberg and Ringo (2021), is to show that this simple model of housing search combined with time-series data on new listings, for-sale inventory, and withdrawals can be used to estimate housing demand.⁵ We use this estimate of buyer demand along with our data and model to provide new results on the contribution of supply and demand to volatility in the housing market.

Our paper is closely related to the empirical work in Ngai and Sheedy (2020) (henceforth NS). Using data from the U.S., NS find that variation in the supply of homes for sale explains essentially all of the volatility in sales volume, whereas we find that supply has very little explanatory power for volatility in sales volume. Motivated by their empirical finding that sales volume dynamics are explained by supply, NS also build a search-and-matching model with endogenous house prices and moving decisions.

⁴Our estimates of the rate semi-elasticity of home buying demand is in line with other estimates in the literature using shocks to the cost of credit. Estimated semi-elasticities of demand for home buying range from 20 for lower-income, credit constrained households to 6 for middle- and upper-income households (Bhutta and Ringo, 2020 and Ringo, 2022). Estimates of the elasticity of the *intensive* margin of mortgage demand—that is, the amount borrowed conditional on a loan being taken out at all—are much smaller. For example, DeFusco and Paciorek (2017) estimate the rate semi-elasticity of intensive margin demand to be about 2.

⁵Concurrent work by Gabrovski and Ortego-Martí (2021) uses a similar model to estimate the pool of active buyers in order to estimate the slope of the Beveridge curve in the housing market. To estimate the pool of active buyers, Gabrovski and Ortego-Martí (2021) use Census data on vacancies and time-to-sell for new home sales.

In Section 6, we show that our results differ from the empirical results in NS for two reasons. First, we make different assumptions about how supply affects the sale hazard rate. In NS, the stock of active listings for sale does not affect the sale hazard rate. In our model the sale hazard depends on the market tightness, meaning it varies endogenously with housing supply. As supply increases (all else equal), the sale hazard goes down and offsets much of the effect of increased supply on sales volume. Our results highlight the importance of modeling market tightness, and its implications for the matching process, when evaluating the relative roles of demand and supply. Second, we use micro data on individual listings that allow us to directly measure supply – i.e. the inflow of new listings for sale. NS use aggregate data and as a result of data limitations, their measure of supply is actually new listings *net of withdrawals*. Because withdrawals are negatively correlated with demand, the NS measure of supply is influenced by demand factors, causing their estimates to overstate the contribution of supply to volatility in sales volumes.⁶

Our finding that demand is important for explaining variation in sales volume is consistent with and related to a number of recent papers providing quasi-experimental evidence that sales volumes are sensitive to demand stimulus. Bhutta and Ringo (2020) and Anenberg and Ringo (2022) find that changes in mortgage rates have important effects on home sales. Berger, Turner and Zwick (2020) find that a national first-time homebuyer tax credit, a similarly demand-side policy, had a meaningful stimulative impact on home sales. Best and Kleven (2017) find that sales volumes in the U.K. are sensitive to transaction taxes. However, while the statutory incidence of this tax falls on the buyer, their paper does not attempt to determine whether the change in volume occurred through a demand or supply response. We are not aware of any quasi-experimental studies of the effect of the stock of for-sale listings on sales volumes.

2 Data and Motivating Empirical Patterns

Our data are MLS records provided by CoreLogic. The data come directly from regional boards of realtors, and cover over 50 percent of property listings in the U.S.

⁶Withdrawals are negatively correlated with demand because sellers often become discouraged after failing to find a buyer for an extended period of time. For additional evidence on the counter-cyclicity of withdrawals, see Carrillo and Williams (2019).

Information on homes listed for sale includes the initial listing date, the withdrawal date if the home is removed from the MLS without a sale and is not subsequently relisted a short time later, the contract date if the home is sold to a buyer, the asking price, and many home characteristics, including the address.⁷ The MLS data have some advantages for our purposes over aggregated listings data, such as those published by the National Association of Realtors (NAR). The data on individual listings allow us to observe the actual inflow of new listings as opposed to inferring it from net changes in total for-sale listings and home sales, a procedure that can lead to mismeasurement caused by withdrawals and homes that list and sell within the same month. Furthermore, the listing-specific data allows us to control for characteristics of the house or listing that could affect the sale hazard (e.g. compositional effects or whether the seller has set an asking price well above or below prevailing prices).

The data run from 2002-2021. From the full sample, we select a subset of 263 counties due to data limitations for some counties. We describe these limitations as well as our procedure for selecting counties in the Appendix.

Figure 1 shows trends in sales volume, new listings, for-sale inventory, and the sale hazard rate over our sample period. The sale hazard rate is calculated as the number of sales contracted each month divided by the number of homes actively listed at some point during the month. The annual sale hazard is the average of the monthly sale hazards, weighted by the number of homes listed for sale each month. Sales volume rises during the early 2000s, and then falls sharply during the Great Recession. Sales volume slowly recovers from its fall and only in recent years has the level of sales volume returned to early 2000s levels.

One might expect sales volume to be closely related to new listings, as homes can only transact if they are put on the market for sale. Remarkably, however, Figure 1 shows that new listings are only weakly correlated with sales volume over our sample period. New listings and home sales both rise during the early 2000s, but then diverge as sales volume declines during the Great Recession and new listings remain elevated. New listings have remained fairly flat over the last decade even as sales volume has been on a strong upward trend. These trends suggest that understanding the behavior of new listings alone is not sufficient for understanding cyclicity in sales volume.

In contrast, the figure shows that sales volume and the sale hazard rate have a

⁷CoreLogic treats properties that are withdrawn from the market and then relisted a within two months as a single, continuous listing.

very high correlation. In addition, Figure 2 shows that the sale hazard is also strongly associated with house price growth. The figure shows the 12-month change in a real house price index and the “months’ supply”.⁸ Months’ supply is the ratio between the number of homes for sale and the number of sales, or the inverse of the monthly sale hazard. Months’ supply alone can explain about 80 percent of the variation in house price growth over our sample period. The strong associations of the sale hazard rate with sales volume and house price growth suggest that understanding the drivers of the sale hazard rate is key for understanding the cyclical nature of the housing market. The importance of the sale hazard rate motivates our model of housing search in Section 3. Our model focuses on this sale hazard rate and allows us to predict how the sale hazard rate changes with supply and demand.

2.1 County-level Evidence

County-level statistics provide additional motivating evidence for the limited role new listings play in explaining short-run variation in sales volume. Table 1 reports regression results of the 12-month growth in sales volume on the 12-month growth in new listings and the sale hazard rate. Each variable is measured at the county level and each regression pools observations across counties. The first column shows that a one percentage point increase in new listings growth is associated with a 0.35 percentage point increase in sales volume growth, but the R^2 is only 0.08. New listings growth alone explains a tiny fraction of the variation in sales volume growth. The second column shows that a one percentage point increase in sale hazard growth is associated with a 0.73 percentage point increase in sales volume growth, and the R^2 from this regression is 0.70. Growth in the sale hazard explains almost ten times more of the variation in sales volume growth than does new listings growth. The remaining columns show that the results are similar when fixed effects are added for county, or for county and year-month.

⁸We compute the quality-adjusted house price index from our micro data using a standard hedonic price regression, as described in the Appendix.

3 Model of Housing Search

To facilitate our decomposition of housing market cyclicalities into demand and supply factors, we use a simple model of housing search. We define supply as the *flow* of homes coming onto the for-sale market each period. Similarly, we define demand as the *flow* of prospective buyers that enter the market to find a home to purchase. In our model, supply and demand affect the housing market equilibrium through their effect on market tightness, θ , which in turn affects the rate at which homes are sold. Market tightness is the ratio of the *stock* of prospective buyers, b , relative to the stock of sellers, s : i.e. $\theta = \frac{b}{s}$.

The stock of sellers (i.e. the inventory of homes for sale), s , in each month t evolves as

$$s_{t+1} = s_t - s_t q_t^s(\theta_t) - s_t(1 - q_t^s(\theta_t))w^s + n_{t+1}^s \quad (1)$$

where q^s is the rate at which homes are sold, w^s is the rate at which unsold homes are withdrawn from the market, and n^s is the inflow of new sellers (i.e. our fundamental measure of supply). Equation 1 expresses the stock next period as the stock this period (first term) minus the outflow arising from sales and withdrawals (middle terms) plus the inflow (final term).

Similarly to the supply of homes for sale, there is a stock of currently-searching buyers, b , that is replenished by an inflow of new buyers, and depleted as buyers purchase a home and exit the market, or drop out without purchasing. The stock of buyers evolves as

$$b_{t+1} = b_t - b_t q_t^b(\theta_t) - b_t(1 - q_t^b(\theta_t))w^b + n_{t+1}^b \quad (2)$$

where q^b is the rate at which a buyer finds a home to buy, w^b is the rate at which buyers leave the market, and n^b is the inflow of new buyers.

