Online Appendix to "The Propagation of Demand Shocks Through Housing Markets"

Elliot Anenberg[*] Daniel Ringo[†]

February 11, 2021

A Details of Matching on Buyer and Seller Name

Each property transaction records a first name and a last name field for up to two buyers (or current owners, if the listed transaction is a refinancing). The first name field often contains a middle name or middle initial. We refer to the most recent names listed on a transaction for a property prior to 2014 as the sellers. Names listed as purchasers of properties in 2014 and 2015 are buyers. Names are listed in the order they appear on the deed.

We first search for all potential buyers that match with (i.e., are potentially the same household as) each seller with a home listed on the MLS sometime in 2014 or 2015. Matches are restricted to occur within a six month window around the period the seller's home was listed for sale. As a first step, we require that the last name of the first listed buyer (buyer 1) be an exact match to the last name of the first listed seller (seller 1). We also require that the new home have a different address than the seller's current home.

We then calculate the Jaro-Winkler distance between the first names of seller 1 and buyer 1. Matches with a distance greater than 0.1 are dropped.

[*]Board of Governors of the Federal Reserve System. Email: elliot.anenberg@frb.gov.
[†]Board of Governors of the Federal Reserve System. Email: daniel.r.ringo@frb.gov.

This fuzzy matching criteria is introduced to allow for nicknames, omitted middle names and typos.

To choose between the remaining possible matches, we then turn to the second listed names (seller 2 and buyer 2). If the Jaro-Winkler distance between the first name of seller 2 and buyer 2 is less than 0.1, then the closest match is kept. Last names of seller and buyer 2 are ignored, as they may change due to marriage and they generally match the last name of seller and buyer 1, respectively.[1] Cases in which seller 2 does not match to buyer 2 are dropped in favor of cases in which no seller 2 or buyer 2 is listed.

To break further ties, the matches in which the purchase date lies closest to the time period when the seller's home was listed on the MLS are kept.

A.1 Assessment of Match Quality

Using this procedure, we can link about 45 percent of households in the listing data who successfully sold their home to another purchase around the same time. This match rate is similar to those found by Anenberg and Bayer (2020) and DeFusco, Nathanson and Zwick (2017). One possible concern is false negatives; that is, does this match rate imply a too-low probability of home buying following a sale? To determine if the match rate is reasonable, we compare this implied probability of purchasing another house around the same time as selling a current one to data from the Panel Survey of Income Dynamics (PSID). From 2011 through 2015, approximately 50 percent of households in the PSID that sold a piece of real estate property in the two years between surveys bought one as well during the same period. This figure includes primary residences but excludes farmland.

One significant difference between our data and the PSID is that the PSID samples households, while our data samples properties. Investors who own multiple properties are thus represented in a greater fraction of our observa-

[1] Inspecting the data, it appears that a male name is listed first and a female name second in the vast majority of cases in which two, recognizably gendered names appear. It also appears that the listed order of names tends to be consistent within couples across transactions - we get very few additional matches when we repeat our matching procedure, attempting to match seller 1 to buyer 2.

tions than in the PSID. In fact, about 10 percent of listed homes for sale in our data have an owner with no listed last name, or a name that contains the strings "TRUST" or "LLC". These homes are not owner-occupied, so their sale doesn't have to coincide with the owner finding another place to live (and hence the purchase of another house). There are likely additional investors who own multiple properties in their own name as well. Given the number of non-owner occupied houses, we think the slightly lower purchase rate in our data relative to the PSID is reasonable.

A further concern is the possibility of false positive matches. Home sellers with common names in particular may be spuriously identified as having purchased another home, due to being matched with a different buyer of the same name. However, having a non-unique name will not necessarily produce a false positive match. A different person with the same name would have had to coincidentally purchase a home within the six month window of the home sale to potentially produce a false positive. Nonetheless, to make sure that our results are not driven by false positive matches, below we show robustness of our results to restricting the estimation sample to the 75 percent of sellers in who have a name that is unique within our sample, and who should therefore be much less likely to generate a false match.

B Robustness, Validity Checks, and Further Results

B.1 Testing for Direct Effects on Current Owner Purchases

Our identifying assumption is that any difference between our treatment and control groups following the MIP cut is due to the change in the relative demand for their homes, rather than a direct effect of the lower premiums on the owners' purchase decisions. A direct effect of the MIP cut on incumbent owner purchase behavior, correlated with Z but not mediated through the demand for their current home, would bias our estimates of the strength of the propogation mechanism and resulting multiplier effects. In this section we show the results of several different tests of a direct effect on incumbent owner purchases.

First, we rerun our estimator on the subsample of the population that did not use an FHA loan for their previous home purchase or refinance. The CoreLogic deeds data we use for our main estimation sample also includes records for whether the property was refinanced or purchased with a mortgage, and, if so, whether the mortgage carried FHA insurance. Homeowners who previously needed an FHA loan will be more likely to use FHA insurance on subsequent purchases, and would thus be the group most plausibly sensitive to a direct effect of the MIP cut. As can be seen in Figure 1, these individuals are (unsurprisingly) more common in neighborhoods and prices ranges with higher values of Z.

