

The Propagation of Demand Shocks Through Housing Markets*

Elliot Anenberg[†] Daniel Ringo[‡]

March 15, 2019

PRELIMINARY AND INCOMPLETE. DO NOT CITE OR CIRCULATE.

Abstract

The presence of incumbent homeowners creates a friction in housing markets, as incumbents may wait to match with a buyer for their current home before buying their next home. As a result, demand stimulus can produce a multiplier effect by freeing up owners attempting to sell their current home, allowing them to re-enter the market as buyers. Exploiting a shock to housing demand caused by the 2015 surprise cut in Federal Housing Administration mortgage insurance premiums, we find that homeowners buy their next home sooner when the probability of their current home selling increases. This effect is especially pronounced in cold housing markets, in which homes take a long time to sell. We build and calibrate a model of the joint buyer-seller decision that explains these findings as a result of homeowners avoiding the cost of owning two homes simultaneously. Simulations of the model demonstrate that stimulus to home buying generates a substantial multiplier effect, particularly in cold housing markets.

*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

The federal government has devoted considerable resources toward stimulating housing demand through a variety of channels—for example, through quantitative easing, first time home buyer tax credits, and subsidies to the Federal Housing Administration (FHA) and Government Sponsored Enterprises (GSEs). A major motivation for these policies is to provide quick stimulus to home sales and economic activity during episodes of weak economic growth. Indeed, home sales are accompanied by sizable purchases of durable goods (Benmelech et al. (2017)) and directly generate income for Realtors, loan officers, and others. Facilitating movement up the housing ladder can also lead to increases in new construction and home ownership (Abel (2018)). In addition, allowing homeowners to sell more easily can help households re-optimize their location and consumption of housing services.

From a welfare perspective, housing demand is a potentially fruitful target for stimulus because of frictions that prevent households from efficiently sorting into their optimal residence. In particular, *incumbent homeowners* who are attempting to move must match on both sides of a search market, as a buyer for their new home and a seller for their current one.¹ To the extent that incumbents wait to enter the market as a buyer until they have sold their current home (e.g. due to the high costs of carrying two homes), search behavior by incumbents creates a mechanism by which stimulus to housing demand can have a multiplier effect. Each policy induced home purchase frees up the incumbent to re-enter the market as a buyer, who can then buy a new home and free *that* home's incumbent to re-enter, and so on.

This multiplier effect—and hence the efficacy of stimulus and subsidies to housing demand—depends crucially on how quickly sellers and buyers in that housing market can expect to find a match. In cold markets where the probability of selling is low (but a home to buy can be found quickly), sellers prefer to wait to find a buyer for their existing home before buying themselves, as buying before selling is likely to put them in the costly position of owning two homes for a lengthy period of time. Therefore, in a cold market, policies that stimulate housing demand can unleash substantial additional demand as existing sellers are freed up to enter the market as buyers themselves. Conversely, in hot markets where it is relatively difficult to find

¹Investors can smooth this friction by acting as a market maker, providing liquidity by both buying and selling homes. However, investors are involved in only a minority of single family home transactions in the United States.

a home to buy (but a home can be sold quickly), sellers may prefer to buy before selling to avoid a potentially lengthy period of “homelessness” (or short-term rental). As a result of this search behavior, the multiplier effect of demand stimulus on sales volume in hot markets should be lower than it is in cold markets.

In this paper, we show that demand stimulus is effective in slow housing markets because it generates substantial multiplier effects. In hot markets, in contrast, we do not find any evidence of a meaningful multiplier effect, suggesting that policy-generated housing demand shocks are restricted to the households directly affected by those policies in such markets, and do not propagate. Estimates of the multiplier effect and an understanding of the mechanisms that generate it are important because they help policy makers design appropriate housing stimulus policy.

We begin the paper by providing evidence suggesting the presence of multiplier effects due to the joint buyer-seller search problem, and that the multiplier effect is likely to vary with market conditions. Using a novel data set that follows individual owners who list their home for sale to see if they buy another home elsewhere within the United States, we show that home purchase activity is sensitive to the ability of existing owners to sell their homes, especially in cold housing markets. To identify this effect, we exploit a change in FHA pricing that provides exogenous variation in the probability that an existing homeowner is able to sell her listed home.² In a cold market, we estimate that selling the listed home is associated with a 16 percentage point increase in the monthly hazard rate of that seller buying another home. However, in a hot market, the estimated effect is less than 3 percentage points and not statistically significant. This finding suggests that a substantial multiplier effect may arise in cold markets from homeowners who wait until their current home gets an offer before agreeing to buy their next home.

We then develop a model of housing search and transactions that we calibrate to match these empirical findings and other moments. The calibrated model allows us to quantify the multiplier effect of stimulus under different market conditions. In the model, homeowners occasionally receive moving shocks, in which case they must choose whether to search the market as a seller first, as a buyer first, or as a buyer and seller simultaneously. Search is random and the matching technology is constant returns to scale. As in Moen et al. (2015), an owner’s optimal strategy depends on

²Bhutta and Ringo (2017) use the same policy shock to show that home buying is highly responsive to interest rates in a large segment of the population.

others' choices. For example, in a buyer's market where homes for sale have a low probability of matching (i.e. the ratio of buyers to sellers, or market tightness, is low), owners tend to choose to sell first to avoid a long period of owning two homes. This behavior reinforces the low market tightness. Consistent with the data, in such a market, the causal effect of selling one's home on buying another home is relatively high, as the inability to sell was preventing many owners from searching to buy in the first place.³

Simulations of the estimated model show that the two-year multiplier associated with a generic shock to first-time home buyer demand is over 3 in cold markets, meaning that each additional transaction by a first-time home buyer stimulates more than two additional transactions within 24 months. In the cold market, owners tend to choose to sell before buying, so the additional inflow of first-time buyers into the market immediately unleashes a significant amount of demand from existing owners. Furthermore, the additional inflow of buyers encourages newly mismatched owners to buy first, which strengthens the multiplier effect. In contrast the estimated multiplier in hot markets is less than one. This is because the incoming first-time home buyers crowd out more repeat home buyers than are encouraged by the multiplier effect. The supply of homes coming onto the market – either from new construction or existing owners deciding to list their homes for sale – is exogenous in our model and thus policy-invariant. Our model therefore delivers a sizable and quick multiplier effect in cold markets simply through dual-search and the endogenous decisions of existing homeowners to buy or sell first.

Our paper is related to Moen et al. (2015) and Anenberg and Bayer (2015). Like our paper, these papers have models predicting that home purchase activity is sensitive to the ability of existing owners to sell their homes, and that the sensitivity varies with market conditions. Our paper contributes by providing direct, empirical support for these predictions using an exogenous source of variation in the propensity of existing owners to sell. In addition, our paper focuses on estimating multipliers on transaction volume while Moen et al. (2015) focuses theoretically on how the joint buyer-seller problem can generate multiple equilibria and Anenberg and Bayer (2015) focus empirically on how the joint buyer-seller problem can amplify price volatility.