Buyers and sellers interact via the search and matching process, which we model as Cobb-Douglas with constant returns to scale. We discuss this choice of functional form and robustness to alternative specifications of the search-and-matching function in the Appendix. Under Cobb-Douglas, the probabilities of buying and selling are:

$$q_t^s(\theta_t) = \theta_t q_t^b(\theta_t) = A_t \theta_t^\eta \quad (3)$$

where $0 < \eta < 1$ is the elasticity of the probability of sale with respect to market tightness and A is a parameter that determines the efficiency of the matching function. We allow A_t to vary over time based on factors exogenous to our model, discussed further in Section 4.1.

Because q^s is an increasing function of market tightness, the more sellers there are in the market, the slower a given house is likely to sell (all else equal). Sellers crowd each other out and create congestion by competing for the stock of available buyers - the more houses there are for sale, the less likely any particular house is to receive an offer. This prediction is consistent with the very strong negative correlation in the data between the number of sellers on the market, s , and the sale hazard rate, q^s shown in Figure 1.

An alternative to the random search model described above would be to model housing search via stock-flow matching. For example, Smith (2020) uses such a stock-flow model with endogenous seller entry to explain hot and cold housing markets. In the Appendix, we show robustness of our main results in a stock-flow model where homes that have just come to market are more efficient searchers than homes that have been on the for-sale market for some time. Another possible modeling choice would be directed search. Albrecht, Gautier and Vroman (2016) develop a directed search model where motivated sellers choose low list prices and relaxed sellers choose high list prices, leading to shorter and longer time-to-sell, respectively.⁹ Our model abstracts from the list price decision, but, as we describe below in Section 4.1, we incorporate list prices by allowing them to affect matching efficiency. As a result, a change in the composition of sellers—for example, from relaxed, high list-price sellers to motivated, low list-price sellers—can affect the sale hazard rate without affecting our estimate of housing demand.

4 Estimation and Calibration

To impute the relative influence of supply and demand on housing market fluctuations, we need estimates of n_t^b and n_t^s . As discussed above, n_t^s is directly observed from our listings data. We estimate n_t^b using our model structure and the listings data. Our approach to estimating n_t^b requires estimates of the sale hazard, q^s ; estimates of the matching efficiency, A ; and calibration of several parameters. We discuss each in

⁹List prices play a similar role in the quantitative directed search model in Hedlund (2016).

turn.

4.1 Estimating sale hazard, q^s , and matching efficiency, A

We estimate sale hazards and matching efficiency using our panel of active listings at a monthly frequency. Houses enter the panel either in the month they are listed for sale, or in January 2002 if the listing was already active at that point. They exit when the house is delisted, and the panel as a whole ends in November 2021. Some homes are delisted because a sale has occurred, others are delisted because the seller has decided to no longer market the home for sale. Homes that are delisted from the market without a sale are treated as censored observations.

Using this sample, we estimate a time period specific sale hazard for each month of the panel. This sale hazard is intended to represent that of a generic listing, affected only by the number of active buyers and sellers, so we need to control for variation in the composition of listings that could affect sale hazards. To accomplish this, we estimate an accelerated failure time model where the hazard rate of sale for house i at time t is parameterized as

$$h_{it} = \exp(\delta_t + \beta^A X_{it}^A) \quad (4)$$

where δ_t denotes a set of month-year fixed effects and X^A is a vector of observables that affect matching efficiency. In X^A , we include characteristics of the home, such as its age and number of bathrooms, as well as the home's list price relative to an expected market sales price.¹⁰ These listing-specific characteristics could affect sale hazards for reasons external to the count-based notions of supply and demand we are concerned with. For example, very old homes may not be suitable matches for many buyers. If the pool of homes for sale happen to be older than is typical, matching efficiency for that time period could be low. A high list price relative to market sales price could proxy for a low seller search intensity or unrealistic expectations, also lowering matching efficiency. A mismatch between seller expectations and buyer willingness to pay was especially relevant during the years of the financial crisis, when homeowners were slow to accept how much the price of their homes had fallen.

The effect of X^A on probability of sale is identified using cross-sectional variation,

¹⁰The expected market sales price is the predicted value from a auxiliary regression of log sales prices on home characteristics and month-year dummy variables.

and the month-year fixed effects (δ_t) capture residual variation in average sale hazard over time that is not related to X^A . We interpret variation in δ_t as variation in the sale hazard rate that is related to variation in market tightness. We use an exponential hazard function because equation 3 implies that A_t has a proportional effect on the sale hazard given the market tightness. In equation 4, the log of the sale hazard is additively separable in the logs of matching efficiency and market tightness, consistent with equation (3).

Our estimate of the sale hazard rate is just the average predicted value from estimating equation 4:

$$\hat{q}_t^s = \frac{1}{N_t} \sum_{i=1}^{N_t} \exp(\hat{\delta}_t + \hat{\beta}^A X_{it}^A) \quad (5)$$

where N_t denotes the number of homes listed on the market in period t .

We estimate A_t as

$$\hat{A}_t = \bar{A} \left(\frac{\hat{q}_t^s}{\exp(\hat{\delta}_t)} \right) \quad (6)$$

which is the estimated sale hazard rate net of the estimated contribution of the month-year fixed effects. Because the month-year fixed effects can only be estimated relative to a baseline period, A_t is identified up to a scale parameter, \bar{A} .

We also estimated a much simpler alternative specification which attributes all variation over time in the sale hazard to market tightness (i.e. do not allow the sale hazard to vary with X^A). This simpler specification where A_t is fixed in every period yields qualitatively similar results to our main specification for A_t described in this subsection.

4.2 Calibration of parameters

As is common in the housing literature with random search, we calibrate the elasticity of the matching function, η , to 0.84 to match the estimate from Genesove and Han (2012).¹¹ Genesove and Han (2012) estimate η using cross-market regressions and survey data on buyer time-on-market, seller time-on market, and number of homes visited by buyers. Subsequent work has validated this estimate using a variety of

¹¹For example, Anenberg and Bayer (2020), Guren and McQuade (2020), and Guren (2018) also calibrate to Genesove and Han (2012).

different identification strategies. Head, Lloyd-Ellis and Sun (2014) calibrate η in their housing search model to target the relative volatility of sales growth to income growth, and arrive at $\eta = 0.86$. Recent work by Grindaker et al. (2021) also arrives at almost an identical estimate to Genesove and Han (2012) using a shift-share shock to market tightness.

Under our matching function, the addition of an extra buyer or seller to the market increases sales volume, but not one-for-one, as the addition of an extra buyer (seller) creates competition or crowd out for other buyers (sellers), lowering the probability of a match. The calibration of $\eta = 0.84$ implies that the addition of an extra buyer to the market has a relatively low crowd out effect on the probability that other buyers in the market match with a for-sale home. Adding an extra seller to the market, however, has a comparatively larger (negative) effect on the probability that other sellers in the market match with a buyer.¹² Genesove and Han (2012) discuss how these crowd-out results could be generated by the MLS. The MLS allows buyers to observe all for-sale listings, but sellers cannot typically observe anything about the pool of potential buyers or take active steps to match with a particular buyer. As a result, buyers can more easily and quickly substitute to other listings when multiple buyers are interested in the same house (i.e. if it sells just before they tour it). Sellers, in contrast, must passively wait for interested buyers to arrive.

We calibrate $\bar{A} = 1.4$ using survey data from the NAR on average search time for buyers in 2019.¹³ Buyers reported searching for 10 weeks on average, and we calibrate \bar{A} so that the median buyer simulated in 2019 matches in this time frame.

For our counterfactual simulations, described below in Section 5, we calibrate $w^s = 0.061$ to match the average monthly withdrawal hazard in our MLS data. In those simulations, we hold the withdrawal hazard fixed to ensure we are isolating the variation in housing market outcomes due exclusively to supply or demand, respectively. As described in Section 6, the withdrawal hazard of for-sale listings is a function of demand conditions, and failing to distinguish changes in the inflow of new listings from changes in the outflow of withdrawals can lead to spurious inference.

¹²To see this, note that $\partial M/\partial b = A\eta\theta^{\eta-1}$ and $\partial M/\partial s = A(1-\eta)\theta^\eta$ where $M = sq^s$ denotes the number of matchings or sales. The relative crowd out effects depend on θ , but except in very tight markets (i.e. those with very large values of θ), $\partial M/\partial b > \partial M/\partial s$ for $\eta = 0.84$. When market tightness is high, the addition of an extra seller does relatively more to stimulate sales than when tightness is low.

¹³Source: National Association of Realtors (2019)

We do not have any data or external estimates to inform w^b , and the dynamics of the model depend on the *net* inflow of potential buyers less withdrawals, rather than the gross inflows and withdrawal outflows individually. Therefore, for simplicity we normalize $w^b = w^s$. Our estimates of the inflow of new buyers, n^b , as described below can therefore be thought of as the fluctuations over time in the net inflow of new buyers less withdrawals, up to a constant scalar determined by our normalization of average w^b .

