

The Propagation of Demand Shocks Through Housing Markets*

Elliot Anenberg[†] Daniel Ringo[‡]

June 7, 2019

Abstract

The presence of incumbent homeowners creates a friction in housing markets, as incumbents wait to match with a buyer for their current home before buying their next home. As a result, demand stimulus produces 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 first-time home buyer 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 can stimulate housing demand through a variety of channels—for example, through quantitative easing, first-time homebuyer tax credits, and subsidies through the Federal Housing Administration (FHA) and Government Sponsored Enterprises (GSEs). Housing demand stimulus can be used to quickly increase home sales and economic activity, which may be especially desirable 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. In addition, allowing homeowners to sell more easily can help households re-optimize their location and consumption of housing services (Karahan and Rhee (2019); Brown and Matsa (2016)), and can increase new construction and homeownership as households move up the housing ladder (Abel (2018)).

Housing demand may also be a fruitful target for stimulus because of the potential for multiplier effects. Multiplier effects can arise because of the large role played in housing markets by *incumbent homeowners* who are attempting to move. These owners must match on both sides of a search market, as a buyer for their new home and a seller for their current one. Many incumbents wait to buy until they have sold their current home—for example, due to the high costs of carrying two homes. Therefore, a policy induced home purchase can immediately free up an 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.¹ Multiple transactions could end up taking place due to the initial, policy induced home sale.

A main contribution of this paper is to show that multiplier effects exist and that, under certain market conditions, they can be very large. A policy implication of our findings is that accounting for the indirect effects of stimulus on home sales is just as important as—and sometimes more important than—accounting for the direct effects when assessing the efficacy of stimulus policy.

We begin the paper with evidence that home purchase activity is sensitive to the ability of existing owners to sell their homes, especially in cold housing markets where the probability of selling is low (but a home to buy can be found more quickly).

¹Search frictions, which prevent an instantaneous and efficient match between buyers and sellers, are a crucial element of the multiplier mechanism. Third party investors 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.

We construct 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. To overcome endogeneity concerns associated with the relationship between selling and purchasing, 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 19 percentage point increase in the monthly hazard rate of that seller buying another home. In a hot market, the estimated effect is a smaller 11 percentage points. These findings suggest that the decision to buy does indeed depend on the ability to sell for many incumbents, especially in cold markets. As a result, stimulus may generate substantial multiplier effects by triggering a chain of transactions.

To quantify the multiplier effect of stimulus under different market conditions, we calibrate a model of housing search and transactions to match our empirical findings and other moments from our micro data. 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 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. Conversely, in hot markets where it is relatively difficult to find a home to buy (but a home can be sold quickly), sellers tend to choose to buy first to avoid a long period of “homelessness” (or short-term rental).

Simulations of the estimated model show that the two-year multiplier associated with a generic shock to first-time home buyer demand is substantial at 2.48 in cold markets, meaning that each additional transaction by a first-time home buyer stimulates one and a half additional transactions, in expectation, within 24 months. In the cold market, owners tend to choose to sell before buying in the model, 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

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.

multiplier effect. The supply of homes coming onto the market – either from new construction or existing owners deciding to move – 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. The estimated multiplier in hot markets is much smaller, although still significantly above 1, as fewer incumbents wait to buy until they have sold under such market conditions.

We close the paper by showing that housing market stimulus can be an effective method of fiscal stimulus due to multiplier effects, especially in cold markets. In the first year following the decrease in FHA premiums we used to calibrate our model, we find that each dollar of foregone revenue by the government directly leads to an additional \$4.25 and \$2.56 in GDP in the cold and hot market, respectively. The fiscal multipliers are large because each policy-induced home sale leads to sizable additional home sales, but the government does not lose any revenue on the additional home sales indirectly generated by the stimulus.³ We assume that home sales increase GDP only through Realtor commissions and spending on furniture, home improvement, and related expenditures that typically accompany a home sale (Benmelech et al. (2017)). Accounting for additional effects, such as the encouragement of new residential investment, would push these estimated fiscal multipliers higher.

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.

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 on sales volume from demand stimulus policies implemented in the wake of the

³We do not make conclusions about the efficacy of the FHA premium cut in particular because its main motivation was likely not fiscal stimulus. We use the variation induced by the premium cut to evaluate the fiscal multiplier from generic housing stimulus.

financial crisis. Notably, both papers find little or no reversal in home sales in the year or two following the expiration of stimulus.⁴ Our results may offer an explanation for the lack of a swift reversal following demand stimulus in the housing market. In our simulations, the marginal first time home buyer continues to induce an elevated overall volume of transactions months and years after making the initial purchase, 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.

Finally, our finding that housing stimulus can be a relatively effective form of fiscal stimulus is consistent with the findings of Best and Kleven (2017). Like us, Best and Kleven (2017) finds a spending multiplier from housing stimulus that is larger than estimates from existing work analyzing the effects of tax rebates on consumer spending (Parker et al. (2013); Johnson et al. (2006); Shapiro and Slemrod (2003); Agarwal et al. (2007)). Interestingly, our finding that stimulus is especially effective in cold markets does not appear to hold generally for other, non-housing focused stimulus. Based on her review of the literature, Ramey (2019) finds no robust evidence of higher multipliers during recessions or times of slack. For taxes, the evidence actually suggests that multipliers are higher during expansions. Therefore, our results suggest that housing stimulus is also relatively effective because it is most effective in slow markets – exactly the times when such stimulus policies are likeliest to be implemented.

The rest of the paper is organized as follows. Section 2 explains the reduced-

⁴As noted by Berger et al. (2016) and Best and Kleven (2017), these results differ from reversal patterns found in the auto market. Mian and Sufi (2012) find quick reversal in auto sales after the Cash-for-Clunkers program expired.