³The effect is attenuated, in both our reduced-form estimates and the model, by owners who do not actually intend to buy another home (e.g. they are moving to the rental market or combining households).

Our finding of a large multiplier effect in cold markets is in concordance with the findings of Berger et al. (2016) and Best and Kleven (2017), who both find large effects from demand stimulus policies implemented in the wake of the financial crisis. Notably, Berger et al. (2016) find that the increase in home sales caused by the First-Time Home Buyer Tax Credit was not offset by a reduction in sales following the credit’s expiration—that is, the stimulated demand was not simply a shift in the timing of future demand. Our results may offer an explanation. In our simulations, the overall volume of transactions is elevated months and years after the rate of entry of first time home buyers has returned to normal, due to the multiplier effects described above.

Our paper contributes to a broader literature that has theorized about the role of the joint buyer-seller decision in housing market dynamics. These include Wheaton (1990), who shows that a search and matching model of home sales with incumbent owners can explain structural vacancy rates, and Rosenthal (1997), who shows that linked chains of buyers and sellers can cause lags in the movement of house prices. Also related is the literature on vacancy chains in housing markets (see e.g. White (1971) and Turner (2008)), which focuses on how prospective buyers must wait for a vacancy to appear before moving into their next residence, creating another vacancy in turn. Ortalo-Magne and Rady (2006) develop a model in which existing homeowners’ demand to move up the housing ladder is a function of the demand for their current home.

The rest of the paper is organized as follows. Section 2 explains the reduced-form estimator we use to identify the effect of a marginal home sale on its owners probability of purchasing a subsequent home. In Section 3 we describe the data used, and in Section 4 present the results. We describe our model of the housing market and the joint buyer-seller decision in Section 5, and the calibration of the model in Section 6. Section 7 contains our simulations of a shock to first-time home buying demand, which we use to calculate the magnitude of the multiplier effect under different market conditions.

2 Estimation

The housing demand multiplier, which we discussed in the Introduction and will explore in depth beginning in Section 5, is produced by the marginal home sale

releasing the incumbent owners to re-enter the market as home buyers themselves. When they buy their next home, this releases a further set of incumbent owners who re-enter the market as buyers, and so on. Multiple transactions could end up taking place due to the initial, marginal home sale. The size of the housing demand multiplier therefore depends crucially on how much the marginal home sale increases the seller's probability of purchasing another home over a given window of time.

Estimating this effect is not trivial, however. There are a number of factors that could bias simple regressions of the probability of an incumbent homeowner purchasing their next home on the sale of their current home. One major concern is reverse causality. We are interested in the degree to which homeowners wait to sell their current home before buying their next one. If some homeowners instead wait until they have found a new home to buy before selling their current one, this could produce a spurious correlation between selling and buying. Another concern is property investors. These individuals own homes that they do not occupy, and so may sell homes without any need to quickly buy another one. If investors transact more frequently than owner-occupiers, their presence in transactions data will bias estimates. A third concern is overall market conditions, which could affect both homeowners' sale and purchase probabilities regardless of the causal relationship between the two.

Over and above these potential sources of bias, the timing of sale agreements presents a major obstacle to estimating the effect of a home sale on its owner's subsequent purchases. Specifically, a buyer and seller may agree on a transaction months before it is actually scheduled to take place (and recorded). Observing only transaction dates, it is possible for a purchase to be caused by a sale that had not occurred yet, if the sale was agreed to before the purchase was.

For all these reasons, we want an exogenous source of variation in the probability a particular home sells to identify how marginal sales affect their owner's purchasing behavior. Such variation is provided by the January 2015 reduction in mortgage insurance premiums (MIP) for Federal Housing Authority (FHA) loans. As shown by Bhutta and Ringo (2017), this surprise 50 basis point reduction in the cost of credit caused an immediate jump in the volume of home buying by populations dependent on the FHA for access to mortgage credit—that is, borrowers with low credit scores and down payments. Overall, the volume of purchase mortgages increased by about 2 percent in response to the MIP cut. The abrupt reaction was due to credit rationing, as households who were on the margin of being denied a mortgage due to high ratios

of debt-service payments to income were able to slip below otherwise binding underwriting thresholds as a result of the reduction in mortgage costs. The MIP cut caused an influx of new buyers that increased the probability a current homeowner gets an offer for their home, but it had essentially no *direct* effect on current homeowner’s purchase probabilities. This is because, as Bhutta and Ringo (2017) find, the increase in home buying came entirely from lower income, highly leveraged FHA borrowers who are almost 90 percent first time home buyers. Any effect of the MIP cut on current homeowner’s purchase behavior came *indirectly* through the cut’s effect on their ability to sell their current home. In the appendix, we show that purchases by current homeowners who were cash buyers, or who had a high credit score or low LTV ratio (and thus did not need for FHA insurance), increased just as much as purchases by their lower credit score, lower down payment counterparts. This finding implies that the MIP cut affected the purchases of current homeowners only through the indirect channel.

Because the direct response to the FHA MIP cut was confined to lower credit score, highly leveraged FHA borrowers, we have cross sectional as well as across time variation in which homes were exposed to the resulting demand shock. In particular, by using geography and price range, we can identify homes that were and were not in the choice set of this responsive population. We define the responsive population to be borrowers with FICO scores below 680 and loan-to-value (LTV) ratios greater than 80 percent, just as in Bhutta and Ringo (2017). Houses in census tracts and prices ranges (divided into \$50,000 buckets) where no responsive borrowers purchased a home in 2013 or 2014 form our control group. Our treatment group is houses in tracts and price ranges where there was some purchase activity by the responsive population. As a first stage, we estimate:

$$S_{it} = \alpha_0 + \alpha_1 Z_i + \alpha_2 Post_t + \alpha_3 Z_i \times Post_t + \mu_{it} \quad (1)$$

where S_{it} is an indicator that house i sells in month t , Z_i is an indicator that i is in the treatment group, and $Post_t$ is an indicator that t is after January 2015. Our first stage is thus a difference-in-differences estimator, comparing the monthly sale probabilities of treatment and control group homes, before and after the January 2015 MIP cut.

Our second stage estimates how the sale of a house affects the monthly probability

that the owner purchases a new home. We estimate:

$$P_{it} = \beta_0 + \beta_1 S_{it} + \beta_2 Z_i + \beta_3 Post_t + \epsilon_{it} \quad (2)$$

where P_{it} is an indicator that the owner of house i purchased a new home in month t . Equation 2 is estimated via 2SLS, with $Z_i \times Post_t$ used as an instrument for S_{it} .

Note that while equation 2 is written as if we are testing for effects of S on P only in the same month, t , the only time variation in the instrument is an indicator for before and after the FHA MIP cut. Therefore, we are really estimating how much the monthly purchase hazard of treatment group homeowners increased relative to the control group after January 2015, scaled by how much the monthly sale hazard of treatment group homeowners increased relative to the control group. In other words, we are using the following Wald estimator:

$$plim_{n \rightarrow \infty} \hat{\beta}_1 = \frac{\Delta \Pr(P|Z = 1) - \Delta \Pr(P|Z = 0)}{\Delta \Pr(S|Z = 1) - \Delta \Pr(S|Z = 0)} \quad (3)$$

where Δ indicates the difference between the periods before and after the MIP cut. Using these broader time windows (essentially each of the full years before and after the MIP cut) allows us to get around the issue of agreement and transaction dates not necessarily lining up.