Dropping previous FHA users from the sample, we reestimate equation 2 of the main manuscript using the same 2SLS strategy, and present results in column 2 of Table 1. Results for the full sample are presented in column 1 for comparison. As can be seen, the estimate of β_1 is similar for both samples considering the standard errors, suggesting that the *total* effect of the MIP cut is strong even in the population for which we would not expect a *direct* effect. Next, we further restrict our sample to homeowners that had not used any mortgage to purchase their current home - that is, cash buyers. Previous cash buyers were very unlikely to use FHA insurance on subsequent purchases (only about 3 percent of our sample went from buying with cash to using an FHA loan) and so should be essentially free of any direct effect of the MIP cut. Results from estimating equation 2 on this subsample are presented in column 3 of Table 1. Again there is evidence of a strong total effect of the MIP cut in a population that should have no direct response. The point estimate is actually higher in this subsample than in the full estimation sample (although the confidence interval is wider).

For our next tests of direct effects, we continue to exploit the fact that among current owners, only certain subsets could realistically show a direct response to a cut in the FHA's premiums. First, owners who do not intend to use a mortgage for their next purchase (cash buyers) are not directly influenced by the price of a particular form of mortgage credit. Second, mortgage borrowers who put down a down payment of 20 percent or more, or who have a high credit score, have lower cost options than FHA insurance. The pricing of FHA insurance should not influence these owners' decisions to buy either. Any direct effect of the MIP cut on the purchase probabilities of current owners (that is stronger in neighborhoods with high *Z)* should therefore appear as a relative increase in the share of purchases by current owners who make use of a mortgage, and who have a low credit score and high LTV ratio, coming from those neighborhoods.

To test for such effects, we make use of additional data from both CoreLogic and McDash Analytics. The CoreLogic deeds data records the amount of any mortgage used, while McDash, which records servicing data for over half of all mortgage originations in the US, provides FICO scores and LTV ratios at origination. We match the McDash data to the deeds by loan amount and purchase price (rounded to the nearest $1,000), month of origination, ZIP code, and indicators for FHA and VA status. We then rerun versions of equation 1, estimating the reduced form effect of the instrument on the probability a home purchase by a current owner makes use of a mortgage (limiting the sample to months with a successful purchase), and conditional on using a mortgage, on the probability the purchaser has a low FICO score and high LTV ratio.

For purposes of comparison, we also estimate the direct effect of the instrument on current owners' monthly purchase probabilities. Results are presented in Table 2. As can be seen in column 1, the reduced form effect of the instrument on purchase probability is a statistically significant 0.002. With a baseline monthly purchase probability of 0.033, this means switching the instrument from zero to one increases the number of current owners who purchase a home each month by over 6 percent. If these purchases were *directly* caused by the MIP cut, we would expect to see the number of owners using a mortgage to buy a home (relative to cash buyers) to increase by a similar amount, in particular the number of mortgage borrowers with low FICO scores and high LTV ratios.

In column 2 of Table 2 we show the estimated reduced form effect of the instrument on the share of homeowners who used a mortgage to purchase their next house. The estimate is not significantly different from zero, and is actually negative. Purchases by current owners using cash were at least as responsive to the MIP cut as purchases making use of a mortgage, suggesting the correlation between any direct effect and Z was negligible relative to the indirect effect. In column 3 we show the estimated reduced form effect of the instrument on the share of low FICO, high LTV ratio borrowers among homeowners using a mortgage to purchase their next house. Although this point estimate is positive, it is not statistically significantly different from zero and its magnitude is too small to explain more than a fraction of the 6 percent increase in purchases caused by the instrument.

The confidence bands around the estimates in Table 2 are tight enough to allow us to put meaningful bounds on the magnitude of any direct effect as well. As mentioned above, we estimate that the effect of changing Z from zero to 1 increases subsequent purchase probability by over 6 percent. In 2014, 75 percent of incumbent owners used a mortgage on their next purchase. Suppose the 6 percent reduced form effect of the instrument was due entirely to direct effects. Then, for every 100 inframarginal purchases in a given month by incumbent owners from neighborhoods with $Z = 1$, 6 marginal purchases would be induced by the MIP cut, all of which must be mortgage borrowers. The fraction of incumbent purchasers using a mortgage would rise from $\frac{75}{100}$ to $\frac{75+6}{100+6}$, an increase of 1.4 percentage points. From column 2 of Table 2, we can calculate that the upper bound of the 95 percent confidence interval of the estimated effect on the fraction using a mortgage is less than 0.6 percentage points, less than half of what would be necessary to explain the increase in incumbent purchases.