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. Section 8 evaluates the fiscal multiplier from housing stimulus.

2 Estimation

As discussed in the Introduction, the size of the housing demand multiplier depends crucially on how much the marginal home sale increases the seller's probability of purchasing another home over a given window of time. In this section, we describe how we address a number of endogeneity concerns in order to convincingly estimate this effect. Our results from this exercise will form the key moments that we use to calibrate our model developed in Section 5.

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 downward. A third concern is overall market conditions, which could affect both homeowners' sale and purchase probabilities regardless of the causal relationship between the two.⁵ On net, the bias in a simple regression could be positive or negative.

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

⁵In our housing search model below, sale and purchase probabilities are negatively correlated as market conditions change.

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. Furthermore, the lag between agreement and transaction can vary significantly across transactions. These timing issues will introduce an additional source of bias in naive estimates.

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.⁶ 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, making the premium cut an ideal source of variation for our research design. 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 FHA insurance), increased just as much as purchases by their lower credit score, lower down payment counterparts. This finding supports our assumption 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 first time homebuyers, we have cross sectional as well as across time variation in which homes were exposed to the resulting demand shock. 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 price ranges (divided into \$50,000 buckets) where no responsive

⁶Overall, 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.

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. The treatment intensity increases with the share of 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 the share of home purchase loans in i 's tract and price range that went to low FICO, high LTV borrowers, and $Post_t$ is an indicator that t is after January 2015. Our first stage is thus similar to 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 somewhere within the U.S. in month t . Equation 2 is estimated via 2SLS, with $Z_i \times Post_t$ used as an instrument for S_{it} .

Note that the only time variation in the instrument is an indicator for before and after the FHA MIP cut. Therefore, we are effectively 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. The estimator could be simplified to:

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

under the additional assumption that $Var(Z | Post = 1) = Var(Z | Post = 0)$. By using these broader time windows for identification (essentially each of the full years before and after the MIP cut), we do not need to take a stand on the precise lead or lag structure through which S affects P . This estimator mitigates bias from misalignment of agreement and transaction dates.

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 somewhere within the U.S., 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 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 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 residences. 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. 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, we show in the appendix that our results are quite similar when we restructure the estimation sample to sellers who have a combination of first and last name that is unique in our data set. These sellers represent about 75 percent of the sellers in our sample and should be much less likely to generate a false positive match.

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. Some homes are delisted because a sale agreement is reached, others are delisted because the seller has decided to no longer market the home for sale. Considering listed homes for sale in a survival analysis framework, homes that delist without selling are implicitly treated as censored observations. 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.⁷

⁷Note that all S and P are defined by the date of transaction (which is recorded for all home sales in our data) rather than the date of agreement.

Summary statistics for the estimation sample are presented in Table 1 for the treatment ($Z > 0$) and control groups ($Z=0$) 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

In Figure 1, we plot OLS estimates of the effect of the instrument Z on the probability a home in the estimation sample sells in a given month, for each month 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 correlation between Z and monthly sale probabilities. Following the MIP cut, however, treatment group homes become significantly more likely to sell than control group homes. Through most of 2015, the estimated effect of the instrument is about one percentage point higher than it was 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. In the appendix, we show that this result is robust to the inclusion of a detailed set of control variables.⁸

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

8

smaller multiplier effects (and hence low values of $\hat{\beta}_1$) in hot markets where homes sell quickly. In these markets, homeowners are less concerned about being stuck holding two properties for an extended period, and so are more willing to wait until they have found a new residence to put their current home up for sale.

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. We then re-estimate equation 2 on each of these three groups separately. Results are presented in Table 2. A marginal home sale increases the homeowner's monthly purchase hazard by about 19 percentage points in cold markets, almost double the strength of the effect found in hot markets. 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 to avoid the cost of owning two properties for an extended period.

5 Model

We now develop a simple model of home sales in a housing market with search frictions. In Section 2, we will calibrate the model to match the reduced-form results just described, and in Section 7, we use the calibrated model to estimate the magnitude of the sales volume multiplier effect under different market conditions.

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. 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. For example, some agents will be very motivated to move 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 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.⁹ House prices are exogenous. In the appendix, we also show robustness of our main results to a modified version of the model where house prices depend on the market tightness.

Renters

We refer to agents who are searching the market to buy a home, but do not own a home, as renters.¹⁰ 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. 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

⁹This assumption is almost not binding because the parameter values we estimate imply that the buyer and seller would endogenously agree to a sale. There is one exception where the assumption is binding, which we discuss below.

¹⁰We 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.

because the price is assumed to be exogenous and the same regardless of which types of agents are transacting.¹¹ With probability $1 - q_b(\theta)$, the renter does not match with a seller and remains a renter.

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

$$V^c = u + \beta[(1 - \rho)(1 - \omega)V^c + \rho(1 - \omega)V^m + \omega V^e] \quad (5)$$

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

¹¹Omitting 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.

$$V^m = 0.5772 + \ln[\exp(V^s) + \exp(V^b) + \exp(V^{sb})] \quad (6)$$

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

$$V^s = u - \chi + \beta[q_s(\theta)V^r + (1 - q_s(\theta))V^s] \quad (7)$$

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

$$V^b = u - \chi + \beta[q_b(\theta)V^d + (1 - q_b(\theta))V^b] \quad (8)$$

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 .¹² The value function associated with being a double owner is therefore

¹²Wlog, 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.

$$V^d = u_2 + \beta[q_s(\theta)V^c + (1 - q_s(\theta))V^d] \quad (9)$$

Note that we are assuming for simplicity 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 happen to buy and sell at the same time). The value function associated with being a seller-buyer is

$$\begin{aligned} V^{sb} = & u - \chi + \beta[q_s(\theta)q_b(\theta)V^c + (1 - q_s(\theta))q_b(\theta)V^d + \dots \\ & \dots + q_s(\theta)(1 - q_b(\theta))V^r + (1 - q_s(\theta))(1 - q_b(\theta))V^{sb}] \end{aligned} \quad (10)$$

