The Existing Home Sales Volatility Puzzle: An Irreversible Construction Explanation*

Joshua Abel†

August 18, 2018

Abstract

The recent housing cycle in the United States saw a large swing not only in home prices but in the number of home sales as well. This paper begins by using a comprehensive dataset on US home sales to investigate two popular explanations for the cyclicality of selling activity: "house lock," which conjectures that falling prices cause down-payment constraints to bind and prevent current homeowners from selling their homes; and nominal loss aversion, which proposes that cognitive frictions prevent homeowners from selling when doing so would not garner a price as high as the one they originally paid for the house. I find that while there is evidence that both of these mechanisms are active at the household level, they explain a fairly small portion of the decline in sales from boom to bust: likely no more than 10%. I then propose a novel mechanism, which is that construction of New homes, which tend to be of high quality, unlocks sales of Existing homes during booms, as there is aggregate movement up the housing ladder. In the bust, this movement freezes, and it does not reverse, as the irreversibility of construction prevents the market from tearing down nice homes and facilitating an aggregate down-size. As a result, sales are high in the boom and low in the bust. I show that the model's predictions are consistent with recent dynamics of aggregate prices and volume, as well as cross-MSA variation in sales. Overall, the model is able to explain up to 30% of the aggregate movements in sales over the previous cycle. I conclude by discussing factors that could increase or decrease this magnitude.

*I give special thanks to my dissertation committee: Edward Glaeser, Paul Willen, Gabriel Chodorow-Reich, and John Campbell. I am also grateful for helpful conversations with Andrew Abel, Ben Austin, Valentin Bolotnyy, Maxim Boycko, John Coglianese, Daniel Cooper, Oren Danielli, Ellora Derenoncourt, Natalia Emanuel, Emmanuel Farhi, Chris Foote, Andreas Fuster, Andrew Garin, Adam Guren, Michael Kincaid, Claire Labonne, David Laibson, Eben Lazarus, Lara Loewenstein, Kurt Lunsford, Libby Mishkin, Ben Moll, Michael Reher, Stuart Rosenthal, Andreas Schaab, Jeremy Stein, Linh To, and Eric Zwick, as well as comments received at Harvard's Macro and Labor/PF lunches. The views expressed in this paper are solely those of the author and not necessarily those of the Federal Reserve Bank of Boston or the Federal Reserve System.

†Harvard University, Department of Economics; Federal Reserve Bank of Boston, Research Department. E-mail: jabel@fas.harvard.edu

1 Introduction

The surge and collapse in home prices in the recent housing cycle coincided with a dramatic rise and fall in the number of homes sold.[1] This dramatic swing—along with the financial panic and global recession it ignited—has renewed interest for researchers and housing market participants into why the volume of homes sold is so volatile and, in particular, why it seems to be so strongly associated with home prices.

In this paper, I first use a comprehensive dataset on US home sales to empirically evaluate the contributions of down-payment constraints and nominal loss aversion to these gyrations in home sales. I find that they explain only a very small share of the variation. I then propose a novel mechanism to explain the volatility of home sales. The key insight is that an asymmetry exists between booms and busts in how households navigate the housing quality ladder. New construction plays a central role. Suppose there is a broad-based increase in the willingness-to-pay for housing, so that many households want to move into higher-quality homes. This can be accommodated by the construction of a new unit with high quality: a fairly wealthy household will move into that unit, freeing up its former home to be sold to a household of lower means. This second household may then sell its former house to a third household, and so on. In principle, then, the construction of one new home can lead to an unbounded number of sales of existing homes. Now consider a period with a broad-based *decline* in the willingness-to-pay for housing. Will the movement up the ladder during the boom be reversed as households move down the ladder in the bust? If housing investment is sufficiently costly to reverse, the answer is "no." Individual households may want to move down in order to substitute away from housing towards other types of consumption, but in aggregate this would require tearing down homes and converting it into that other consumption—and the costly reversibility of housing investment implies that this cannot be done efficiently. So, rather than a large procession of households down the ladder, home prices drop to deter households from trading down, and the movement along the housing ladder freezes.[2] As a result, rising housing valuations coincide with high sales volume as the churn of households up the quality ladder proceeds, while this churn freezes in periods of falling valuations, and housing market activity diminishes.

Understanding movements in home sales provides an interesting challenge for economists. Our models, even the simplest supply-and-demand diagrams taught to students studying the field for the first time, make stark and compelling predictions about prices and quantities. Yet in many foundational models, trading volume is indeterminate, as buying 1 share of an asset is equivalent to buying 2 and selling 1. In these models, trading volume is indeterminate because it is irrelevant. Agents do not care how much paper they shuffled, they only care about their final positions.

[1] See Figure 1.

[2] Note that I am certainly not claiming that households never trade down. Individual households can trade down, but this is coupled with a different household trading up, as in Ortalo-Magne and Rady (2006). The point is that movement down the ladder cannot exceed movement up the ladder—individuals can trade down to lower quality homes but in aggregate the market does not decline in quality. This is in contrast to movement up the ladder, where the construction of new, high-quality homes does allow for improvements to the aggregate housing stock and therefore a net upward movement of households along the quality ladder.

The housing market is an interesting setting to think about sales volume because this irrelevance argument does not hold: a household that sells a home and buys another will end up with the same number of houses as before, but changing their place of residence likely has a great impact on their consumption of housing services and/or local amenities. Far from shuffling paper, most households probably consider the decision of whether to move one of great importance and consideration in their lives.

This paper makes two broad contributions in moving us closer to an understanding of this phenomenon. In the first part of the paper, I empirically evaluate the quantitative power of two popular explanations for the volatility of home sales, "house lock" (or equity constraints) and nominal loss aversion. I show that while the household-level predictions of these models are borne out in the data, they are able to account for only a small share of the aggregate decline in sales from boom to bust. While a household's hazard of sale does depend on its loan-to-value ratio (LTV), as the theory predicts, the housing bust generated only a relatively modest movement in the LTV distribution. Sales fell not because of this movement but because of a downward shift in the hazard function across the LTV support. In total, house lock can account for only about 2% of the decline in sales. A similar approach suggests that nominal loss aversion cannot account for more than 8%. I go on to show that while these factors may cause homeowners to leave their homes on the market somewhat longer once they are listed, as has been emphasized in the previous literature, the speed of sale is a quantitatively trivial determinant of sales volume. So, while it is quite common to see these mechanisms used to explain the decline in sales, the vast majority of the decline is accounted for by homeowners who do not appear to have been facing equity constraints and who had not suffered nominal losses on their housing investments.

In the second part of the paper, I present a dynamic equilibrium model of a housing market with homes of different qualities. Importantly, the model features "filtering"—construction of new homes allows their buyers to sell their old homes to new occupants. This creates a link between construction of New homes and sales of Existing homes that, while long-studied in the affordable housing literature, has not been considered as a driver of time series variation in the turnover of the existing housing stock. I model the housing cycle as a temporary increase in expectations of future housing valuations. This, unsurprisingly, leads to a construction boom, but what's novel is that the model demonstrates how, due to the filtering process, this leads to a boom in sales of existing homes as well. When expectations revert, the building boom disappears and in fact reverses, as an overhang of housing suppresses construction. Once again, due to the filtering process, existing sales follow suit, collapsing below pre-boom levels.

The mechanism leaves a number of footprints in the data. First, I use microdata on home sales to document the existence of sale chains. The key to the mechanism is that people who buy new homes often simultaneously sell their previous homes. I document that this is true for most buyers of new construction. In fact, the chain often extends beyond the sale of one existing home, as the household that fills the home vacated by the buyer of new construction oftentimes sells a previous

home of their own. These sale chains provide some key micro-evidence that the mechanism is operative in the real world.

I also provide evidence of a key cross-MSA empirical relationship predicted by my model that other models of the housing market cannot account for: US cities with high elasticities of housing supply (as estimated by Saiz (2010)) also exhibit higher rates of turnover of their existing housing stocks. As the housing supply elasticity serves as a source of exogenous variation in construction of New homes, I interpret this as a causal effect of new construction on Existing home sales. This is at the heart of my model with filtering. Note that while many explanations of volatility in home sales rely on price growth to generate surges in sales, they make the counterfactual prediction that less elastic cities would have more sales during the boom, as these areas had stronger price growth. By highlighting the dependence of Existing home sales on New home sales, I generate a theory that is consistent with this novel empirical result.

I also show that the pattern of Existing home sales across US cities is consistent with my model in which construction of New homes causes sales of Existing homes. In particular, I find that sales of New and Existing home sales are tightly linked in a panel of MSAs, and if the reader is concerned that this is being driven by an omitted third factor, I use the estimates of housing supply elasticity in Saiz (2010) as an exogenous source of cross-MSA variation in new construction and find that cities where it is easier to build new homes exhibit higher turnover of the existing housing stock. Not only is this a novel empirical result, but many models of home sale volatility have a lot of trouble accounting for it—and in fact, they would predict the opposite relationship. This is because mechanisms that use price growth to generate a surge in sales, then, would predict the opposite effect of housing supply elasticity, as elastic cities experienced less price growth. This empirical test, therefore, provides a novel result that is predicted by my theory but not others.

I then return to the aggregate data for additional validation of the model. A strong positive correlation between home price growth and sales volume has been well-documented and is one of the empirical facts motivating the home sale volatility literature. I show that while the relationship between those two variables was extremely tight in the years leading up to the financial crisis, it has changed substantially since then. In particular, whereas years of high price growth were very reliably periods of high sales volume before 2007, the years since have had consistent, depressed sales volume despite enormous variation in home price growth, from -15% in 2008 to +10% in 2013. My model predicts precisely this change in the relationship. After the decline in housing valuations in the bust, the housing market is left with substantial overhang from the construction boom. As a result, further movements in valuations do not elicit much supply response. Therefore, the filtering process continues to be limited, generating few sales of existing homes, while the price of housing is the main margin of adjustment. So as in the data, the post-bust period can exhibit above-average price growth without a rebound in sales.

Finally, I show that within the framework of the model, the quantitative impact of the mechanism on Existing home sales boils down to the product of three sufficient statistics. The first is how

much new construction moved. Since the argument of the paper is that New homes generate sales of Existing homes through the filtering process, more variation in new construction will lead to a more powerful mechanism. Second, among the people who buy the marginal new homes (those in excess of normal levels of new construction), the share that sells existing homes is crucial, as this tells us to what extent the construction of new homes actually leads to filtering. These first two sufficient statistics are essentially observable in the data. The third sufficient statistic is the share of the above group that would not have moved in the absence of the housing market shock. If all the people who bought (marginal) New homes were going to move anyway and were merely induced to buy New rather than Existing homes, then the mechanism would not impact the sale of Existing homes. On the other hand, if these households were not planning to move but then decided it made sense to do so in order to upgrade, this mechanism can be strong. This final sufficient statistic is difficult to observe in data, so I show that as this share tends toward 1, the mechanism is able to explain about 1/3 of the variation in home sales over the cycle. I then discuss how richer models would likely predict larger impacts.

The previous literature has proposed a variety of explanations to explain volatility in home sales. Many use some kind of friction to generate a relationship between prices and the incentives of home sellers to transact. House lock and nominal loss aversion, as in Stein (1995) and Genesove and Mayer (2001), are two straightforward and popular examples. House lock proposes that financial market frictions make it difficult for low-equity homeowners to sell their homes, so price declines deter borrowers from selling their homes as their equity gets wiped out. Similarly, nominal loss aversion is a *cognitive* friction that makes homeowners uncomfortable selling their homes at prices below those at which they bought them. So as prices fall and nominal gains turn into nominal losses, sales could dry up for this reason. I show that while these mechanisms are appealing at the household level, they fail to account for much of the aggregate decline in sales.

Search frictions have also become extremely popular in housing models (e.g. Krainer (2001), Novy-Marx (2009), and Head et al. (2014)). With search frictions, sellers' payoffs are 2-dimensional: they want to sell at a high price, and they want to sell quickly. Demand shocks then provide a windfall to home sellers, who respond by adjusting reservation prices to generate higher sale prices and quicker selling speeds. As homes spend less time on the market, sales volume increases. These models provide valuable insights into short-run dynamics of housing markets,[3] but as I show below, time on market is a nearly trivial factor in determining aggregate sales volume. If one wants to explain why home sales are so volatile, one must explain why there is so much variation in homeowners' decision to sell in the first place, not the ease with which they do so.

A growing literature has also emphasized the role of investors/speculators in driving sales volume up in the boom. Haughwout et al. (2011) show that in localities that experienced a very dramatic housing cycles, a large share of purchase mortgages during the height of the boom were taken out

[3]For instance, the presence and dynamics of housing vacancies seems to require search frictions, as in Wheaton (1990). Guren (2018) and Head et al. (2014) show how models with search frictions can generate short-term price momentum, an interesting and important feature of housing markets.

by borrowers with first mortgages on other properties, suggesting an important role for investors. Bayer et al. (2017a) and Bayer et al. (2017b) find that there was a substantial rise in participation in the market by amateur speculators in the boom and that they did not perform well, as few seem to have foreseen the coming decline in prices. Mian and Sufi (2018) argue that the acceleration of the private label securitization market at the height of the boom expanded credit and facilitated a large increase in activity from buyers they identify as "flippers," who engaged in a large number of transactions in a short period of time. DeFusco et al. (2018) provide a model of speculative behavior in which the short-term predictability of home prices endogenously entices agents with short expected holding times to enter the market in a boom.[4] Because they do not hold on to the homes for very long, a surge in sales volume results. The authors also provide evidence that there was a large increase in the number of people selling homes within just a few years of purchase during the height of the boom, a key prediction of their model.

The work on speculation is a rich and important strand of the literature, and it seems likely that speculative motives were an important contributor to the variation in sales volume. Nonetheless, these papers focus largely on the peak of the boom, when speculation drove sales volume (and perhaps prices) to new heights, but they have less to say about the longer-term rise in sales that preceded the most frenzied days of the boom, as well as the collapse during the bust that brought sales below pre-boom levels. There is also evidence in Anenberg and Bayer (2016) that an important source of the cyclicality of sales volume comes from households buying and selling (presumably as they move) within a market. These households are unlikely to be speculators, as their transactions do not change their exposure to home price growth. Furthermore, the boom and bust in home sales coincided with a rise and fall in household mobility, as shown in Bachmann and Cooper (2014). Household formation also tends to be pro-cyclical, as shown in Lee and Painter (2013), and it rose and fell in this cycle as well, which also suggests that a portion of sales volatility came from residential decisions of households and not just speculative trading from investors. Additionally, Loewenstein (2018) uses the PSID to show that during the boom, continuing homeowners became more likely to increase their housing consumption by trading up into more expensive homes. This means that despite a potentially important role for investors during this cycle, it is still important to model housing transactions as a means by which households change their consumption of housing services and/or locations, as I do in this paper.

The two theoretical papers to which I am closest are Ortalo-Magne and Rady (2006) and Ngai and Sheedy (2017). Ortalo-Magne and Rady (2006) study a housing market with distinct tiers of housing, as I do, but they abstract from new construction and focus on dynamic steady states of an overlapping generations model. Sales in the model are generated by old households with low housing needs but nice houses trading with young households in the reverse situation. They show that when young households enter the economy with more income, they bid up the price of nicer

[4] Other models of speculation highlight disagreement between agents, with recent examples being Nathanson and Zwick (Forthcoming) and Burnside et al. (2016). The focus of this work is typically in explaining variation in prices moreso than sales volume.

houses, inducing more old homeowners to trade down and cash out the large difference in prices, thus generating more sales. While the model relies on a housing ladder to generate variation in home sales, as mine does, it uses an entirely distinct mechanism. In my model, the increase in home sales depends crucially on construction of new homes, because sales are generated by an aggregate increase in home quality, as individual households churn up the ladder. With their fixed housing stock, Ortalo-Magne and Rady (2006) get variation in sales not from an aggregate upgrading but rather from an increase in mismatch between who owns nicer homes and who wants them. Both phenomena are likely occurring in real housing markets.

Ngai and Sheedy (2017) model moving as a way for a household to invest in better idiosyncratic match utility from its home, which they trade off against a moving cost. They show that if there is a reduction in the (effective) moving cost—perhaps from increasing wages or the introduction of internet-based housing search—more household will bear that cost and move. While this has parallels to the mechanism I explore, it is distinct. As in Ortalo-Magne and Rady (2006), their model has a fixed housing stock, but there is variation over time in the return to moving. My model below features both idiosyncratic match utilities between agents and homes, as in their paper, *and* objective quality differences in homes. In the baseline model, the incentive to move *within* a quality to find a better idiosyncratic match is constant over time,[5] and the movement in home sales comes from a common increase in the taste for quality. For this to result in an increase in home sales, new construction is necessary, as it facilitates the aggregate upgrading that is required for households to move up the ladder. So again, my paper should be viewed as distinct from and quite complementary to Ngai and Sheedy (2017), as with Ortalo-Magne and Rady (2006).

The paper proceeds as follows. In Section 2 I describe my main data source. Section 3 uses that data to evaluate the models of house lock and nominal loss aversion described above. Section 4 presents a model of a housing ladder with new construction and explores its implications, while Section 5 shows evidence from the data supporting its relevance. In Section 6, I discuss the quantitative implications of the model, and I conclude in Section 7.

2 Deeds Records

My primary dataset comprises property-level public records on home sales and mortgages. This data is provided by CoreLogic, which scraped electronic records from county registers of deeds from across the country. The data covers the years 2000-14, and its geographic coverage is strong and improves over the sample period. In 2000, about 50% of U.S. counties are included, covering over 85% of the U.S. population. By 2004, counties with around 90% of the population are covered, and by the end of the sample, counties including over 99% of the population are included.

When a home is sold, the address of the property being transferred is observed, as well as the names of (up to two) buyers and sellers involved in the transaction, the price and date of the

[5] I show an extension that allows for market thickness effects, which increase the return from moving in periods when a high number of other households are also moving.

transaction, and the terms of any mortgage taken out by the buyer. In addition to transactions, the data records mortgage refinances, when mortgage terms change but there is no transfer of ownership.

2.1 Primary Estimation Sample

My main empirical results come from a property-by-month panel constructed from a 1% random sample of the dataset described above. Specifically, I randomly select 1% of properties in the CoreLogic dataset. Next, I drop sales that are classified by CoreLogic as nominal (or "non-arm's-length") transactions, which is about 1/3 of all sale records. I then generate an observation for each property for each month between January 2000 and December 2014. From this panel of properties, I generate "ownership spells," which begin with the sale of a property and end with its subsequent sale. Foreclosures require some further care. When a lender seizes a house from a homeowner, a record is generated in the deeds data, but I do not count this as a sale. Rather, the sale occurs when the lender (or a subsequent lender to whom the property has been transferred) makes an arm's-length sale sale of the property.

One final sample selection criterion is important to discuss. In order to compute a homeowner's LTV, I need to know the value of the house, which requires knowing the previous sale price. This means that, while the dataset contains nearly the universe of transactions, the panel used in estimation will not contain the universe of properties. For instance, if a property sells for the first time in, say, June 2004, it cannot be included in the sample until July 2004, because its LTV in the months January 2000-June 2004 is unknown. This means the sample size will increase over time. As a result, I will not be discussing the *count* of sales but rather the *hazard* of a sale occurring. Furthermore, I begin the analysis in 2004 rather than 2000, as this allows the full sample to have a more representative set of properties,[6] both because of the issue detailed in this paragraph as well as the improved coverage of the dataset over time as discussed earlier.

Figure 2 shows that the remaining sample, while not fully representative due to the necessary sample restriction discussed above,[7] still exhibits a collapse in sales during the housing bust—and slow recovery thereafter—that is extremely similar to the aggregate data provided by the National Association of Realtors. This offers reassurance that subsequent analysis of this sample is likely to apply to the broader population of properties.

All told, this sample has 26,498,330 property-month observations spanning 265,832 properties

[6]For instance, in the early years of the sample, the LTV distribution is very tightly packed around 80% and 100%, since the sample is comprised of people who have only very recently purchased their homes and who therefore have not experienced home movement in local home prices (or amortized much of their mortgage debt). By waiting a handful of years, I allow the sample to look more like the population: as discussed below, the LTV distribution in my sample in later years compares favorably with the distribution calculated from a separate dataset that does not have this same sample selection issue.