Similar findings apply to the other test. In 2014, 9 percent of incumbent purchasers who used a mortgage belonged to the low FICO, high LTV group identified in Bhutta and Ringo (2020) as being responsive to the MIP cut. If the entire reduced form effect of Z on purchases happened through direct effects, the fraction of incumbent purchase borrowers from the responsive group would rise from $\frac{9\times 0.75}{100\times 0.75}$ to $\frac{9\times 0.75+6}{100\times 0.75+6}$, or 6.7 percentage points, adjusting for the fact that only 75 percent of incumbent purchasers used a mortgage. The upper bound of the 95 percent confidence interval calculated from column 3 in Table 2 is only 2.5 percentage points.

All of these tests reject the possibility that direct effects can explain more than a minority of the effect of Z on incumbent owner purchases. The point estimates suggest that the most likely magnitude of the correlation between Z and any direct effect is negligible or zero. Overall, we do not find any compelling evidence that the instrument affected the purchase probability of current homeowners except through a demand effect for their current homes.

B.2 Restricting Estimation Sample to Unique Names

Our matching procedure identifies sellers as having purchased another home if we can find a home buyer with the same name as them in a certain time window somewhere in the United States. Some names are quite common, however, so this procedure runs the risk of producing false positive matches. However, in our sample, approximately 75 percent of sellers have a unique combination of first and last name for the first individual listed on the property. While this certainly doesn't guarantee that these names are globally unique, this subset should be much less susceptible to the false positive problem.

As a test for whether false positive matches are biasing our results, we rerun the estimator on the subsample with unique names. Results are presented in Table 3. Results are quite similar to the main estimation sample. This test suggest false positive matches are not materially biasing our main estimates.

B.3 Robustness to the Inclusion of Control Variables

Our main results are robust to the inclusion of a wide range of detailed control variables. These include census tract and month fixed effects, as well as the original listed asking price of the home. To clear out any seasonal differences in the selling and buying behavior of homeowners in the treatment versus the control group, we also include month-of-the-year by treatment group status fixed effects. Results are presented in Table 3. The estimated effect with the additional controls is very similar to that using our main specification.

B.4 Effects of MIP Cut on Home Listings

One alternative interpretation of our main reduced form result is that the MIP cut increased current homeowner purchase hazards because it increased the for-sale inventory by drawing more sellers onto the market. To test whether the MIP cut elicited a significant listing response, we regress the treatment measure against "Post", an indicator for whether the listing first went onto the market after the MIP cut. If treatment neighborhood owners responded to the MIP cut by listing their homes, the average value of "treatment" of new listings should increase after the cut because a greater fraction of listings come from high treatment neighborhoods.

Table 4 shows that we actually see a negative coefficient on "Post". The estimate is small, representing a change of about ¼ of a percent of the standard deviation of the treatment measure, but standard errors are tight meaning we can rule out an increase in listings in response to the MIP cut. The small response of new listings combined with the fact that the flow of listings onto the market is small relative to the stock of listings at any point in time suggests that changes in listing behavior are unlikely to explain our main results. In the longer run, listing behavior may play a more important role in the housing market's response to stimulus, but exploring long-run effects is beyond the scope of this paper.

B.5 Effects of MIP Cut on House Prices

The FHA MIP cut caused a demand shock at the low end of the market, so the price of the average home sold actually fell immediately following the premium cut due to sample selection effects. To test whether the cut had an effect on the price current homeowners received for the homes conditional on quality, we take the initial listed price as given and test if homes sold for a higher amount conditional on that price. First, we calculate a discount = ln(sale price) – ln(asking price). We regress this discount against the treatment measure, and "Post", an indicator for the sale taking place after the MIP cut, and the interaction. The coefficient on the interaction shows how much more (or less) sellers in treatment group neighborhoods received for a given home following the MIP cut.

Table 4 reports the results. The MIP cut appears to have a small but statistically significant increase on the sale price, as would be expected given the shorter time-on-market. The effect of increasing the treatment from 0 to 1 – its minimum to its maximum value – is to increase the sale-to-asking price by 1.4 percent. In Figure 1 of the main manuscript, we found that the comparable effect on the sale hazard was 7 percent, which is much larger elasticity compared to the sale price response.

B.6 Robustness of Sales Volume Multipliers to Endogenous House Prices

We add prices to the baseline model. At the time of every transaction, we assume buyers pay a price $p(\theta)$ to the seller. Agents who receive exit shocks receive the price once their home sells. We compute the multiplier from stimulus under various assumptions about the relationship between the price and market tightness.

To operationalize this model, we first need to re-calibrate it. We calibrate the model using the same procedure used for the baseline model and we set the steady-state price in each market equal to our estimates of $V^c - V^{s2}$ under the baseline model. The rationale for this price level is that the difference in utility associated with being contented relative to owning two homes is roughly equal to the utility of the price that the double owner would receive from selling one of her homes. We verified that our results are not sensitive to alternative values for the pre-stimulus steady-state price level. The model fit for this calibrated version of the model is almost identical to the baseline model fit. The parameter estimates adjust somewhat to account for the price level that is added to some of the value functions and subtracted from others.