Here we note that our assumption that all matchings lead to transactions becomes binding. Under our calibrated parameters, a seller-buyer who matches with a seller but not a buyer prefers not to buy and become a double owner. However, allowing seller-buyers to make transaction decisions complicates the model. Therefore, we choose to simplify the model by assuming that all matchings lead to transaction.¹³

Equilibrium and Discussion

An equilibrium in the housing market consists of value functions and a market-tightness θ that satisfies equations (4) through (10). In the calibration of the model, we will focus on the steady state equilibrium. We allow for an inflow of agents into the renter pool to balance out the outflow of agents who receive exit shocks, ω . For example, this inflow could reflect the formation of new households who enter the housing market.

¹³A 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.

The only decision that agents face in our model is whether to search as sellers, buyers, or seller-buyers upon receiving a mismatch shock. This decision is irreversible. Under our type 1 extreme value assumption, the probability of choosing each strategy has the following closed form

$$Pr(i) = \frac{\exp(V^i)}{\exp(V^b) + \exp(V^s) + \exp(V^{sb})}, \text{ for each } i \in \{b, s, sb\} \quad (11)$$

Since the value functions depend on the market tightness, the equilibrium choice probabilities do as well. This feature of the model creates interesting feedback effects between market tightness and choice probabilities. An agent's optimal strategy affects the market tightness, and the optimal strategy depends on the market tightness. For example, consider a buyer's market (low θ) where homes for sale have a low probability of matching. Furthermore, suppose that u_2 – the flow utility of holding two homes – is very low. In such a market, mismatched owners will tend to choose to sell first to avoid a long and costly 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. A given set of parameters could support both a cold market equilibrium (where the market tightness is low and agents choose to sell first) and a hot market equilibrium (where the market tightness is high and agents choose to buy first).

6 Calibration

We assume an urn-ball matching technology, so that

$$q_s(\theta) = 1 - \exp(-A\theta) \quad (12)$$

$$q_b(\theta) = \frac{1 - \exp(-A\theta)}{\theta} \quad (13)$$

where A is a technology parameter that determines the efficiency of the market. We chose this functional form because we found that it provided the best fit to our data, as we describe further below.

The parameters of the model – $u, u_0, u_2, \chi, A, \omega, \gamma, \beta$ – are summarized in Table 3. We normalize $u = 0$ and set $\beta = 0.95^{1/12}$ so that each model period can be thought of as a month. We also set $\omega = \gamma = 0.0035$ implying an expected value of being in the contented state of about 12 years. The assumption that $\omega = \gamma$ implies that the share of sales by exiters is roughly equal to the share of sales by internal movers, consistent with what we observe in our data for both the hot and cold markets.

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 in equations 12 and 13 to take on two different values, A_L and A_H . One interpretation of A is that it measures the fraction of buyers who are suitable matches for a randomly selected home for sale (see Petrongolo and Pissarides (2001)). A could be higher in some markets than others due to factors outside of the model, such as differences in the housing stock or in buyer tastes across markets.

We construct the data moments using the micro data discussed in Section 3. The data moments we use are shown in Table 4 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.¹⁴ 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 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

¹⁴Moen 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 $\theta < 1$.

buys is $q_b(\theta)$ regardless of whether they are able to sell or not. Therefore, we can write

$$\beta_1 = q_b(\theta) \frac{s}{s+d+e+sb} \quad (14)$$

where s, d, e, sb denote the steady state number of agents in the seller, double-owner, exiter, and sell-buyer pools, respectively.¹⁵

6.1 Identification

The parameters of the matching function, A_L and A_H , are largely identified by the probabilities of buying and selling (moments 2 and 4). Note that the value of the market tightness (moment 6) is implied by the probabilities of buying and selling given our matching technology described in equations 12 and 13.¹⁶ The three flow utility parameters, u_2, u_0, χ , are largely identified by the probability of choosing seller, buyer, and seller-buyer. We have six moments (moments 1, 3, 5 in each type of market) related to these choice probabilities to identify these three parameters.

6.2 Results

The parameter estimates in Table 3 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, consistent with our intuition that double ownership and short-term rentership are costly states to be in. Double ownership is costly because agents can typically only live in one house at a time and credit and rental frictions make the carrying costs associated with owning two homes high. Short-term rentership is costly because, put simply, it is inconvenient. The estimate of u_2 is less than u_0 , implying that double ownership is more costly than short-term rentership, all else being equal.

Our estimates of A imply that $|\partial q_b / \partial \theta|$ – the sensitivity of the probability of buying to the market tightness – is low. This implies that the addition of an extra

¹⁵To see this even more clearly, note that $\beta_1 = (q_b(\theta) - 0) \frac{s}{s+d+e+sb} + (0 - 0) \frac{d}{s+d+e+sb} + (0 - 0) \frac{e}{s+d+e+sb} + (q_b(\theta) - q_b(\theta)) \frac{sb}{s+d+e+sb} = q_b(\theta) \frac{s}{s+d+e+sb}$.
¹⁶ $\frac{q_s}{q_b} = \frac{1 - \exp(-A\theta)}{(1 - \exp(-A\theta))/\theta} = \theta$.

buyer to the market does not have a large crowd out effect on the probability that other buyers in the market match with a for-sale home. This result is consistent with Genesove and Han (2012) who also find $|\partial q_b/\partial\theta| < |\partial q_s/\partial\theta|$ using survey data on buyer time-on-market, seller time-on market, and number of homes visited by buyers.¹⁷

Table 4 shows that the model fit is very good. Our urn-ball matching function can fit the buy and sell probabilities exactly in both types of markets. The model does a good job of matching our IV-estimate of β_1 – the fit is almost perfect in the hot market. The model fit is poorest for the share of double owners relative to total sellers in the cold market. The fit of this moment could be improved by increasing u_2 so that internal movers are more likely to become double owners. However, increasing u_2 would also lower the model-implied estimate of the causal effect of selling on buying, which is already slightly below the estimate in the data.