3 Data

We use a number of different sources to put together the data set for our estimation. Our primary requirement is the ability to observe households who are attempting to sell their home, whether they succeed, and when (and if) they purchase another home. In addition, the instrument, described in Section 2, requires information on the location and price range of the home.

The data set is built around Multiple Listing Service (MLS) records provided by CoreLogic. The data comes directly from regional boards of Realtors, and covers over 50 percent of property listings nationwide. Information on homes listed for sale includes the dates the listing was opened and closed, whether the home actually sold, the asking price and location. Our main estimation sample is restricted to single-family homes that had an active listing some time in the years 2014 and 2015. This leaves

us with just under 6 million properties with a listing in this period.

To track the home purchase behavior of the owners of these listed homes, we turn to property transaction data, also provided by CoreLogic. Sourced from county deeds records offices, this data covers over 98 percent of the U.S. population. This source give us the name(s) of the owner(s) listed on properties that transacted or were refinanced. A unique property ID allows an exact match of these transactions to the listings in the MLS data.

To construct the instrument, Z , we use mortgage records collected under the Home Mortgage Disclosure Act (HMDA) merged with rate lock data provided by Optimal Blue. The HMDA data contain individual loan records for the vast majority of residential mortgage loans originated each year, including information on loan amount, purpose, property location, borrower income and whether the loan carried FHA insurance. Optimal Blue provides underwriting data, including FICO scores and LTV ratios, for approximately one quarter of the mortgage market. From the merged data, we can observe the fraction of home purchase loans in each census tract and \$50,000 purchase price range that went to a borrower with a low FICO score and high LTV ratio in the years around the MIP cut.

3.1 Tracking households between homes

We track individual households between the sale of house i and their purchase of the next house using the named owners on the deed. To get the names of the current owners of i , we match deeds records to the MLS records using the unique property ID. The CoreLogic deeds records extend back only to the year 2003, so our sample is limited to houses that transacted or were refinanced between 2003 and 2013, inclusive. This leaves us with just over 3 million total properties listed for sale between 2014 and 2015 matched to the names of the sellers.

To determine if and when these sellers purchased another house, we match these names to the the names of buyers of single family homes over the the 2014-2015 period. We use an exact match on last names and a fuzzy match on the first and middle names, to allow for abbreviations, dropped initials, nicknames or other misspellings. Details of the matching procedure are available in the appendix. Matches are required to fall within a 6 month window of the period in which the seller's home was listed in the MLS.

3.1.1 Assessing 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 (2015) and DeFusco et al. (2017). One concern this raises 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 observations 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.

Who are these owners that sell a house without buying another one? In addition to investors who own multiple properties, they include owner-occupiers who are moving into other living arrangements. These could be people moving from owning to renting or into institutionalized living arrangements. They could be people combining households, through marriage or moving in with family. Homeowners who emigrate or die also leave a home to be sold, without showing up as having purchased another.

A further concern is the possibility of false positive matches. Names are not unique identifiers of persons—home sellers with common names in particular may be identified as having purchased another home, due to being matched with a different buyer of the same name. To address this issue, we note that among our sample of home sellers, about 75 percent have a combination of first and last name that is unique in our data set. These owners should be much less likely to generate a false

positive match. As we show in the appendix, all our results are quite similar when we restrict the estimation sample to just these uniquely named owners. We therefore do not think false positive matches are a meaningful source of bias.

3.2 Creating the panel

The final step of building our estimation sample is to construct a panel based on the dates of listing, delisting and sale of each listed house, as well as the purchase date if the owners bought another house. We create pseudo observations at the monthly level to produce an unbalanced panel. Houses enter the panel either when they are listed for sale, or in January 2014 if the listing was already active at that point. They exit when the house is delisted for good, regardless of whether it sold. Unsold listings are therefore implicitly treated as censored. Each month the house is in the panel, the dummy variable S is set to one if the house sold that month, and P is set to one if the owners bought another house that month, and are zero otherwise.

Summary statistics for the estimation sample are presented in Table 1 for the treatment and control groups separately. Treatment group homes are somewhat less expensive on average, as would be expected given that they are in the price range of lower-income FHA borrowers. The two groups had similar hazard rates of selling and buying new homes.

4 Results

Bhutta and Ringo (2017) showed that the 2015 FHA MIP cut increased the number of home purchase mortgages originated. To confirm that our instrument based on the MIP cut is relevant, we need to show that this increase in originations translated to a measurable increase in the probability of sale for an identifiable set of homes. Furthermore, we would like to support the claim that any relative increase in the sale probability of treatment group homes after the MIP cut was due to the cut itself, rather than a prevailing difference in secular trends.

In Figure 1, we plot OLS estimates of the effect of treatment group status on the probability a home in the estimation sample sells in a given month, for each month in from 2012 through 2015. The dashed lines mark the 95 percent confidence interval, using standard errors robust to clustering at the tract level. Through 2014, there is no

clear trend in the difference between treatment and control group sale probabilities. Following the MIP cut, however, treatment group homes become significantly more likely to sell than control group homes. Through most of 2015, treatment group homes are about 1 percentage point more likely to sell each month than they were in 2014—approximately a 7 percent increase in sale hazard. This fits the pattern demonstrated in Bhutta and Ringo (2017) of an immediate and apparently sustained jump in treatment group sales following the MIP cut.

Turning to the second stage, we estimate equation 2 on the main estimation sample. Results are shown in Table 2, with naive OLS estimates shown as well for comparison. The IV results suggest that selling one’s home increases the seller’s monthly purchase hazard by over 8 percentage points. The F-statistic indicates that the IV easily passes weak-instrument thresholds. We therefore conclude that marginal home sales do indeed produce a multiplier effect, spurring further home sales as they release the incumbent owner to enter the market as a buyer.

As we described in the introduction, however, this average treatment effect conceals substantial heterogeneity across market conditions. In particular, we would expect a stronger multiplier effect (and hence a larger $\hat{\beta}_1$) in cold housing markets, where homes take a long time to sell. Homeowners in these markets have an incentive to find a buyer for their current home before buying a new one, or they may be stuck with the carrying costs of two homes for a long time. In contrast, we would expect little or no multiplier effect (and hence low values of $\hat{\beta}_1$) in hot markets where homes sell quickly. In these markets, homeowners are not concerned about being stuck holding two properties for an extended period, and so can wait until they have found a new residence to put their current home up for sale. These homeowner’s purchase of their next home would therefore be unaffected by the sale of their current one.