4.3 Inferring demand, n^b

In this section, we present the equation that expresses n^b in terms of variables and parameters that we can observe or estimate using our data. First, note that by inverting equation (3), we can express the number of buyers in any period, which is unobserved in our data, as

$$b_t = s_t \left(\frac{q_t^s}{A_t} \right)^{\frac{1}{\eta}} \quad (7)$$

This equation provides an estimate of b_t because s_t is observed in our data, η is a parameter that we calibrate, and equations 5 and 6 give estimates of q^s and A .¹⁴ Second, by plugging equation (3) into equation (2) and rearranging, we can express n^b as

$$n_t^b = b_t - b_{t-1} + s_{t-1} q_{t-1}^s + b_{t-1} \left(1 - \frac{s_{t-1} q_{t-1}^s}{b_{t-1}} \right) w^b \quad (8)$$

Given equation (7), the right-hand side of equation (8) depends only on variables that can be observed or estimated, allowing us to estimate n_t^b .

¹⁴The estimate of b_t is essentially a residual and so we would attribute any unobserved shocks that affect matching efficiency to variation in b_t . As described in Section 4.1, however, observable characteristics of listings that we can control for in X^A do very little to explain the time-series variation in sale hazards. We therefore believe that variation in sale hazards, conditional on the number of listings and their observed characteristics, is largely due to variation in demand and that our estimates of b_t are a good proxy for the number of actively searching buyers.

5 Partialling Out the Effects of Demand and Supply

Our next step is to infer how much of the variation in housing market activity is driven by n_t^s and \hat{n}_t^b , observed supply and estimated demand, respectively.

Our first counterfactual scenario isolates the effect of supply by holding demand constant and letting supply follow its observed course. That is, we set the simulated inflow of potential buyers \widetilde{n}_b^t equal to the estimated sample mean, \bar{n}_b , in every period. The simulated inflow of new listings, \widetilde{n}_s^t , is set equal to the observed n_s^t . We initialize the market using the observed values of the state variables in January 2002, and simulate forward following equations 1, 2, and 3 of our model. Our second scenario isolates the effect of demand by allowing demand to follow its estimated course, while holding supply constant. We set the simulated inflow of buyers equal to its estimated monthly values, $\widetilde{n}_b^t = \hat{n}_b^t$. The simulated inflow of new listings is instead set to equal its sample mean ($\widetilde{n}_s^t = \bar{n}_s$). Again, the market is initialized at January 2002 levels and simulated forward. In both scenarios, A_t , which captures changes in observable characteristics of listings that affect the sale hazard, is held fixed at its sample mean. Withdrawal hazards for both buyers and sellers, w^b and w^s , are likewise fixed at a constant in both simulations to isolate the effect of new entrants.

It is worth noting that many home sellers are builders, investors, or households leaving the owner-occupied housing market. Conceptually, supply from these sellers can vary while demand remains fixed as in our counterfactual simulations. Similarly, many buyers (such as first time home buyers) are not simultaneously trying to sell another home, so demand from these buyers can vary while supply remains fixed. However, many households are attempting to move from one owner-occupied home to another, and so enter the market as both buyers and sellers. For this reason, fixing the entirety of either the supply or demand side constant while letting the other vary fully may not be a realistic economic counterfactual for the housing market. Still, these joint buyer-sellers make choices about their dual-search problem, including whether to enter the market first as a buyer, first as a seller, or as a buyer and seller simultaneously. Our simulations help us understand which side of the market their potential entry in matters more.

5.1 Sales Volume

Figure 3 displays the monthly volume of home sales recorded in our MLS data, alongside counterfactual values from the constant-demand and constant-supply simulations described above. The simulated volumes with constant demand and time-varying supply are nearly flat and only weakly correlated with observed sales. The implication is that variation in supply has essentially no explanatory power over sales volumes.

In contrast, the simulated volumes with constant supply but time-varying demand match the realized sales data very well. There is some small deviation between the series during the Great Recession period, when sellers were particularly likely to set asking prices well above what buyers were willing to pay, but variation in demand explains the overwhelming majority of variation in sales.

5.1.1 Model Choice and Crowd-Out

An important parameter in our model is the elasticity of the matching function, calibrated to $\eta = 0.84$. The closer to 1 is η , the more crowd-out each listing creates. That is, a marginal new listing is more likely to poach a buyer from a different listing, rather than create a new sale. Marginal potential buyers, on the other hand, create little crowd out and so overall sales volume should be quite responsive to variation in the number of buyers. Variation in supply could thus be more important for smaller values of η .

As discussed above, the empirical literature without exception finds evidence for a high value of η . To gauge the sensitivity of our results to lower values of η , however, we show results for $\eta = 0.16$. This alternative calibration is symmetric to the baseline calibration, except buyers, and not sellers, create more crowd out. Results from these alternative simulations are shown in Figure 4. The reduced supply-side crowd out is clearly visible in the greater variation in sales generated by the constant-demand, varying-supply simulations. However, this simulation still makes a poor fit to the actual data. A regression of the sales volume data on simulated sales has an R^2 of essentially zero. The varying-demand, constant-supply simulations, on the hand, still do much better in matching the data. While the fit is not as good as in our preferred specification, a regression of sales volume data on these simulated sales has an R^2 of 0.78.

The impotence of supply (and hence dominance of demand) follows from the data

and is essentially a requirement of the observation that total for-sale listings and sales volumes are not well correlated. As can be seen in Figure 1, over the time period we study the correlation is actually substantially negative. If the supply of homes for sale were the major determinant of sales volumes, total listings and sales should generally move together. Under a calibration with more buyer than seller crowd out, the logic of the search model implies an even greater pro-cyclicality of demand to match the observed negative correlation between listings and sales than it does in our preferred calibration. External evidence that crowd-out is actually greater on the supply than on the demand side (motivating a high value of η) only reinforces the relative importance of demand over supply.

5.2 Months' Supply and House Prices

With the various counterfactual series of b and s simulated as described above, we can also simulate counterfactual paths of the months' supply of homes for sale, shown in Figure 5. Compared to sales volumes, months' supply is more responsive to variation in new listings. While variation in new listings has little effect on the numerator of the months' supply ratio (sales volume), it is an important determinant of the denominator (total active listings). The figure shows that the vary-supply counterfactual rises and falls between 2002 and 2012, which is consistent with the general pattern in the data over this time period. However, the counterfactual simulation shows an increase in months' supply from 2012 onward whereas in the data, months' supply declined over this time period. Overall, the R^2 from a regression of the true months' supply on the vary-supply simulated months' supply is 0.26.

As with the sales volume counterfactuals, we find that the vary-demand simulation explains a larger share of the variation. The vary-demand simulation follows the true path of months' supply closely, though it cannot account for the full rise in months' supply during the Great Recession. The elevated level of new listings over this time period contributed significantly to the rise in months' supply. The R^2 from a regression of the months' supply on the vary-demand simulation is 0.8, three times as large as the vary-supply R^2 .

Months' supply of homes for sale is a figure of particular interest because of its close connection with house prices. As shown in Figure 2, there is a very tight negative correlation between months' supply and house price growth, a relationship

sometimes referred to as the housing Phillips curve. The R^2 of this relationship is about 0.80. Caplin and Leahy (2011) and Guren (2018) discuss how a negative relationship between the *level* of months' supply and *changes* in house prices is difficult to explain in a model with full information and rationality. While our housing search model does not take a stand on the specific house price formation process, one way to motivate this tight connection is with a model in which sellers have only limited information about the demand for their homes, but can observe the recent experience of other for-sale listings. A tight market, as evidenced by the rapid sale of recent listings, informs sellers that they can raise prices. Slower sales, indicating that the market is not so tight, would inform them that they may need to lower prices to make a timely sale. Price adjustments to rebalance the number of prospective buyers and sellers willing to transact at current price levels (i.e. market tightness) cause the negative correlation between price growth and months' supply. We describe such a model in the Appendix. Similar intuition for the housing Phillips curve can be found in the models of Carrillo, de Wit and Larson (2015) and Guren (2018).¹⁵

Taking the reduced-form relationship in Figure 2 as given and feeding counterfactual months' supply from Figure 5 into that relationship, we find that price fluctuations in the short term mostly depend on demand. Variation in supply also has a meaningful influence on months' supply and hence price growth, though its influence is much smaller than that of demand.¹⁶

¹⁵In Carrillo, de Wit and Larson (2015), sellers are slow to realize when there has been a shock to the number of buyers searching on the market, and that their bargaining power has consequently changed. Market tightness improves the seller's relative bargaining position and hence increases sales prices. As individuals learn about the new level of market tightness, they slowly adjust their positions. Market tightness thus predicts price growth over multiple future periods. Guren (2018) models some fraction of sellers as using a backward-looking heuristic for house prices. This behavior generates momentum in house prices and corresponding inventory volatility, exacerbated by other, forward-looking households strategically timing their market entry to take advantage of the predictability of price growth.