7 Estimates of the Multiplier from Stimulus

Using our parameter estimates recovered in the previous section, we explore how sales volume in our model economy responds to 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 exogenously stimulate demand by permanently increasing the inflow into the renter pool. For example, inflow may increase in response to a first-time home buyer tax credit or a decrease in mortgage rates, but we do not actually model the response of inflow to policy. Our results focus on the size of the total sales volume response relative to the direct sales volume response caused by the inflow of renters, which is the sales volume multiplier from stimulus. The inflow shock changes the equilibrium market tightness and optimal choice probabilities. In the Appendix, we describe how we solve for the equilibrium of the model at each period along the transition path to a new steady state.

We consider a small shock to the inflow. In unreported results, we show that the multiplier is not sensitive to the size of the inflow shock for small values. Focusing on small shocks alleviates concerns about multiple equilibria in our model. We think

¹⁷In our data, 14% and 25% of transactions occur above the list price in the cold and hot market, respectively. If we assume that such transactions proxy for cases where multiple buyers are matched with a single seller, then these data also suggest that buyer crowd out is relatively low.

it is reasonable to assume that for a small policy shock, the housing market does not switch to a new equilibrium with a drastically different market tightness.

The left panel of Figure 2 illustrates the transition dynamics of sales volume for the cold market in the first 100 months following the stimulus.¹⁸ Sales in each period are reported as changes relative to their steady state level prior to the stimulus.

The black line simply shows the permanent impulse to inflow that is the stimulus to the housing market in our simulations, which we set at $1e-7$. The blue line shows that as the number of first-time buyers entering the housing market increases, the number of sales to first-time buyers also increases, and eventually to a level that almost equals the first-time buyer inflow. The response of first-time homebuyer sales is a measure of the direct response to the stimulus. First-time 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 so remain in the buyer pool.¹⁹

The red line shows the response of total sales, which includes first-time buyer sales and all others. The main result from Figure 2 is that the response of total sales significantly exceeds the response of first-time homebuyer sales. The multiplier from stimulus equals:

$$multiplier = \frac{\Delta TotalSales}{\Delta First-timeBuyerSales} \quad (15)$$

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 Figure 2, the multiplier is equal to the area under the red line divided by the area under the blue line. Table 5 shows that the multiplier for the cold market over two years is sizable at 2.48. Each first-time homebuyer sale generated by the stimulus leads to 2.48 total sales, or to an additional 1.48 total sales over and above the sale directly generated by the stimulus.

The right panel of Figure 2 illustrates the transition dynamics for the hot market. Qualitatively, the responses are similar as in the cold market, however the magnitude of the multiplier is much smaller. Table 5 shows that the multiplier for the hot market over two years is 1.48.

¹⁸The transition to the new steady state takes more than 100 months, but we show only the first 100 months in the figure because our focus is on the short-run response to stimulus.

¹⁹First-time homebuyers are a subset of renters. Renters include some previous homeowners who chose to be “sellers” or “seller-buyers” and are not first-time homebuyers.

There are two main mechanisms in the model that generate the multiplier effect. First, the stimulus helps to clear for-sale inventory, allowing some of the sellers of those homes to become buyers themselves, creating an endogenous increase in internal demand. The existence of agents who are waiting to buy until they sell is key for this result. To emphasize this point, Table 5 shows that when when we set u_0 equal to a large negative number, which implies that all mismatched agents choose to buy first, the multiplier estimates are close to one. 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 total sales volume. We call the second mechanism the “switching effect”.²⁰

Both mechanisms contribute to a larger multiplier in a cold market than a hot market. In a cold market, for-sale inventory and the share of mismatched owners choosing sell first is relatively high, and so there is more internal demand for stimulus to unleash. In addition, in a cold market where buyers are relatively scarce, the marginal effect of an additional buyer on sales volume is larger, and so the increase in internal demand from the switching effect increases sales volume to a greater extent.

To gauge the quantitative importance of the two mechanism, Figure 3 plots the transition dynamics assuming that the probability of choosing seller, buyer, and seller-buyer upon mismatch remain fixed at their pre-stimulus steady state levels. This simulation shuts down the switching effect and isolates the effect of releasing pent-up demand of sell-first owners. Comparing Figures 2 and 3, we see that without the switching effect, the response of total sales volume to the stimulus is much lower in the cold market and somewhat lower in the hot market. In both markets, the effect on first-time homebuyer sales is similar to the baseline simulation. Table 5 shows that the multiplier is 1.23 and 1.50 in the hot and cold market, respectively.

These results suggest that the switching effect increases the multiplier effect, and substantially so in the cold market. Even with the choice probabilities fixed, however, the multiplier effects are still sizable and remain larger in the cold market than in the hot market. For a given increase in the number of first-time homebuyer sales, the total sales volume would increase about 22 percent more in cold markets than in hot markets purely through releasing the pent-up demand of sell-first owners. When

²⁰In both the cold and hot market calibrations, the probability of matching as a buyer is much larger than the probability of matching as a seller. Therefore, the switching effect has a positive effect on total sales volume under both hot and cold market conditions.

mismatched owners are allowed to change their strategy in response to the demand shock, the difference in overall sales is almost 70 percent.