[7]It is also not fully representative because, as discussed, some counties are omitted. Furthermore, there are instances when a sale occurs but a price is not listed. LTV cannot be computed in these cases, so they are also not included in the estimation.

and 392,074 ownership spells. There are 127,334 sales in this sample,[8] generating a monthly hazard of 48bp. Figure 3 shows that, during the crash, there was a very large increase in the share of these sales that were "distressed." In 2004, over 95% percent of sales were non-distressed—sales by homeowners who were able to pay back the mortgage in full. In the height of the crisis, nearly 50% of sales were either short sales or sales[9] out of foreclosure. While this number declined throughout the recovery, it was still close to 20% in 2014, well above the historical norm.

2.1.1 Computing LTV

A homeowner's LTV is a key variable of interest in Section 3, so I now describe in detail how this measure is calculated.

The numerator is the mortgage debt being collateralized by the house. As discussed above, the deeds records provide information about mortgages taken out when a home is sold, as well as subsequent refinances. This means that at the time of purchase, the loan balance is directly observed (inclusive of "piggyback" second liens). In subsequent months when the property is not observed in the data, I amortize the loan balance according to the interest rate and mortgage term listed,[10] assuming that the payments are being made.

Some ambiguity arises when a household refinances. While the balance and terms of the new loan are listed, there is not explicit information about what has happened with the previous loan. Specifically, it is not known whether the new loan is a replacement of the prior loan or an additional loan on top of it. In the results that follow, I assume that if the new loan has a balance that is less than 25% of the outstanding balance on the old loan, then it is a cashout refinance and I add the two balances together. If it is more than 25% of the old loan's outstanding balance, I assume it is a replacement, so that the new level of total mortgage debt is just the balance of the new loan.[11]

The denominator of the LTV is the market price of the home. This is of course only observed when the home sells. To get an estimate of the LTV in months between sales, I update the sale price using local[12] home price appreciation, as measured by CoreLogic's repeat sale home price index, which is standard practice in this literature. Specifically, if home i in geography g sells in months t_0 and $t_1 > t_0$, I compute the market price in month $t_0 < t \leq t_1$ as:

[8] The number of sales is nearly equal to the difference between the number of ownership spells and the number of properties because each sale begins a new ownership spell. There are 1,092 "extra" sales because this is precisely the number that occurred in December 2014, the last month of the sample. As a result, these sales do not begin new ownership spells.

[9] Short sales refer to cases where the proceeds from selling the home are insufficient to pay off the mortgage, but the lender agrees to allow the borrower to sell the house and not pay back the entire loan, rather than go through the foreclosure process.

[10] If the term is missing, I assume it is a 30-year mortgage. If the interest rate is missing, I assume it is the interest rate reported by Fannie Mae as the average contract rate on a conforming loan for that month.

[11] I have run the analysis using cutoffs of 20% and 0% and while the shape of the LTV distribution is somewhat sensitive to this, the results later in the paper are hardly affected at all.

[12] For properties in ZIP codes covered by CoreLogic's index, I use ZIP code-level appreciation. If that is not possible, I use county-level appreciation, followed by MSA and finally state.

$$\hat{P}_{i,t} = P_{i,t_0} \frac{HPI_{g,t}}{HPI_{g,t_0}}. \tag{1}$$

Figure 4 shows the LTV distribution in this sample for selected years (excluding the roughly 30% of properties that have no mortgage at any given time). I discuss the evolution of the LTV distribution in much greater detail below, but for now I note that it compares very favorably with Figure 5, which shows the LTV distribution computed in a separate dataset.[13] Specifically, Figure 5 comes from a 0.1% random sample of the Credit Risk Insights Servicing McDash (CRISM) dataset, which is a monthly mortgage loan performance dataset (from Lender Processing Services) with borrowers' credit records (from Equifax) merged in. While the price has to be imputed in a similar way (using local price indexes), the mortgage balance is directly observed each month, precluding any need to impute it based on the loan's terms. Furthermore, the credit bureau data lists the borrower's outstanding debt on home equity lines of credit (HELOC) and closed-end second mortgage balances.[14] That the LTV distributions look so similar when computed in these two very different datasets is important validation for the construction of this key measure.[15]

2.1.2 Computing Nominal Gains/Losses

Another important measure for my analysis is nominal gains/losses that a homeowner faces if she sells her home. I will focus only on nominal gains/losses coming from local home price appreciation. If home i in geography g sold in month t_0 and sells again in month t_1, then for $t_0 < t \leq t_1$:

$$NomGain_{i,t} = \frac{HPI_{g,t}}{HPI_{g,t_0}} - 1. \tag{2}$$

So a homeowner faces nominal losses if and only if local home prices have declined since the time of purchase. This assumes that when the household initially bought the home, the price did not include any discount or premium. Some papers studying nominal loss aversion (e.g. Genesove and Mayer (2001), Anenberg (2011)) relax this assumption and use home characteristics to estimate a "fair" market price using a hedonic model. However, given the time period I am studying, the

[13] This alternative dataset, CRISM, is not available until the middle of 2005, which is why I do not show a comparison for 2004.

[14] Note that CRISM is not perfect for this exercise either. The reason is that the second mortgage balances are given at the borrower level, so for borrowers with multiple properties, I could be including second mortgage debt in the LTV when it really is collateralized by a different property. However, this problem is quite distinct from the issues with the deeds data discussed in the text, so the corroboration of the two datasets is still quite reassuring.

[15] Two differences between the datasets are apparent. First, because CRISM is based on mortgage data, it does not contain the 30% of properties that do not have mortgages. Therefore, the densities are a bit higher in general in the CRISM data. Secondly, the deeds data shows much tighter bunching around notable cutoffs like 80%. This is because the measure of value being used in the deeds data is the price, whereas the available measure in CRISM is the appraisal. Since many borrowers take out loans for exactly 80% of the purchase price, there is sharp bunching in the deeds records. The bunching is mitigated somewhat in CRISM to the extent that appraisals and sale prices differ.

preponderance of variation in nominal gains/losses will come from market-level movements in prices, so I choose to focus on market-level variation, as in Engelhardt (2003) and Bracke and Tenreyro (2018).

Figure 6 shows the distribution of this variable in my sample for selected years. I will discuss this in greater detail below, but the reader can immediately note that, as expected, this distribution deteriorates badly as the price decline unfolds.

2.2 100% Seattle Sample

In addition to the results from the 1% national sample, I provide estimates coming from the universe of transactions in the Seattle-Tacoma-Bellevue MSA. Results from this sample mostly serve as robustness checks for the main results—showing that what's true across MSAs is also true within a particular MSA—and so will be largely confined to Appendix A.1. However, some of the analysis below will leverage Multiple Listing Service (MLS) data on home listings, which is linked to the deeds data at the property level. The coverage and quality of that data is heterogeneous across MSAs, and the data quality in Seattle is very high, so it is a sensible place to focus in on. Furthermore, the size of the price boom and bust in Seattle was roughly equal to the national experience, further bolstering its appropriateness for closer inspection.

The Seattle sample has 47,518,278 property-month observations from 446,286 properties and 662,170 ownership spells. There were 217,698 sales. Figure A-1 shows the evolution of the sale hazard in the Seattle sample, which is quite similar to the national series, while Figures A-2 and A-3 show the distributions in selected years of LTV and nominal gains/losses, respectively. Comparing these with the national distributions in Figures 4 and 6, we see some of the geographic heterogeneity of the housing cycle. Because prices continued to grow in Seattle for an extra year (compared to national measures of prices), the 2009 distributions of both LTV and nominal gains/losses are fatter in both tails than in 2004, as households who bought early in the sample period got to experience an additional year of appreciation rather than depreciation.

3 The Role of House Lock and Nominal Loss Aversion

This section evaluates two common explanations for why falling house prices coincide with declining housing market turnover: house lock and nominal loss aversion. House lock attributes the decline in sales during times of declining home prices to the fact that existing homeowners lose their home equity and so face frictions when trying to sell. In a similar way, if households are averse to realizing nominal losses, then a price decline may decrease the volume of home sales because homeowners do not want to realize a loss on their investments. I will confirm that these descriptions of household decision-making are borne out in the data, but I will also show that they are able to explain only a very small portion of the aggregate decline in home sales from boom to bust.

House lock was first proposed as a determinant of home sales by Stein (1995). Genesove and

Mayer (1997) then used data on listings of condos in the Boston area to show that home sellers with high LTVs set higher list prices, wait longer to sell their homes, and receive higher prices. Anenberg (2011), in a more recent study of the San Francisco area, also finds that high LTVs are associated with higher sale prices. The evidence on the effect on sales has been more mixed. Ferreira et al. (2010) and Ferreira et al. (2012) use the American Housing Survey (AHS) to argue that mobility is reduced by negative equity, though Schulhofer-Wohl (2012) casts doubt on these findings. Engelhardt (2003) has trouble finding an effect on mobility in the National Longitudinal Survey of Youth (NLSY). More recent work focusing on home sales and using approaches similar to what I will use below do find an effect in Florida (Andersson and Mayock (2014)) and the United Kingdom (Bracke and Tenreyro (2018)).

The seminal paper on on nominal loss aversion in the housing market was Genesove and Mayer (2001), which found that home sellers facing nominal losses set higher list prices, though the evidence on whether they received higher sale prices or waited longer to sell their homes was mixed. Since then, Anenberg (2011) confirmed in his San Francisco sample that nominal losses do seem to lead to higher sale prices, and Engelhardt (2003) and Bracke and Tenreyro (2018) have argued that nominal losses reduce measures of turnover (mobility in the NLSY in the former, sales in the UK in the latter).

I contribute to this literature by using a dataset that comprises a large share of the universe of US home sales (and corresponding mortgage information) during a period of tremendous interest in the housing market. The sheer size of the dataset allows me to estimate models that are quite a bit more flexible than what is standard in the literature. Most importantly, I am able to use my estimates to make credible statements about the impact of these mechanisms on the aggregate time series of home sales.

3.1 House Lock

I provide a model of house lock in Appendix A.2, which is an extension of Stein (1995), but here I briefly summarize its intuition and predictions, which are fairly straightforward. Prospective home buyers must make a down-payment in order to purchase a house. For current homeowners, an important source of cash for this down-payment is often the equity in their current home. Therefore, homeowners with low equity (high LTV) will not be able to recover a large down-payment by selling their current house, meaning that moving could require a substantial down-sizing. As a result, the model predicts that there is some $L\tilde{T}V$ above which the homeowner will not be able to put together a down-payment for her optimal house. As LTV begins to rise above $L\tilde{T}V$, the model then predicts that the likelihood of sale will fall, as the household chooses not to move rather than move into a small house. Therefore, when home prices fall, the loss in home equity among existing homeowners could decrease their incentive to move, causing a decline in overall sale activity.

There is a second potential effect of mortgage debt, however, resulting from homeowners' option to default. Once a homeowner has $LTV > 1$, selling the home no longer raises any money, so further

increases in the LTV do not have any impact on the down-payment incentive described above.[16] Instead, further increases in LTV heighten the household's incentive to default on their mortgage and shed this debt.[17] As a result, once LTV rises above 1, the model predicts that the likelihood of a home selling will begin increasing, as households default and lenders sell off their homes. This default channel suggests that if prices decline sufficiently and LTVs increase dramatically enough, we could in fact see an *increase* in sales.

In summary, the house lock model has stark predictions about household behavior. In particular, there should be some region where $LTV < 1$ and the probability of a home sale is a decreasing function of LTV. Once LTV exceeds 1, the probability of sale should increase. The ability of this model to explain the aggregate movement in sales volume will then depend on the strength of these effects and, crucially, how the LTV distribution evolved over time.

3.1.1 Household Behavior

My workhorse econometric framework is a linear probability model.[18] I estimate the equation:

$$Sale_{i,t} = \mathbf{LTV}_{i,t} \cdot \gamma_1 + \mathbf{X}_{i,t} \cdot \gamma_2 + u_{i,t}, \tag{3}$$

where $Sale_{i,t}$ is an indicator function for whether property i was sold in month t; $\mathbf{LTV}_{i,t}$ is a set of K indicator functions for whether property i has an LTV in month t in one of K regions of the LTV support; $\mathbf{X}_{i,t}$ is a set of controls that varies by specification, as described below; and $u_{i,t}$ is a residual. In particular, I partition the LTV support into 24 bins,[19] with the partition being particularly fine in the range $LTV \in [0.7, 1.1]$, as this region is both well-populated throughout the sample period and likely to be the region where the down-payment and mortgage default incentives described above are quite active. This specification allows the effect of LTV to vary arbitrarily across these bins, so I can estimate a very flexible relationship. In the reported results that follow, the omitted category will be those observations with $LTV_{i,t} \in (0, 0.2]$, so all effects are relative to being in that group.

I will show results from 4 specifications, which differ in what is included in $\mathbf{X}_{i,t}$. All specifications include flexible controls for months since purchase ("duration"), calendar month (to soak out any effects of seasonality), and nominal gains/losses. The "Baseline" specification controls for

[16] To the extent that there are transaction costs, this threshold may in fact be below 1.

[17] In the static model presented in Appendix A.2, this is simply due to an increase in net worth associated with liquidating a debt and an asset when the former is larger. The higher is the LTV, the larger is the increase in net worth from doing so. In a richer dynamic model with home price uncertainty, owning an underwater home has option value, but this diminishes as the LTV increases and the option becomes further and further out-of-the-money, making default more likely, as in Campbell and Cocco (2015).

[18] I have also run the analysis using a Cox proportional hazard model, and results are hardly changed. I present the results using the linear model to allow for a more exhaustive set of controls while maintaining computational tractability.

[19] The 24 bins are as follows: [0], (0,0.2], (0.2,0.4], (0.4,0.6], (0.6,0.7], (0.7,0.75], (0.75,0.8], (0.8,0.82], (0.82,0.84], (0.84,0.86], (0.86,0.88], (0.88,0.9], (0.9,0.92], (0.92,0.94], (0.94,0.96], (0.96,0.98], (0.98,1], (1,1.05], (1.05,1.1], (1.1,1.2], (1.2,1.3], (1.3,1.4], (1.4,1.6], (1.6,∞).

only these. The specification "+ Covariates" further controls for county-level economic conditions (county unemployment rate 6 months prior, change in the county unemployment rate from 12 to 6 months prior, the gross number of jobs created in the county in that year [expressed as a percentage of employment]), ZIP code-level measures of financial health (average FICO score on purchase mortgages, average household income, and share of adults with at least a college—all measured soon before the start of the sample period), and some property level controls (initial down-payment and percentile rank of property's value within MSA). The specification "+ MSA-,year-FEs" adds in indicators for each MSA and year in the sample, while the final specification, "+ MSA-by-year FEs" adds in indicators for each MSA in each year. Standard errors are clustered at the MSA-year level.

Figure 7 shows the results of this estimation on the 1% national sample. The probability of sale is quite stable for values of LTV between 0.2 and 0.8, but it falls monotonically and dramatically between 0.8 and 1. Depending slightly on the specification, the monthly probability of a sale is about 15bp lower when $LTV = 1$ than when $LTV = 0.8$, which is about 30% of the sample average (48bp). This confirms the primary prediction of the house lock model, which is that high LTVs (less than 1) deter homeowners from selling their properties.

The results also confirm the model's secondary prediction, which is that the likelihood of sale is an *increasing* function of LTV when $LTV > 1$, as homeowners become increasingly likely to default. The estimates suggest that while this effect begins to take hold immediately at $LTV = 1$, it does not begin to outweigh the house lock effect until the LTV is quite high, around 1.4. Homes with the highest LTVs (in the range of 1.6 and above) are in fact the most likely to be sold.

In addition to confirming the household predictions of the house lock model and laying the groundwork for evaluating its quantitative implications for aggregate movements in home sales (see below), these results are consistent with a growing body of work concluding that borrowers are not inclined to default on a mortgage in the absence of a sharp cash flow shock until their LTVs are quite high. The estimates suggest that this "strategic" incentive[20] does not outweigh the other consequences of default (including being forced to move and perhaps down-size, having diminished access to credit markets, and moral qualms associated with not paying back debts as in Guiso et al. (2013)) until the mortgage debt reaches about 150% of the home's value, a number similar to that estimated by Bhutta et al. (2017).

[20]The distinction between "strategic" and "liquidity-based" default is a bit contrived. As Campbell and Cocco (2015) show, if a homeowner can save a sufficient amount on monthly payments by defaulting (say, if local rents are less than the mortgage payments), then there will be some threshold LTV above which she will default. Liquidity shocks can move that threshold around, so ultimately every default decision is based on both liquidity and LTV. Nonetheless, despite this imprecision, the term "strategic default" is useful shorthand for describing the role that LTV has in determining households' incentives to default.

3.1.2 Aggregate Impact on Home Sales

The fact that the sale hazard depends on LTV allows for the possibility that changes in the distribution of LTV over time may be responsible for changes in sales volume. Figure 4 shows the LTV distribution's rightward shift during the housing bust, as home prices fell. Whereas very few borrowers were underwater on their mortgages in 2004—near the end of a long period of increasing home prices—by 2009 around 20% found themselves in this position. I will now assess the magnitude of the effect of this shift on overall selling activity. In particular, I will look at the decline in the sale hazard from the period 2004-6 ("boom") to the period 2007-11 ("bust").

First, I will treat the hazard function[21] in Figure 7 as constant across time. Let s_k be the hazard coefficient of a home in bin k, and let $w_{k,t}$ be the percent of homes in period t that are in bin k. We can then compute the change in the sale hazard generated by a shift in the LTV distribution as:

$$\Delta_{t,t-1} = \sum_{k=1}^{K} s_k \cdot (w_{k,t} - w_{k,t-1}). \tag{4}$$

To the extent that the shift in the distribution moved mass from bins with high hazard rates (e.g. $LTV \in (0, 0.2]$) to bins with low hazard rates (e.g. $LTV \in (0.98, 1]$), the average hazard rate will decline. This exercise is depicted visually in Figure 8, where I show the hazard function plotted against the difference in LTV distributions between the two periods.

The headline result is that $\Delta_{Bust,Boom}$ = -0.5bp. This is compared to the actual decline in the hazard rate of 32.5bp, meaning this channel is able to account for about 1.5% of the decline in sales. Figure A-4 shows that the story is very similar if one focuses exclusively on a single MSA, as the house lock effect can account for 3.3% of the decline in sales in the Seattle area.

Why is the effect so small? One might be tempted to attribute it to the offsetting impact of two competing effects: yes, sales declined due to the lock-in effect of the down-payment constraint as borrowers shifted into the high-LTV region, but this was offset by the increase in sales coming from defaulting households who had been pushed into the *very*-high-LTV region. And while this is true, it is only a small part of the story. If we shut down the default channel by assuming the hazard function is constant for $LTV > 1$, where that constant is the depressed level for the group with $LTV \in (0.98, 1]$, we would still only get a decline of 2.0bp, or just 6.2% of the actual decline.

To get a better sense of why the house lock story is quantitatively insufficient to explain the drop in sales, consider what would happen if all homeowners had been in the reference bin ($LTV \in (0, 0.2]$) in the boom and had all moved to the $LTV \in (0.98, 1]$ bin in the bust, so that the sale hazard was as depressed as possible. This would generate a decline in the sale hazard of 20bp,

[21] Specifically, the hazard function corresponding to the fourth specification, which includes MSA-by-year fixed effects.

about 60% of the true decline. In other words, it would have required a massive shift in the LTV distribution to generate movement in the overall sale hazard on a comparable scale to what happened to the true hazard rate. While the LTV distribution did of course deteriorate badly, it was not nearly on that scale.

Instead, there was a downward shift in the hazard function. To see this, I re-estimate the model allowing the effect of LTV to vary over time:

$$Sale_{i,t} = (\mathbf{Period}_{i,t} \times \mathbf{LTV}_{i,t}) \cdot \gamma_1 + \mathbf{X}_{i,t} \cdot \gamma_2 + u_{i,t}, \tag{5}$$

where **Period** is a set of indicators for whether the observation is occurring in the "boom" (2004-6), "bust" (2007-11), or "recovery" (2012-4). Figure 9 shows the results of this estimation, where each coefficient is interpreted as the basis point difference in the sale hazard relative to borrowers with $LTV \in (0, 0.2]$ *in the boom.*

The results are stark: the hazard function for homes with LTVs between 0 and 0.8 fell dramatically—by 30-35bp. As this is where the large majority of homeowners are positioned in the distribution (about 80% in the boom, 70% in the bust), it is this decline—in a region where house lock and mortgage default are irrelevant—that is mostly responsible for the 32.5bp decline in sales. Most of the remainder comes from downward shifts in the hazard function in other parts of the support, with only a very small part (again, 1.5%) coming from shifts in the distribution along the hazard function. Figure A-5 shows that this conclusion remains true when focusing exclusively on Seattle.