With the re-calibrated model, we conduct the same exercise presented in Section 7 to see how sales volume responds to stimulus in this version of the model. Because the model continues to abstract from price determination, we assume that the price elasticity with respect to market tightness is equal to a multiple of the sale probability elasticity with respect to market tightness. We consider several values of the multiple.

Table 6 reports the sales volume multipliers for this version of the model. As prices become more responsive to market tightness, the multiplier estimates decrease, but not by much. Existing evidence suggests that the responsiveness of price to market tightness is significantly less than the responsiveness of sale probability. For example, in a model with search frictions and endogenous prices (but without a joint buyer-seller problem), Anenberg and Kung (2018) find that the elasticity of house prices is 1/3rd as large as the elasticity of sale probability in response to an interest rate shock. Diaz and Jerez (2013) find that in the data, the volatility of prices is 1/4th the volatility of time-on-market. Even when we conservatively assume that the elasticity of house prices is equal to the elasticity of sale probability, Table 6 shows that stimulus still leads to large sales volume multipliers of 2.22 and 1.43 in the cold and hot markets, respectively – only slightly less than our baseline estimates.

B.7 Sales Volume Multipliers Under Temporary Stimulus

Our baseline simulations assume that the stimulus is permanent. However, our model can deliver sizable multipliers from temporary stimulus as well. Figure 2 shows impulse responses when the stimulus is in place for one period and then is immediately withdrawn. Mechanically, the inflow into the renter pool is increased for one period and after that period the inflow returns to its pre-stimulus steady state level. The estimated multiplier in the hot market is 1.59, which is very similar to the baseline multiplier reported in Table 5 of the main manuscript. The estimated multiplier in the cold market is 2.33, which is somewhat lower than the baseline estimate, but is still sizable.

B.8 Steady State Multiplier

The transition to the new steady state after stimulus takes decades, as mismatch shocks occur very infrequently in our calibration. Our model is likely too simple to explain such long-run effects of stimulus. For example, supply is exogenous in our model but in the long-run supply may respond to stimulus. For completeness, Table 9 shows the multiplier in the new steady state following the permanent stimulus. As in the short-run, the steady state multiplier is larger in the cold market than the hot market. The steady state multipliers are larger than the short-run multipliers shown in Table 7 of the main manuscript.

The multipliers are above one in the steady state because increasing the population with a fixed housing stock increases the market tightness and efficiency of the matching process. Mismatched households cycle back to the contented state faster, leading to stimulated transaction volume in steady state. During the transition to the new steady state, market tightness is increasing. Therefore, the matching efficiency benefits are the highest in the new steady state. Table 9 also shows the multipliers are much lower when we fix choice probabilities at pre-stimulus levels, suggesting that the switching effect has a strong effect on the steady state multiplier as well.

C Model Details

C.1 Laws of Motion

The pool sizes evolve according the following equations:

$$b' = (1-\omega)[(1-q_b(\theta))b + \rho c * \sigma(b)] \tag{1}$$

$$d' = (1-\omega)[(1-q_s(\theta))d + q_b(\theta)b + q_b(\theta)(1-q_s(\theta))sb] \quad (2)$$

$$s' = (1-\omega)[(1-q_s(\theta))s + \rho c * \sigma(s)] \quad (3)$$

$$sb' = (1-\omega)[(1-q_s(\theta))(1-q_b(\theta))sb + \rho c * \sigma(sb)] \quad (4)$$

$$r' = (1-\omega)[(1-q_b(\theta))r + q_s(\theta)s + q_s(\theta)(1-q_b(\theta))sb] + inflow \quad (5)$$

$$e' = (1-q_s(\theta))e + \omega(c + b + s + sb + 2d) \quad (6)$$

$$c' = 1 - b - s - sb - 2d - e \quad (7)$$

$$\theta = \frac{b + r + sb}{d + s + sb + e} \quad (8)$$

where $inflow$ is an exogenous inflow into the renter pool and $\sigma(i)$ is defined in Equation 11 of the main manuscript.

C.2 Details on Model Calibration

We first note that the steady state market tightness is implied by the observed probabilities of buying and selling: $\frac{q_s}{q_b} = \frac{1-exp(-A\theta)}{(1-exp(-A\theta))/\theta} = \theta$. Using the values of

q_s and q_b reported in Table 4 of the main manuscript implies $\theta_L = 0.25$ and $\theta_H = 0.55$.

Given these values of θ, we choose A_L and A_H to match our estimates of $q_s(\theta)$ and $q_b(\theta)$ from the data shown in Table 3 of the main manuscript. The table shows that we are able to match these moments almost exactly for $A_L = 0.51$ and $A_H = 0.57$.