7.1 Robustness

Endogenous Prices

Our baseline model abstracts from house prices and so house prices do not change in response to stimulus in the simulations just described. Fully endogenizing house prices significantly complicates the model, as shown in Moen et al. (2015).

In the Appendix, we show robustness of our sales volume multipliers to a modified version of our baseline model that allows the house price to vary with market tightness through a reduced-form relationship. In this version of the model, the house price rises as the market tightens following stimulus. A higher house price lowers the sales volume multipliers, but only slightly. The multipliers in the model with endogenous prices are quite similar to the estimates presented in Table 5.

Price varying with tightness produces slightly smaller estimates of the multiplier because of a discounting effect. When the house price rises, the incentive to sell first increases, as selling first allows a mismatched owner to receive the higher price sooner (and pay the higher price later). The discounting effect counteracts some of the switching effect described above, resulting in a slightly lower multiplier.

Because this exercise suggests that the level of house prices is not as important as time-to-sell and time-to-buy for explaining the optimal transaction sequence for internal movers, we choose to abstract from prices in the baseline model, which keeps the model parsimonious.

Temporary Shocks

Our baseline simulations assume that the stimulus is permanent. However, our model can deliver sizable multipliers from temporary stimulus as well. Table 5 shows estimates of the multiplier when the stimulus is in place for one period and then is immediately withdrawn.²¹ The estimated multiplier in the hot market is similar to

²¹Mechanically, 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. If instead the inflow returns to a level below its pre-stimulus steady state level, the multiplier could be substantially reduced due to a reversal effect.

the baseline. The estimated multiplier in the cold market is lower than in the baseline, but still well above 2.

8 Fiscal Multiplier from Housing Stimulus

In this final section, we use the sales volume multipliers recovered in the previous section to evaluate the fiscal multiplier from housing stimulus. As a case study, we consider the same cut in FHA premiums that we used to calibrate our model. The fiscal multiplier is equal to the total economic activity generated by the premium cut relative to the expenditure (or foregone revenue) by the government. Our calculation of the multiplier is back-of-the-envelope and focuses only on partial-equilibrium, direct, and short-run effects of stimulus.²²

Bhutta and Ringo (2017) estimate that the FHA premium cut caused first time home buyer volumes to increase by about 72,000 total purchases per year. They find little difference in the direct effect of the rate cut across market conditions, so suppose that first time home buying increased by approximately 24,000 per year in both the hottest and coldest thirds of the market. Multiplying these estimates by our short-run sales volume multipliers from Table 5, the effect of the FHA premium cut on total sales volume is 58,000 and 34,000 for the cold and hot market, respectively. Home sales contribute most directly to GDP through Realtor commissions and lender fees, and spending on furniture, home improvement, and related expenditures that typically accompany a home sale. We assume that each sale generates 5.5 percent of the sale price in fee income, and \$5000 in complementary spending (Benmelech et al. (2017)). In the year of the the premium cut, the average sale price associated with homes financed with FHA loans was about \$190k according to HMDA data. Therefore, the premium contributes \$896 million and \$529 million to GDP in cold and hot markets, respectively.

The FHA premium cut cost the government 50 basis points on all inframarginal FHA borrowers. According to HMDA data, about 650,000 FHA loans were originated at an average loan amount of \$190k in the year of the premium cut. Averaging the lost revenue across the different market types, the premium cut reduced the government's

²²For example, we do not consider the effects of the FHA premium cut on refinancing or substitution into FHA loans from other types of loans. As our model abstracts from changes in house prices in response to the stimulus, our fiscal multiplier calculation does not account for any changes in spending due to changes in house prices.

revenue by \$206 million in each market ($217,000 \times 190,000 \times 0.005$). Our calculations imply sizable fiscal multipliers of \$4.35 and \$2.56 per dollar of government spending in the cold markets and hot market, respectively.²³

Over a longer time horizon, the revenue costs of the rate cut increase, however. By lowering its premiums, the FHA commits to reduced income over the life of the loan. The median FHA loan defaults or prepays approximately 7 years after origination (see data presented in Castelli et al. (2014)), so the short term bump in expenditures should be weighed against the foregone revenue over this extended period.

Yet a further consideration is the direct stimulative effect of the reduced payments on the consumption of inframarginal FHA borrowers. For them, the reduced premiums are functionally equivalent to a tax credit. Estimates in the literature of the marginal propensity to consume (MPC) vary (see, for example, Shapiro and Slemrod (2003), Johnson et al. (2006), Agarwal et al. (2007), Parker et al. (2013) and Jappelli and Pistaferri (2014)), but many of these studies find an MPC of 50 percent or more. Applying a 5 percent annual discount rate over 7 years, the net present value of the foregone revenue from the premium cut is about \$1.25 billion in each market. At an MPC of 50 percent, the net present value of the additional consumption is \$675 million. All told, the fiscal multiplier reduces to \$1.26 in cold markets and \$0.96 in hot markets in net present value terms.

Using estimates of GDP per home sale that are similar to ours, Best and Kleven (2017) estimate a multiplier of around 1 in response to a transaction tax cut in the UK and Berger et al. (2016) estimate a multiplier of less than one-half in response to the first-time home buyer tax credit in the U.S. One potential explanation for our somewhat larger estimate is that Best and Kleven (2017) and Berger et al. (2016) estimate the response of total sales volume to stimulus directly from the data using treatment and control groups. As Best and Kleven (2017) and Berger et al. (2016) acknowledge, their design may understate the effects of stimulus because some homeowners who sell their home in a treatment area in response to the stimulus also buy a home in the control area. Such spillovers between treatment and control groups will tend to attenuate their estimates of the treatment effect, as the indirect effect of the stimulus bleeds into the control group. In contrast, we estimate the multiplier effect

²³This calculation ignores the increased revenue from the marginal FHA borrowers, as well as the off-balance sheet costs of any future insurance claims on those marginal loans. If FHA insurance pricing was actuarially fair, these two factors should offset.

using a calibrated model, and apply that to estimates of the direct effect.