Therefore, despite its success in explaining an element of household behavior, the house lock story is incapable of explaining the aggregate decline in home sales—in an accounting sense, this decline was driven by homeowners who do not face constraints from their mortgages and yet became far less likely to sell in the bust than they were in the boom. An explanation for this development must come from elsewhere.

3.2 Nominal Loss Aversion

Nominal loss aversion provides another mechanism through which falling home prices can lead to decreased sales activity. If households get disutility from realizing a loss on an investment—or more generally, if utility is an increasing function of realized gains on investments, independent of any impact on lifetime wealth or consumption—then homeowners will become less likely to sell after home prices fall because realized gains will be lower. As this mechanism only applies to homeowners and not their lenders if and when they foreclose on a home, I will focus in this subsection on non-distressed sales. Therefore, I will censor an ownership spell when a lender forecloses on the home or when it is sold via short sale. Figure 10 shows the non-distressed sale hazard in this censored sample. As expected, the decline in this measure of sales is more dramatic than the overall decline in sales, as the increase in distressed sales is not present to buffer the collapse.

I use the same empirical framework to estimate the impact of nominal loss aversion as I did in the previous subsection for house lock. I partition the support of nominal gains/losses into 8 bins: ($-\infty$,-20%], (-20%,-10%], (-10%,-5%], (-5%,0%], (0%,10%], (10%,20%], (20%,40%], (40%,∞). In the reporting of results that follows, the omitted group will be those in the lowest bin with nominal gains in the range ($-\infty$,-20%].

Figure 11 shows the results of this estimation. First note the sensitivity of the estimates to the set of controls. When no controls are included for time period, the effects appear extremely large. This is essentially a restatement of this paper's motivating fact, which is that periods of booming prices have high levels of home sales. When MSA-by-year indicators are included,[22] the effect is weakened substantially, so that homeowners experiencing nominal gains in excess of 20% are about 12bp more likely to sell than homeowners experiencing nominal losses of over 20%, compared to an average non-distressed sale probability of 38bp in the sample.

To assess the aggregate impact, I perform the same sort of calculation as in the previous subsection on house lock:

$$\tilde{\Delta}_{t,t-1} = \sum_{k=1}^{\tilde{K}} \tilde{s}_k \cdot (\tilde{w}_{k,t} - \tilde{w}_{k,t-1}), \tag{6}$$

where $\tilde{s}_k$ is the sale hazard coefficient for bin k and $\tilde{w}_{k,t}$ is the share of nominal gain/loss distribution in period t that is in bin k. Figure 12 shows the components of this calculation: the estimated hazard coefficients and the change in the nominal gain/loss distribution between the boom and bust.

Performing this calculation, the change in the nominal gain/loss distribution can account for a 3.5bp decline in the non-distressed sale hazard between boom and bust, or about 7.9% of the actual decline. While this is hardly trivial, it suggests that, as in the case of house lock investigated earlier, other explanations are required to explain the dramatic decline in home sales experienced during the crash.

The reason nominal loss aversion fails to account for a large share of the decline in aggregate sales is somewhat different than why house lock is not capable of doing so. Whereas, as discussed previously, only about 10% of the LTV distribution shifted from being below 0.8 to above 0.8, about 50% of the nominal gain distribution shifted from being positive to negative between boom and bust. However, the effects on household behavior are not strong enough to convert that large change in the distribution into a steep decline in sales. Even if the entire distribution where shifted from the bin with highest gains into the bin with highest losses, this would generate only about 30% of the observed decline in non-distressed sales. Given the more modest shift, the true number appears to be quite a bit smaller than that.

[22] Further stringency in the set of controls seems to impact the results very little. For instance, when including ZIP code-by-month FEs rather than MSA-by-year FEs, the coefficients are hardly changed.

3.3 Time on Market and Sales Volume

I conclude this section by discussing how the time a listed home spends on the market relates to this broader discussion of sales volume. This is relevant both because, as discussed in Section 1, many papers have used search models to explain time variation in home sales with some mechanism that changes the speed at which listed homes sell, and because some prominent empirical work has focused on this part of the selling process. Seminal work in Genesove and Mayer (1997) and Genesove and Mayer (2001) used listings data to demonstrate that house lock and nominal loss aversion affect the amount of time a house spends on the market after being listed for sale. Table A-1 uses the Seattle (non-REO[23]) listings data to verify the results of those studies. In particular, I find that listed homes becomes less likely to be sold as their LTVs increase past 0.8, consistent with the house lock story. Homes that have experienced more appreciation since they were last purchased are also more likely to sell, with this effect being stronger in the region where homeowners have experienced losses than the region with gains. This is consistent with the literature's results with regards to nominal loss aversion.

While this suffices to qualitatively demonstrate that homeowners are affected by these channels when selling their homes, it tells us very little about their quantitative impact—even at the household level. The reason is that the speed at which homes turn over is not very dependent on how quickly listed homes sell. Rather, what is critical is the frequency with which they are listed for sale in the first place, a point emphasized by Ngai and Sheedy (2017).

To see why, let n be the hazard that an unlisted home is put on the market and s be the hazard of a home selling, conditional on being listed. In steady state, then, the unconditional hazard of sale is given by:

$$V = \frac{sn}{s+n}. \tag{7}$$

This allows us to decompose variation in the unconditional sale hazard into a part coming from variation in the probability of being listed and a part coming from the probability of a listed home selling:[24]

$$dV \approx \frac{n^2}{(n+s)^2} ds + \frac{s^2}{(n+s)^2} dn. \tag{8}$$

The contribution of each hazard is scaled by the square of the other hazard. Intuitively, if n is small and s is large (as is the case), the set of listed homes will be a fairly low share of the housing stock. As a result, changes in the rate at which they sell will not have a large impact on overall sales. An increase in n, however, will get a lot more homes out on the market, and since they sell

[23]This stands for "real estate owned," and refers to properties owned by lenders.

[24]Technically, this variation has be to between steady states, but the high rate at which listed homes sell makes the transition period relatively brief.

quickly, this well have large effect on overall sales. In the listings data available for Seattle, $n \approx \frac{1}{100}$ and $s \approx \frac{1}{4}$, so the variation in V comes almost entirely from movement in n, not s.[25] Figure 13 shows this explicitly—almost all of the variation in V comes from movements in n. As a result, while studying listings data can be informative about how housing markets function, ultimately time-on-market is not a quantitatively important determinant of the frequency with which homes sell. The results I presented above, then, show that the mechanisms uncovered by Genesove and Mayer (1997) and Genesove and Mayer (2001) extend to the more important decision of whether to list the home in the first place. In this sense, then, these results extend and bolster that older literature. Nonetheless, while the household results are compelling, they simply cannot account for much of the decline in sales, as demonstrated.

4 A Novel Explanation: Filtering and (Irreversible) Construction

In this section, I present a novel mechanism to explain the cyclicality of home sales. The model allows for new construction and homes of differing discrete qualities—a housing "ladder." When there is a boom in the housing market, agents want to live in nicer homes. This aggregate upgrading is enabled by the construction of high-quality homes. Agents with the highest willingness-to-pay for housing move into those homes, and they sell their old homes to people with lower willingness-to-pay. As a result, sales of existing homes are high during the boom. But the irreversibility of housing construction means that there is no aggregate downgrading in a bust—agents simply stay in the homes they have, and sales of existing homes dry up.

This process by which the construction of a new home allows for improved home quality for agents beyond the one moving into that home has long been known as "filtering" in the affordable housing literature. The seminal work of Lansing et al. (1969) started with 1,000 pieces of new construction, and interviewers asked the occupants where they had moved from. They then went to that address and did the same, following each chain until it concluded with a demolished home or a new household. They found that 1 newly built home was associated with around 2.5 *additional* moves and 1.4 additional sales (they tracked renters as well), with a strong tendency for the new occupant of a home to have lower income than its previous occupant. Rosenthal (2014) provides a more modern look at this question and argues that the construction of new homes for relatively wealthy households is an important source of housing for low-income households, precisely because of this filtering process. In this paper, I argue that the filtering process, combined with

[25] This is not to say that that time-on-market does not matter for anything. First of all, as the search process is costly—a pillar of the search-and-matching literature—periods of extended time-on-market can be harmful to home sellers. Second of all, while it is true that home sellers waiting a bit longer to sell their homes does not impact sales volume, that extra time could be quite important for them to garner higher bids and thus sell at higher prices. Finally, in terms of aggregate variables, inventory—the stock of listed homes—does depend quite a bit on both n and s. Defining L to be the share of homes that are listed, an analogous analysis to the one above yields $dL \approx -\frac{n}{(s+n)^2}ds + \frac{s}{(s+n)^2}dn$. Now each hazard's contribution is no longer scaled by the *square* of the other, which allows s to be an important determinant of inventory. Intuitively, if listed homes sell very quickly, not many homes will be listed at any given point in time. So while overall sales are determined almost entirely by the flow of homes onto the market, the stock of homes on the market does depend on both flows.

the cyclicality of construction, induces cyclicality in sales of Existing homes.

The model presented in this section begins by establishing what I will call a "constant construction path" (CCP) with construction of new homes, formation of new households,[26] and filtering. The CCP is not a true steady state because the housing stock is growing. It is also not a "balanced growth path" because all variables are not growing at constant rates. Rather, the amount of new construction is constant and thus is declining relative to the size of the housing stock, while sales of existing homes increase over time (linearly, not exponentially). Note that this is consistent with the long-run data in the US. Since 1968, the *number* of homes built has had essentially no trend, while the housing stock and the number of existing sales has grown in such a way that share of existing home being sold has also essentially not had a trend (see Figure 12). I will then show that when the market is disturbed off of the CCP and then converges back to it, the filtering process generates dynamics for sales volume that are consistent with the data we observe in housing cycles.

The model presented here is stripped down to its essentials in order to make the intuition for the mechanism as clear as possible. I do not include mortgages, and I not explicitly discuss a rental market, among other omissions. Nonetheless, I will argue in later sections that the model is not only useful for making qualitative points about the mechanism, but it is also a useful guide for what to look for in the data to assess the quantitative power of the mechanism.

4.1 Setup

In each period, agents choose from 3 types of housing arrangements, which offer different flows of housing utility. In ascending order of desirability, these options are:

1. Non-household ($\emptyset$)
2. Existing home (E)
3. New home (N)

In the exposition to come, the term "agents" refers to all decision-makers, while "household" refers only to those who have selected housing of type E or N. Agents choose between these options by weighing costs against their taste for housing quality. Taste for housing quality has a common component, Ω_t, and an idiosyncratic component, ω_i. I will assume that Ω_t grows over

[26] One way to generate a steady state with construction is to have depreciation, so that new homes are built to replace those that crumble. And indeed, the filtering literature often focuses on the depreciation of homes as what drives them down the income distribution (e.g. Bond and Coulson (1989)). However, in my model, I will not have depreciation. Instead, the constant stream of construction will be driven by a secular increase in the willingness-to-pay for housing (think rising income), so that homes travel down the distribution not because they are becoming worse but because the population is getting richer. This is both consistent with a growing housing stock, as in the real world, and it avoids some pitfalls that a model with depreciation runs into. Since home qualities are discrete, exponential depreciation is not practical, while random and complete depreciation generates a perverse (and counter-cyclical) source of sales, as it scrambles the relationship between households' tastes for housing and the qualities of their homes. As a result, I abstract from depreciation.

time, generating construction of New homes and household formation. Idiosyncratic tastes (ω_i) are unbounded below and, without loss of generality, I will assume that the household with highest taste has $\omega_i = 0$. The measure of households at each level ω is given by $M(\omega)$.

Flow utility has 2 components:

$$u_{i,t} = HousingServices_{i,t} + b \cdot CashFlow_{i,t}. \tag{9}$$

Higher quality homes will tend to give higher housing services but, in equilibrium, they will be costlier. This will hurt the agent's cash flow, generating the key tradeoff.

Define $a_{i,t}$ to be the quality of the home that agent i *enters* period t owning, and let $h_{i,t}$ be the quality of the home in which agent i chooses to *live* in period t. Then:

$$CashFlow_{i,t} = \begin{cases} 0 & \text{if no move} \\ \phi_t^{a_{i,t}} - \phi_t^{h_{i,t}} - \mu & \text{if move} \end{cases}. \tag{10}$$

μ is a resource cost of moving that is paid to no one, and $\boldsymbol{\phi} = [\phi^{\emptyset}, \phi^E, \phi^N]$ is the prices of the different housing qualities in a given period. For notational convenience, I will normalize the price of being a non-household to 0 ($\phi_t^{\emptyset} = 0$), and I will denote the price of an Existing home to be P ($\phi_t^E = P_t$) and the price of a New home to be C ($\phi_t^N = C_t$).

Turning to the utility derived from housing, $HousingServices_{i,t}$ itself has 2 components:

$$HousingServices_{i,t} = Quality_{i,t} + Match_{i,t}. \tag{11}$$

$Quality_{i,t}$ is the component of housing services that comes from the objective quality of the home. It is a function of the home's quality as well as the agent's taste for quality:

$$Quality_{i,t} = \begin{cases} \chi & \text{if } h_{i,t} = \emptyset \\ (\Omega_t + \omega_i) & \text{if } h_{i,t} = E \\ (1+q)(\Omega_t + \omega_i) & \text{if } h_{i,t} = N. \end{cases} \tag{12}$$

The parameter q is crucial, as $q > 0$ means that New homes are preferred to Existing homes. The fact that New homes increase utility in a way that is increasing in ω (taste for quality) will tend to sort households with higher tastes into New homes and households with lower tastes into Existing homes.

The other component of the housing service flow, $Match_{i,t}$, is an idiosyncratic piece that the household receives from the house that is independent of its objective quality. This is a standard feature of housing models with searching and matching and is meant to capture the idiosyncracies

that homes offer as well as the idiosyncratic tastes of the people who live in them. For example, in the context of this model, it could capture location. If an agent receives good job offers in one area of the country and not another, her match to homes in the former area will be better than in the latter.

At the beginning of each period, households draw a match to their *current* home, $\epsilon_{i,t}^{stay} \sim G$. If they stay in their current home, they receive this value for their match flow. If they decide to move, then they will perform a search process. For now, assume that this search process yields $Match_{i,t} = \eta_t^{h_{i,t}}$ on average. (Below, I will present a simple search framework that fits tractably into the surrounding model while providing micro-foundations for $\eta_t^{h_{i,t}}$ and motivating why it could be time-varying. However, I defer that until after more discussion of the rest of the model.) Therefore:

$$Match_{i,t} = \begin{cases} 0 & \text{if } h_{i,t} = \emptyset \\ \epsilon_{i,t}^{stay} & \text{if no move and } h_{i,t} \neq \emptyset \\ \eta_t^E & \text{if move and } h_{i,t} = E \\ \eta_t^N & \text{if move and } h_{i,t} = N \end{cases}. \tag{13}$$

4.2 Household Solution

We can now write the household's value function. The aggregate state is given by $\mathbf{X} \equiv [\Omega, \boldsymbol{\phi}, \boldsymbol{\eta}]$—the common component of taste for housing, the prices of different qualities of homes, and the average match utilities from moving. The individual-level state variables are $\mathbf{Z} \equiv [a, \epsilon^{stay}, \omega]$—the quality of the home the agent owns entering the period, the match utility from staying in that home, and the idiosyncratic piece of the taste for housing quality. The value function is then given by:

$$V(\mathbf{Z}, \mathbf{X}) = \max_{h,y} \mathbb{E}\Big[u(h, y; \mathbf{Z}, \mathbf{X}) + \beta V(\mathbf{Z}_+, \mathbf{X}_+)|h\Big], \tag{14}$$

where β is the agent's discount factor and y is an indicator for whether the agent moved. The decision problem in 14 is a discrete-choice problem. The household chooses one of two values for y (move or not) and, if moving, one of three values for h ($\emptyset$, E, or N). This leaves four combinations of these two variables corresponding to the cases listed in Equation 13, and households will choose among these four options depending on their idiosyncratic values of ω_i and $\epsilon_{i,t}^{stay}$ (and aggregate variables). Figure 15 displays the idiosyncratic taste for housing quality (ω_i) on the horizontal axis and the match utility from the current home ($\epsilon_{i,t}^{stay}$) on the vertical axis. Agents with low taste for housing quality, on the left side of the graph, will not be households. Households with high values of match utility in their current homes, in the top portion of the graph, will elect to stay in their homes. Households with low match utility in their current homes, in the bottom part of the graph, will choose to move. Among this group, those with the highest taste for housing quality will buy

newly-constructed homes (N) and those with lower tastes for quality will buy Existing homes (E). The borders between these adjacent regions in Figure 15 can be explicitly represented in terms of cutoff rules, which directly indicate the various factors that lead households to locate in one region or another. I discuss each of these four cutoffs now.

The first cutoff rule determines who forms households in a given period. Note that these "marginal households"—agents who did not want to form households until the current period—will have relatively low housing tastes, which means (as demonstrated below) they will move into Existing homes rather than New homes, as they are unwilling to pay the higher cost associated with the higher quality of New homes. The cutoff household is thus characterized by indifference between moving into an Existing home in period t versus waiting until period $t+1$, when the common component of the taste for quality will have increased.[27] This yields:

Result 1 *Agents will be households in period t if and only if $\omega_i > \hat{\omega}_t$, where*

$$\hat{\omega}_t \equiv \chi - \Omega_t - \eta_t^E + b\mu + bP_t - \beta(bP_{t+1} + J_{t+1}), \tag{15}$$

and $J_t \equiv \int_{\eta_t^E - b\mu}^{\infty} (\epsilon - (\eta_t^E - b\mu))dG(\epsilon) - b\mu.$

There are a few takeaways from this result. First, as Ω_t grows over time, this will put downward pressure on $\hat{\omega}_t$, generating household formation and new construction, as additional homes are required in order for the housing stock to accommodate the increased number of households.[28] Second, the threshold is increasing in P_t. This simply says that demand for housing is downward-sloping: all else equal, higher prices lead fewer agents to want to form households. Third, the cutoff is a decreasing function of P_{t+1}. Since Ω_t is increasing, all agents will eventually form households. The marginal household, then, has to determine whether she wants to form in period t or period $t+1$. So as next period's price increases, it will be more appealing to do so in the current period.[29]

The second cutoff describes, conditional on moving, who will choose a New versus an Existing home. Unsurprisingly, it is agents with high idiosyncratic housing tastes who choose New homes:

Result 2 *Conditional on moving, a household will choose N instead of E if and only if $\omega_i > \tilde{\omega}_t$, where*

[27] One caveat for this result is that it only holds for equilibria with positive construction. If there were no construction, then there might be no marginal household, so that all non-households are sufficiently far away from being indifferent that even when Ω_t increases, they will not form households. However, since new construction is always positive empirically, I will simply assume this to be the case and focus on areas of the parameter space where it is.

[28] Crucially, while there are New homes and new households, the new households do not move into the New homes. Rather, they move into Existing homes, and it is current homeowners who move into the New homes.

[29] Another way of saying this is that higher price growth lowers the user cost of ownership and so encourages household formation.

$$\tilde{\omega}_t \equiv \frac{b(C_t - P_t) + (\eta_t^E - \eta_t^N)}{q} - \Omega_t. \tag{16}$$

The key takeaway from this equation is that $\tilde{\omega}_t$ is an increasing function of $C_t - P_t$, ceteris paribus. Intuitively, as the price of Existing homes rises relative to the price of New homes, households suffer less of a negative cash flow hit when moving from an Existing home to a New home. This increases the set of households that would choose a New home upon moving—a group I will refer to as "conditional Ns." I will refer to the group with $\omega_i \in [\hat{\omega}_t, \tilde{\omega}_t]$ as "conditional Es."

Now that we know who will move into what type of home, *conditional* on moving, we can establish the final two cutoff rules, which describe which current households choose to move.