To test for this differential effect across markets, we divide our sample into three groups of approximately equal numbers of listed homes. Groups are defined by how hot the housing market is in the county that the house is located in. The "Cold" group includes the third of listed homes located in the slowest paced markets, where active listings have a monthly probability of sale of just under 10 percent, on average. The "Hot" group includes the third of homes in the fastest markets, with an average monthly probability of sale of 21 percent. The middle, "Moderate", group has a monthly sale probability of 16 percent. We then re-estimate equation 2 on each of

these three groups separately. Results are presented in Table 3.⁴

A marginal home sale increases the homeowner’s monthly purchase hazard by about 16 percentage points in cold markets, almost double the strength of the effect found for the overall market shown in Table 2. In contrast, the estimated effect in hot markets is small (less than 3 percentage points) and not statistically significant. This finding supports the story that the multiplier effect arises from homeowners who wait until their current home gets an offer before agreeing to buy their next home. Note that while there is essentially no evidence for a multiplier effect in hot markets, the instrument is still quite strong for this group (as it is for the other two groups). This means that there is a substantial effect of the FHA MIP cut on first time home buying in all market types, but the effect on repeat buyers only appears where would expect to see a multiplier effect.

5 Model

We consider a simple model of home sales in a housing market with search frictions. Time is discrete and agents discount the future at rate β . There is a fixed stock of homes normalized to have measure one.

Most of the time, homeowners are *contented* in their homes, which means that they receive the flow utility u from owning the home. Occasionally, however, contented owners receive exogenous mismatch shocks, in which case their flow utility of living in the home drops to $u - \chi$. These mismatched owners then enter the search market with the goal of moving house. That is, they will try to sell the home they are currently mismatched with and buy a different home that puts them back in the contented state.

The key decision is whether to enter the market as a buyer first (“buyers”), a seller first (sellers”), or as a buyer and seller simultaneously (“seller-buyers”). Market conditions will endogenously affect this decision. However, even for a given set of market conditions, agents in the model will choose different strategies because of exogenous idiosyncratic preferences. Some agents will be very motivated to move (i.e. low flow utility of being mismatched) and so will not want to wait to sell until they buy. Some agents will have high costs of holding two homes, and so will wait to

⁴As we show in the appendix, these results are robust to the inclusion of a detailed set of control variables.

buy until they sell.

Buyers meet sellers via a frictional matching process. The matching function simply depends on the total stock of buyers and sellers searching, and is assumed to be constant returns to scale. Let $\theta = b/s$ be the ratio of buyers to sellers in the market, often referred to as *market tightness*. Then, the probability that a buyer meets a seller is $q_b(\theta)$ and the probability that a seller meets a buyer is $q_s(\theta) = \theta q_b(\theta)$. If a buyer and a seller are matched, we assume that the matching results in a sale.⁵

House prices are exogenous and there is no aggregate uncertainty.⁶

Renters

We refer to agents who are searching the market to buy a home, but do not own a home, as renters.⁷ The net flow utility associated with being a renter and searching the market to buy is u_0 . Renters include agents who are entering the housing market for the first time as well as previously contented agents who have sold their old home and are looking to buy a different one, but to solve the model we do not need to distinguish between these types. The value function associated with being a renter is therefore

$$V^r = u_0 + \beta[q_b(\theta)V^c + (1 - q_b(\theta))V^r] \quad (4)$$

With probability $q_b(\theta)$, the renter matches with a seller and becomes contented. We omit the transfer of a price, p , from the buyer to the seller in the value functions because the price is assumed to be exogenous and the same regardless of which types of agents are transacting.⁸ With probability $1 - q_b(\theta)$, the renter does not match with a seller and remains a renter.

⁵For most matchings, this assumption is not binding because the parameter values we estimate imply that the buyer and seller would endogenously agree to a sale.

⁶Moen et al. (2015) show that in a model similar to ours with endogenous prices, the transaction sequence of joint buyer-sellers is still key for differences in sales activity across hot and cold markets. We choose to treat prices as exogenous because we don't expect that endogenizing them would change our main results and the model becomes more complicated with endogenous prices.

⁷We do not call them buyers because some agents in our model who are searching to buy a home also own a home, and we want to distinguish between these types.

⁸Omitting the price is wlog if we assume that all homes are financed with 100 percent LTV, interest-only mortgages. The interest payments on the loan simply get subsumed by the flow utility parameters.

We assume that there is free entry and exit into the renter pool. Let k denote the utility associated with not entering the renter pool. k can be thought of as the outside option of living somewhere else, or of renting and not searching to buy forever. In equilibrium,

$$V^r = k \tag{5}$$

as agents will freely enter the market until the point in which the marginal buyer is indifferent between entry or exit.

Contented Owners

Contented owners receive the flow utility u , until they receive either of two exogenous shocks. With probability ω , contented agents become mismatched with the housing stock altogether, in which case they will try to sell their home and exit our model economy upon sale. For example, these could be death shocks or emigration shocks. We introduce these shocks because in our data, not every seller goes on to buy another home. With probability ρ , contented agents become mismatched with their current home and want to move into a different home. The value function associated with being a contented owner is simply

$$V^c = u + \beta[(1 - \rho)(1 - \omega)V^c + \rho(1 - \omega)V^m + \omega V^e] \tag{6}$$

where V^m and V^e denote the value functions associated with being mismatched and exiting, respectively. We normalize the utility associated with selling and exiting to zero, so $V^e = \frac{u - \chi}{1 - \beta(1 - q_s(\theta))}$.

Mismatched Owners

With probability ρ , contented homeowners become mismatched and can follow one of three strategies: (1) search the market as a seller first, then search as a buyer once their house has sold (2) search the market as a buyer first, then search as a seller once they have bought a new home (3) search as a buyer and seller simultaneously.

We denote these agents “sellers”, “buyers”, and “seller-buyers”, respectively. The value functions associated with each of the three strategies are V^s , V^b , V^{sb} . We assume

that each strategy is associated with a type 1 extreme value shock, so that we can write the expected value function associated with being mismatched as

$$V^m = 0.5772 + \ln[\exp(V^s) + \exp(V^b) + \exp(V^{sb})] \quad (7)$$

where 0.5772 is Euler's constant.

Sellers

Mismatched owners who choose to sell first receive a flow utility $u - \chi$. Upon finding a buyer for their home, which occurs with probability $q_s(\theta)$, sellers enter the renter pool, as they will no longer own a home. The value function associated with being a seller is therefore

$$\begin{aligned} V^s &= u - \chi + \beta[q_s(\theta)V^r + (1 - q_s(\theta))V^s] \\ &= u - \chi + \beta[q_s(\theta)k + (1 - q_s(\theta))V^s] \end{aligned} \quad (8)$$

where the second line of equation 8 follows from substituting in equation 5.

Buyers

Like sellers, mismatched owners who choose buy first receive a flow utility $u - \chi$. However, upon finding a home to buy, which occurs with probability $q_b(\theta)$, these mismatched owners will own two homes. The value function associated with being a buyer is therefore

$$V^b = u - \chi + \beta[q_b(\theta)V^d + (1 - q_b(\theta))V^b] \quad (9)$$

where V^d is the value function associated with being a "double owner" (i.e. owning two homes).