¹⁶This inference about the determinants of price growth assumes that the observed relationship between market tightness and prices is causal, or mechanically linked as in the model described in the Appendix. However, it is possible that the effect of a shock to market tightness on price growth would not be as strong as the tight correlation apparent in Figure 2 might suggest. Prices and months' supply could both be influenced independently by some third factor. Lacking a clean source of quasi-experimental variation in demand and supply, we cannot be certain that the true effect is as strong as the observed correlation. Nonetheless, given the clear theoretical connection and that tight correlation, it seems very likely that market tightness (and, consequently, demand) is the primary driver of short-run house price growth.

5.3 Explaining Cross-Sectional Variation in Housing Markets

The analysis shown thus far has been exploiting and explaining time-series variation in the U.S. housing market in the aggregate. Housing market dynamics differ widely by locality within the country, however. In the Appendix we show that, across counties, variation in sales growth can also be well explained by variation in demand but not at all by variation in supply. Similarly to the results from the aggregate time series, this inference does not rely on a calibration in which there is more crowd out on the supply side than on the demand side (although it is reinforced by such a calibration). As in the aggregate results, supply plays a larger role in explaining cross-sectional variation in months' supply than in explaining sales volumes.

6 Contrast with a Reduced Form Approach

An alternative, more reduced-form approach to partialling out the relative importance of supply and demand factors is to simulate various counterfactuals, varying or holding constant the terms in the accounting identity described in equation 1. This is the method used by Ngai and Sheedy (2020) (NS). In this section we replicate their motivating empirics, including a key element of their data construction. As described in Section 2, the aggregate NAR data NS use doesn't allow for measuring the number of new listings distinct from the number of withdrawals. We show why the reduced-form approach, along with the conflation of new listings and withdrawals, leads to opposite inferences about the importance of the supply side in determining sales volumes.

NS perform counterfactual simulations similar to those described in our Section 5, holding some determinants of sales fixed while allowing others to follow their observed path. The principal conceptual difference relative to our approach is that, rather than modelling the matching process and inferring the number of active buyers, they treat the sale hazard (q^s , in equation 1) as a model primitive. Depending on the counterfactual being considered, this q^s is fixed as a constant or set following its observed historical path. The inflow of new listings, n^s , is treated similarly as in our approach. However, instead of using actual new listings and holding the withdrawal hazard constant as we do, given their aggregate data they must infer new listings as

the monthly difference in for-sale inventories, net of sales. This construction is actually *net* new listings because it does not distinguish between an increase (decrease) in new listings and a decrease (increase) in withdrawals.

We follow their basic approach, once again initializing the market at the values observed in January 2002. We use equation 1 to simulate two counterfactual paths of active listings and sales. In the first, the sale hazard q^s is fixed at its sample mean, while the inflow of new listings n^s follows its observed path. To replicate their data construction, in this section we measure n^s as monthly new listings minus the monthly number of withdrawals and otherwise set the withdrawal hazard, w^s , equal to zero. In the second simulation, q^s is allowed to vary as it does in the data, while n^s is fixed at its sample mean. Simulated sales volumes under the two reduced-form scenarios, along with actual sales, are presented in Figure 6. In contrast to the simulations using our full model (shown in Figure 3), the reduced-form counterfactuals suggest the inflow of new listings has a powerful effect on the number of sales, and can explain much of the time-series variation. Unlike our simulated inflow of new buyers, variation in the sale hazard by itself does a poorer job explaining the number of sales. While there are some differences due to time period and additional data construction issues, these reduced-form findings qualitatively match those of NS.

Why does fixing the sale hazard as in NS lead to such different conclusions from the method we used in Sections 3 through 5? Recall from our model that the sale hazard, q^s , is a function of supply as well as demand: the more houses there are for sale, the more actively searching buyers are needed to maintain a particular sale hazard. This modeling of the sale hazard is consistent with the negative empirical correlation between for-sale inventory and the sale hazard shown in Figure 1. Given the substantial observed time variation in the supply of new for-sale listings, a substantial amount of variation in the number of active buyers would be necessary to have kept the sale hazard fixed at a constant. Implicitly, the NS simulations in which listings vary but sale hazard is fixed involve considerable variation in demand.

This can be seen by taking the counterfactual sales data from the reduced-form simulations with varying supply and fixed sale hazard, and backing out the implied time series of demand (b_t) using the model described in Section 3. We apply the method described in Section 4.3 to this simulated data, and present the imputed level of counterfactual demand in Figure 7. For comparison, we show the inferred time series of b_t based on our estimates from Section 4. As can be seen, the two series

are quite similar. Given the actual inflow of new listings, a fixed sales hazard implies a time series of demand that follows the true historical demand series closely, with a small delay. The variation in sales volumes in NS’s “vary supply” counterfactual mostly comes from variation in demand, even though the sales hazard is fixed.

Including withdrawals in the “vary-supply” counterfactual contributes to the tight fit between the simulated and true data. As can be seen in Figure 1, for example, during the years 2006-2008 new listings were still coming on the market at an elevated pace while sales volumes and sale hazard rates were falling and the housing boom turned to bust. Yet, the reduced-form “vary-supply” counterfactual sales volume takes a downturn at almost the same time as the true data do (see Figure 6). This is possible because as shown in Appendix Figure 10, withdrawal rates rose as the sale hazard fell, growing about 40 percent from 2005 to their peak in 2008. The surge in withdrawals (and an unchanging sale hazard) depletes the counterfactual stock of for-sale listings faster than even the elevated level of new listings could replenish it, causing sales volumes in this simulation to fall as well. The combination of the reduced-form approach (effectively allowing zero supply-side crowd-out) and this conflation of withdrawals and new listings allows the “vary-supply” simulations in this section and NS to fit the true sales data so well.¹⁷

The results of this section highlight the importance of taking market tightness, and its implications for the matching process, into consideration when evaluating the relative roles of demand and supply. The reduced-form results would suggest the supply of new listings, rather than demand for homes, is the most important factor in determining sales volumes. Our full set of results suggest that the opposite is true.

7 Implications

7.1 COVID-19 Housing Boom

Figure 2 shows that during the COVID-19 pandemic, the housing market tightened considerably. After a brief dip at the onset of the pandemic, the sale hazard rate surged to record levels and house price growth also moved up to record highs. In this

¹⁷In addition to their headline counterfactuals, NS do attempt simulations that partial out the effects of withdrawals. However, lacking observations of individual listings, they are forced to assume an elasticity of withdrawals relative to sales hazard. This exercise does weaken the power of listings to explain sales in their paper, but does not eliminate it.

section, we use our model to decompose the tightening of the housing market during the pandemic into supply or demand factors.

A priori, the recent observed tightening in the housing market could be due to reduced supply or increased demand, or both. On the demand side, lower interest rates and widespread telework may have induced more buyers into the market. On the supply side, homeowners could be reluctant to list their home for sale during a pandemic, which could have reduced the for-sale supply. Generous mortgage forbearance programs and the foreclosure moratorium may also have reduced supply. Indeed, new listings plummeted at the onset of the pandemic.

Figure 8 shows counterfactual months supply using our model under (i) fixed demand and true supply and (ii) true demand and fixed supply. When demand or supply is fixed, we set it at average 2019 (pre-pandemic) levels. At the very beginning of the pandemic, the vary-supply simulation drops below the true months' supply while the vary-demand simulation rises above the true months' supply, showing that some of the initial decrease in months' supply is driven by a decrease in new listings. As the pandemic progresses, however, the figure shows that stronger demand overtakes lower supply as the main factor behind the observed decrease in months' supply. By the middle of 2021, the contribution of reduced supply has disappeared and higher demand can explain essentially all of the decrease in months' supply since March 2020. We conclude that, outside of a brief shock at the beginning of the pandemic, reduction of supply was a minor factor relative to increased demand in explaining the tightening of housing markets.

We can also use our model to estimate how much additional supply would be needed to keep house prices on their pre-pandemic trend, given the observed increase in demand. Figure 9 shows counterfactual house prices in which demand (n_t^b) is set at its actual estimated levels, but supply (n_t^s) is set at some multiplier, x , of average 2019 (pre-pandemic) levels. We find that a value of $x = 1.3$ or greater is necessary to bring the counterfactual house price back to its pre-pandemic trend by November 2021. This means that a 30% increase in the monthly number of homes coming on to the market would have been necessary to keep up with the pandemic-era surge in demand. This is a very large increase in supply. Since new construction typically accounts for about 15% of supply, our estimates imply that new construction would have had to increase by roughly 300% to absorb the pandemic-era surge in demand. One implication of this result is that policies targeted at increasing supply, for example

construction subsidies or zoning reforms, would have done little to cool the pandemic house price boom in the short-run.