Result 3 *A household with $\omega_i > \tilde{\omega}_t$ will move if and only if $\epsilon_{i,t}^{stay} < \bar{\epsilon}_t^N(\omega_i)$, where*

$$\bar{\epsilon}_t^N(\omega_i) \equiv b(P_t - C_t - \mu) + q(\Omega_t + \omega_i) + \eta_t^N. \tag{17}$$

Households will move when they get low realizations of $\epsilon_{i,t}^{stay}$—when the idiosyncratic characteristics of their current home do not conform well to their idiosyncratic tastes. For these conditional Ns, market prices play an important role. When the price of Existing homes is far below that of New homes ($P_t - C_t$ very negative), trading up requires a large reduction in cash flow, so the cutoff value for $\epsilon_{i,t}^{stay}$ will be low, as only those households with very bad matches will move. As the price gap narrows, more of these conditional Ns will be induced to move to gain the extra utility associated with living in a New home.[30]

Result 4 *A household with $\omega_i \leq \tilde{\omega}_t$ will move if and only if $\epsilon_{i,t}^{stay} < \bar{\epsilon}_t^E$, where*

$$\bar{\epsilon}_t^E \equiv \eta_t^E - b\mu. \tag{18}$$

Again, households who find themselves with bad matches to their current homes will move. Interestingly, prices are irrelevant for this margin. As conditional Es are moving from an Existing home into an Existing home, they are on both sides of the market, so the price does not factor into their cash flow or, by extension, their decision. This group, then, demonstrates a powerful intuition: if individual agents are both buying and selling homes in the same market, sales activity will not

[30] One somewhat awkward feature of this solution is that this cutoff rule depends on ω_i—in particular, households with higher tastes will have higher thresholds. Intuitively, they get more of a utility boost from living in a New home, so they will be willing to move out of their Existing homes even with some match realizations that others would not consider very bad. This is a natural feature of the model—a consequence of the necessity of making the difference in utility between E and N an increasing function of ω, which (imperfectly) sorts higher ωs into Ns. However, it predicts that, if all agents draw $\epsilon_{i,t}^{stay}$ from the same distribution, then those with high ω will move more frequently. As this is neither an important prediction of the model nor one that has strong empirical backing, I will assume that $\epsilon_{i,t}^{stay}$ is drawn from a distribution that depends on ω in such a way that, in steady state, all households are equally likely to move. Qualitatively, this will mean agents with high ω draw $\epsilon_{i,t}^{stay}$ from better distributions.

vary with price without some kind of friction. In Stein (1995), this friction is a down-payment constraint, which prevents some households from swapping one home for another. In models with search frictions (e.g. Krainer (2001), Novy-Marx (2009), Wheaton (1990)), sellers trade off price with time-on-market, which determines the volume of homes selling in a period. In my model, in contrast, there is no friction. Rather, as I will show below, the volatility in sales comes from the fact that, while there is a group of households that are insensitive to the price, there is a different set of households that are sensitive to prices because, in a sense, they are operating in 2 markets:[31] those for Existing and New homes. As described above, they care about prices because they are buying in one market and selling in the other. As a result, they will be the drivers of the interesting dynamics to come.

4.2.1 The Search Process

Now that I have described the basic structure of the household's decision, let me return to the question of how moving households search for their homes. Households first choose which market they will look in, New or Existing. The market is segmented in the sense that all who choose Existing will indeed up in Existing homes, and same for New homes. Once they have made this choice, they sell their old home for the market price and pay the market price for the housing quality they have chosen to search within. The search process then begins.

Let's first focus on the market for Existing homes. Following Ngai and Tenreyro (2014), I assume that the arrival of homes to be viewed by a searching household is a stochastic Poisson process with arrival rate λ, meaning the number of homes viewed by a household is drawn from a Poisson(λ) distribution. Letting v_t be share homes of that are listed listed and b_t the share of households that are searching, the arrival rate is given by:

$$\lambda_t = \frac{m(b_t, v_t)}{b_t}, \tag{19}$$

where m is a matching function that is homogeneous of degree $\alpha \geq 1$. As discussed in greater detail below, this is a model where markets clear period-by-period, which means $b_t = v_t$.[32] Therefore:

$$\lambda_t = m v_t^{\alpha-1}, \tag{20}$$

[31] The model of Ortalo-Magne and Rady (2006) also fits this description. Though they emphasize the importance of a down-payment constraint, similar to Stein (1995), this is actually not necessary in order for them to generate volatility in home sales. Interestingly, that paper also features a housing ladder, and so there are agents engaged in 2 markets simultaneously, as in mine, and they are sensitive to prices. As discussed in the main text, their mechanism is totally distinct from mine—mine requires new construction (see below), whereas theirs occurs in a model with a fixed housing stock—yet it shares this feature that prices matter to agents because some of them are buying and selling different types of homes that have different prices, meaning they are not insulated from price movements.

[32] For those movers going between one Existing home and another, they clearly contribute equally to b and v. There are new households that form who search for homes without listing any, but they are equal in measure to the group of households that put Existing homes up for sale and then go looking for New homes. Thus, $v_t = b_t$.

where $m \equiv m(1,1)$. Suppose that, upon being matched to a home, the household observes what would be their match utility from living in that home, $\epsilon \sim H$. I will assume that the household selects the home that they viewed that delivered the highest match utility. Those who are not matched to any homes will then be put in a house at random. If H is Uniform[0,1][33] and we define $\epsilon_{i,t}^{max}$ to be the maximum match utility observed in period t by a searching household i, it can be shown that:

$$E[\epsilon_{i,t}^{max}] = \frac{\lambda_t^2 - \lambda_t + 1 - e^{-\lambda_t}}{\lambda_t^2}. \tag{21}$$

This provides a micro-foundation of η_t^E, the average match utility from moving. Note that it is an increasing function of λ: when households are able to view more homes on average, the match utility from the home they ultimately choose will also be higher on average, meaning their incentive to move will be as well.

Another important observation is that, when $\alpha > 1$, λ_t is an increasing function of sales volume ($v = b$). This captures the idea of "market thickness"—when more homes are available on the market, searching households are more likely to find matches. As a result, they will on average get better match utilities. This generates a feedback cycle that can amplify any initial increase in sales volume—as more homes are put on the market, households will have greater incentive to put their own homes on the market and search, as they will get better matches. This will then increase the quality of the matches, incentivizing more households to enter, and so on. This is an issue I will discuss when evaluating the quantitative power of the model.

So far, I have focused on η_t^E, the match utility from search for an Existing home. There are different ways one could think about how to derive η_t^N, the match utility from New homes. If we are thinking of these homes as being custom-built, then perhaps η_t^N would be very high (and constant). Or perhaps developers build a set of homes and people search among them in much the same way as described above for Existing homes. As there is not much empirical basis to distinguish between these—or other—approaches, I will typically just assume $\eta_t^N = \eta_t^E$, which is quite reasonable in this model because (as discussed below) changes in the sale hazard of Existing homes are tightly linked to changes in the sales of New homes, so to the extent that one market is thick, the other one will be, too.

4.3 Market Clearing

Two markets must clear in this model. The first is the market for New homes. This of course requires a construction sector. I will work with a simple one. The construction sector is made up

[33] This assumption is made for simplicity and the comparative statics discussed below do not depend on it. It is easy to show that the distribution of $\epsilon_{i,t}^{max}$ with high λ_t first-order stochastic dominates the distribution with low λ_t, as emphasized by Ngai and Tenreyro (2014). I make this distributional assumption in order to produced a simple closed-form solution for the mean of the distribution. In keeping with it, I will assume the $\epsilon_{i,t}^{stay}$ is also drawn from a Uniform distribution when solving the model below.

of a continuum of infinitesimal firms, each with constant marginal cost in a given period. While firms take the cost as given, it is increasing in both the size of the housing stock and the aggregate flow of construction. I assume a linear structure to this marginal cost:

$$\zeta_t = c_0 + c_1 \cdot \int_{\hat{\omega}_t}^{0} M(\omega) d\omega + c_2 \cdot \int_{\tilde{\omega}_t}^{0} \Big(\int_{\epsilon_{min}(\omega)}^{\bar{\epsilon}^N(\omega)} dG(\epsilon) \Big) M(\omega) d\omega. \tag{22}$$

This representation of the cost curve already embeds the demand side of the model, as I have written the stock of housing and the flow of construction as functions of the households' cutoff rules.[34] Market clearing then just requires us to ensure that the firms are optimizing. Because firms have no market power, they will set the price of a New home equal to the cost of building it:

$$C_t = \zeta_t. \tag{23}$$

The other market that must clear is that for Existing homes. Supply comes from current households who receive bad match realizations to their homes, prompting sales by both Conditional Es (who move into Existing homes) and Conditional Ns (who move into New homes). Demand comes from those same Conditional Es, as well as new households. Since Conditional Es are both buyers and sellers, they do not contribute any net demand. Therefore, market clearing requires that the supply coming from households moving into New homes equals the demand from newly-formed households:

$$\int_{\tilde{\omega}_t}^{0} \Big(\int_{\epsilon_{min}(\omega)}^{\bar{\epsilon}^N(\omega)} dG(\epsilon) \Big) M(\omega) d\omega = \int_{\hat{\omega}_{t-1}}^{\hat{\omega}_t} M(\omega) d\omega. \tag{24}$$

This means that construction of New homes will be equal to household formation, but critically, the new households do not move into the New homes. Rather, they move into Existing homes that are being vacated by continuing households. This generates a key link between construction of New homes and sales of Existing homes, as the expansion/improvement of the aggregate housing stock is facilitated by turnover of the existing housing stock at the individual level.

Another important assumption is embedded in this market clearing condition: the irreversibility of housing construction. Households who move are able to get rid of their old homes only by selling them to other households. This means the measure of households cannot fall, as this would require some agents to sell homes without buying any, which would violate this condition. One could imagine a different model in which households have the option of simply scrapping their homes and consuming the materials. In this case, households could dissolve and scrap their homes, maintaining

[34] The first integral represents the stock of housing, as it takes the measure of all agents with tastes for housing about the threshold required for household formation ($\hat{\omega}_t$). The double integral gives the flow of housing construction, because it represents all agents who both have tastes high enough to make them conditional Ns ($\omega_i > \tilde{\omega}_t$) and received low enough values of $\epsilon_{i,t}^{stay}$ that they want to leave their current homes.

the market clearing condition while shrinking the housing stock. Implicitly, then, I am assuming that the technology for turning homes into consumption of a different sort is sufficiently bad to preclude its use.

4.4 Equilibrium

The previous subsections have described the behavior of optimizing agents and market clearing. To pin down the equilibrium, all that is left to do is to describe how agents forecast prices. As mentioned earlier, this is critical because agents must choose when to form households, and so they must have beliefs about what the cost of doing so will be at all points in time. I will now show that a "constant construction path" (CCP) exists in which households forecast prices correctly and construction is constant in all periods. On this CCP, each period's prices clear the housing markets in that period in addition to setting forecasts correctly in all previous periods to clear those markets as well.

This CCP requires an additional assumption, which is that the measure of agents is constant over the entire support:

$$M(\omega) = \begin{cases} 1 \text{ if } \omega \leq 0 \\ 0 \text{ if } \omega > 0 \end{cases} \tag{25}$$

With this in place, the following can be shown:

Result 5 *If $\Omega_t = \Omega_{t-1} + d^{\Omega}\ \forall t$, then $P_t = P_{t-1} + d^P\ \forall t$, $C_t = C_{t-1} + d^C\ \forall t$, and $\hat{\omega}_t = \hat{\omega}_{t-1} + d^{\hat{\omega}}\ \forall t$, where d^P, d^C, and $d^{\hat{\omega}}$ are constants given below:*

$$d^P = \frac{b \cdot c_1 - q}{b\big(1 + b \cdot c_1(1-\beta)\big)} d^{\Omega}, \tag{26}$$

$$d^C = c_1 \frac{1 + q(1-\beta)}{1 + b \cdot c_1(1-\beta)} d^{\Omega}, \tag{27}$$

$$d^{\hat{\omega}} = -\frac{1 + q(1-\beta)}{1 + b \cdot c_1(1-\beta)} d^{\Omega}. \tag{28}$$

This is shown by asserting that $\tilde{\omega}_t$ is constant and then, given d^{Ω}, finding the changes in P_t and C_t required to make it so, while still obeying the various optimization conditions above.

If $d^{\Omega} = 0$, then there is a true steady state, with constant prices and no construction or household formation. If $d^{\Omega} > 0$, then unformed households will want to form as time progresses and continuing households will be more willing to bear the cost of building New homes. Therefore, there will be construction and household formation. This will be tempered by the fact that the

cost of New homes will rise, assuming $c_1 > 0$. The model does not make a strong prediction about whether the price of Existing homes will rise on the CCP. If $q > b \cdot c_1$, households with high tastes for housing quality like New homes so much that they are willing to purchase them even as the price of their Existing homes falls, meaning they have to cover more of the expense. However, given that, empirically, home prices typically increase, I will always assume $b \cdot c_1 > q$.

The CCP is thus characterized by equations: 15 (marginal household), 16 (marginal New home buyer, conditional on moving), 17 (marginal movers among Conditional Ns), 18 (marginal movers among Conditional Es), 20 (arrival rate of homes for searching households), 21 (expected match utility for searching households), 22 (marginal cost curve), 23 (market clearing for New homes), 24 (market clearing for Existing homes), 26 (CCP change in price of Existing home), 27 (CCP change in price of New home), and 28 (CCP household formation/new construction).

4.5 A Housing Cycle: Shocking the CCP

I turn now to the main objective of the model, which is to demonstrate a mechanism to explain the cyclicality of Existing home sales. While the CCP is consistent with long-run features of the housing market—a constant level of construction and a constant proportion of Existing homes selling—the far more interesting part of the model is what happens when the market deviates from the CCP.

To this end, I will study an unanticipated, temporary shock to beliefs about the future of the housing market, which will turn out to be incorrect. In particular, in the "boom" period, agents will believe that the common component of housing tastes will begin to grow in the next period and all periods thereafter at some speed $\tilde{d}^{\Omega} > d^{\Omega}$. The following period, agents' beliefs are corrected, returning to the initial d^{Ω}.

Why do I use this shock? This choice is motivated by a growing literature that has argued that an increase in expectations of future home prices was a prominent feature—and perhaps the driving force—of the housing boom. Case and Shiller (2003) and Case et al. (2012) provide direct evidence based on household surveys that during the housing boom, expectations of home price growth—particularly over long horizons—were very bullish. Perhaps equally importantly, Gerardi et al. (2008) and Foote et al. (2012) provide evidence that financial market participants had similarly optimistic beliefs about the future of home prices. Glaeser et al. (2013) argue theoretically and empirically that credit market conditions, and mortgage interest rates in particular, are not so tightly linked to home prices and are unlikely to be able to account for a substantial portion of the price movements observed during the cycle. In a similar vein, Adelino et al. (2016), Foote et al. (2016), and Albanesi et al. (2017) collectively argue that the cross-sectional patterns of credit growth and mortgage default across the income and credit score distributions are not predicted by theories centered on a credit boom fueled by subprime lending, a popular alternative hypothesis, and thus suggest that optimistic beliefs likely played an important role. Furthermore, Kaplan et al. (2017) estimate an overlapping-generations, incomplete markets model with realistic mortgages and

shocks to labor productivity, credit constraints, and expectations of future home prices. They find that a shock to these beliefs is an important driver of the cycle.[35] Having said this, the model's mechanism does not depend crucially on this particular shock. The key is that a temporary boom in construction will lead to a temporary boom in Existing sales, so even if that boom were generated by, for instance, a loosening of credit standards for marginal households, the mechanism would still go through.[36] However, this expectations shock has become increasingly well-supported in the literature and it fits into the model in a straightforward way, so I proceed with it.

The top panel of Figure 16 shows the response of the model to this shock. In the boom, marginal households respond to the perceived high future price of housing by deciding to buy homes in the present period, before this large increase in prices. This pushes up the price of Existing homes in the boom. This in turn decreases the net expense of trading up from an Existing home to a New home, which a larger share of high-taste households will thus choose to do. This period then has the hallmarks of a housing boom: increasing prices, high construction of New homes, and high turnover of Existing homes, as the aggregate expansion of the housing stock is facilitated by more individuals selling off their old homes to move into New homes. This intuition is shown diagramatically in the top panel of Figure 17.

The following period, agents' beliefs are corrected, returning to the initial d^{Ω}. The market, however, does not return to the CCP. The building boom generated a glut of housing, resulting in overhang that would not have existed had the boom not occurred. Due to the oversupply of houses and the inefficiency in destroying them, the price must collapse below its CCP level to induce households to fill all of them. Construction similarly collapses, though not to zero. Even with the crash in prices, there are some households whose tastes for quality is sufficiently high—and whose matches to their current homes are sufficiently bad—that they move out of their old homes and build New homes, covering a large gap between C_{bust} and P_{bust}. With construction in its depressed state, the filtering process slows down from its CCP level, so sales of Existing homes decline, too.

[35] Even if the reader is accepting of the role of beliefs in driving the housing cycle, the unanticipated nature of the boom and its subsequent reversal may seem peculiar. To be sure, one benefit of this formulation is the tractability it provides. However, while extreme, this is a context where an "MIT shock"—one that occurs despite agents placing no probability weight on it—may actually approximate reality. Gerardi et al. (2008) and Foote et al. (2012) discuss a report from Lehman Brothers in 2005 about the firm's exposure to the housing market. The report provides different scenarios of home price growth and, conditional on those scenarios, they calculate gains/losses that are actually quite clear-eyed. The problem is the probability weights they assign to these scenarios: the "Meltdown" scenario assumes home price declines of 5% per year for 3 years and then grow 5% per year after that, and even this is only given a probability weight of 0.05. For this market participant, then, the actual 30% decline in home prices truly was unanticipated. Furthermore, the quantitative model of Kaplan et al. (2017) requires transition probabilities from (essentially) a calm to a boom and a boom to a bust. They calibrate the model to match the average lengths of boom and bust episodes historically. In doing so, they end up with a 4% (biennial) probability of entering a boom and an 8% probability of transitioning into a bust from that boom. While these probabilities are not 0%, they suggest that what unfolded truly was a "tail event" assigned very little weight by market participants.

[36] In fact, Kaplan et al. (2017) find that while the expectations shock is the key quantitative driver of the home price movements, a relaxation of credit constraints was important for allowing these beliefs to translate into an increase in borrowing by marginal lenders, an element of the cycle strongly emphasized by Mian and Sufi (2009) and Mian and Sufi (2017). In my model, since there are no credit constraints, a relaxation of credit constraints is not necessary to allow the expectations boom to induce a surge in household formation.

The bottom panel of Figure 17 gives diagrammatic intuition for this.[37]

The recovery from the bust continues to be characterized by overhang, so construction and sales of Existing homes remain below their CCP levels. Notably, the price of Existing homes increases faster during the recovery than it did on the CCP. The overhang generated by the building boom mutes the supply response to the increasing taste for housing (recall, Ω_t is rising throughout this period), so the price response is now a more active margin.

In all then, this model is able to generate the key qualitative features of the recent housing cycle in terms of prices, quantities, and turnover: home price growth, construction of New homes, and sales of Existing homes rise and fall together, followed by a sluggish recovery of construction and sales but strong price growth. This was all done with a single, simple shock. I also demonstrate the possibility of amplification of the cycle with market thickness effects, in the spirit of Ngai and Tenreyro (2014). Recall that when $\alpha > 1$, searching households tend to view more homes when searching, which increases the incentive to move. This generates a positive feedback cycle as the initial impulse that increases sales then induces more homeowners to sell their homes, and so on. As the response of Existing home sales demonstrates, this can amplify the model's ability to generate volatility in sales.

The middle and bottom panels of Figure 16 shed further light on the model's key mechanism. The middle panel shows how the model responds if there is a fixed housing stock, with no construction possible. While there are still distinct tiers of housing ($\emptyset$ and E), the boom has no effect on sales of Existing homes. Since there can be no supply response, the boom in expectations simply raises the price, which ensures that the marginal household remains marginal and no agents are induced to form households. As far as the homeowners are concerned, this boom is irrelevant—while the price surges, any move by them requires both the sale and purchase of an Existing home, so they are perfectly hedged against the price fluctuation. As a result, they continue to sell their homes at the same rate as always. This again demonstrates the difficulty of getting sales to be volatile in a model where buyers are also sellers. Notice also that in the bust, the price simply falls back to the CCP—there is no over-shooting because there was no building boom and thus no overhang.

This is probably where it is clearest that the model in this paper is illustrating a distinct mechanism from that of Ortalo-Magne and Rady (2006), a prominent paper that uses a housing ladder to explain housing market volatility. Their model generates time-varying sales even though the housing stock is fixed. In that model, sales occur because there is partial mismatch between housing taste and housing quality: some young households with high housing needs own small homes, and some old households with low housing needs own large homes. Prices induce some of these households to swap. Their cycle is generated by a spike in the income of young households, which increases their willingness-to-pay for nicer homes. This generates a higher price for the large

[37] Note that in this case, the upward pressure on $\hat{\omega}_t$ does not result in its upward shift: this is impossible due to the irreversibility of housing construction.

homes and induces more old households to trade down, as they can cash out more consumption. Thus, sales rise. In my model, this mismatch does not exist, as those with Existing homes at the beginning of a period all have higher tastes than all non-households.[38] Rather, the increased sales volume comes from the fact that there can be an aggregate up-sizing from construction, which requires churn from individual households as they move up the ladder.