The simple calculations presented in this section do not take into account any effects on house prices and the resultant wealth effects on consumption, nor the potential for additional productivity by allowing households to better sort into their optimal labor market. Accounting for these effects would likely increase our multiplier estimates.

9 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 more muted. Overall, the takeaway is that housing demand shocks can have large effects on sales volume and economic activity through multiplier effects, and so considering only the direct effect of stimulus policies on home buying misses much of the economic response.

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. Finally, our results suggest that policies that reduce housing demand – for example, increases in FHA insurance premiums or guarantee fees charged by the GSEs – can have sizable negative effects on sales volume, as the mechanism that generates propagation in our model responds symmetrically to positive and negative demand shocks.

References

Abel, Joshua, “Why do Housing Markets Freeze? Investigating House Lock, Nominal Loss Aversion, and Irreversible Construction,” *mimeo Harvard University*, 2018.

- Agarwal, Sumit, Chunlin Liu, and Nicholas S Souleles**, “The reaction of consumer spending and debt to tax rebates: evidence from consumer credit data,” *Journal of political Economy*, 2007, 115 (6), 986–1019.
- Anenberg, Elliot and Edward Kung**, “Interest Rates and Housing Market Dynamics in a Housing Search Model,” *working paper*, 2018.
- **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.
- Brown, Jennifer and David A Matsa**, “Locked in by leverage: Job search during the housing crisis,” Technical Report, National Bureau of Economic Research 2016.
- Castelli, Francesca, Damien Moore, Gabriel Ehrlich, and Jeffrey Perry**, “Modeling the Budgetary Costs of FHA’s Single Family Mortgage Insurance Program,” *Presentation. Retrieved from Congressional Budget Office website <https://www.cbo.gov/publication/45730>*, 2014.
- 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.
- Diaz, Antonia and Belen Jerez**, “House Prices, Sales, and Time on the Market: A Search-Theoretic Framework,” *International Economic Review*, 2013, 54 (3), 837–872.

- Genesove, David and Lu Han**, “Search and matching in the housing market,” *Journal of urban economics*, 2012, 72 (1), 31–45.
- Jappelli, Tullio and Luigi Pistaferri**, “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics*, October 2014, 6 (4), 107–36.
- Johnson, David S, Jonathan A Parker, and Nicholas S Souleles**, “Household expenditure and the income tax rebates of 2001,” *American Economic Review*, 2006, 96 (5), 1589–1610.
- Karahan, Fatih and Serena Rhee**, “Geographic reallocation and unemployment during the Great Recession: The role of the housing bust,” *Journal of Economic Dynamics and Control*, 2019, 100, 47 – 69.
- Mian, Atif and Amir Sufi**, “The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program,” *The Quarterly journal of economics*, 2012, 127 (3), 1107–1142.
- 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.
- Parker, Jonathan A, Nicholas S Souleles, David S Johnson, and Robert McClelland**, “Consumer spending and the economic stimulus payments of 2008,” *American Economic Review*, 2013, 103 (6), 2530–53.
- Petrongolo, Barbara and Christopher A Pissarides**, “Looking into the black box: A survey of the matching function,” *Journal of Economic literature*, 2001, 39 (2), 390–431.
- Ramey, Valerie A.**, “Ten Years after the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?,” *Journal of Economic Perspectives*, May 2019, 33 (2), 89–114.
- Rosenthal, Leslie**, “Chain-formation in the Owner-Occupied Housing Market,” *The Economic Journal*, 1997, 107 (441), 475–488.

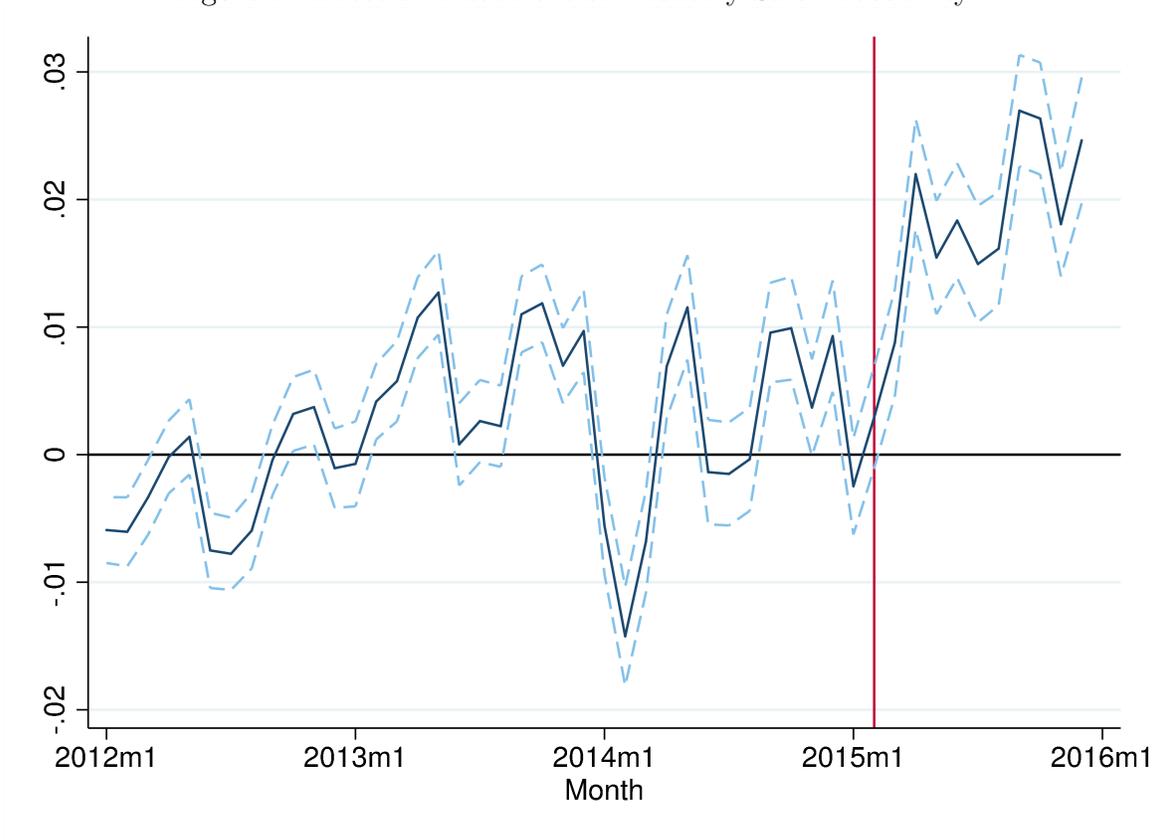
Shapiro, Matthew D and Joel Slemrod, “Consumer response to tax rebates,” *American Economic Review*, 2003, *93* (1), 381–396.