To calibrate the remaining parameters χ, u_0, u_2, we minimize the distance between model and data moments using the following procedure. For a guess of the parameter values, we first iterate on the following loop until convergence

1. Compute V^s under θ using equation 7 in the main manuscript
2. Compute V^b under θ using equation 8 in the main manuscript
3. Compute V^d under θ using equation 9 in the main manuscript
4. Compute V^{sb} under θ using equation 10 in the main manuscript
5. Compute V^r under θ using equation 4 in the main manuscript
6. Compute V^c under θ using equation 5 in the main manuscript

where θ is 0.25 and 0.55 for the loose and tight markets, respectively. After convergence, solve for the steady state values of the pool sizes by guessing at the pool sizes, including the exogenous inflow, and forward-simulating the economy until the pool sizes converge using the laws of motion described in equations 1 through 8. The converged pool sizes sizes must also satisfy $\theta_L = 0.25$ and $\theta_H = 0.55$.

Once the pool sizes converged, use the steady state pool sizes and value functions to compute the moments 1, 3, and 5 shown Table 4 of the main manuscript. Evaluate the squared distance between data and model moments and repeat until parameter values are found that minimizes the distance.

C.3 Details on Moments for Calibration

To calibrate the model's parameters, we match 12 moments from the data (6 in each of the hot and cold markets, respectively) listed in Table 4 of the main manuscript. The first moment is the effect of selling a home on the current homeowner's monthly probability of purchasing another home. The empirical counterpart of this moment is estimated in Section 4, as described in Section 6.

The second moment is the monthly hazard rate of selling for listed homes. In the data, this is the simple average probability a listing open in a given month closes with a sale that month. The third moment is the fraction of all open listings for which the seller is a double-owner. This is calibrated to the fraction of open listings per month for which we observe a purchase by the same owner in a prior month.

The fourth moment is the monthly hazard rate of purchase for households searching the market as a buyer. Finding a counterpart in the data for the purchase hazard is somewhat more complicated than for the sale hazard, because we do not have data directly on households searching, as we do for houses listed for sale. Instead, we infer that incumbent owner households that have already sold a home (and are thus not waiting to find a buyer before searching as buyers themselves) and who we do see eventually purchase a home (and are thus not exiters) are actively searching as buyers every month between the dates of sale and purchase. The estimated purchase hazard rate is the average probability of such households completing a purchase in one of these months.

The fifth moment is the fraction of mismatched households that choose the strategy "buy first". Restricting to all listed homes for which we see the owner purchase another home (to exclude exiters), this moment is calibrated to the fraction that bought prior to the initial listing date.

The sixth moment is θ, the market tightness. Because each match consists of one buyer and one seller, θ is simply the ratio of the monthly sale hazard to the monthly purchase hazard.

Each of these above moments is calculated separately for listings appearing in the coldest and hottest thirds of the country to provide different moments to match in the cold and hot markets.

C.4 Identification and Sensitivity

Figure 3 shows how each moment changes in response to a 0.01 increase in each parameter. Blue bars represent the moments for the cold market, and red bars represent the moments for the hot market. First, note that moments 2 and 4 from Table 4 of the main manuscript are only affected by A_L and A_H. Therefore, moments 2 and 4 for the cold market can identify A_L while moments 2 and 4 from the hot market can identify A_H. High values of u_0 increase the causal effect of selling on buying (moment 1) and decrease the share of double owners (moment 3) and the probability that a newly mismatched household searches as a buyer (moment 5). Higher values of χ and u_2 have opposite effects on these moments. χ and u_2 affect the moments in similar directions, but by different magnitudes. For example, increasing u_2 has a relatively large affect on moment 5 in both the cold and hot markets, whereas increasing χ has a relatively large effect in the cold market but not in the hot market.

C.5 Details on Model Simulation

To solve for the transition path to the new steady state following the stimulus shock, we follow the steps below. First, we solve for the new, post-stimulus steady state. To do this, we take an initial, iteration 0, guess at the steady state θ_0, compute the value functions at the guess of θ_0, solve for the steady state θ_1 implied by the value functions, and iterate on θ_n until convergence. With the new steady state θ in hand, we next iterate on the following loop until convergence:

1. Guess at a transition path for θ to the new steady state level.
2. Solve for the value functions along the transition path for the guess of the transition path for θ using backwards recursion from the new steady state.
3. Simulate the pool sizes implied by the value functions from step 2 according to equations 1 through 8.

4. Check if the guess of θ from step 1 equals the θ implied by the pool sizes from step 3 for every period along the transition path. If not, update the guess of θ toward the transition path implied by the pool sizes in step 3, and return to step 1.

References

Anenberg, Elliot, and Edward Kung. 2018. "Interest Rates and Housing Market Dynamics in a Housing Search Model." *working paper.*

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.