The bottom panel of Figure 16 shows what happens if we allow for new construction but assume that New homes deliver the same flow of utility as Existing homes ($q = 0$). In this case, we get a housing cycle that matches the qualitative facts for prices and construction, but the likelihood of an Existing home selling is not affected. Here, unlike the case with no construction, the boom does induce household formation and construction, but (effectively) the new households move into the New homes,[39] and the owners of Existing homes are once again unaffected by the boom. In this variant of the model, $C_t = P_t$ because the two types of homes are perfect substitutes, so again, the owners of Existing homes are perfectly hedged against price movements: whether they end up moving into Ns or Es, any change in the price of a home they move into will be paid for by a change in the price of the home they sell, and so their incentives for moving are not changed. Therefore, getting the model to produce time-variation in the sale of Existing homes requires not just construction, but a quality ladder, as well.

5 Evidence

This section provides three pieces of evidence that the mechanism highlighted in Section 4 is operative and potentially large. I start at the micro level and document the presence of sale chains—a buyer of a new home selling a different home to another buyer, and that buyer potentially doing the same. This is at the heart of the model, and I show that it is a prominent feature of the data. I then zoom out to the city level and look at a panel of MSAs. I find that localities with more new construction also have more sales of Existing homes, and this holds true even if restricting to variation in new homes coming from plausibly exogenous variation in housing supply elasticity. As the model argues that new construction is an important driver of sales of Existing homes, this is a compelling piece of evidence. Finally, I zoom all the way out to the macro data again and show that the relationship between home price growth and sales has changed since the onset of the housing bust, and it did so in a way that the model predicts quite naturally.

[38] Another way of expressing this distinction is that the result in Ortalo-Magne and Rady (2006) depends on the income shock favoring the owners of the lower tier of housing, so they bid up the price and increase turnover. If the increased willingness-to-pay for housing had come from occupants of the upper tier, sales would *fall* as some young households are not able to induce old households to down-size, as the latter's reservation price has increased. In my model, it does not matter who the shock hits: as long as there is an increase in the willingness-to-pay for housing, it will induce more construction and sales of Existing homes.

[39] Technically, some of the New homes may be moved into by continuing homeowners, who sell off Existing homes. But the key is that their likelihood of moving/incentive to move has not been altered. It is an issue of bookkeeping to say who moves into the New versus the Existing homes, and since the 2 are perfect substitutes, this is not actually pinned down by the model. But while we do not know who exactly moves into the Exisiting homes, we do know how many are sold and that it does not change (in proportion to the housing stock) over time, even during the housing cycle.

5.1 Micro-Evidence on Sale Chains

The mechanism described above relies on households being induced to move into newly-built homes to enjoy the extra quality they deliver, and at the same time selling their previous home. Figure 18 shows that, indeed, New homes are of higher quality than Existing homes. To construct the figure, I take a 10% sample of the deeds data and residualize the (log) prices of those sales on a full set of MSA-by-month indicators. I then plot the distribution of those residuals, split by New and Existing homes. The distribution for New homes is essentially a rightward shift of that for Existing homes, meaning that New homes sold in a given MSA in a given month tend to be more highly-valued than their Existing counterparts. In an reported regression, I find that New homes sell at a premium of about 17%.[40] All of this is hardly surprising, but as the high quality of New homes is an essential feature of the model, this is an important piece of evidence in its favor.

Now, I perform an exercise in the spirit of the seminal filtering study by Lansing et al. (1969) discussed earlier. In particular, I randomly select 200 newly constructed homes in each year from 2001-14. I then look at the name(s) of the buyer(s) of these homes in a fashion similar to Anenberg and Bayer (2016) and DeFusco et al. (2018) and look to see whether they can be found selling a home at roughly the same time. If I do not find them as sellers of another home, then the chain ends. If I do find that they sold a different house, I repeat the process with the buyers of that house and follow the chain until it ends.

Consider a newly-built home ("H_1") that appears in my randomly selected sample. I look to see whether a home ("H_2") is sold within 365 days before or after the sale of H_1 in which the *seller(s)* of H_2 have first and last names that match the buyer(s) of H_1. If there are multiple buyers of H_1, then there must be multiple sellers of H_2 or I do not consider it a match, and vice versa. Furthermore, if middle initials are provided in the records for both H_1 and H_2, then those must match as well. If there are multiple sales that match all of these characteristics, I select the one that is geographically closest to H_1 and determine that that is H_2. Note that, unfortunately, this admits both Type I and Type II errors. In some cases, I will miss matches in which a person has her name recorded differently in different transactions due to a name change, use of a nickname, or a transcription error. On the other hand, I will sometimes incorrectly determine that a match exists when it does not because names do not uniquely identify people, as social security numbers do. Given the limited identifying information in the data, this is a hard issue to solve, though as I will discuss, the results do seem reasonable based on some external evidence. Where possible, I err on the side of Type I errors (missing matches instead of making incorrect matches). For instance, I do not aggressively search for nicknames to match people and, as mentioned above, I do not allow for someone to sell a house by themselves and buy a different home with a partner (perhaps after

[40] Similar results obtain if we use ZIP-by-month indicators, rather than MSA-by-month indicators, though the differences are a bit smaller: the New home premium falls to 13%. This suggests that part of the appeal of New homes is that they are being built in more desirable parts of the MSA. If I do not control for geography at all and instead only residualize on time indicators, there is not much difference from the MSA results—the premium rises to 18%.

a marriage), or vice versa.

Figure 19 shows the distribution of sales in these chains, split by period: 2001-7 versus 2008-14. In all, a bit more than 50% of New homes are bought by people who are (roughly) simultaneously selling an Existing home. These subsequent Existing homes are then often sold to buyers who are also selling off Existing homes (also around 50%). On average, a chain contains 1.2 sales of Existing homes, which is actually quite similar to the 1.4 found in the study by Lansing et al. (1969) back in the 1960s. In a more recent snapshot, the The National Association of Realtors reports that approximately 1/3 of homes are bought by first-time buyers.[41] Given that some buyers are able to buy a home while holding onto one they already own, my results are fairly consistent with this evidence. Furthermore, Anenberg and Bayer (2016) study the Los Angeles metro area specifically and find that around 30% of home buyers are also selling homes in that market. As I do not restrict to within-MSA moves, this also suggests that my matching algorithm is reasonable. Of course, in the model presented Section 4, 100% of New homes were sold to buyers who were selling Existing homes. However, as I discuss in the next section, the fact that many New homes are bought by people not selling Existing homes can accommodated by adding some simple taste shocks to the model.[42]

Further corroborating the mechanism, Figure 20 shows sale prices of homes at different points on the chain. Link 1 refers to the New home that starts the chain, and links 2+ refer to the Existing homes sold as part of the chain. The figure shows that, as expected, the New home sells at a high price, so on average the buyer of this home is forced to cover a shortfall even after selling off the Existing home, as in the model. Interestingly, there is not much difference in the sale prices of subsequent Existing homes in the chain, though the price does tend to be lower for links further down the chain.

5.2 Panel of MSAs

As the model predicts that sales of New homes allow for sales of Existing homes, I now show that when looking across US cities (and city-years), markets with more construction also have more sales of Existing homes, and this holds true when using plausibly exogenous variation in new construction. I focus on the 75 largest MSAs in the country from 2000-14, and all sale counts are normalized by the MSA population in 2000.[43]

I start by looking at within-city variation in sales of New homes across years and find that it is highly predictive of Existing home sales. Figure 21 shows a binned scatterplot where I residualize

[41] See National Association of Realtors (2017).

[42] Another reason that New homes may be bought by someone who is not selling an Existing home is that some New homes are not of high quality. This does not fit into the model so easily, but it is surely true. And in fact, in an unreported regression, I find that New homes with higher sale prices are more likely to be sold to buyers who are selling Existing homes, consistent with this notion.

[43] Note that this part of the analysis—and only this part—CoreLogic's aggregated MSA time series for sales, rather than the microdata. As I am not interested in conditioning on LTV or anything like that, the microdata is not necessary.

MSA-year New and Existing home sales on MSA FEs and year FEs to take out variation that is attributable purely to differences in cities or the stages of the national housing cycles. The relationship is very tight ($R^2 = 0.84$) and fairly strong: one New home sale is associated with 1.43 sales of Existing homes. In other words, in a given city, years with lots of New homes selling also see many sales of Existing homes, and this is not simply driven by high levels of both across the country in the boom.

While this is suggestive of a role for new construction in driving sales of Existing homes, the possibility still exists that a third factor is driving both. For instance, perhaps credit conditions vary over time within city, and years with looser credit make it easier for new homeowners to buy newly built homes and current homeowners to move between Existing homes. As discussed in Section 4, there would be no filtering in this case and therefore the mechanism would not be active, but the relationship in Figure 21 would still appear.

To address this and move toward a more plausibly causal estimate of the effect of new construction on Existing home sales, I exploit variation in the ease of building across cities. In particular, I use the measure of housing supply elasticity estimated in Saiz (2010).[44] The top panel of Figure 22 shows a binned scatterplot of New home sales against elasticity and finds that, indeed, areas with more elastic supply saw more construction.[45] If we think of the elasticity as an instrument for new construction, this shows a relevant first stage. The bottom panel shows that a very similar relationship exists between elasticity and *Existing home sales*, which suggests that indeed, sales of New homes are a factor in determining Existing home sales. Given the slopes of the two lines, it appears that 1 New home sale generates about 2.5 sales of Existing homes.

The first two columns of Table 1 show the formal regressions associated with Figures 21 and 22. Importantly, the standard errors (which are clustered in two ways, by year and MSA) are quite tight, so the relationships discussed above are quite unlikely to be statistical noise. The final three columns of the table introduce January temperature, which Glaeser et al. (2001) found to be a strong predictor of changes in housing demand in the US around this time. I show robustness to using this demand shifter as another source of variation in construction and also as a control to purge the elasticity measure of any differences in demand it might be picking up. Column (3) first controls for January temperature and finds no impact on the results—the effect is still around 2.5 Existing homes sold for 1 New one, and the standard error is reduced a bit. Column (4) uses January temperature as an excluded instrument rather than an included one, and while the point estimate shrinks to around 2.1, the standard error is reduced sharply, so it still supports the existence of a large effect. The causal interpretation of this regression assumes that high demand only affects

[44] This measure has been used to instrument for house price changes during the housing cycle (e.g. Mian and Sufi (2009)). My approach is essentially to look at the other axis of a supply-and-demand chart: if it is a valid instrument for prices, as others have argued, then it should be valid for new construction.

[45] Note that in this part of the analysis, MSA FEs cannot be included, as the elasticity is constant for a given MSA. Year FEs can be included but are irrelevant, as each MSA is included in each year and so there is of course no relationship between year and elasticity. Figure 22 pools all years together, but the relationships are quite stable from year to year, too, as shown in Figure 23.

Existing home sales through its effect on New home sales, whereas Column (3) allowed for demand to affect sales through other channels. Column (5) makes a more subtle assumption, as I allow for January temperature to affect Existing home sales through channels other than construction (i.e. use it as an included instrument), but I use the interaction of January temperature and housing supply elasticity as an excluded instrument. This means that to the extent that increased demand causes more of an increase in Existing home sales in more elastic places, I will attribute that to the effect of construction. The effect falls a bit further to 1.9, again with tight standard errors. Finally, Figure 23 shows what happens if we run this final regression one year at a time. Reassuringly, the estimates are relatively stable from year to year, with the estimates being a bit larger in the second part of the sample, an issue I will take up below.

On the whole, then, the MSA evidence is quite supportive of the effect of new construction in determining Existing home sales. While this reduced form evidence cannot pinpoint the mechanism, the filtering process I have highlighted in this paper is a straightforward and appealing one, especially given the evidence on sale chains shown in the previous subsection.

Before concluding this subsection, I want to emphasize the novelty of these results. Note that housing supply elasticity *negatively* predicts price growth in the boom, as emphasized by Mian and Sufi (2009), so in fact this exogenous measure of new construction is negatively predictive of the price boom. This means that explanations that link the boom and bust in home sales to the boom and bust in prices, such as the literature on speculators detailed above, would predict that, at least in the boom years, the relationship between elasticity and sales would be *negative*, as opposed to the positive relationship in the data and predicted by the filtering model presented in the Section 4. Interestingly, this could be why the relationship strengthens from boom to bust, as the price growth in low elasticity areas evaporates and the sales volume being driven by speculation goes with it. So in fact, one could interpret the movement of the coefficients in Figure 23 to be some evidence in favor of the role of speculation in determining sales volume. Regardless of whether the reader is convinced of this last piece, the first-order fact that elastic supply is predictive of turnover of Existing homes is a robust fact that provides strong support for the model in Section 4.

5.3 The Changing Relationship Between Prices and Volume

In this subsection, I return to the aggregate data and show that the relationship between home prices and home sales has evolved in a way that is consistent with the model presented in this paper. In particular, the model in Section 4 delivers the prediction that the presence of overhang in the housing stock alters the co-movement of prices and sales volume. This is driven by the possibility (and irreversibility) of new construction. When there is no overhang (e.g. before the housing collapse), increases in housing demand lead to more homes being built, so while prices can still increase, part of the increase in the willingness-to-pay for housing goes into new construction. Through the filtering process, this drives sales of Existing homes up as well. So price growth, new construction, and sales of Existing homes all move together. When there is overhang, though, the

construction margin is far less active, and so movements in the willingness-to-pay for housing are felt moreso on the price margin rather than in quantities. Due to the filtering process, this means sales of Existing homes will be sluggish, too. As Figure 16 demonstrates, the model's recovery exhibits fairly low sales volume while having price growth that is higher than it was on the CCP.

Figure 24 shows that the US housing market has followed this path in the recent housing cycle. The top panel shows the relationship between housing starts and price growth, and for the years 1991-2006, there is a very tight, positive relationship between the two, as years with high price growth also had high levels of construction. Following the collapse, however, this changed drastically. Since the crisis, we have seen years of sharp price declines ($\approx -17\%$ in 2008) and increases ($\approx 8\%$ in 2013) and years with more modest increases and declines, but housing starts have not moved much at all in that period. Only in the past few years has new construction begun returning to levels considered normal in the pre-boom period. This is precisely what one would expect in the presence of overhang.

The bottom panel shows that a similar pattern has played out when looking at the sales of Existing homes, though admittedly to a less extreme degree. Again, the years 1991-2006 show a remarkably tight, positive relationship between price growth and sales, and this has changed starkly in the years since, which have seen very different movements in home prices but stable (and low) sales of Existing homes. While there was strong price growth in 2013-4, on par with what was witnessed in the boom years (with the exception of the very peak in 2004-5), sales of Existing homes were quite a bit lower—in line with the more typical levels seen pre-boom in the mid-1990s. Interestingly, as new construction has normalized in the past few years, price growth and sales of Existing homes appear to be returning to their pre-crisis relationship.

The idea that a building boom would generate overhang that suppresses construction in a subsequent bust—and that this can lead to increased price volatility—is not a novel point being made in this paper. It is in fact a fairly direct prediction of basic theory and one that has been used to explain housing market phenomena previously.[46] What is novel is to point out that there is a mechanism through which this overhang also suppresses the sale of Existing homes. And while this analysis focuses on the role of overhang and low construction in keeping home sales low, the other side of the coin is that high sales volume in the boom could have been driven in large part by the boom in construction in those years.

[46] For example, Glaeser and Gyourko (2005) argue that the irreversibility of housing construction can explain the nature of urban decline: as demand in a city declines, there is not a surge in out-migration, but rather a severe drop in home prices, which works to retain residents and keep the homes filled. Liu et al. (2016) argue that new construction during booms helps to keep price growth stable across different tiers of housing, as construction firms respond to differences in prices. During busts, however, when building largely ceases, there can be a good deal more volatility in the relative prices of different tiers of housing.

6 How Large Is the Effect?

In this section, I will use the model in Section 4 as a guide to assess the magnitude of the mechanism explored in this paper for driving variation in sales volume. As the model was intended to highlight the intuition of the mechanism, it is quite parsimonious, and so I defer a more precisely-tuned quantitative exercise to future work. Nonetheless, even with its simplicity, the model is able to deliver significant insight in helping us gauge the magnitude of the effect being studied. The argument in the paper is that construction of New homes causes sales of Existing homes, as households move into the New homes and sell off their previous homes. The model points to 3 key factors that determine how much of the boom and bust in Existing home sales can be explained by the boom and bust in New home sales.

The first is the size of the boom/bust in New home sales compared to that of Existing sales. The larger the relative size of the swings in New home sales, the more plausible it is as a driver of movements in Existing home sales. To assess this, Figure 25 shows the deviations of New and Existing home sales from their trends established during the years 1968-1994. Unsurprisingly, large deviations in one of the series from its trend are typically associated with large deviations for the other series. The recent cycle is no exception. While the movements in Existing sales are larger than New sales, the movement in the latter is substantial. In particular, in the boom years, the deviation of New home sales from its trend is about 30% of that of Existing sales. This number is more like 70% in the post-crisis years. These relatively large magnitudes suggest that, if indeed the filtering process is an important source of home sales, then new construction may be driving a significant share of volatility in sales of Existing homes.

The second key determinant is the share of "marginal" buyers of new construction who simultaneously sell Existing homes. By "marginal," I mean the people who are buying the excess homes that push the New home sale series above its trend. These people are buying New homes *because of* the boom—they would not be moving into New homes if the market were on its CCP. (When New home sales are below trend, the marginal buyers are those who are *not* buying New homes but would have been if the market were on its CCP.) In the model presented in Section 4, this was 100%, as only those with high housing tastes—and thus, current homeowners—would buy New homes. As we saw in Figure 19, this is not the case, as only roughly 50% of buyers of New homes appear to be selling Existing homes (55% in 2001-7, 47% in 2008-14). In Appendix **??**, I make some simple extensions to the model to allow for exogenous entry/exit by households, along with taste shocks that induce some marginal households—those forming for the first time ($\omega_i \approx \hat{\omega}_t$)—to buy New homes. If these taste shocks arrive with probability $(1 - \sigma)$, then σ represents the share of marginal buyers of New homes that are selling Existing homes.

I compute σ in the following way. On average, the annualized rate of New home sales fell by about 650 (thousand) from the period 2001-7 to 2008-14. The share of New home buyers who also sell Existing homes fell from 55% to 47%, so the *number* of New home buyers who also sell Existing homes fell from 571 to 183, or a reduction of 388 (thousand). Therefore, $\sigma = \frac{388}{650} = 0.60$. Note that

if one repeats this exercise at a higher frequency, from year $t-1$ to year t over the course of the sample period, the conclusion does not change much, as most years give estimates of σ near 0.6.

The third and final determinant of home sales in the model is also the most difficult to glean from the data. This is the share of the previous group (marginal buyers of New homes who also sold Existing homes) who would not have moved in the absence of the boom (or when considering the bust, those who do not move that would have moved, if the market were on the CCP). Of the people who were induced to buy New homes by the boom, some would have moved anyway because they had poor match utility to their existing home ($\epsilon_{i,t}^{stay}$), and the only difference the boom caused was that they chose to move into New rather than Existing homes. These households do not contribute to the boom in Existing home sales, as they would have sold Existing homes anyway. On the other hand, some households had match utility that was not bad enough to induce them to move, had the market stayed on the CCP. However, because the boom pushed up the price of their current home and made a New home more affordable, they were induced to sell their Existing home. This group is the heart of the mechanism, as they are selling their Existing homes *because* they have the opportunity to move into New homes. Denote this share by ρ. Figure 26 shows this issue diagramatically. The trapezoid "A" shows that indeed some people who move into New homes have match qualities (ϵ^{stay}) that were not low enough to induce them to move initially, but because of the increase in P_t during the boom, they are better able to afford the cost of a New home, and so they move. These are sales of Existing homes that would not have occurred in the absence of the boom. On the other hand, the rectangle "B" shows people with sufficiently bad match utilities that they were going to move anyway: but during the boom, they move into New rather than Existing homes. As a result, $\rho = \frac{A}{A+B}$. As I alluded to above, there is not much in the data that is capable of informing what ρ should be. As we know $\rho \in [0, 1]$, I will simply show how the magnitude of the mechanism depends on different values of ρ.