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.

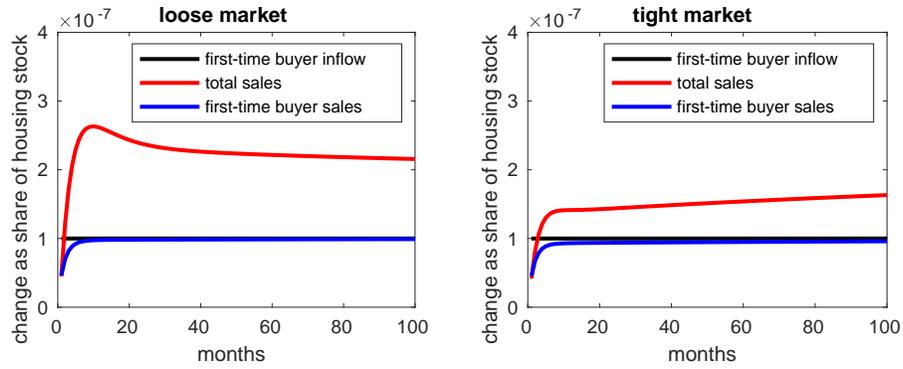
White, Harrison C, “Multipliers, vacancy chains, and filtering in housing,” *Journal of the American institute of planners*, 1971, *37* (2), 88–94.

Figure 1: Effect of Treatment on Monthly Sale Probability



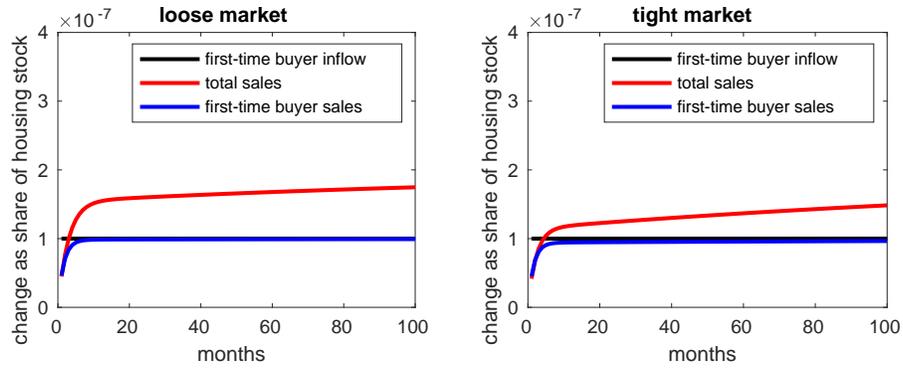
Note: Figure shows the estimated effect, by month, of the instrument Z on the probability a home listed for sale closes in that month. Treatment group sales in February 2015 and later are potentially affected by the reduction in FHA insurance premiums. Point estimates and the 95 percent confidence interval,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

Figure 2: Sales volume response to a demand shock



At time 0, stimulus is introduced by increasing the first-time homebuyer inflow by the amount shown in the black line. 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: Sales volume response to a demand shock, no change in choice probabilities after stimulus



At time 0, stimulus is introduced by increasing the first-time homebuyer inflow by the amount shown in the black line. After the stimulus, all agents continue to make decisions using the pre-stimulus optimal policy functions so that there is no change in the probability of choosing seller, buyer, or seller-buyer. 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.

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 | | |
|-------------------|---------------------|---------------------|--------------------|---------------------|
| | | All Markets | Cold | Hot |
| Sold | 0.041** (0.0002) | 0.117** (0.022) | 0.192** (0.060) | 0.115** (0.019) |
| Z | | 0.002** (0.0003) | -0.0004 (0.001) | 0.008** (0.0005) |
| Post January 2015 | | 0.005** (0.0005) | 0.0025* (0.001) | 0.005** (0.0005) |
| $N \cdot T$ | | 16,778,818 | 6,789,714 | 4,256,171 |
| F-stat | | 597.90 | 103.38 | 708.52 |

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

**p < 0.01

*p < 0.05

Table 3: Parameter Estimates

| parameter | Description | Value |
|-----------|-----------------------------------|---------|
| u | contented flow utility | 0 |
| u_0 | renter flow utility | -0.1480 |
| u_2 | double owner flow utility | -0.3788 |
| χ | mismatch flow utility penalty | 0.0965 |
| A_L | matching efficiency, loose market | 0.5100 |
| A_H | matching efficiency, tight market | 0.5700 |
| ρ | probability of mismatch | 0.0035 |
| ω | probability of death | 0.0035 |
| β | monthly discount factor | 0.9957 |