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.*

DeFusco, Anthony A, Charles G Nathanson, and Eric Zwick. 2017. "Speculative Dynamics of Prices and Volume." National Bureau of Economic Research Working Paper 23449.

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.

Figure 1: Fraction of Owners with Previous FHA use, by Z

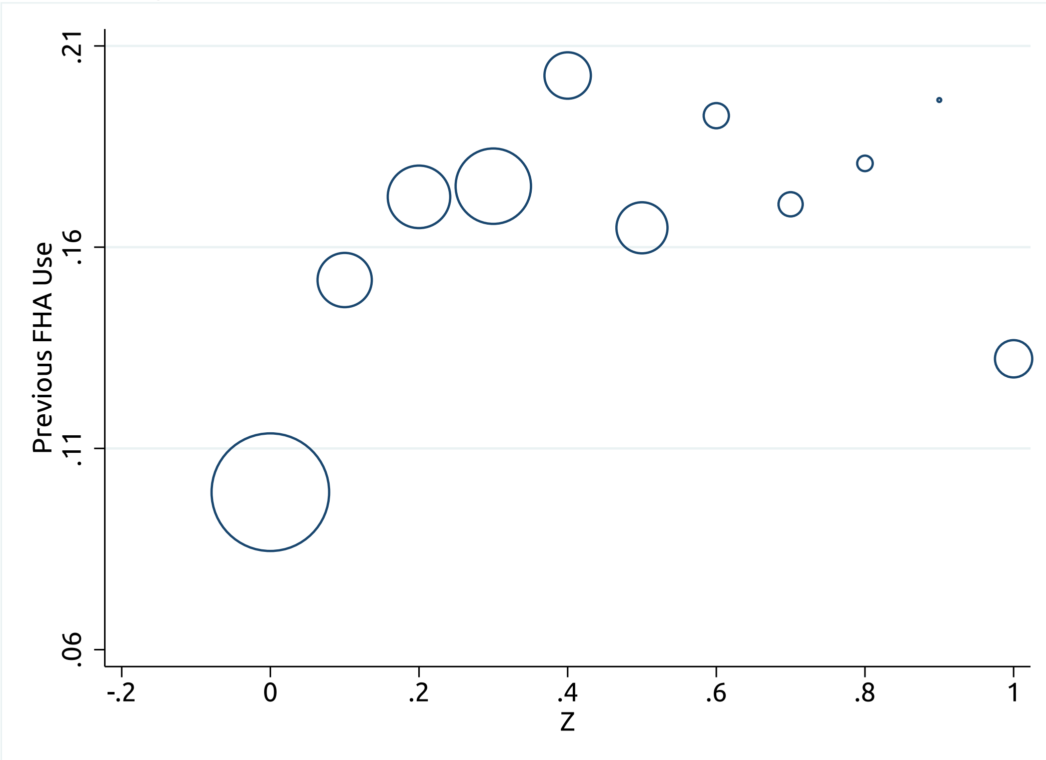

Note: Figure displays the fraction of incumbent owners with a listed home for sale whose most recent purchase or refinance of that home was financed with an FHA loan. Circle size indicates the mass of the distribution for each value of Z rounded to the nearest 0.1. Z is the fraction of home-purchase mortgages in the neighborhood and price range of house i that went to low FICO, high LTV buyers in the years prior to the FHA premium cut.

Figure 2: Sales Volume Response to a Temporary Demand Shock

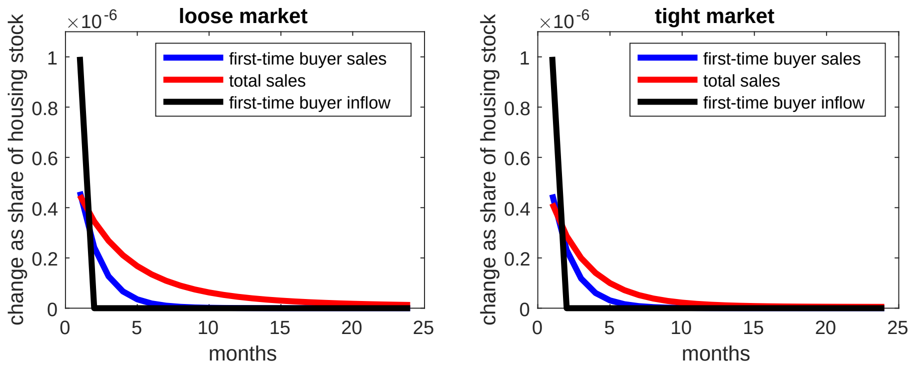

Note: At time 0, stimulus is introduced by increasing the first-time homebuyer inflow by 1e-6. At time 1, stimulus is permanently removed so that the first-time homebuyer inflow equals its pre-stimulus steady state level. The black line shows the path of stimulus. First-time homebuyers are agents searching to buy a home who have not previously owned a home. Changes shown are relative to the steady state prior to the stimulus.