If we define $\hat{ES}_t$ and $\hat{NS}_t$ to be the deviation of Existing and New home sales from their trends, respectively, then the model says:

$$\hat{ES}_t = \hat{NS}_t \cdot \sigma \cdot \rho. \tag{29}$$

Figure 27 shows the explanatory power of the model for Existing home sales under different assumed values of ρ. Of course, when $\rho = 0$, so no current homeowners are induced to move by the boom (or deterred by the bust), the model can generate no movement in Existing sales. This is not a particularly plausible case, given the MSA analysis earlier. Furthermore, this is a limiting case of the model when $q = 0$, but it also implies that $C_t = P_t$, and we know that in fact New homes are of higher quality (and more expensive) than Existing homes. Still, it is a useful benchmark. As ρ increases, the ability of the model to explain Existing home sales' deviation from trend also increases, peaking at $\rho = 1$ around 20% for the boom years and 40% for the bust years.

6.1 Amplification

The previous discussion only concerned the initial impulse from the mechanism; to the extent that there are amplification mechanisms active in the housing market, the variation in new construction may be able to generate larger swings in Existing home sales. In Section 4, I described how the search process has the ability to amplify movements in home sales. When $\alpha > 1$, the matching function exhibits increasing returns to scale, which can generate market thickness effects, as explored in Ngai and Tenreyro (2014). Periods with with an initial increase in sales volume endogenously garner yet more sales, as homeowners see that there is more variety out on the market and so expect to get better matches from searching. In the terminology of the model, this means that Conditional Es, who do not care about market conditions when $\alpha = 1$ because they are insulated from price movements due to the fact that they are both buyers and sellers, will now be responsive, as their return from searching increases in the boom and decreases in the bust. As shown in Figure 16, this will increase the volatility of sales. In reality, it could also be the case that homeowners have a strong interest in a particular house (or set of houses), and that once they see it go on the market, they too will sell their homes, as they can now move into their dream home. Finally, there's also the distinct possibility that the quality ladder has more rungs and so construction of homes on higher rungs may generate multiple sales of Existing homes, as opposed to the one assumed in the model.[47]

If we then define θ to be a reduced-form multiplier capturing these effects, we can amend Equation 29 to be the following:

$$\hat{ES}_t = \hat{NS}_t \cdot \sigma \cdot \rho \cdot \theta. \tag{30}$$

This then leaves us with 2 parameters, ρ and θ that are difficult to ascertain empirically. Figure 28 shows the share of movements in Existing home sales that can be explained by the model as a function of ρ and θ. Of course, when ρ is low, the model cannot generate much volatility, as discussed before. But when ρ is relatively large, then with amplification, the model can generate a substantial amount of variation, again particularly in the bust, when the collapse in construction was particularly severe. For instance, with $\theta = 2$, the model can account for over 80% of the low sales of Existing homes in the bust. Of course, it is hard to know if $\theta = 2$ is reasonable. Ngai and Tenreyro (2014) study a seasonal steady state and, in their calibrated model, find that if the exogenous arrival of moving shocks is 25% higher in the summer than the winter (their proposed explanation for this is the school calendar), market thickness effects can make sales volume 120% higher in the summer than the winter, as agents prefer to transact in the thicker summer months.[48]

[47] Note that this is different than chain length. The model generates chains that have more than 1 Existing home sale, but the *causal effect* of 1 New home sale is 1 Existing home sale. The additional home sales in the chain are by people moving within a tier (Existing homes), and so while this can go on without bound, these people would have been moving anyway and so are not part of the causal impact of new construction.

[48] While the "120%" number is not reported explicitly in Ngai and Tenreyro (2014), I back it out. They report that

So in that context, market thickness effects amplify the initial impact by a factor of about 5. Of course, that is quite a different context—moving shocks arrive exogenously and sellers have to decide whether to list in the winter or summer—and it is reasonable to think that the effect is more potent at the seasonal frequency (when the next season is entirely predictable and will arrive relatively quickly) than at the housing cycle frequency. Nonetheless, that paper provides evidence that market thickness effects are likely quite real and potentially powerful, potentially enabling the modest-sized impact of my mechanism on transactions volume to in fact be an important driver of home sale volatility at the housing cycle frequency. As described above, there are other reasons to believe that the true effect of the mechanism may be stronger than what the parsimonious model generates, such as a longer ladder. What I have demonstrated in this section is that, so long as ρ is not very low—so long as a significant set of people is induced to trade up during the boom—the impact of movements in new construction on movements in sales of Existing homes can be substantial.

7 Concluding Remarks

At the heart of this paper is a fairly obvious observation: due to the cyclicality of housing construction, the housing stock changes relatively rapidly in booms, while it is far more static in busts. While this may be obvious, I believe this is the first paper to propose it as a mechanism for explaining the closely related housing market fact that turnover of Existing homes is very cyclical, as well. I borrow the idea of filtering from the affordable housing literature as a way to link aggregate changes in the housing stock to turnover of individual homes: new homes tend to be of high quality, so as more are built, more households are able to trade up into nicer homes, which then of course generates more sales of Existing homes. I argued that the boom and bust in construction was dramatic enough that this is likely an important source of the volatility of Existing home sales.

I do not claim that this mechanism is the *only* driver of sales volume volatility, as many mechanisms have been proposed, and it is likely that a combination of a number of these is necessary to generate the surge and collapse in home sales that we witnessed. In the first half of this paper, I investigated two of these: house lock and nominal loss aversion. I found that while both are consistent with household behavior, the former can account for almost none of the decline in aggregate sales volume from boom to bust, while the latter can account for a bit more, but still not very much.

Looking at all of these results, explanations beyond the three explored in this paper are probably necessary in order to explain the movements in sales volume, particularly the enormous rise in the boom. An important part of this gap is likely filled by speculators—investors who were tempted by rising home prices to enter the market and make a quick profit by simply allowing their new asset to appreciate. As detailed in Section 1, a bevy of papers has found evidence of people doing

"seasonality," which they define as $\left(\frac{N_S}{N_W}-1\right)^2-\left(\frac{N_W}{N_S}-1\right)^2$, where N_S is the number of transactions in summer and N_W is the number of transactions in winter, is equal to 115%. This then implies that $\frac{N_S}{N_W}=2.20$.

just that, particularly at the height of the boom in the mid-2000s. In a different vein, Ngai and Sheedy (2017) argue that the introduction of internet housing search tools reduced search costs during the boom and that rising labor productivity (i.e. wages) and falling interest rates increased households' willingness to bear those costs, both of which increase selling activity. They argue that these factors could be quantitatively important. My hope is that the reader will be convinced that the filtering mechanism, which is related to but distinct from mechanisms in the related literature, is deserving of a meaningful place in this constellation of explanations.

Given the importance I have ascribed to new construction in this paper for the functioning of housing markets, understanding its recent movements is an important avenue for future study. The model in the present paper showed that given the realized path of construction, it was likely a driving force behind the gyrations in overall sales volume. But the recent behavior of new construction has been quite interesting, as it remained depressed for years after the financial crisis and is still subdued today. As many regions in the US grapple with shortages in affordable housing, this lack of construction—and the associated filtering that goes along with it— in the face of steeply rising rents is an issue of great policy concern and may be a fruitful path for future research.

Bibliography

Adelino, M., A. Schoar, and F. Severino (2016): "Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class," *Review of Financial Studies*, 29, 1635–1670.

Albanesi, S., G. De Giorgi, and J. Nosal (2017): "Credit Growth and the Financial Crisis: A New Narrative," NBER Working Paper 23740.

Andersson, F. and T. Mayock (2014): "How Does Home Equity Affect Mobility?" *Journal of Urban Economics*, 84, 23–39.

Anenberg, E. (2011): "Loss Aversion, Equity Constraints and Seller Behavior in the Real Estate Market," *Regional Science and Urban Economics*, 41, 67–76.

Anenberg, E. and P. Bayer (2016): "Endogenous Sources of Volatility and Momentum in Housing Markets: The Joint Buyer-Seller Problem," Nber working paper 18980.

Bachmann, R. and D. Cooper (2014): "The Ins and Arounds in the U.S. Housing Market," Research department working paper series no. 14-03, Federal Reserve Bank of Boston.

Bayer, P., C. Geissler, K. Mangum, and J. Roberts (2017a): "Speculators and Middlement: The Strategy and Performance of Invesetors in the Housing Market," Nber working paper 16784.

Bayer, P., K. Mangum, and J. Roberts (2017b): "Speculative Fever: Investor Contagion in the Housing Boom," Nber working paper 22065.

Bhutta, N., J. Dokko, and H. Shan (2017): "Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust," *The Journal of Finance*, 72, 2433–2466.

Bond, E. and N. E. Coulson (1989): "Externalities, Filtering, and Neighborhood Change," *Journal of Urban Economics*, 26, 231–249.

Bracke, P. and S. Tenreyro (2018): "History Dependence in the Housing Market," *Bank of England Staff Working Paper (No. 630).*

Burnside, C., M. Eichenbaum, and S. Rebelo (2016): "Understanding Booms and Busts in Housing Markets," *Journal of Political Economy*, 124, 1088–1147.

Campbell, J. and J. Cocco (2015): "A Model of Mortgage Default," *Journal of Finance*, 70, 1495–1554.

Case, K. and R. Shiller (2003): "Is There a Bubble in the Housing Market?" *Brookings Papers on Economic Activity*, 2, 299–342.

Case, K., R. Shiller, and A. Thompson (2012): "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets," *Brookings Papers on Economic Activity*, 45, 265–298.

DeFusco, A., C. Nathanson, and E. Zwick (2018): "Speculative Dynamics of Prices and Volume," Nber working paper 23449.

Engelhardt, G. (2003): "Nominal Loss Aversion, Housing Equity Constraints, and Household Mobility: Evidence from the United States," *Journal of Urban Economics*, 53.

Ferreira, F., J. Gyourko, and J. Tracy (2010): "Housing Busts and Household Mobility," *Journal of Urban Economics*, 68, 34–45.

——— (2012): "Housing Busts and Household Mobility: An Update," *FRBNY Economic Policy Review.*

Foote, C., K. Gerardi, and P. Willen (2012): "Why Did So Many People Make So Many Ex Post Bad Decisions: The Causes of the Foreclosure Crisis," in *Rethinking the Financial Crisis*, ed. by A. Blinder, A. Lo, and R. Solow, New York, NY and Washington, DC: The Russell Sage Foundation and the Century Foundation, chap. 6, 136–186.

Foote, C., L. Loewenstein, and P. Willen (2016): "Cross-Sectional Patterns of Mortgage Debt during the Housing Boom: Evidence and Implications," Federal Reserve Bank of Boston Working Paper No. 16-12.

Genesove, D. and C. Mayer (1997): "Equity and Time to Sale in the Real Estate Market," *American Economic Review*, 87, 255–269.

——— (2001): "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *The Quarterly Journal of Economics*, 116, 1233–1260.

Gerardi, K., A. Lehnert, S. Sherlund, and P. Willen (2008): "Making Sense of the Subprime Crisis," *Brookings Papers on Economic Activity*, 39, 69–159.

Glaeser, E., J. Gottlieb, and J. Gyourko (2013): "Can Cheap Credit Explain the Housing Boom?" in *Housing and the Financial Crisis*, ed. by E. Glaeser and T. Sinai, Chicago: University of Chicago Press, chap. 7, 301–359.

Glaeser, E. and J. Gyourko (2005): "Urban Decline and Durable Housing," *Journal of Political Economy*, 113, 345–375.

Glaeser, E., J. Kolko, and A. Saiz (2001): "Consumer City," *Journal of Economic Geography*, 1, 27–50.

Guiso, L., P. Sapienza, and L. Zingales (2013): "The Determinants of Attitudes toward Strategic Defaults on Mortgages," *Journal of Finance*, 68, 1473–1515.

Guren, A. (2018): "House Price Momentum and Strategic Complementarity," *Journal of Political Economy*, 126, 1172–1218.

Haughwout, A., D. Lee, J. Tracy, and W. Van der Klaauw (2011): "Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis," Staff report 514, Federal Reserve Bank of New York.

Head, A., H. Lloyd-Ellis, and H. Sun (2014): "Search, Liquidity, and the Dynamics of House Prices and Construction," *American Economic Review*, 104, 1172–1210.

Kaplan, G., K. Mitman, and G. Violante (2017): "The Housing Boom and Bust: Model Meets Evidence," NBER Working Paper 23694.

Krainer, J. (2001): "A Theory of Liquidity in Residential Real Estate Markets," *Journal of Urban Economics*, 49, 32–53.

Lansing, J., C. Clifton, and J. Morgan (1969): *New Homes and Poor People: A Study of Chains of Moves*, Ann Arbor, MI: Institute for Social Research, University of Michigan.

Lee, K. O. and G. Painter (2013): "What Happens to Household Formation in a Recession?" *Journal of Urban Economics*, 76, 93–109.

Liu, C., A. Nowack, and S. Rosenthal (2016): "Housing Price Bubbles, New Supply, and Within-City Dynamics," *Journal of Urban Economics*, 96, 55–72.

Loewenstein, L. (2018): "Consumption of Housing During the 2000s Boom: Theory and Evidence," Working Paper.

Mian, A. and A. Sufi (2009): "The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis," *The Quarterly Journal of Economics*, 124, 1449–1496.

——— (2017): "Household Debt and Defaults from 2000 to 2010: The Credit Supply View," in *Evidence and Innovation in Housing Law and Policy*, ed. by L. A. Fennell and B. Keys, Cambridge: Cambridge University Press, chap. 11, 257–289.

——— (2018): "Credit Supply and Housing Speculation," Nber working paper 24823.

Nathanson, C. and E. Zwick (Forthcoming): "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market," *The Journal of Finance*.

National Association of Realtors (2017): "2017 Profile of Home Buyers and Sellers," Tech. rep.

Ngai, L. R. and K. Sheedy (2017): "The Decision to Move House and Aggregate Housing-Market Dynamics," Working Paper.

Ngai, L. R. and S. Tenreyro (2014): "Hot and Cold Seasons in the Housing Market," *American Economic Review*, 104, 3991–4026.

Novy-Marx, R. (2009): "Hot and Cold Markets," *Real Estate Economics*, 37, 1–22.

Ortalo-Magne, F. and S. Rady (2006): "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints," *The Review of Economic Studies*, 73, 459–485.

Rosenthal, S. (2014): "Are Private Markets and Filtering a Viable Source of Low-Income Housing? Estimates from a 'Repeat Income' Model," *American Economic Review*, 104, 687–706.

Saiz, A. (2010): "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125, 1253–1296.

Schulhofer-Wohl, S. (2012): "Negative Equity Does Not Reduce Homeowners' Mobility," *Federal Reserve Bank of Minneapolis Quarterly Review*, 35.

Stein, J. (1995): "Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects," *The Quarterly Journal of Economics*, 110, 379–406.

Wheaton, W. (1990): "Vacancy, Search, and Prices in a Housing Market Matching Model," *Journal of Political Economy*, 98, 1270–1292.

Figures

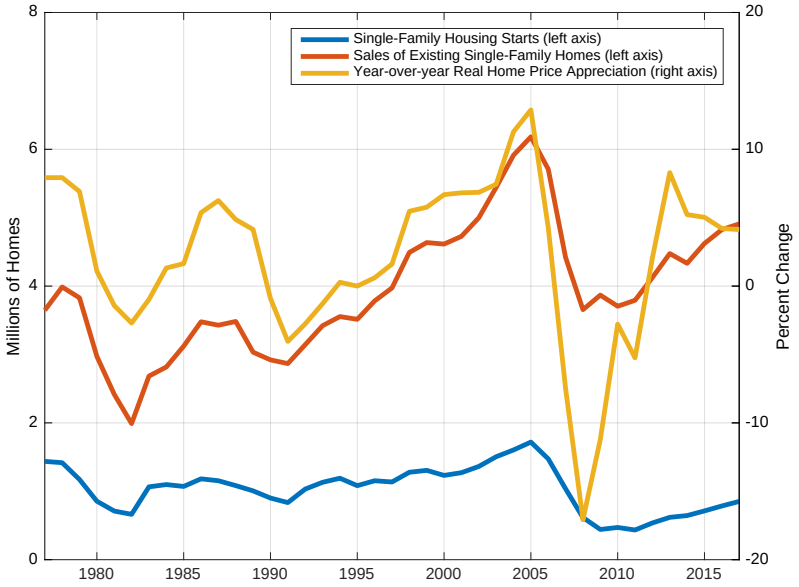

Figure 1: Housing starts for single-family homes (from Census Bureau) and sales of existing single-family homes (from National Assocation of Realtors) are plotted against the lefthand axis, and year-over-year growth of CoreLogic's national home price index is shown on the righthand axis.

	(1)	(2)	(3)	(4)	(5)
New Home Sales	1.43	2.59	2.54	2.05	1.91
(SE)	(0.27)***	(0.54)***	(0.43)***	(0.26)***	(0.33)***
Elasticity	No	Ex.	Ex.	Ex.	Ex.
Temperature	No	No	In.	Ex.	In.
Elas·Temp	No	No	No	Ex.	Ex.

Table 1: Regressions of Existing home sales on New home sales across 75 MSAs, from 2000-2014. Sales are normalized by MSA population in 2000. Column (1) uses OLS. Columns (2)-(5) use 2SLS, with a varied set of instruments. "Ex." means the variable is an excluded instrument (i.e. appears only in the first stage), and "In." means it is an included instrument (i.e. appears in both stages). "Elasticity" refers to the housing supply elasticity measure reported in Saiz (2010). "Temperature" refers to average January temperature. Year FEs are included, and standard errors use two-way clustering at the year and MSA levels. *** indicates significance at the 1% level.

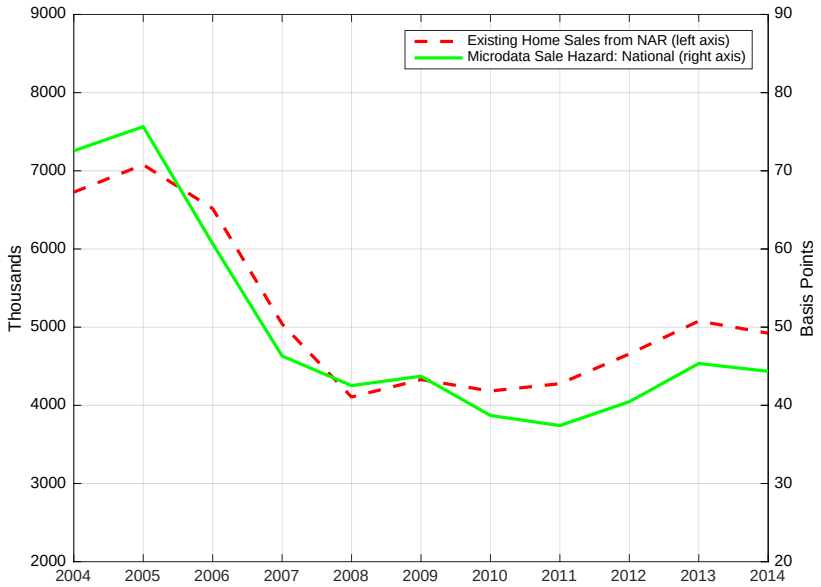

Figure 2: Comparison of National Association of Realtor's count of existing sales with the sale hazard from the final sample of microdata.

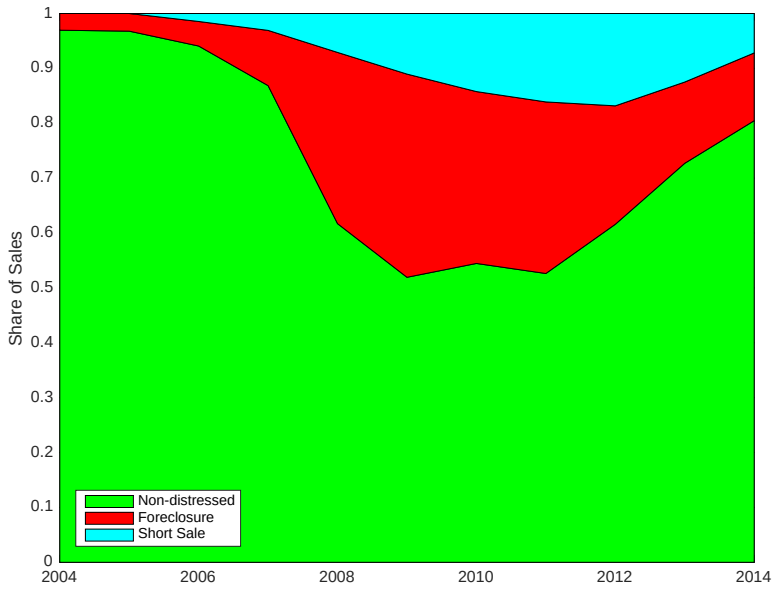

Figure 3: Sales in the deeds data are either non-distressed, short sales, or sales out of foreclosure. I show how the share of sales in each category evolves over time.