Table 4: Model Fit

| Moment | Description | Tight Market | | Loose Market | |
|-------------------------------------|------------------------------------|--------------|--------|--------------|--------|
| | | Data | Model | Data | Model |
| 1. $q_b(\theta) \frac{s}{s+d+e+sb}$ | causal effect of selling on buying | 0.1160 | 0.1165 | 0.1930 | 0.1788 |
| 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 | 0.1895 | 0.22 | 0.1073 |
| 4. $q_b(\theta)$ | buy probability | 0.49 | 0.4893 | 0.48 | 0.4788 |
| 5. $Pr(b)$ | probability of searching as buyer | 0.16 | 0.1891 | 0.12 | 0.0861 |
| 6. θ | market tightness | 0.55 | 0.5500 | 0.25 | 0.2500 |

Table 5: Sales Volume Multiplier Estimates from Stimulus

| Assumptions | Cold Market | Hot Market |
|--|-------------|------------|
| Baseline model | 2.48 | 1.48 |
| All mismatched agents choose buy first | 1.04 | 1.00 |
| No change in choice probabilities following stimulus | 1.50 | 1.23 |

Model implied multiplier estimates. The multiplier is $\frac{\Delta TotalSales}{\Delta First-timeBuyerSales}$ 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.

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.²⁴ 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.

²⁴Inspecting 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 6. The estimated effect with the additional controls is very similar to that using our main specification.

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. Results are presented in Table 6. Results are quite similar to the main estimation sample. 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 7. 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 7 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.

B.4 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. We adjust the value functions to account for this transfer. 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 presented in Table 4. 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.²⁵

²⁵In the simulations, we assume that the price level immediately adjusts to its new steady-state

Table 8 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 conservatively assume that the elasticity of house prices is equal to the elasticity of sale probability, Table 8 shows that stimulus still leads to large sales volume multipliers of 2.33 and 1.42 in the cold and hot markets, respectively – only slightly less than our baseline estimates.

level after the stimulus is imposed. When we alternatively allow for the price to gradually adjust to its steady state level along with the gradual adjustment in the market tightness and sale probability, we find slightly stronger sales volume multiplier estimates, as gradual price increases increase the incentive to buy first.

Table 6: Effect of Home Sale on Owner’s Monthly Purchase Hazard, Robustness Checks

| | Main Specification | Additional Controls | Sample with Unique Names |
|-------------|--------------------|---------------------|--------------------------|
| Sold | 0.117** (0.022) | 0.121** (0.021) | 0.080** (0.025) |
| $N \cdot T$ | 16,778,818 | 16,765,134 | 12,459,383 |
| F-stat | 597.90 | 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. Standard errors adjusted for clustering at the census tract level.

** $p < 0.01$

* $p < 0.05$

Table 7: 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) | Note: |
| $Post_t$ | 0.007** (0.0001) | 0.025** (0.002) | 0.003 (0.002) | |
| $N \cdot T$ | 16,804,476 | 563,836 | 158,207 | |

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

Table 8: Sales Volume Multiplier Estimates from Stimulus, Endogenous Prices

| Assumptions | Cold Market | Hot Market |
|---|-------------|------------|
| $\frac{\partial \ln p}{\partial \ln \theta} = 0$ | 2.44 | 1.47 |
| $\frac{\partial \ln p}{\partial \ln \theta} = 0.5 * \frac{\partial \ln q_s}{\partial \ln \theta}$ | 2.38 | 1.45 |
| $\frac{\partial \ln p}{\partial \ln \theta} = \frac{\partial \ln q_s}{\partial \ln \theta}$ | 2.33 | 1.42 |
| $\frac{\partial \ln p}{\partial \ln \theta} = 2 * \frac{\partial \ln q_s}{\partial \ln \theta}$ | 2.22 | 1.38 |

Model implied multiplier estimates. The multiplier is $\frac{\Delta TotalSales}{\Delta First-timeBuyerSales}$ 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.

C Model Details

C.1 Details on Model Calibration

We first note that the steady state market tightness can be inferred from the data for each type of market, as shown in Table 4. Denote this tightness as θ . For a guess of the parameters values, we first iterate on the following loop until convergence

1. Compute V^s under θ using (7)
2. Compute V^b under θ using (8)
3. Compute V^d under θ using (9)
4. Compute V^{sb} under θ using (10)
5. Compute V^r under θ using (4)
6. Compute V^c under θ using (5)

After convergence, 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:

$$b' = (1 - q_b(\theta))b + \rho(1 - \omega)c \frac{\exp(V^b)}{\exp(V^b) + \exp(V^s) + \exp(V^{sb})} \quad (16)$$

$$d' = (1 - q_s(\theta))d + q_b(\theta)b + q_b(\theta)(1 - q_s(\theta))sb \quad (17)$$

$$s' = (1 - q_s(\theta))s + \rho(1 - \omega)c \frac{\exp(V^s)}{\exp(V^b) + \exp(V^s) + \exp(V^{sb})} \quad (18)$$

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

$$e' = (1 - q_s(\theta))e + \omega c \quad (20)$$

$$c' = 1 - b - s - sb - 2d - e \quad (21)$$

where c denotes the mass of contented owners. Once the pool sizes converged, use the pool sizes and value functions to compute the moments shown in Table 4. Evaluate the objective function and repeat until parameter values are found that minimize the objective function. Once the parameter values have been found, we can easily solve for the steady state inflow into the renter pool that rationalizes θ as an equilibrium.

C.2 Details on Model Simulation

To solve for the transition path to the new steady state following the stimulus shock, we follow the following steps. First, we solve for the new, post-stimulus steady state. To do this, we guess at the steady state θ , compute the value functions at the guess of θ , solve for the steady state θ implied by the value functions, and iterate on θ 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 16-21.
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.