7.2 Interest Rate Elasticity

This section compares the sensitivity of housing demand and supply to changes in interest rates, which is an important channel through which policy makers can influence the housing market.

We estimate the regression:

$$y_t - y_{t-12} = \alpha_0 + \alpha_1(frmt_t - frmt_{t-12}) + \epsilon_t \quad (9)$$

where y is the housing market variable of interest in month t and frm is the average monthly 30-year fixed mortgage rate in percentage points as reported in the Freddie Mac primary mortgage market survey. We estimate the regressions using our monthly sample between January 2002-November 2021.

Table 2 reports estimates of α_1 for different outcome variables. The first column shows that higher mortgage rates have a strong negative effect on buyer demand, n_t^b . A negative effect is expected because higher mortgage rates increase the cost of owning a home, which should decrease demand all else equal. A one percentage point increase in the mortgage rate is associated with about 9,000 fewer buyers entering the market in our sample counties. Relative to the sample average value of buyer demand, this is a decrease of about 10.4 percent, or a semi-elasticity of 10.4.

Column 2 shows that home sales are also negatively associated with mortgage rates, but the magnitude of the effect is much smaller. The semi-elasticity of home sales to mortgage rates is estimated to be 6, about one-half the estimate of the demand semi-elasticity. Why are home sales much less mortgage rate sensitive than our estimate of buyer demand? One explanation is search frictions. Because it takes time for buyers to transact, home sales today reflect demand from a mix of periods in the past. This mixture effectively smooths the response of home sales to demand shocks, leading to attenuated estimates. A second reason is that, as column 3 shows, there is a small, positive association between mortgage rates and new listings.¹⁸ Higher sup-

¹⁸One potential explanation for a positive association between mortgage rates and new listings is investor activity. Investors may be more inclined to sell homes when interest rates rise and the expected returns from investing in assets other than housing become relatively more attractive.

ply results in higher sales volume all else equal, so the negative relationship between demand and mortgage rates is somewhat offset by the positive relationship between supply and mortgage rates. An implication of these results is that housing demand is much more responsive to mortgage rates than simple regressions based on observables imply.

One potential issue is that changes in mortgage rates could be correlated with unobservables that also influence demand, supply, and sales. To address this endogeneity concern, the final three columns of the table show results where we instrument for the change in mortgage rates using a monthly series of monetary policy surprises estimated in Bu, Rogers and Wu (2021). Contractionary (expansionary) monetary policy tends to raise (lower) mortgage rates, and Bu, Rogers and Wu (2021) develop an estimation procedure that extracts any component of monetary policy that is unrelated to economic fundamentals.¹⁹ For demand and sales, the table shows that the IV estimates of the semi-elasticity are larger than the OLS estimates. Larger IV estimates are expected because increases in mortgage rates are typically associated with an improved economic outlook, which likely increases buyer demand and biases the OLS estimate towards zero. Consistent with the OLS estimates, buyer demand is more rate sensitive than home sales in the IV estimates. The semi-elasticity is 21 for demand and 15 for sales volume, though the demand estimate is somewhat imprecise and is only marginally significant. The semi-elasticity for supply is no longer statistically significant when the IV is used.

8 Conclusion

We use a housing search model to decompose fluctuations in home sales and prices into supply or demand factors. Simulations of the estimated model show that housing demand drives short-run fluctuations in home sales and prices.

For longer-run changes in the housing market, supply may play a much larger role. For example, new supply today also increases supply in the future as today's buyer eventually sells her new home. Our simulations do not account for such a response as we are focused on the short run, but the accumulation of new supply (including new

¹⁹Bu, Rogers and Wu (2021) also show that their monetary policy surprise measure contains no significant central bank information effect. The measure is available from the authors' website for our full sample period.

construction) likely explains more of the variation in sales volume over long horizons. Similarly, long-run levels of house prices may not be as closely connected to market tightness as the short-run price growth we consider in this paper. Understanding the relative importance of supply and demand and other factors for longer-run changes in the housing market remains a topic for future research.

References


- Aladangady, Aditya.** 2017. “Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata.” *American Economic Review*, 107(11): 3415–46.
- Albrecht, James, Pieter A Gautier, and Susan Vroman.** 2016. “Directed search in the housing market.” *Review of Economic Dynamics*, 19: 218–231.
- Andrews, Martyn J, Steve Bradley, Dave Stott, and Richard Upward.** 2013. “Estimating the stock-flow matching model using micro data.” *Journal of the European Economic Association*, 11(5): 1153–1177.
- Anenberg, Elliot, and Daniel Ringo.** 2021. “Housing Market Tightness During COVID-19: Increased Demand or Reduced Supply?” *Federal Reserve Note*.
- Anenberg, Elliot, and Daniel Ringo.** 2022. “The propagation of demand shocks through housing markets.” *American Economic Journal: Macroeconomics*, 14(3): 481–507.
- Anenberg, Elliot, and Patrick Bayer.** 2020. “Endogenous Sources Of Volatility In Housing Markets: The Joint Buyer–Seller Problem.” *International Economic Review*, 61(3): 1195–1228.
- Benmelech, Efraim, Adam Guren, and Brian T Melzer.** forthcoming. “Making the house a home: the stimulative effect of home purchases on consumption and investment.” *Review of Financial Studies*.
- Berger, David, Nicholas Turner, and Eric Zwick.** 2020. “Stimulating housing markets.” *The Journal of Finance*, 75(1): 277–321.

- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra.** 2017. “House Prices and Consumer Spending.” *The Review of Economic Studies*, 85(3): 1502–1542.
- Best, Michael Carlos, and Henrik Jacobsen Kleven.** 2017. “Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK.” *The Review of Economic Studies*, 85(1): 157–193.
- Bhutta, Neil, and Daniel Ringo.** 2020. “The effect of interest rates on home buying: Evidence from a shock to mortgage insurance premiums.” *Journal of Monetary Economics*.
- Bu, Chunya, John Rogers, and Wenbin Wu.** 2021. “A unified measure of Fed monetary policy shocks.” *Journal of Monetary Economics*, 118: 331–349.
- Caplin, Andrew, and John Leahy.** 2011. “Trading frictions and house price dynamics.” *Journal of Money, Credit and Banking*, 43: 283–303.
- Carrillo, Paul E, and Benjamin Williams.** 2019. “The repeat time-on-the-market index.” *Journal of Urban Economics*, 112: 33–49.
- Carrillo, Paul E, Eric R de Wit, and William Larson.** 2015. “Can tightness in the housing market help predict subsequent home price appreciation? Evidence from the United States and the Netherlands.” *Real Estate Economics*, 43(3): 609–651.
- DeFusco, Anthony A, and Andrew Paciorek.** 2017. “The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit.” *American Economic Journal: Economic Policy*, 9(1): 210–40.
- Diaz, Antonia, and Belen Jerez.** 2013. “House Prices, Sales, and Time on the Market: A Search-Theoretic Framework.” *International Economic Review*, 54(3): 837–872.
- Gabrovski, Miroslav, and Victor Ortego-Marti.** 2019. “The cyclical behavior of the Beveridge Curve in the housing market.” *Journal of Economic Theory*, 181: 361–381.

- Gabrovski, Miroslav, and Victor Ortego-Marti.** 2021. “On the Slope of the Beveridge Curve in the Housing Market.” *working paper*.
- Genesove, David, and Lu Han.** 2012. “Search and matching in the housing market.” *Journal of urban economics*, 72(1): 31–45.
- Grindaker, Morten, Artashes Karapetyan, Espen R Moen, and Plamen Nenov.** 2021. “Transaction Sequencing and House Price Pressures.”
- Guren, Adam M.** 2018. “House price momentum and strategic complementarity.” *Journal of Political Economy*, 126(3): 1172–1218.
- Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson.** 2021. “Housing wealth effects: The long view.” *The Review of Economic Studies*, 88(2): 669–707.
- Guren, Adam M, and Timothy J McQuade.** 2020. “How do foreclosures exacerbate housing downturns?” *The Review of Economic Studies*, 87(3): 1331–1364.
- Han, Lu, and William C Strange.** 2015. “The microstructure of housing markets: Search, bargaining, and brokerage.” In *Handbook of regional and urban economics*. Vol. 5, 813–886. Elsevier.
- Head, Allen, Huw Lloyd-Ellis, and Hongfei Sun.** 2014. “Search, liquidity, and the dynamics of house prices and construction.” *American Economic Review*, 104(4): 1172–1210.
- Hedlund, Aaron.** 2016. “Illiquidity and its discontents: Trading delays and foreclosures in the housing market.” *Journal of Monetary Economics*, 83: 1–13.
- Karahan, Fatih, and Serena Rhee.** 2019. “Geographic reallocation and unemployment during the Great Recession: The role of the housing bust.” *Journal of Economic Dynamics and Control*, 100: 47 – 69.
- Krainer, John.** 2001. “A Theory of Liquidity in Residential Real Estate Markets.” *Journal of Urban Economics*, 49(1): 32–53.
- Mian, Atif, and Amir Sufi.** 2011. “House prices, home equity-based borrowing, and the US household leverage crisis.” *American Economic Review*, 101(5): 2132–56.