Figure 3: Parameter Sensitivity Analysis

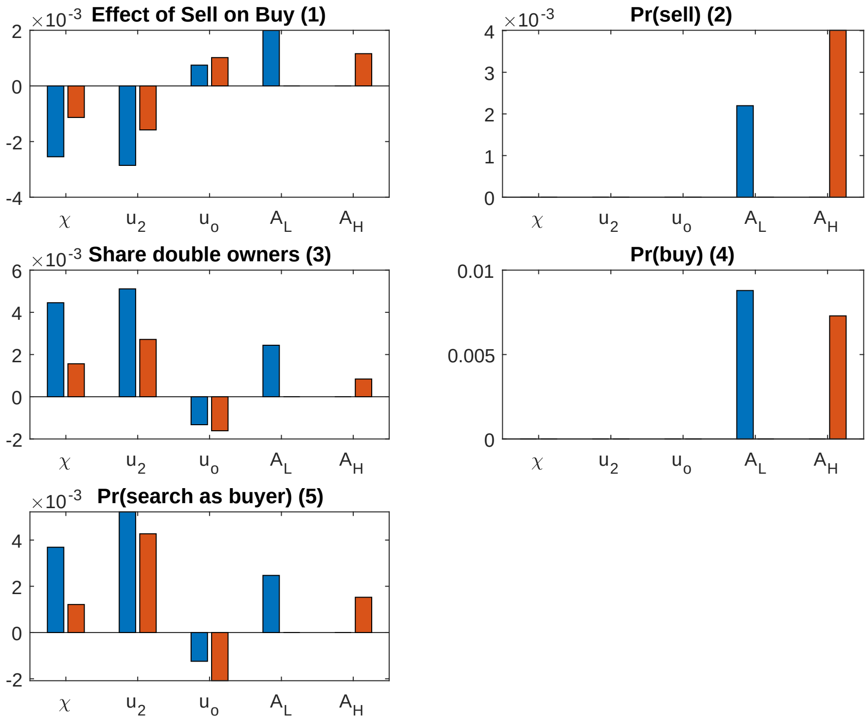

Note: Figure shows how each moment changes in response to a 0.01 increase in each parameter. Blue bars represent the moments for the cold market, and red bars represent the moments for the hot market. Numbers in parenthesis refer to the moment numbers in Table 4 of main manuscript. See the table for a more complete description of moments. Parameters tested are χ, the mismatch shock to flow utility, u_2 and u_0, the flow utilities of being a double owner and renter, and A_L and A_H, the matching technologies in hot and cold markets.

Table 1: Effect of Home Sale on Owner's Monthly Purchase Hazard, Restricted Samples

	Full Sample (1)	Non-FHA (2)	Cash Buyers (3)
Sold	0.119	0.095	0.198
	(0.022)	(0.033)	(0.099)
Z_i	0.002	0.002	0.004
	(0.0003)	(0.0004)	(0.001)
$Post_t$	0.005	0.0098	0.0065
	(0.0005)	(0.0007)	(0.0008)
$N \cdot T$	16,765,134	12,190,670	4,740,844
F-stat	598.2	380.04	38.29

Note: Z is the fraction of home-purchase mortgages in the neighborhood and price range of house i that went to low FICO, high LTV buyers in the years prior to the FHA premium cut. Column 1 shows 2SLS estimates of equation 2 of the main manuscript for the full estimation sample. Column 2 restricts the sample to owners whose previous purchase or refinance was not financed by an FHA loan. Column 3 is further restricted to owners who bought their current home without using a mortgage. Standard errors adjusted for clustering at the census tract level.

Table 2: Testing for Direct Effect of the Instrument

	Bought (1)	Used a Mortgage (2)	Low FICO, High LTV Ratio (3)
$Z_i \cdot Post_t$	0.002	-0.005	0.011
	(0.0004)	(0.005)	(0.007)
Z_i	0.003	0.02	0.058
	(0.0002)	(0.004)	(0.005)
$Post_t$	0.007	0.025	0.003
	(0.0001)	(0.002)	(0.002)
$N \cdot T$	16,804,476	563,836	158,207

Note: Column 1 shows the estimated reduced form effect of the instrument on the monthly purchase probability. Column 2 restricts the sample to months in which a purchase occurred, and shows the estimated reduced form effect of the instrument on the probability a mortgage was used to purchase the house. Column 3 further restricts the sample to purchases with a mortgage that were matched to the McDash data, and shows the estimated reduced form effect of the instrument on the probability the borrower had a FICO score below 680 and an LTV ratio greater than 80. Standard errors adjusted for clustering at the census tract level.