Double Owners

The total flow utility associated with owning two homes is u_2 .⁹ The value function associated with being a double owner is therefore

$$V^d = u_2 + \beta[q_s(\theta)V^c + (1 - q_s(\theta))V^d] \quad (10)$$

Note that we are assuming that double owners do not receive mismatch shocks.

Seller-Buyers

Seller-buyers can transition directly into renters (if they sell first), double owners (if they buy first), or contented owners (if they buy and sell at the same time). The value function associated with being a seller-buyer is

$$\begin{aligned} V^{sb} = & u - \chi + \beta[q_s(\theta)q_b(\theta)V^c + (1 - q_s(\theta))q_b(\theta)V^d + \dots \\ & \dots + q_s(\theta)(1 - q_b(\theta))k + (1 - q_s(\theta))(1 - q_b(\theta))V^{sb}] \end{aligned} \quad (11)$$

Here we note that our assumption that all matchings lead to transactions has some bite. For example, a seller-buyer who matches with a seller but not a buyer may not actually wish to buy if owning two homes is especially costly and so V^d is low. However, allowing seller-buyers to make transaction decisions significantly complicates the model.¹⁰ Therefore, we choose to simplify the model by assuming that all matchings lead to transaction.¹¹

⁹Wlog, one could also write the flow utility as $u + u - \chi - u_2$ where u_2 captures the effects of frictions that prevent homeowners from realizing the consumption benefits of owning two homes simultaneously.

¹⁰For example, to avoid this complication, Moen et al. (2015) assume that seller-buyers allocate a fraction of their time to buying and the remaining fraction to selling.

¹¹A motivation for this assumption is that realtors put pressure on their clients to transact because they are only compensated if a transaction occurs. Therefore, in reality one reason why all mismatched owners do not choose to search to buy and sell at the same time may be that they do not want to be pressured to buy if they match with a home that seems like a plausible fit before they are able to sell. Our model captures this disincentive to being a seller-buyer.

Equilibrium and Discussion

An equilibrium in the housing market consists of value functions and a market-tightness θ that satisfies equations (4) through (11). Note that because of our assumption of free entry into the renter pool, the equilibrium market tightness must satisfy

$$\theta = q_b^{-1} \left(\frac{(1 - \beta)k - u_0}{\beta(V^c - k)} \right) \quad (12)$$

which we obtain by plugging equation 5 into equation 4 and rearranging to isolate θ . Equation 12 shows that more renters will want to enter the housing market, and θ will be higher, when (1) the utility of owning a home, V^c , is large relative to their outside option k , and (2) when the costs of searching for a home to buy, u_0 , are low relative to the flow utility they receive from not entering the market and searching to buy, $(1 - \beta)k$.

From equation 12, we can see that the equilibrium market tightness does not depend on the distribution of agents in the economy as long as the value function associated with being contented does not depend on the distribution. Without equation 12, the equilibrium of our economy would generally depend on the distribution of agents across the various pools (e.g. the number of seller-buyers, double owners, etc). To solve for the equilibrium of our model both in and out of steady state, we do not need to keep track of this heterogeneity, which significantly simplifies our computational work. The Appendix describes how we solve the model.

An interesting feature of the equilibrium is that mismatched owners' strategies depend on others' choices. For example, in a buyer's market (low θ), homes for sale have a low probability of matching. As a result, mismatched owners will tend to choose to sell first to avoid a long period of double ownership. The decision to sell first reinforces the low market tightness. As shown in Moen et al. (2015), this strategic complementarity in the transaction sequence may lead to multiple equilibria.

6 Calibration

To calibrate the model, we must first parameterize the matching technology. We assume a Cobb-Douglas matching technology, so that

$$q_s(\theta) = A\theta^{1-\gamma} \tag{13}$$

$$q_b(\theta) = A\theta^{-\gamma} \tag{14}$$

where A is a technology parameter that determines the efficiency of the market.

The parameters of the model are summarized in Table 4. The parameters are $u, u_0, u_2, \chi, k, A, \omega, \gamma, \beta$. Wlog, we normalize $u = 0$. We set $\beta = 0.95^{1/12}$ so that each model period can be thought of as a month. We also set $\omega = \gamma = 0.002$ implying an expected value of being in the contented state of 21 years.

We calibrate the remaining parameters by matching data moments from a hot and a cold market. To generate hot and cold markets in our model, we allow A to take on two different values, A_L and A_H [add interpretation].

We construct the data moments using the micro data discussed in Section 3. The data moments we use are shown in Table 5 and we describe how we compute each data moment in the Appendix. We assume that the markets from which we create our data moments are in steady state. Therefore, to fit the data moments, we assume that the model economy is also in steady state.¹² The mapping of the steady state of the model to the data moments is straightforward, except in one important case, which we elaborate on in the next paragraph.

A key data moment is the IV-estimate of β_1 from equation 2. Recall that β_1 measures the causal effect of selling one's home on the monthly probability of purchasing another home. What is β_1 according to the model? Of the four types of agents with homes on the market for sale in our model (seller-buyers, double owners, exiters, and sellers), the ability to sell only affects the purchase behavior of the seller-types. Seller-types simply cannot buy until they have sold. Therefore, selling increases the probability they buy in the next period by $q_b(\theta)$. Double-owners and exiters are simply not in the market to buy, so selling generates no change in the probability that these types buy a home. Seller-buyers are in the market to buy, but

¹²Moen et al. (2015) show that under certain parameter values, a similar model will produce multiple stable equilibria, one with $\theta < 1$ and one with $\theta > 1$. In our data the match rate of buyers is always higher than sellers, implying that $\theta < 1$. We therefore confine our equilibrium selection to instances in which $\theta < 1$.

they are searching to buy while they are searching to sell, so selling also generates no change in the probability that a seller-buyer buys. The probability a seller-buyer buys is $q_b(\theta)$ regardless of whether they are able to sell or not. Therefore, we can write

$$\beta_1 = q_b(\theta) \frac{s}{s+d+e+sb} \quad (15)$$

where s, d, e, sb denote the steady state number of agents in the seller, double-owner, exiter, and sell-buyer pools, respectively.¹³

6.1 Identification

The model is nonlinear and so almost all parameters affect all outcomes. Nonetheless, here we provide a discussion of the main features of the data that identify each of our parameters.

We have six parameters that largely affect the probability that a buyer and seller transact. These are $A_L, A_H, k_L, k_H, \gamma, u_0$. The effect of A and γ on transaction probabilities can be seen directly from equations 13 and 14. The effects of k and u_0 can be seen from equation 12. This equation shows that k and u_0 are key for pinning down the equilibrium θ , which then helps determine the transaction probabilities from equations 13 and 14. The six data moments that help to identify these six parameters are the market tightness in both the cold and hot markets, the probability of selling in both the cold and hot markets, and the probability of being in both the cold and hot markets.

The remaining parameters, χ and u_2 , are largely identified by the probability of searching as a buyer and the probability of searching as a seller.¹⁴ For example, mismatched owners will be less likely to choose to be “buyers” or “seller-buyers” when u_2 is very low. When the mismatch utility penalty is high, sellers want to get rid of their mismatched home as soon as possible, increasing the probability that they choose “seller” or “seller-buyer”.