- Mian, Atif, Kamalesh Rao, and Amir Sufi.** 2013. “Household Balance Sheets, Consumption, and the Economic Slump*.” *The Quarterly Journal of Economics*, 128(4): 1687–1726.
- National Association of Realtors.** 2019. “2019 Profile of Home Buyers and Sellers.”
- Ngai, L Rachel, and Kevin D Sheedy.** 2020. “The decision to move house and aggregate housing-market dynamics.” *Journal of the European Economic Association*, 18(5): 2487–2531.
- Ngai, L Rachel, and Silvana Tenreyro.** 2014. “Hot and cold seasons in the housing market.” *American Economic Review*, 104(12): 3991–4026.
- Novy-Marx, Robert.** 2009. “Hot and cold markets.” *Real Estate Economics*, 37(1): 1–22.
- Ortalo-Magne, Francois, and Sven Rady.** 2006. “Housing market dynamics: On the contribution of income shocks and credit constraints.” *The Review of Economic Studies*, 73(2): 459–485.
- Petrongolo, Barbara, and Christopher A Pissarides.** 2001. “Looking into the black box: A survey of the matching function.” *Journal of Economic literature*, 39(2): 390–431.
- Piazzesi, Monika, and Martin Schneider.** 2009. “Momentum Traders in the Housing Market: Survey Evidence and a Search Model.” *American Economic Review*, 99(2): 406–11.
- Ringo, Daniel.** 2022. “Monetary Policy and Home Buying Inequality.” *Available at SSRN 4220206*.
- Smith, Eric.** 2020. “High and low activity spells in housing markets.” *Review of Economic Dynamics*, 36: 1–28.

Figure 1: Annual sales volume, new listings, inventory, and sale hazard rate




Notes: All series are indexed to 2002 values. The sale hazard rate is calculated as the number of sales contracted each month divided by the number of homes actively listed for sale at some point in the month. The annual sale hazard is the average of the monthly sale hazards, weighted by the number of homes listed for sale each month.

Figure 2: Months supply and house price growth




Notes: Months supply is equal to the inverse of the sale hazard rate. House price index shows the estimates of the time dummies from a regression of log house prices on time dummies, home characteristics, and zipcode fixed effects. House price index is adjusted for inflation using the consumer price index excluding shelter.

Figure 3: Sales Volumes, Observed and Counterfactual




Notes: “True” is the actual sales volume in the data. “Vary supply” is the counterfactual sales volume according to our model when demand is held fixed at its sample mean, but supply varies as in the data. “Vary demand” is the counterfactual sales volume according to our model when supply is held fixed at its sample mean, but demand varies as in our estimates.

Figure 4: Sales Volumes, Alternative Crowd-Out




Notes: Shows simulated sales volume for an alternative calibration of the elasticity of the matching function: $\eta = 0.16$. “True” is the actual sales volume in the data. “Vary supply” is the counterfactual sales volume according to this model when demand is held fixed at its sample mean, but supply varies as in the data. “Vary demand” is the counterfactual sales volume according to this model when supply is held fixed at its sample mean, but demand varies as in our estimates.

Figure 5: Months' Supply, Observed and Counterfactual




Notes: Months' supply is equal to the inverse of the monthly sale hazard rate. "True" is the actual months supply in the data. "Vary supply" is the counterfactual months supply according to our model when demand is held fixed at its sample mean, but supply varies as in the data. "Vary demand" is the counterfactual months supply according to our model when supply is held fixed at its sample mean, but demand varies as in our estimates.

Figure 6: Sales Volume, Reduced Form




Notes: “True” is the actual sales volume in the data. “Vary supply” is the counterfactual sales volume when the sale hazard is held fixed at its sample mean, but supply varies as in the data. “Vary sale hazard” is the counterfactual sales volume when supply is held fixed at its sample mean, but the sale hazard varies as in the data.

Figure 7: Estimated Number of Active Buyers




Note: “From Observed Sales” shows the estimated number of active buyers implied by our model and the observed time series of listings and sales. “From Counterfactual Sales (Fixed Sale Hazard)” shows the number of active buyers implied by our model using a counterfactual sales volume series generated when the sale hazard is held fixed at its sample mean, but supply varies as in the data.

Figure 8: Months' Supply during COVID-19, observed and counterfactual



Notes: Months' supply is equal to the inverse of the monthly sale hazard rate. "True" is the actual months supply in the data. "Vary supply" is the counterfactual months' supply according to our model when demand is held fixed at pre-pandemic (2019) levels, but supply varies as in the data. "Vary demand" is the counterfactual months' supply according to our model when supply is held fixed at pre-pandemic (2019) levels, but demand varies as in the data.

Figure 9: Real house price during COVID-19, observed and counterfactual



Notes: Shows log real house price under counterfactual supply. “+X” is a counterfactual where supply is set at a multiplier, X, of pre-pandemic (2019) supply levels for each month from March 2020 onward.

Table 1: County-level Growth in Sales, New Listings and Sale Hazard

	Sales Growth	Sales Growth	Sales Growth	Sales Growth	Sales Growth	Sales Growth
New Listings Growth	0.346*** (0.00496)		0.330*** (0.00499)		0.244*** (0.00511)	
Sale Hazard Growth		0.732*** (0.00209)		0.731*** (0.00208)		0.835*** (0.00209)
County FE			X	X	X	X
Month-Year FE					X	X
R-squared	0.0827	0.695	0.0955	0.704	0.192	0.788
N	54026	54026	54026	54026	54026	54026

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: OLS regression results of the 12-month growth in sales volume on the 12-month growth in new listings and the sale hazard rate. Each variable is measured at the county-month level and each regression pools observations across counties.

Table 2: Mortgage Rate Elasticity

	OLS			IV		
	Demand	Supply	Sales	Demand	Supply	Sales
30-yr Fixed Mortgage Rate	-8856.7*** (2185.4)	2837.7* (1599.9)	-4292.2*** (1450.0)	-18609.4* (10674.3)	6531.6 (9069.8)	-10442.6** (4698.9)
Observations	226	226	226	226	226	226

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: All variables are in 12-month changes. Demand is housing demand, as implied by the housing search model. Supply is new listings and Sales is sales volume. The latter three columns show results when we instrument for the change in mortgage rates with a monthly series of monetary policy surprises estimated in Bu, Rogers and Wu (2021). Newey-West standard errors with optimal lag-selection algorithm are shown.

A Sample restrictions

There are 95 million listings in our MLS data over our sample period and these listings are associated with 3,067 unique counties. We drop counties from our estimation sample for one of two reasons. First, not all counties in our sample record the contract date if a sale occurs. The contract date measures when a buyer and seller agree on a sales price, and is a better measure of when a property sells than the sale closing date. The sale closing date, which is always recorded in our data for sales and is part of the public record, measures when property ownership is transferred from the seller to the buyer. The lag between sale agreement and sale closing varies across sales largely due to idiosyncratic factors, such as the buyer move date preference or the processing time for the mortgage lender, and not due to housing market tightness. We therefore drop counties for which the sale contract date is missing for at least two percent of sales in any year.²⁰

Second, the coverage of the MLS data increases over time for certain geographic areas. We drop counties for which there appears to be large changes in coverage during our sample period. We drop counties where (i) the annual change in sales for any year is greater than 75% (either positive or negative) of the previous year's sales or (ii) the ratio of total sales between 2014 and 2019 is over 2.5 times as large as total sales between 2000 and 2005. We also drop counties where (iii) the number of sales over the full sample period is less than 400 or (iv) in any year, the number of sales is less than 10.

After our sample restrictions, we are left with 263 counties and 38 million listings.

B Search-and-Matching Function

Our baseline parameterization of the search and matching function is Cobb-Douglas with constant returns to scale. Although this matching function does not have a clear micro-foundation, its advantage is that it has a free parameter, η , that dictates

²⁰For the small number of remaining sales with missing sales contract date, we set the sales contract date equal to the sales closing date minus 35 days, which is the average delay between sale agreement and sale closing in our data.

the elasticity of the probability of sale with respect to market tightness. This is an advantage because as we discuss in Section 4, there are a few consistent and credible external estimates of this parameter that we can calibrate to. We can also test robustness to alternative values of the parameter. In the next three subsections, we discuss alternatives to our baseline, constant returns to scale Cobb-Douglas matching function.