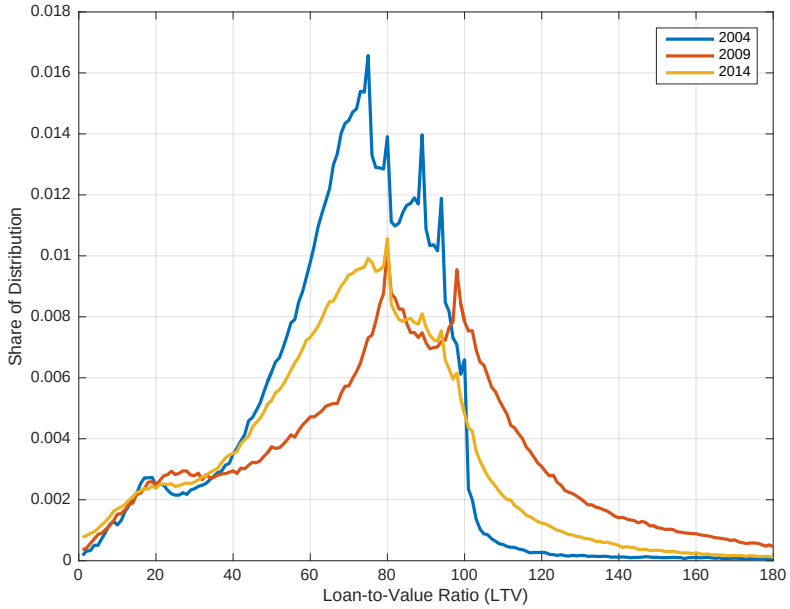

Figure 4: LTV distribution in estimation sample for selected years.

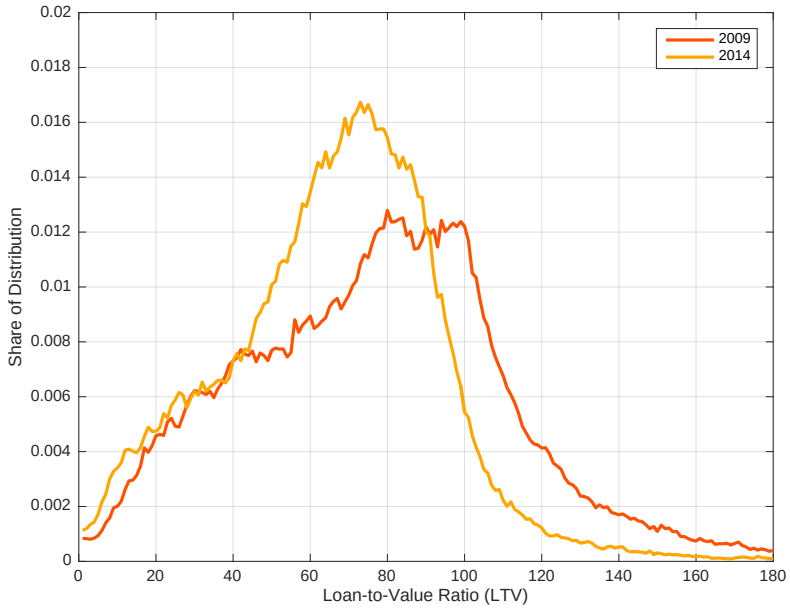

Figure 5: LTV distribution in CRISM dataset for selected years.

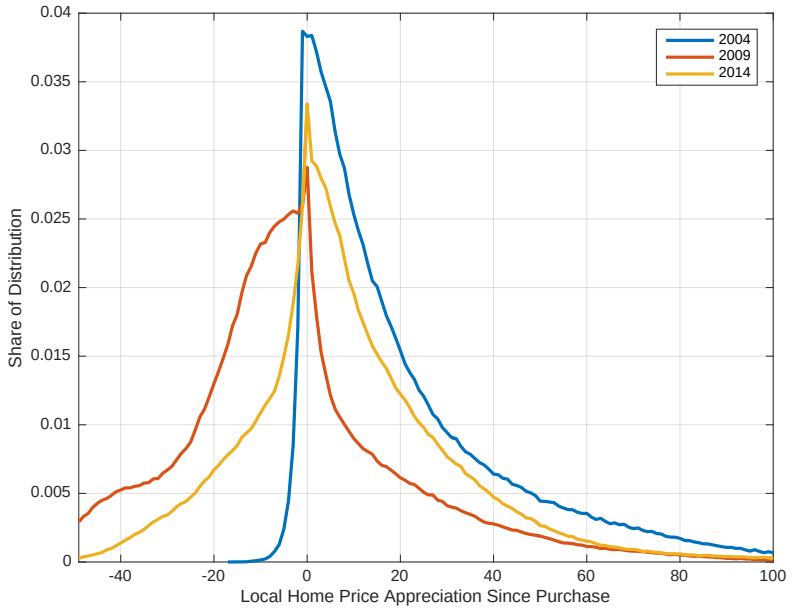

Figure 6: Distribution of local home price appreciation since purchase in estimation sample for selected years.

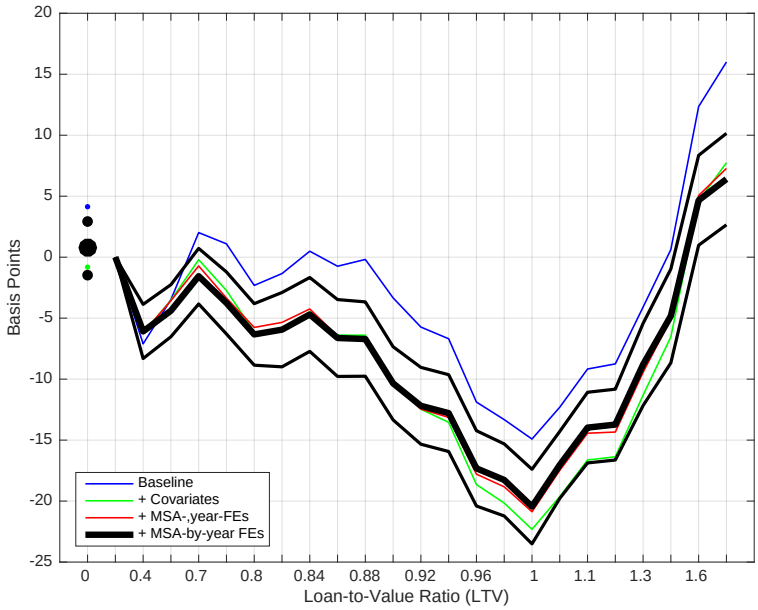

Figure 7: Relationship between sale probability and LTV, 1% national sample. Results come from a linear probability model. All specifications include months since purchase ("duration"), calendar month (to soak out any effects of seasonality), and a flexible step function for nominal gains/losses. The "Baseline" specification controls for only these. The specification "+ Covariates" further controls for county-level economic conditions (county unemployment rate 6 months prior, change in the county unemployment rate from 12 to 6 months prior, the gross number of jobs created in the county in that year [expressed as a percentage of employment]), ZIP code-level measures of financial health (average FICO score on purchase mortgages in 2003, average household income in 2004, and share of adults with at least a college degree in the 2000 Census), and some property level controls (initial down-payment and percentile rank within MSA of property's value). The specification "+ MSA-,year-FEs" adds in indicators for each MSA and year in the sample, while the final specification, "+ MSA-by-year FEs" adds in indicators for each MSA in each year. Standard errors are shown for the final specification, and they are clustered at the MSA-level level. There are 8 bins for nominal gains/losses: [0], (0,0.2], (0.2,0.4], (0.4,0.6], (0.6,0.7], (0.7,0.75], (0.75,0.8], (0.8,0.82], (0.82,0.84], (0.84,0.86], (0.86,0.88], (0.88,0.9], (0.9,0.92], (0.92,0.94], (0.94,0.96], (0.96,0.98], (0.98,1], (1,1.05], (1.05,1.1], (1.1,1.2], (1.2,1.3], (1.3,1.4], (1.4,1.6], (1.6,∞).

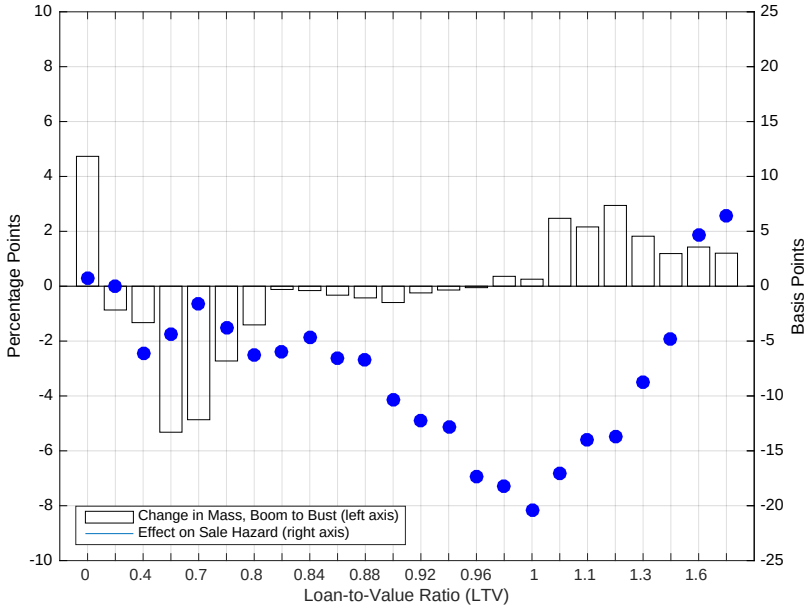

Figure 8: The relationship between the probability of sale and LTV is shown in the blue dots. The white bars show the change from 2004-6 to 2007-11 in the percent of borrowers in each bin.

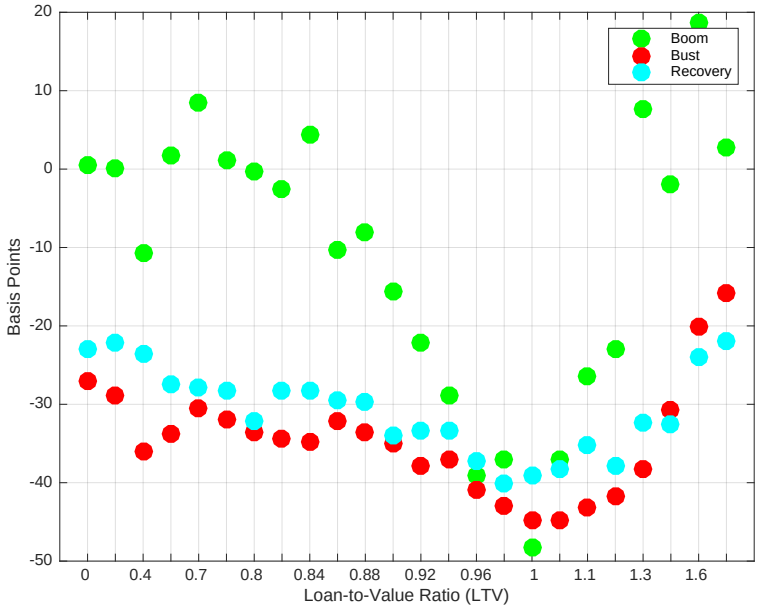

Figure 9: The time-varying relationship between the probability of sale and LTV. The "boom" (2004-6) is shown in green dots, the "bust" (2007-11) is shown in red dots, and the "recovery" (2012-4) is shown in blue dots. The estimates come from a regression with all covariates described in the caption of Figure 7 and MSA fixed effects).

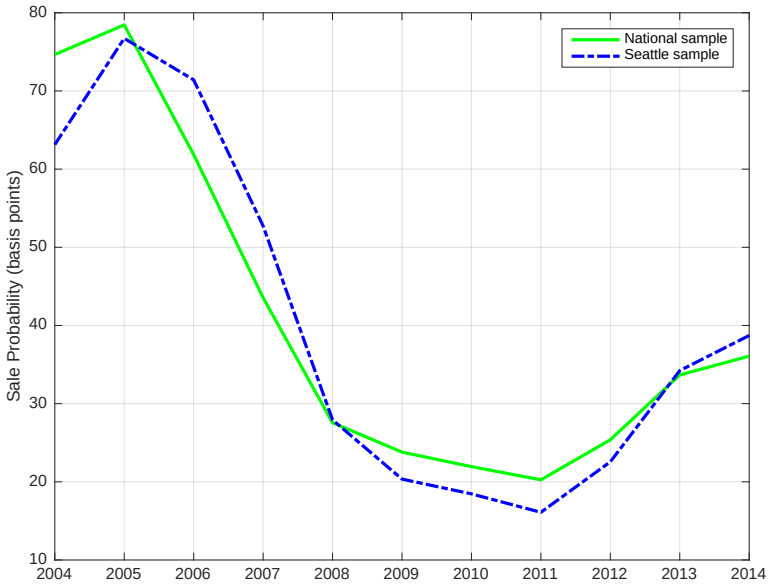

Figure 10: Hazard of a non-distressed sale in the deeds microdata, for both the 1% U.S. sample and the 100% Seattle-Tacoma-Bellevue sample.

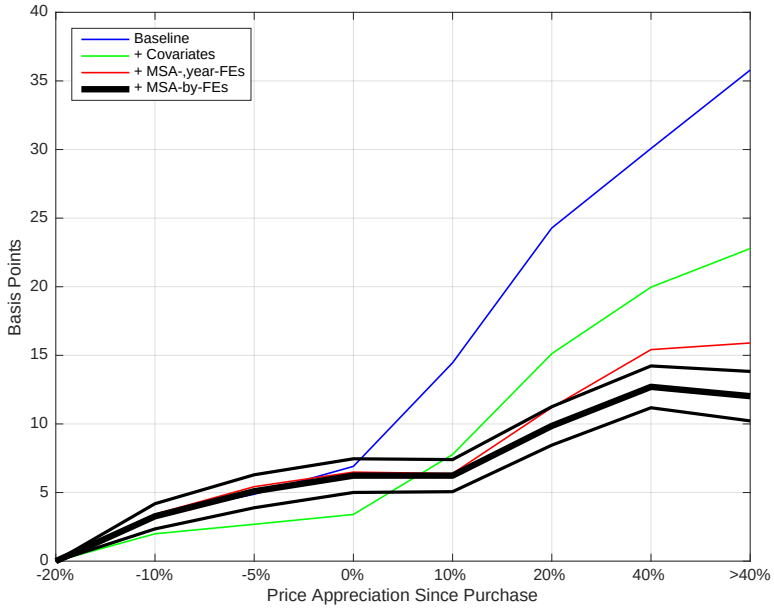

Figure 11: Relationship between non-distressed sale probability and nominal gains/losses, 1% national sample. Results come from a linear probability model. All specifications include months since purchase ("duration"), calendar month (to soak out any effects of seasonality), and a flexible step function for LTV. The "Baseline" specification controls for only these. The specification "+ Covariates" further controls for county-level economic conditions (county unemployment rate 6 months prior, change in the county unemployment rate from 12 to 6 months prior, the gross number of jobs created in the county in that year [expressed as a percentage of employment]), ZIP code-level measures of financial health (average FICO score on purchase mortgages in 2003, average household income in 2004, and share of adults with at least a college degree in the 2000 Census), and some property level controls (initial down-payment and percentile rank within MSA of property's value). The specification "+ MSA-,year-FEs" adds in indicators for each MSA and year in the sample, while the final specification, "+ MSA-by-year FEs" adds in indicators for each MSA in each year. Standard errors are shown for the final specification, and they are clustered at the MSA-level level. There are 8 bins for nominal gains/losses: ($-\infty$,-20%], (-20%,-10%], (-10%,-5%], (-5%,0%], (0%,10%], (10%,20%], (20%,40%], (40%,∞).

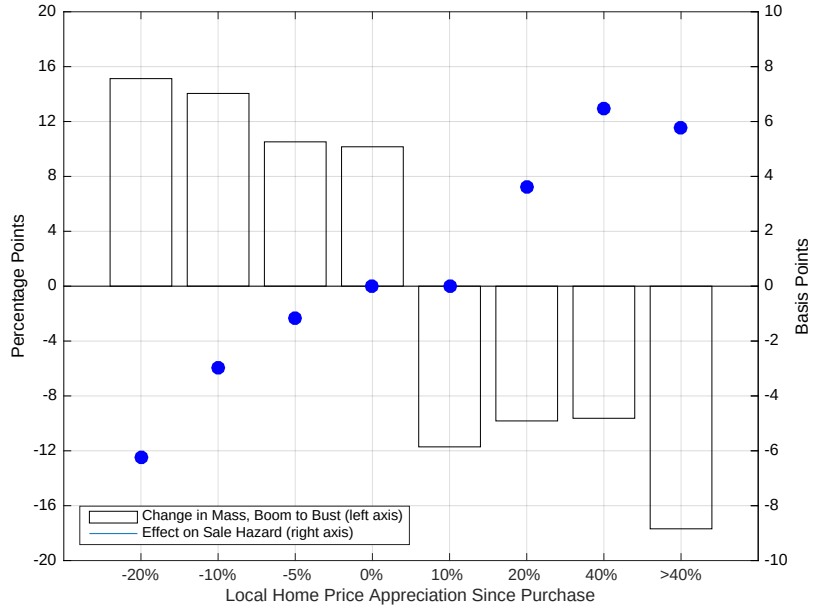

Figure 12: The relationship between the probability of sale and nominal gains/losses is shown in the blue dots. The white bars show the change from 2004-6 to 2007-11 in the percent of borrowers in each bin.

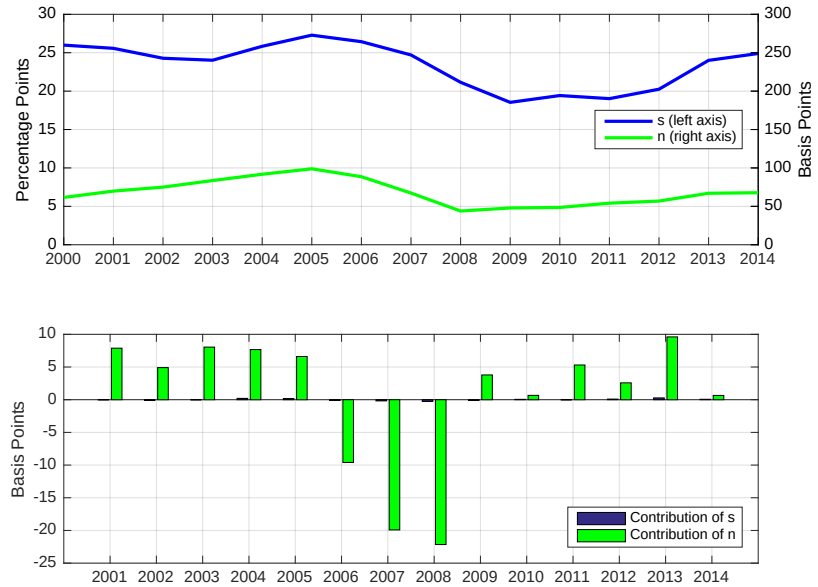

Figure 13: The top panel shows the variation in the probability in the Seattle sample that a listed home sells (s) and that an unlisted home is listed (n). (Note that the 2 series are plotted on different axes.) The bottom panel shows the share of the change in the overall sale hazard that can be attributed to changes in each s and n.

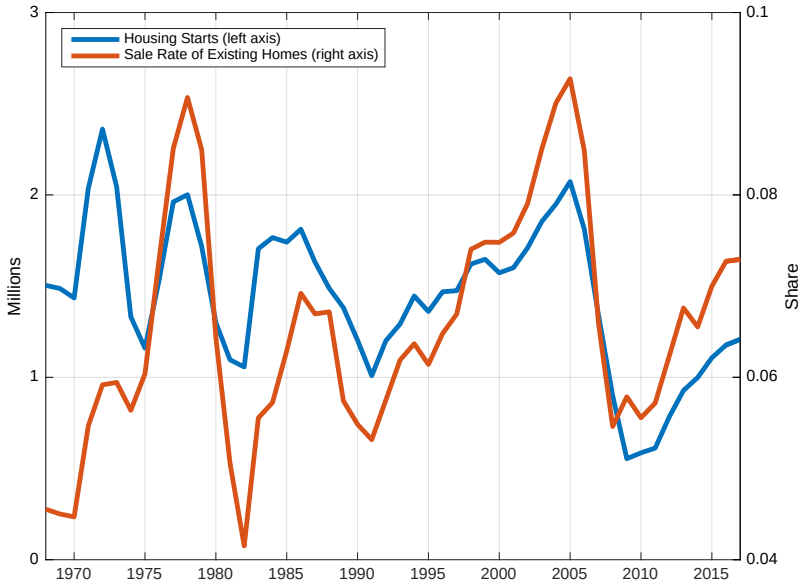

Figure 14: Housing starts and turnover rate for existing homes. Housing starts include single-family and multi-family homes, as reported by the Census Bureau. For the years 1999-2017, turnover of existing homes is measured as sales of single-family and multi-family homes from the National Association of Realtors (NAR) divided by the number of owner-occupied homes, as reported by the Census Bureau. The series on multi-family home sales is absent for 1968-1998, so I scale up the single-family series based on the relationship between the two series for 1999-2017.