¹³To see this even more clearly, note that $\beta_1 = (q_b(\theta) - 0) \frac{s}{s+d+e+sb} + (0 - 0) \frac{d}{s+d+e+sb} + (0 - 0) \frac{e}{s+d+e+sb} + (q_b(\theta) - q_b(\theta)) \frac{sb}{s+d+e+sb} = q_b(\theta) \frac{s}{s+d+e+sb}$.

¹⁴We do not directly observe the probability of searching as a seller, but we observe $q_b(\theta) \frac{s}{s+d+e+sb}$, which is correlated with it.

6.2 Results

Table 4 show the parameter estimates. The parameter estimates show that the flow utility associated with mismatch is larger than the flow utility associated with being a double owner and the flow utility of being a renter. These estimates are consistent with our intuition that double ownership and short-term rentership are costly states to be in.

Table 5 shows the model fit.

7 Counterfactuals

Using our parameter estimates recovered in the previous section, we now explore how sales volume in our model economy responds to stimulus. We consider stimulus in both the cold and hot markets, corresponding to $A = A_L$ and $A = A_H$, respectively. We initialize the two markets at their respective steady states, and then stimulate demand by increasing u_0 , which is the flow utility of being a renter, by a small amount.¹⁵ Increasing u_0 increases the equilibrium market tightness, θ , because more renters will want to enter the housing market until congestion in the buyer pool lowers the value of being a renter sufficiently.¹⁶ To solve for the new market tightness and value functions, we simply follow the steps outlined in Appendix A under the new value of u_0 . With the new market tightness and value functions in hand, we then simulate the model forward until the market eventually converges to a new steady state.

The left panel of Figure 2 illustrates the transition dynamics of sales volume for the cold market. The orange line shows the change in total sales volume relative to sales volume in the steady state prior to the stimulus. Stimulus increases sales volume immediately, as more renters flow into the market, increasing the transaction rate of homes on the market at the time of the stimulus. Although the stimulus to u_0 is assumed to be permanent, its stimulative effect on sales volume is temporary. In our model, the housing stock is fixed. The main driver of steady state sales volume are the mismatch shocks, which determine the supply of homes coming onto the market.

¹⁵For the results shown here, we set the change in u_0 equal to 0.001. This increases the equilibrium market tightness to 0.2461 from 0.2419 in the hot market, and to 0.5097 from 0.5094 in the cold market.

¹⁶Using equation 12, it can be shown that $d\theta/du_0 > 0$ (to do).

Mismatch shocks are exogenous and remain fixed in our counterfactuals.

The blue line shows the change in sales volume to “newborn” households. Newborn households are the subset of households among the pool of agents searching to buy a home who have not previously owned a home. Newborn buyers include those who are drawn into the market because of the stimulus, as well as new entrants from previous periods that have not yet bought a home and exited the buyer pool. Initially, the stimulus increases newborn sales as stimulus draws newborns into the market to buy from the large existing for-sale inventory. For-sale inventory is large in the cold market because the probability of selling, $q_s(\theta)$, is relatively low. In the initial periods after the stimulus, newborn sales actually exceed the difference in total sales volume relative to the initial steady state. This means that on net, newborn sales are crowding out sales that would have gone to agents who were already in the buyer pool at the time of the stimulus.

After the initial periods, newborn sales quickly taper off while total sales volume remains stimulated. In the cold market, a relatively large number of households are waiting to search to buy until they sell. The initial increase in newborn sales therefore unleashes a significant amount of additional demand as sellers are freed up to search as buyers themselves. Furthermore, because the probability of sale is concave in market tightness as can be seen in equation 14, this increase in the number of buyers is happening in a market where sales volume is especially responsive to an increase in the number of buyers. The increase in internal demand therefore leads to sizable further increases in total sales volume. Newborn sales fall off because the increase in internal demand, which increases the market tightness, makes it less attractive for newborns to enter. Newborns also become crowded out by households who receive mismatch shocks after the stimulus goes into effect. These newly mismatch households are now more likely to choose to search to buy at the time of mismatch as a result of the tighter market.

We compute the multiplier from stimulus according to the following equation:

$$multiplier = \frac{\Delta TotalSales}{\Delta NewbornSales} \quad (16)$$

where the change is relative to the pre-stimulus steady state and sales volume is summed over the two years following the implementation of the stimulus. In the context of Figure 2, the multiplier is the the area under the orange line divided by the area under the blue line. Table 6 shows that the multiplier for the cold market over

two years is 3.84. Each newborn sale generated by the stimulus leads to 3.84 total sales, or to an additional 2.84 total sales over and above the sale directly generated by the stimulus.

The right panel of Figure 2 shows the transition dynamics for the hot market. Total sales volume is much less responsive to the stimulus. In the hot market, homes are selling relatively quickly and so for-sale inventory that can be cleared by the stimulus is relatively low. Furthermore, since the market is already tight, adding additional buyers to the market – from the newborn households directly or from internal movers who are freed up to buy – has less of an effect on the probability of sale. Table 6 shows that the multiplier in the cold market is less than one at 0.78. Stimulus has a crowd out effect that dominates the muted multiplier effect on total sales in the hot market.

To summarize, there are two main mechanisms in the model that generate the large multiplier in the loose market. First, the stimulus helps to clear the large for-sale inventory in the cold market, allowing the sellers of those homes to become buyers themselves, creating an endogenous increase in internal demand. Second, because the stimulus increases the market tightness, newly mismatched owners are subsequently more likely to choose to search as buyers, which further increases internal demand and crowds out newborns from entering the market. We call the second mechanism the “switching effect”.

To gauge the quantitative importance of each mechanism, we plot the transition dynamics assuming that the probability of choosing seller, buyer, and seller-buyer upon mismatch does not change after the stimulus goes into effect. Figure 3 shows that the response of total sales volume to the stimulus in both the cold and hot markets is similar to the baseline model. However, the multiplier in the cold market is much lower than in the baseline model, as shown in Table 6. These results suggest that the switching effect is not necessary to stimulate total sales volume in the cold market. Without the switching effect, more newborns enter the housing market, and newborn sales account for almost all of the stimulus to total sales volume. However, with the switching effect, we can achieve the same stimulus to total sales with a much smaller contribution from newborn sales in the cold market. In other words, the switching effect is key for allowing newborn sales to stimulate significant additional demand from internal movers.

Even with the choice probabilities fixed, however, note that the multiplier in the

cold market is substantially larger than in the hot market. This means that for a given increase in the number of first time home buyers, the total sales volume would increase about 50 percent more in cold markets than in hot markets purely through releasing the pent-up demand of sell-first owners. When mismatched owners are allowed to change their strategy in response to the demand shock, the difference in overall sales is almost 400 percent.