Table 3: Effect of Home Sale on Owner's Monthly Purchase Hazard, Robustness Checks

	Main Specification	Additional Controls	Sample with Unique Names
Sold	0.119	0.121	0.080
	(0.022)	(0.021)	(0.025)
$N \cdot T$	16,765,134	16,765,134	12,459,383
F-stat	598.2	260.03	427.42

Note: The main specification column shows results of the IV regression of monthly home purchase hazard on an indicator for whether the current home has sold. Regression controls for the share of purchase mortgages in the listed home's tract and price range that went to a low FICO, high LTV borrower (Z) and an indicator for the listed month being after January 2015. In the "Additional Controls" specification, regression additionally controls for tract and month fixed effects, interactions between month-of-the-year fixed effects and Z, and the original listed asking price. In the "Sample with Unique Names" column, estimation sample restricted to sellers with combinations of first and last name that are unique in the data set. Standard errors adjusted for clustering at the census tract level. Regression controls for Z and an indicator for the listed month being after January 2015.

Table 4: Effect of the FHA MIP Cut on Prices and New Listings

	Log Price Discount		Treatment Measure	
	(1)	(2)	(3)	(4)
$Z_i \cdot Post_t$	0.014	0.014		
	(0.001)	(0.001)		
Z_i	-0.024	-0.024		
	(0.001)	(0.001)		
$Post_t$	-0.001	0.009	-0.001	-0.002
	(0.001)	(0.001)	(0.0004)	(0.0004)
Month-of-the-Year FEs		X		X
$N \cdot T$	2,712,977		4,077,417	2,719,366

Note: Columns 1 and 2 show the estimated reduced form effect of the instrument on the log difference between purchase price and initial listed asking price. Columns 3 and 4 show the estimated change in the average value of the treatment measure after the MIP cut. "Post" refers to sales that occurred after the MIP cut. Columns 2 and 4 control for month-of-the-year fixed effects.

Table 5: Model Summary

Agent Type	Searching as	# Homes Owned	Contented with Homes Owned?	Flow Utility	Transition to
Renters	Buyer	0	Not applicable	u_o	Contented owner
Contented owners	Not Searching	1	Yes	u	Seller, buyer, or seller-buyer
Sellers	Seller	1	No	$u-\chi$	Renter
Buyers	Buyer	1	No	$u-\chi$	Double owner
Seller-buyers	Seller and buyer	1	No	$u-\chi$	Contented owner, double owner, or renter
Double owners	Seller	2	1 Yes, 1 No	u_2	Contented owner

Note: this table shows the various states agents in the model can occupy, as well as their search behavior, ownership status, flow utility, and possible transitions to other states. In addition, all agents transition out of the model economy if they receive the exit shock, ω.

Table 6: Sales Volume Multiplier Estimates from Stimulus, Endogenous Prices

Assumptions	Cold Market	Hot Market
$\frac{\partial lnp}{\partial ln\theta} = 0$	2.46	1.51
$\frac{\partial lnp}{\partial ln\theta} = 0.5 * \frac{\partial lnq_s}{\partial ln\theta}$	2.34	1.47
$\frac{\partial lnp}{\partial ln\theta} = \frac{\partial lnq_s}{\partial ln\theta}$	2.22	1.43
$\frac{\partial lnp}{\partial ln\theta} = 2 * \frac{\partial lnq_s}{\partial ln\theta}$	2.00	1.35

Note: Model implied multiplier estimates. The multiplier is $\frac{\Delta TotalSales}{\Delta First\text{-}timeBuyerSales}$ where the change is with respect to the pre-stimulus steady state and sales volume for both total sales and first-time buyer sales is summed over the two year period following the stimulus.

Table 7: Model Fit, Alternative Calibration

		Tight Market		Loose Market	
Moment	Description	Data	Model	Data	Model
1. $q_b(\theta)\frac{s}{s+d+e+sb}$	causal effect of selling on buying	0.1160	0.0940	0.1930	0.1222
2. $q_s(\theta)$	sell probability	0.27	0.2691	0.12	0.1197
3. $\frac{d}{s+d+e+sb}$	double owners / total sellers	0.22	0.2047	0.22	0.1891
4. $q_b(\theta)$	buy probability	0.49	0.4893	0.48	0.4788
5. $Pr(b)$	probability of searching as buyer	0.16	0.1643	0.12	0.1217
6. θ	market tightness	0.55	0.5500	0.25	0.2500

Table 8: Sales Volume Multiplier Estimates from Stimulus, Alternative Calibration

Assumptions	Cold Market	Hot Market
Baseline model	1.71	1.28
Choice probabilities fixed at pre-stimulus levels	1.32	1.18

Note: Model implied multiplier estimates under alternative calibration shown in Table 7. The multiplier is $\frac{\Delta TotalSales}{\Delta First\text{-}timeBuyerSales}$ where the change is with respect to the pre-stimulus steady state and sales volume for both total sales and first-time buyer sales is summed over the two year period following the stimulus.

Table 9: Sales Volume Multiplier Estimates in Steady State

Assumptions	Cold Market	Hot Market
Baseline model	4.39	3.17
Choice probabilities fixed at pre-stimulus levels	3.20	2.73