B.1 Urn-ball matching function

An alternative search-and-matching function to consider is the urn-ball matching function, where

$$q^s(\theta) = \theta q^b(\theta) = 1 - \exp(-A\theta)$$

This matching function has only one instead of two free parameters as in Cobb-Douglas, but its advantage is that it has a micro-foundation (see Petrongolo and Pissarides (2001)). Each period, sellers post their vacancies, and each buyer randomly visits a seller. Buyers do not coordinate on their visits. One interpretation of A is that it measures the fraction of buyers who are suitable matches for a randomly selected home for sale. A transaction occurs when a seller is matched with at least one suitable buyer. For a large number of sellers, $1 - \exp(-A\theta)$ is a good approximation to the sale hazard rate. Figures 11 and 12 show our results are little changed when we use the urn-ball matching function instead of Cobb-Douglas.

B.2 Stock-flow model

In this section, we consider an extension of our housing search model where new listings are potentially more efficient searchers than old listings. This model can account for the fact that in the data, the sale hazard rate for new listings is much higher than the sale hazard rate for old listings. The model builds on Andrews et al. (2013). The notation is the same as in the model presented in the main text with one exception. We now use s_t to denote the number of old listings on the market in period t . We continue to denote the number of new listings by n^s . New listings become old listings if they do not sell (and are not withdrawn from the market) in the first period after listing. In this model, we modify the matching function to be

$$M_t = A_t b_t^\eta (s_t + a_t n_t^s)^{1-\eta} \quad (10)$$

where M denotes the number of transactions and a_t denotes the search efficiency of new listings relative to old listings. Since we include the aggregate efficiency term A_t , we can normalize the search efficiency on old listings to one. Assuming that matches are divided between new listings and old listings with weights equal to s_t and $a_t n_t^s$ respectively, we can express the sale probability for new listings and old listings as

$$q_{old}^s = \frac{q_{new}^s}{a_t} = A_t \left(\frac{b_t}{s_t + a_t n_t^s} \right)^\eta \quad (11)$$

The probability of buying is simply

$$q^b = \frac{M_t}{b_t} = A_t \left(\frac{b_t}{s_t + a_t n_t^s} \right)^{\eta-1} \quad (12)$$

The stock of old sellers evolves as

$$s_{t+1} = s_t - s_t q_{old}^s - s_t (1 - q_{old}^s) w_{old}^s + n_t^s (1 - q_{new}^s) (1 - w_{new}^s) \quad (13)$$

where w_{new}^s and w_{old}^s are the withdrawal rates for new and old sellers, respectively. The stock of old sellers next period equals the stock this period minus the outflow (old sellers who sell or withdraw) plus the inflow (new sellers who do not sell and do not withdraw).

The stock of buyers evolves as:

$$b_{t+1} = b_t - b_t q^b - b_t (1 - q^b) w^b + n_{t+1}^b \quad (14)$$

We estimate the stock-flow model using the same approach described in Section 4. Sale and withdrawal hazards are estimated separately for new and old listings. The parameters that are new to this model, a_t , can simply be estimated by taking the ratio of the estimated sale hazards for new and old listings, as shown in equation 11.

As in the baseline model, we calibrate $\eta = 0.84$ and we calibrate \bar{A} to match survey data from the NAR on average search time for buyers in 2019. We calibrate $w_{new}^s = 0.052$ and $w_{old}^s = 0.113$ to match the average monthly withdrawal hazard for new and old listings in our MLS data, respectively. We set $w^b = 0.08$, which is the

average seller withdrawal hazard across both new and old listings.

Consistent with our results from the baseline model, Figure 13 shows that demand explains essentially all of the variation in sales volume in the stock-flow search model. The relative role of demand and supply is also comparable to our baseline model for months supply, as shown in Figure 14.

B.3 Increasing returns to scale

It is common in the housing search literature to model the matching function as constant returns to scale. An exception is Ngai and Tenreyro (2014), who use a matching model with increasing returns to scale (IRS) to explain seasonality in the housing market. In their model, there is a “thick market effect”: when there are more homes available for sale, a buyer is more likely to find a match. Market tightness plays no role in their model and the probability of a match is increasing in the stock of houses available for sale (s using our model notation).

While thick market effects may exist, they are unlikely to be a first-order consideration for the matching technology because as we discussed earlier, Figure 1 shows a strong negative correlation between the stock of houses available for sale and sales volume.²¹ As a robustness, we consider simulations for the IRS matching technology:

$$M = Ab^\eta s^\alpha \tag{15}$$

for $\eta = 0.84$, which is our baseline calibration of the exponent for the stock of buyers, and $\alpha = 0.30$, which is an arbitrary value that is greater than 0.16 so that the matching function exhibits IRS. Results with IRS are similar to our main results.

C Construction of house price index and expected house price

To construct a house price index, we estimate the hedonic regression

$$p_{it} = \delta_t + \beta X_{it} + \epsilon \tag{16}$$

²¹Looking specifically at seasonality, Figure 20 shows that there is a positive correlation between seasonality in inventory and sales, but it is far from one.

where p_{it} is the log transaction price of house i that goes under contract in month-year t , δ_t is a set of month-year dummy variables for the contract date, and X is a vector of house characteristics including: home age, its square, and its cube; number of bedrooms and its square; number of bathrooms; the ratio of bedrooms to bathrooms; a dummy for new construction; a dummy for property type; and zipcode fixed effects.²² We drop transactions associated with sale prices below \$10k or above \$10 million. The estimates of the month-year dummy variables are used as the quality-adjusted nominal house price index. To get the real house price index, we deflate using the CPI-less-shelter index from the BLS.

The expected log house price for home i at time t is simply the predicted value from equation 16, \hat{p}_{it} . The expected house price can be computed for any home listed for sale, not just homes that ultimately transact. $p_i^L - \hat{p}_{it}$ is the list price premium, which is included in the variables that affect the sale hazard in equation 5, X^A . Our MLS data only record the initial list price and the final list price associated with each listing, and not the time of any list price changes. We set p_i^L equal to the original log list price associated with listing i . For the stock-flow model described in Appendix Section B, for new listings, we use the original log list price as p_i^L , and for old listings, we use the final log list price as p_i^L .

D County-level findings

Using location data from individual listings, we construct a county-level panel of listings and sales and use these data to back out the implied number of buyers in each county and month. We then simulate counterfactual time series of sales and inventory in each county, varying only either demand or supply as described in Section 5. To determine how well supply and demand individually explain housing market dynamics in the cross section, we compare counterfactual growth in sales and months' supply to the true data over the years 2003-2005, when the U.S. housing market was experiencing a boom with a large amount of geographical variation.

County level sales growth is measured as the percentage difference in total annual sales between 2003 and 2005. The counterfactual values of this sales growth under the vary-demand and vary-supply simulations are shown in a scatter plot against the true

²²We do not use the home's square footage or lot size because these characteristics are frequently missing in our data.

values of sales growth on the horizontal axis in Figure 17. The vary-demand simulated observations are tightly clustered around the 45 degree line, suggesting cross-sectional variation in demand does a very good job explaining differences across counties in sales growth during the mid-2000’s housing boom. The vary-supply simulated observations show no particular correlation with true sales growth, however. Variation in supply cannot explain different counties’ experiences over this time period.

Just as with the aggregate time series, we may be concerned that this result is the consequence of a modeling choice, rather than an implication of the underlying data. We therefore perform a similar exercise as in Section 5.1.1, and rerun the simulations calibrating $\eta = 0.16$. This calibration causes more crowd out on the demand side than the supply side, potentially allowing supply to play greater role in driving sales volumes.

Results from counterfactual simulations under the alternative calibration are shown in Figure 18. A low value of η weakens but does not eliminate the fit of the vary-demand counterfactual observations to the true data, reducing the correlation between the simulated and true data to 0.43. It also allows for supply to drive more variation in simulated sales volumes, evident from the greater vertical spread in the vary-supply counterfactual observations relative to Figure 17. However, the fit of the vary-supply simulations to the true data is not any improved—the correlation is essentially zero. Just as we found from the aggregate time series, supply’s inability to explain cross-sectional variation in sales growth is driven by the data rather, not by our choice of calibration.

In Figure 19 we show counterfactual simulations of months’ supply relative to the true data. We find that supply does a better job explaining months’ supply than it does sales volumes. Across counties, the vary-demand and vary-supply simulated observations have a similar correlation to the true data.

E House prices and market tightness

In this section, we describe a theoretical justification for the tight empirical relationship we find between market tightness and house price growth. For simplicity of exposition, we only consider demand-side shocks.

Time is discrete. Let there be a measure 1 of *Sellers*, who each live a single period. They list a home for sale and choose a take-it-or-leave-it asking price, P . For

simplicity, assume P can take only one of two values: High ($P = H$) or Low ($P = L$), where $H > L$. A Seller's utility is the expected proceeds from a sale, i.e. the asking price times the probability a sale occurs.