8 Conclusion

Incumbent homeowners' desire to avoid prolonged stretches owning either two homes at once, or no home at all, creates frictions in housing markets that complicate the overall response to demand shocks. We show in this paper that in cold housing markets, the direct effect of stimulus to housing demand can be substantially smaller than the indirect effect which propagates due to homeowners' strategic behavior. In contrast, in hot markets the weak propagation mechanism and crowd-out effects can lead to an overall response that is actually smaller than the direct effect.

These results imply that stimulus to housing markets is more effective in periods when markets are slow—exactly the times when such stimulus policies are likeliest to be implemented. The presence of substantial frictions in cold housing markets also suggests that the equilibrium is far from efficient, so stimulus policies may be justified on a welfare enhancing basis. Overall, the takeaway is that considering only the direct effect of stimulus policies on home buying misses much of the economic response.

References

- Abel, Joshua**, “Why do Housing Markets Freeze? Investigating House Lock, Nominal Loss Aversion, and Irreversible Construction,” *mimeo Harvard University*, 2018.
- Anenberg, Elliot and Patrick Bayer**, “Endogenous Sources of Volatility in Housing Markets: the Joint Buyer-Seller Problem,” *mimeo Duke University*, 2015.
- Benmelech, Efraim, Adam Guren, and Brian T Melzer**, “Making the house a home: the stimulative effect of home purchases on consumption and investment,” Technical Report, National Bureau of Economic Research 2017.

- Berger, David, Nicholas Turner, and Eric Zwick**, “Stimulating housing markets,” Technical Report, National Bureau of Economic Research 2016.
- Best, Michael Carlos and Henrik Jacobsen Kleven**, “Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK,” *The Review of Economic Studies*, 2017, 85 (1), 157–193.
- Bhutta, Neil and Daniel Ringo**, “The Effect of Interest Rates on Home Buying: Evidence from a Discontinuity in Mortgage Insurance Premiums,” *working paper*, 2017.
- DeFusco, Anthony A, Charles G Nathanson, and Eric Zwick**, “Speculative Dynamics of Prices and Volume,” Working Paper 23449, National Bureau of Economic Research May 2017.
- Moen, Espen R, Plamen Nenov, and Florian Sniekers**, “Buying first or selling first in housing markets,” 2015.
- Ortalo-Magne, Francois and Sven Rady**, “Housing market dynamics: On the contribution of income shocks and credit constraints,” *The Review of Economic Studies*, 2006, 73 (2), 459–485.
- Rosenthal, Leslie**, “Chain-formation in the Owner-Occupied Housing Market,” *The Economic Journal*, 1997, 107 (441), 475–488.
- Turner, Lena Magnusson**, “Who gets what and why? Vacancy chains in Stockholm’s housing market,” *European Journal of Housing Policy*, 2008, 8 (1), 1–19.
- Wheaton, William C**, “Vacancy, search, and prices in a housing market matching model,” *Journal of Political Economy*, 1990, 98 (6), 1270–1292.
- White, Harrison C**, “Multipliers, vacancy chains, and filtering in housing,” *Journal of the American institute of planners*, 1971, 37 (2), 88–94.

Figure 1: Effect of Treatment on Monthly Sale Probability



Note: Figure shows the estimated effect, by month, of treatment group status on the probability a home listed for sale closes in that month. Treatment group sales in February 2015 and later are potentially affected by the reduction in FHA insurance premiums. Point estimates and the 95 percent confidence interval, based on standard errors clustered at the tract level, shown.

Figure 2: Sales volume response to a demand shock

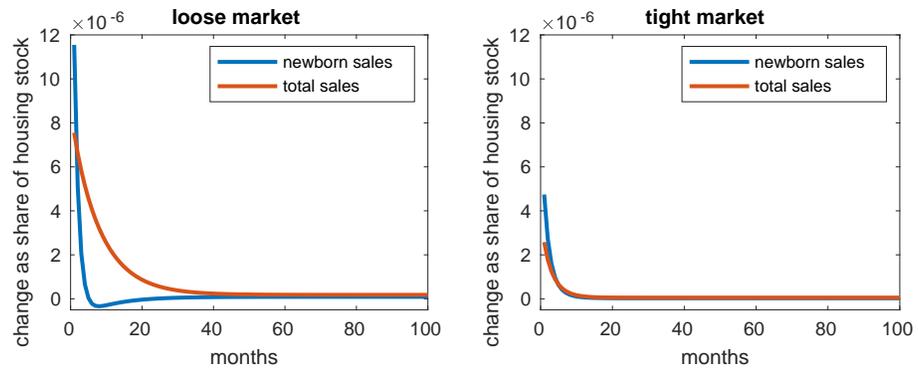


Figure 3: Sales volume response to a demand shock, choice probabilities fixed

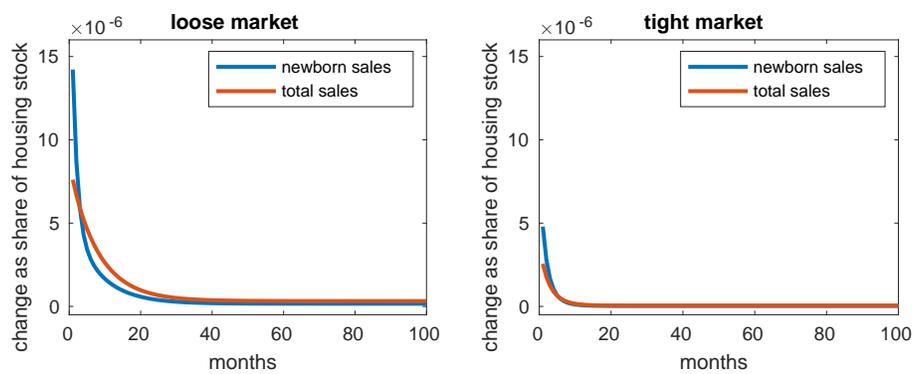


Table 1: Summary Statistics

Variable	Statistic	Treatment Group	Control Group
Initial Listing Price	<i>Median</i>	175	219
	<i>Std. Dev.</i>	(58)	(88)
Days on Market	<i>Median</i>	91	85
	<i>Std. Dev.</i>	(108)	(110)
<i>S</i>	<i>Mean</i>	0.145	0.147
<i>P</i>	<i>Mean</i>	0.034	0.033
	<i>N</i>	526,414	3,431,025
	<i>N × T</i>	2,303,584	14,500,892

Note: Prices listed in \$1,000s. *S* is the monthly hazard rate of the listed property selling. *P* is the monthly hazard rate of the owner of the listed property buying another house.

Table 2: Effect of Home Sale on Owner's Monthly Purchase Hazard

	OLS	IV
Sold	0.041** (0.0002)	0.084* (0.033)
Treatment group		0.001** (0.0002)
Post January 2015		0.004** (0.0006)
$N \cdot T$	16,778,818	
F-stat	260.03	

Note: Standard errors adjusted for clustering at the census tract level.

**p < 0.01

*p < 0.05

Table 3: Effect of Home Sale on Owner’s Monthly Purchase Hazard, by Local Market Conditions

	Cold	Moderate	Hot
Sold	0.161** (0.060)	0.140* (0.059)	0.027 (0.032)
Treatment group	0.001** (0.0002)	0.0002 (0.0007)	0.002** (0.0004)
Post January 2015	0.002* (0.0009)	0.003* (0.001)	0.006** (0.0006)
$N \cdot T$	6,838,723	5,613,310	4,326,785
F-stat	103.4	69.9	220.2

Note: Standard errors adjusted for clustering at the census tract level.

**p < 0.01

*p < 0.05

Table 4: Parameter Estimates

parameter	Description	Value
u	contented flow utility	0
u_0	renter flow utility	-1.2415
u_2	double owner flow utility	-1.0071
χ	mismatch flow utility penalty	0.6478
k_L	renters outside option, loose market	-6.1169
k_H	renters outside option, tight market	-4.1065
A_L	matching efficiency, loose market	0.1988
A_H	matching efficiency, tight market	0.3357
γ	elasticity of matching function	0.6357
ρ	probability of mismatch	0.002
ω	probability of death	0.002
β	monthly discount factor	0.9957

Table 5: Model Fit

Moment	Description	Tight Market		Loose Market	
		Data	Model	Data	Model
$q_b(\theta) \frac{s}{s+d+e+sb}$	causal effect of selling on buying	0.03	0.0465	0.16	0.1464
$q_s(\theta)$	sell probability	0.27	0.2717	0.12	0.1137
$\frac{s+d+e+sb}{c}$	mismatched owners / contented owners	0.008	0.0147	0.014	0.0351
$\frac{e}{s+d+e+sb}$	exiters / mismatched owners	0.5556	0.5005	0.5556	0.5005
$\frac{d}{s+d+e+sb}$	double owners / mismatched owners	0.22	0.2661	0.22	0.1861
$q_b(\theta)$	buy probability	0.49	0.4856	0.48	0.5270
$\frac{\exp(V^b)}{\exp(V^b)+\exp(V^s)+\exp(V^{sb})}$	probability of searching as buyer	0.16	0.1740	0.12	0.0791
θ	market tightness	0.55	0.5595	0.25	0.2158

Table 6: Sales Volume Multiplier Estimates

Economy	Estimate
Hot market	0.78
Cold market	3.84
Hot market, no change in choice probabilities	0.75
Cold market, no change in choice probabilities	1.11

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. 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.¹⁷ 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.

¹⁷Inspecting 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.

B Robustness and Validity Checks

B.1 Robustness to the Inclusion of Control Variables

Our main results, described in Section 4, 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 7. We see a similar pattern as in Table 3, with strong effects of home sales on purchase hazard in cold markets, and no statistically significant effect in hot markets.

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 re-run the estimator on the subsample with unique names, for cold, moderate and hot markets separately. Results are presented in Table 8. Comparing Tables 8 and 3, we can see that the results are quite similar, with selling causing an (approximately) 16 percentage point increase in the purchase hazard in cold markets, and no evidence of an effect in hot markets. This test suggests false positive matches are not materially biasing our main estimates.

B.3 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. We can test for such direct effects by noting that among current owners, not all households would be equally responsive to a cut in the FHA's premiums. Owners who do not intend to use a mortgage (cash buyers) are not directly influenced by the price of a particular form of mortgage credit. Similarly, 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 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.

To test for such effects, we make use of additional data from both CoreLogic and McDash Analytics. The CoreLogic deeds data we use for our main estimation sample also includes records for whether the property was purchased with a mortgage, and the mortgage amount. 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 on the probability the purchaser has a low FICO score and high LTV ratio (among the further subset that made use of a mortgage, and for which we found a match in the McDash data).

For purposes of comparison, we also estimate the direct effect of the instrument on current owner's monthly purchase probability. Results are presented in Table 9. 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 9 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 any direct effect 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. 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.

Table 7: Effect of Home Sale on Owner’s Monthly Purchase Hazard, with Additional Controls

	Cold	Moderate	Hot
Sold	0.214* (0.103)	0.181** (0.065)	0.045 (0.037)
$N \cdot T$	6,838,723	5,613,310	4,326,785
F-stat	35.2	57.8	177.7

Note: Regressions additionally control for tract and month fixed effects, interactions between month-of-the-year fixed effects and a treatment group indicator, and the original listed asking price. Standard errors adjusted for clustering at the census tract level.

**p < 0.01

*p < 0.05

Table 8: Effect of Home Sale on Owner’s Monthly Purchase Hazard, for Sample with Unique Names

	Cold	Moderate	Hot
Sold	0.165** (0.058)	0.114 (0.070)	-0.007 (0.036)
Treatment group	0.001** (0.0003)	-0.0007 (0.0009)	0.001* (0.0005)
Post January 2015	0.001 (0.001)	0.002 (0.002)	0.006** (0.0007)
$N \cdot T$	5,240,316	4,125,971	3,103,563
F-stat	93.1	54.0	168.4

Note: 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.

**p < 0.01

*p < 0.05

Table 9: Testing for Direct Effect of the Instrument

	Bought (1)	Used a Mortgage (2)	Low FICO, High LTV Ratio (3)
$Z_i \cdot Post_i$	0.002** (0.0004)	-0.005 (0.005)	0.011 (0.007)
Z_i	0.003** (0.0002)	0.02** (0.004)	0.058** (0.005)
$Post_t$	0.007** (0.0001)	0.025** (0.002)	0.003 (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.

**p < 0.01

*p < 0.05

C Details on Solving the Model

To solve the model, we iterate on the following loop until convergence. Given an initial guess of the value functions,

1. Compute θ using (12)
2. Compute V^s using (8)
3. Compute V^b using (9)
4. Compute V^d using (10)
5. Compute V^{sb} using (11)
6. Compute V^c using (6)

After convergence, we solve for the steady state values of the pool sizes by guessing at the pool sizes and forward-simulating the economy until the pool sizes converge. The pool sizes evolve according the following equations:

$$b' = (1 - q_b(\theta))b + \rho(1 - \omega)c \frac{\exp(V^b)}{\exp(V^b) + \exp(V^s) + \exp(V^{sb})} \quad (17)$$

$$d' = (1 - q_s(\theta))d + q_b(\theta)b + q_b(\theta)(1 - q_s(\theta))sb \quad (18)$$

$$s' = (1 - q_s(\theta))s + \rho(1 - \omega)c \frac{\exp(V^s)}{\exp(V^b) + \exp(V^s) + \exp(V^{sb})} \quad (19)$$

$$sb' = (1 - q_s(\theta))(1 - q_b(\theta))sb + \rho(1 - \omega)c \frac{\exp(V^{sb})}{\exp(V^b) + \exp(V^s) + \exp(V^{sb})} \quad (20)$$

$$e' = (1 - q_s(\theta))e + \omega c \quad (21)$$

$$c' = 1 - b - s - sb - 2d - e \tag{22}$$

$$renters' = \theta(s + sb + e + d) - (b + sb) \tag{23}$$

where c denotes the mass of contented owners. Note that while the pool sizes will change period-to-period when the economy is not in steady state, the market tightness, θ , will not. θ remains constant both in and out of steady state at the level that is pinned down by equation (12).