Buyers search the market for homes listed for sale. Like Sellers, let Buyers live for a single period. There are two types of Buyers: High Demand and Low Demand. Assume there is a measure 1 of both types of Buyers. High Demand Buyers always have a higher willingness to pay to buy a home than Low Demand Buyers.

Buyers' willingness to pay is determined by their type and by the state of the world, which can be *Good* or *Bad*. In the Good state, High Demand Buyers are willing to pay H and Low Demand Buyers are willing to pay L for a home. In the Bad state, High Demand Buyers are willing to pay L and Low Demand Buyers are not willing to pay anything. Each period, the state of the world stays the same as the last period with probability S , and changes (from Bad to Good, or from Good to Bad) with probability $1 - S$.

Sellers cannot observe the current state of the world—they do not know precisely how in-demand their home is. However, they can observe prices and sale activity from the previous period, and, inferring what the state of the world was at that time, form an expectation about its current state. Using this information, sellers set an asking price to maximize expected utility. Buyers then observe their own willingness to pay and choose to either:

1. Search among the high-priced listings (if any are on offer);
2. Search among the low-price listings (if any are on offer); or,
3. Abstain from the market.

Within each price tier P , the sale hazard is a function q of market tightness, θ_P , the ratio of Buyers searching in price P to sellers asking for P . Following the matching function of Section 3, assume this sale hazard takes the form:

$$q(\theta_P) = A\theta_P^\eta \tag{17}$$

Each match results in a transaction. The utility from these transactions is realized, and the next period begins with a new crop of Buyers and Sellers.

Under certain parameterizations, the following strategy is a pooling equilibrium for every seller: set $P = H$ if the previous period state was Good, and set $P = L$

if the previous period state was Bad. In this equilibrium, all sellers choose the same price in a given period and so the directed search model collapses to random search, as there is only one active price tier for Buyers to search in at any time.

To see that all Sellers following this strategy is an equilibrium, consider first the utility a Seller derives from this strategy when the previous state was Good. The Seller sets $P = H$, as do all the other sellers.

- With probability S , the current state is Good. High Demand Buyers, willing to pay H , search in this price tier as there are no homes priced at L for sale. With all homes priced above their willingness to pay, Low Demand Buyers abstain from the market. $\theta_{H,Good} = 1$, so $q(\theta_{H,Good}) = A$.
- With probability $1 - S$, the current state is Bad. No Buyers are willing to pay H , so everyone abstains. $\theta_{H,Bad} = 0$, so $q(\theta_{H,Bad}) = 0$ and no sales occur.

The Seller's expected utility from following this strategy is the product of the probability of each state times the expected utility in that state, summed over both states, or

$$U = Sq(\theta_{H,Good})H + (1 - S)q(\theta_{H,Bad})H = SAH \quad (18)$$

Now consider a Seller deviating from this strategy, and setting $P = L$ if the previous state was good. They could increase their sale probability, depending on Buyer strategy, but the expected utility cannot rise above L . Therefore, every Seller setting $P = H$ when the previous state was Good is an equilibrium if:

$$L < SAH \quad (19)$$

Next, consider the utility a Seller derives from following the strategy when the previous state was Bad. The Seller sets $P = L$, as do all the other sellers.

- With probability S , the current state is Bad. High Demand Buyers, willing to pay L , search in this price tier where all the homes for sale are. With all homes priced above their willingness to pay, Low Demand Buyers abstain from the market. $\theta_{L,Bad} = 1$, so $q(\theta_{L,Bad}) = A$.
- With probability $1 - S$, the current state is Good. Both High Demand and Low Demand Buyers are willing to pay L , so everyone searches in the low price tier. $\theta_{L,Good} = 2$, so $q(\theta_{L,Good}) = A2^n$.

The Seller's expected utility from following this strategy is

$$U = Sq(\theta_{L,Bad})L + (1 - S)q(\theta_{L,Good})L = SAL + (1 - S)A2^nL \quad (20)$$

Now consider a Seller deviating from this strategy and setting $P = H$ when the previous state was Bad. They could increase their reward if the current state was Good, depending on Buyer strategy, but their reward in the Bad state is certain to be zero (as no Buyer is willing to pay H in the bad state). Their expected utility cannot rise above $(1 - S)H$. Therefore, every seller setting $P = L$ when the previous state was Bad is an equilibrium if:

$$(1 - S)H < SAL + (1 - S)A2^nL \quad (21)$$


For $AH > L$, the inequalities 19 and 21 can be satisfied, and this pooling equilibrium can arise, as long as S is sufficiently close to 1. That is, when recent housing market demand is sufficiently informative about the current expected level of demand, Sellers will price high when demand was high and low when demand was low.

The state of the world in the previous period ($t - 1$) can be inferred from observing its market tightness, θ_{t-1} , and prices, P_{t-1} . Under the equilibrium strategy, there are four possible sets of observed values:

1. $\theta_{t-1} = 0$ and $P_{t-1} = H$. Sellers had set prices high but the world ended up in the Bad state, so no Buyers entered. As the previous state was Bad, Sellers will set $P_t = L$. Price growth from $t - 1$ to t is therefore negative, $L - H$.
2. $\theta_{t-1} = 1$ and $P_{t-1} = L$. Sellers had set prices low and the world ended up in the Bad state, so only High Demand Buyers entered. As the previous state was Bad, Sellers will set $P_t = L$. Price growth from $t - 1$ to t is therefore zero, $L - L$.
3. $\theta_{t-1} = 1$ and $P_{t-1} = H$. Sellers had set prices high and the world ended up in the Good state, so only High Demand Buyers entered. As the previous state was Good, Sellers will set $P_t = H$. Price growth from $t - 1$ to t is therefore zero, $H - H$.
4. $\theta_{t-1} = 2$ and $P_{t-1} = L$. Sellers had set prices low but the world ended up in the Good state, so both High Demand and Low Demand Buyers entered. As the previous state was Good, Sellers set $P_t = H$. Price growth from $t - 1$ to t is therefore positive, $H - L$.

When market tightness is low ($\theta_{t-1} = 0$) price growth is low ($L - H$). When market tightness is moderate ($\theta_{t-1} = 1$) price growth is moderate (0). When market tightness is high ($\theta_{t-1} = 2$) price growth is high ($H - L$). Therefore, price growth is perfectly correlated with, and a function of, market tightness. Intuitively, very high or low market tightness represent an imbalance between supply and demand at a given price point. Subsequent Sellers observe this imbalance, and adjust asking prices up or down to bring demand into alignment with supply. This adjustment does not happen instantaneously because Sellers do not observe shocks to demand directly: they must first observe the market tighten or loosen and then react.

Figure 10: Withdrawal Hazard



Notes: Shows the average monthly probability, by year, that a listing is withdrawn from the market without sale.

Figure 11: Sales Volumes, Observed and Counterfactual from Urn-ball Model




Figure 12: Months' Supply, Observed and Counterfactual from Urn-ball Model




Figure 13: Sales Volumes, Observed and Counterfactual from Stock-Flow Model




Figure 14: Months' Supply, Observed and Counterfactual from Stock-Flow Model




Figure 15: Sales Volumes, Observed and Counterfactual from IRS Model




Figure 16: Months' Supply, Observed and Counterfactual from IRS Model




Figure 17: Sales Volumes Counterfactuals across Counties




Notes: Shows simulated sales volume growth from 2003 to 2005, plotted against actual sales volume growth over this period. Observations represent simulations for individual counties. “Vary supply” is the counterfactual sales volume according to our model when demand is held fixed at its sample mean, but supply varies as in the data. “Vary demand” is the counterfactual sales volume according to our model when supply is held fixed at its sample mean, but demand varies as in our estimates. “True” represents the 45 degree line that a perfect simulation of the true underlying data would lie along.

Figure 18: Sales Volumes Counterfactuals across Counties, Alternative Crowd-Out



Notes: Shows simulated sales volume growth from 2003 to 2005 for an alternative calibration of the elasticity of the matching function: $\eta = 0.16$, plotted against actual sales volume growth over this period. Observations represent simulations for individual counties. “Vary supply” is the counterfactual sales volume according to this model when demand is held fixed at its sample mean, but supply varies as in the data. “Vary demand” is the counterfactual sales volume according to this model when supply is held fixed at its sample mean, but demand varies as in our estimates. “True” represents the 45 degree line that a perfect simulation of the true underlying data would lie along.

Figure 19: Months' Supply Counterfactuals across Counties



Notes: Shows simulated growth in the months' supply of homes for sale from 2003 to 2005, plotted against actual months' supply growth over this period. Observations represent simulations for individual counties. "Vary supply" is the counterfactual sales volume according to our model when demand is held fixed at its sample mean, but supply varies as in the data. "Vary demand" is the counterfactual sales volume according to our model when supply is held fixed at its sample mean, but demand varies as in our estimates. "True" represents the 45 degree line that a perfect simulation of the true underlying data would lie along.

Figure 20: Seasonality in sales, new listings, inventory, and sale hazard rate