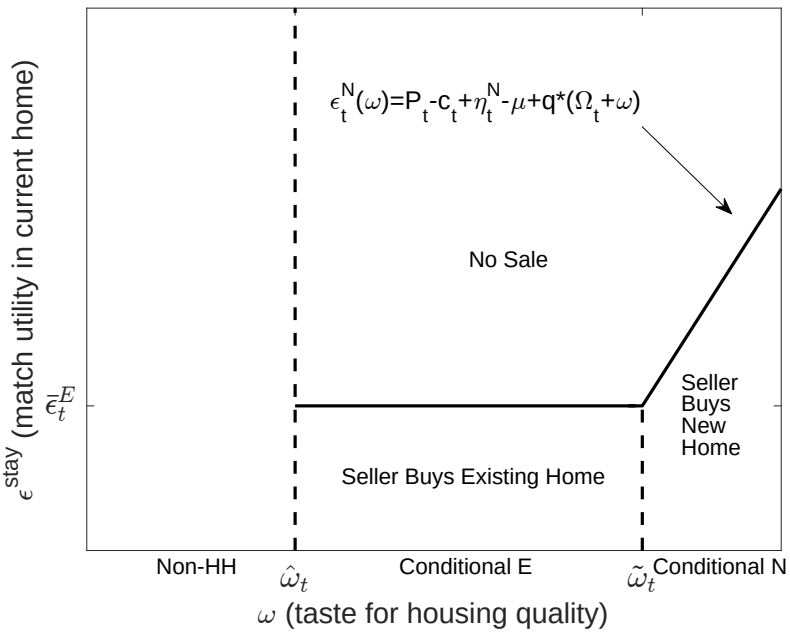

Figure 15: Graphical representation of the household's solution. The horizontal axis represents taste for housing quality and the vertical axis represents match utility from current home. The 4 cutoff rules are represented by $\hat{\omega}_t$, $\tilde{\omega}_t$, $\hat{\epsilon}_t^E$, and $\hat{\epsilon}_t^N$.

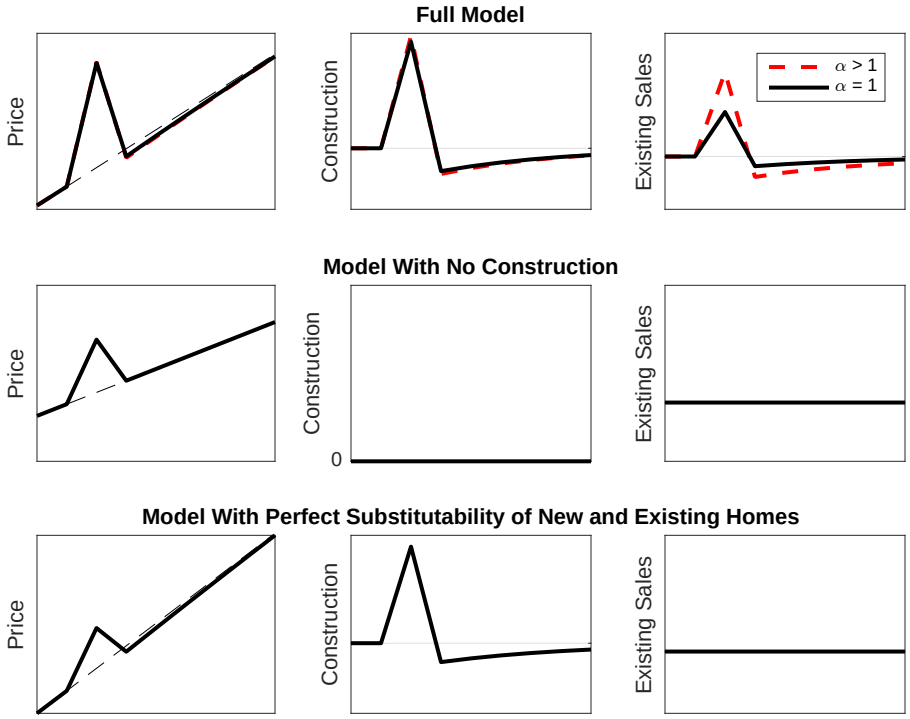

Figure 16: Market responses to a temporary expectations shock in the model. The top panel shows the full model (with construction of New, high-quality homes) under 2 scenarios: $\alpha = 1$ uses the baseline search model with no market thickness effects; $\alpha > 1$ allows for thick market effects. The second panel shows how the model responds when the housing stock is fixed. The bottom panel shows how the model responds when construction is allowed but New homes and Existing homes deliver the same flow of utility. The bottom 2 panels assume $\alpha = 1$.

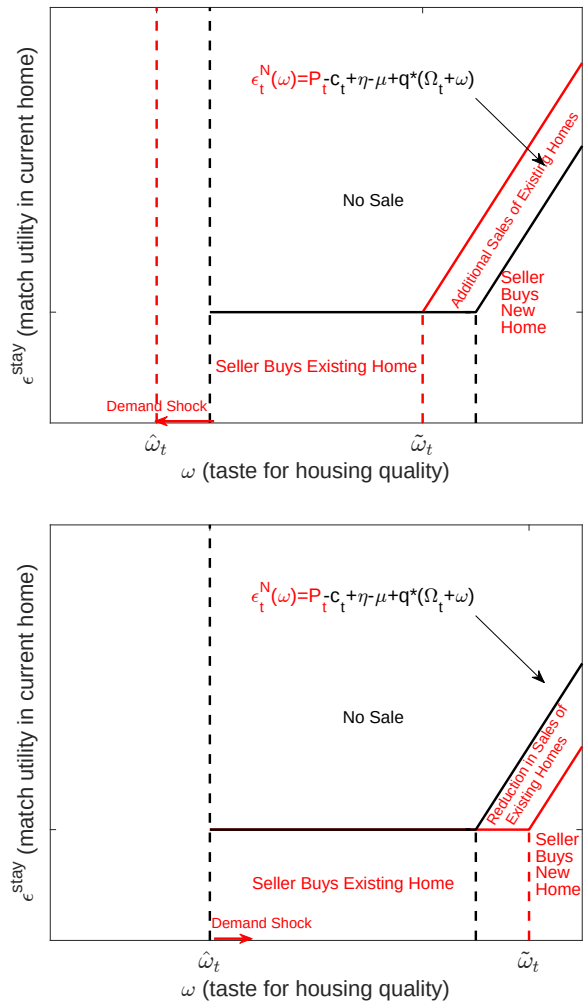

Figure 17: A diagrammatic depiction of the model's responses to a positive (top panel) and negative (bottom panel) demand shock.

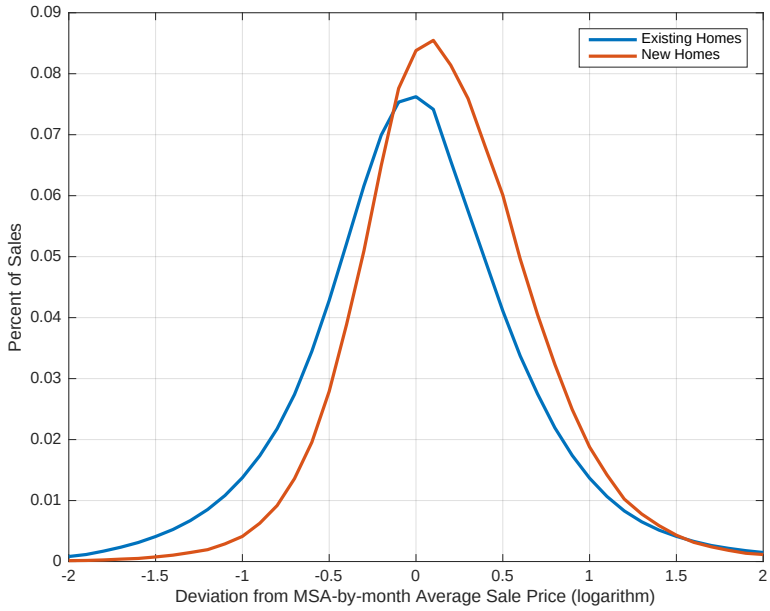

Figure 18: Distributions of log prices of New and Existing home sales, residualized on MSA-by-month indicators. This uses a 10% random sample of deeds records.

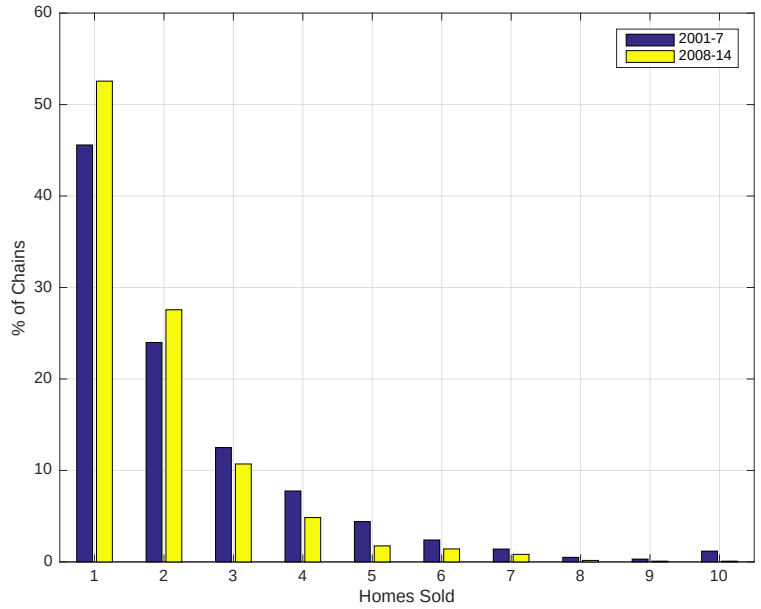

Figure 19: Distribution of homes sold in sale chains constructed from micro-data. For each year from 2001-2014, I select 200 random newly-built homes and then look to see if their buyers were selling other homes contemporaneously. If so, look at the buyers of those homes, and repeat the process. The horizontal axis shows the number of homes sold in these chains, *including the newly-built home that starts the chain.* Each chain is weighted by the number of new homes sold in the year that its new home sold in, according to the National Association of Realtors series.

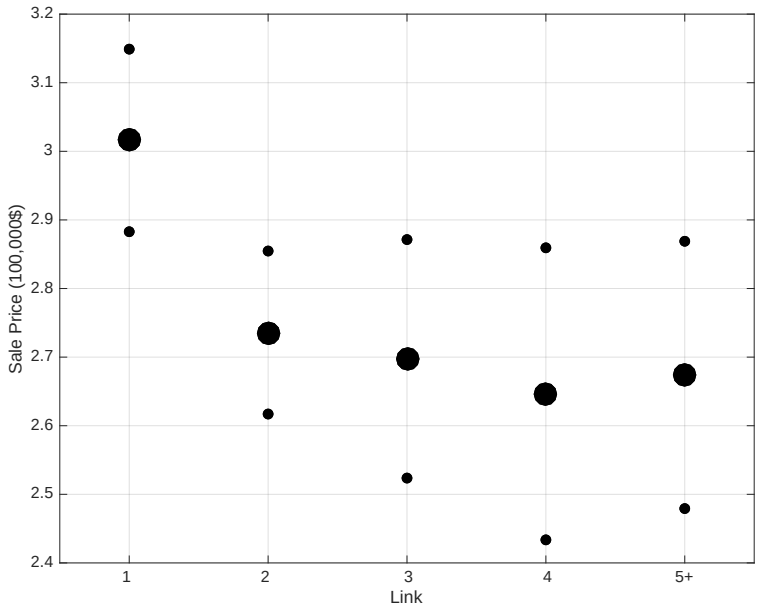

Figure 20: Sale price of homes in sale chains with at least 2 links. Link 1 refers to the first home, which is newly built, while subsequent links refer to Existing homes. 95% confidence intervals are shown.

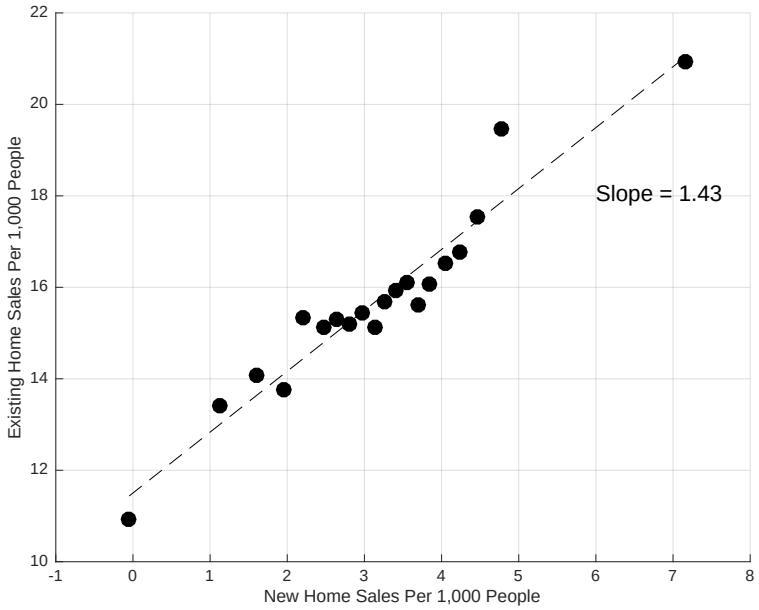

Figure 21: Relationship between sales of New homes and sales of Existing homes across 75 MSAs, from 2000-2014. Sales are normalized by MSA population in 2000. All observations are residualized on MSA FEs and year FEs. Observations are then put into 20 equally-sized bins based on horizontal axis, and the average of value of New and Existing homes is shown for each bin, with the mean of each added back in. (This is the commonly used "binned scatterplot.") All regressions and averages use weight MSAs by their population in 2000.

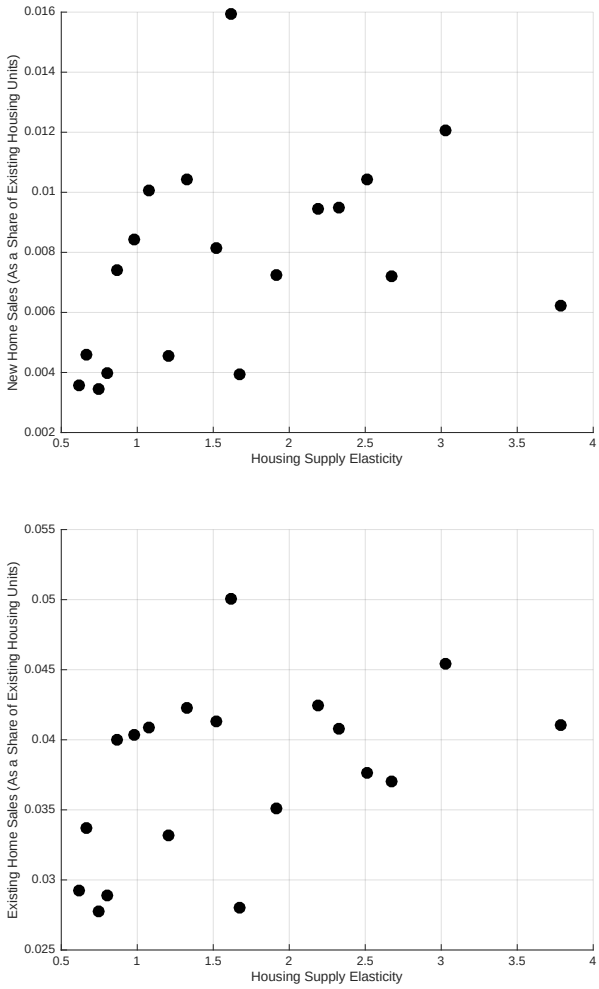

Figure 22: Relationship between housing supply elasticity as reported in Saiz (2010) and sales of New homes (top panel) and sales of Existing homes (bottom panel) across 75 MSAs, from 2000-2014. Sales are normalized by MSA population in 2000. All observations are residualized on year FEs. Observations are then put into 20 equally-sized bins based on horizontal axis, and the average of value of elasticity and the measure of sales is shown for each bin, with the mean of each added back in. (This is the commonly used "binned scatterplot.") All regressions and averages use weight MSAs by their population in 2000.

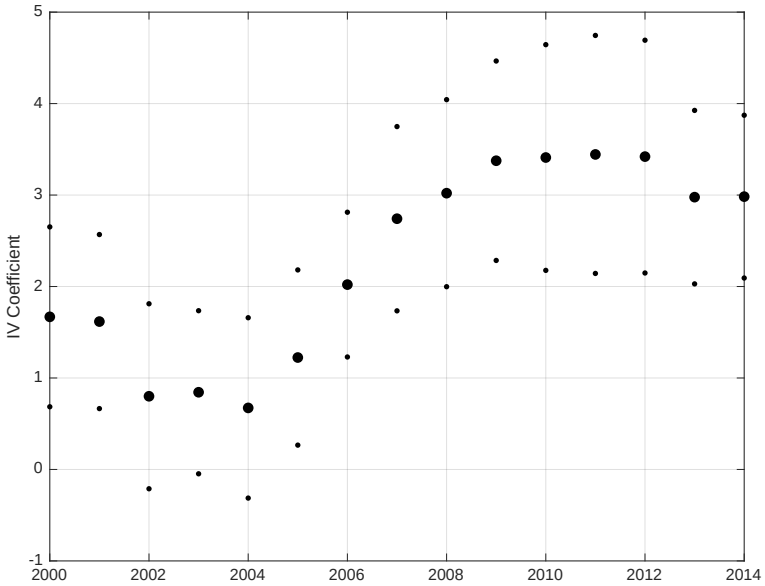

Figure 23: Coefficients from IV regressions of Existing home sales on New home sales across 75 MSAs, run one year at a time. Results correspond to specification (5) in Table 1. 95% confidence intervals are shown.

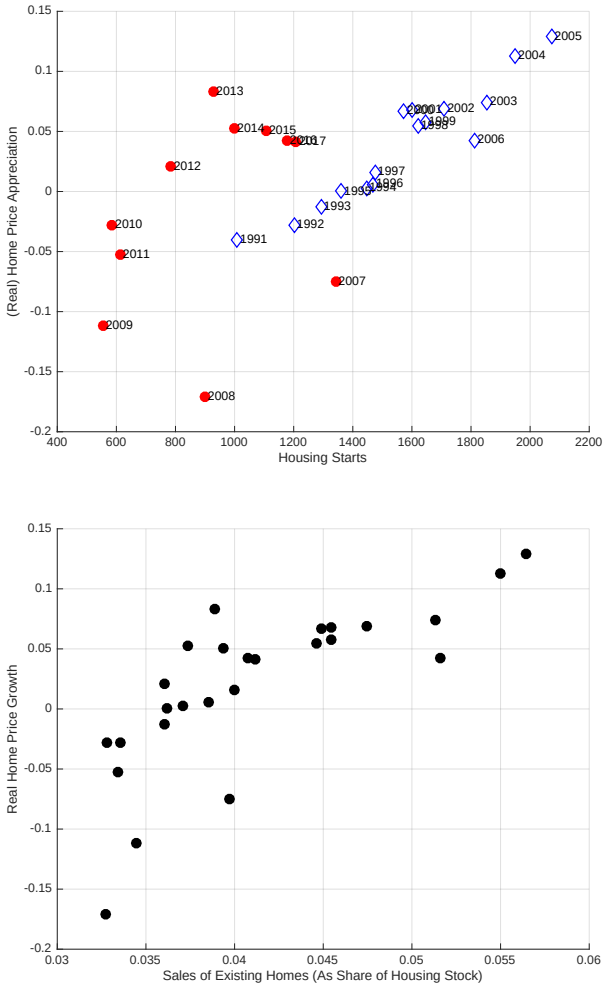

Figure 24: Scatterplot of measures of housing market activity and real home price growth, 1991-2017. The top panel shows housing starts on the horizontal axis, while the bottom panel shows the share of Existing homes sold. I split the sample into a period before the start of the national decline in home prices (1991-2006) and after (2007-2017).

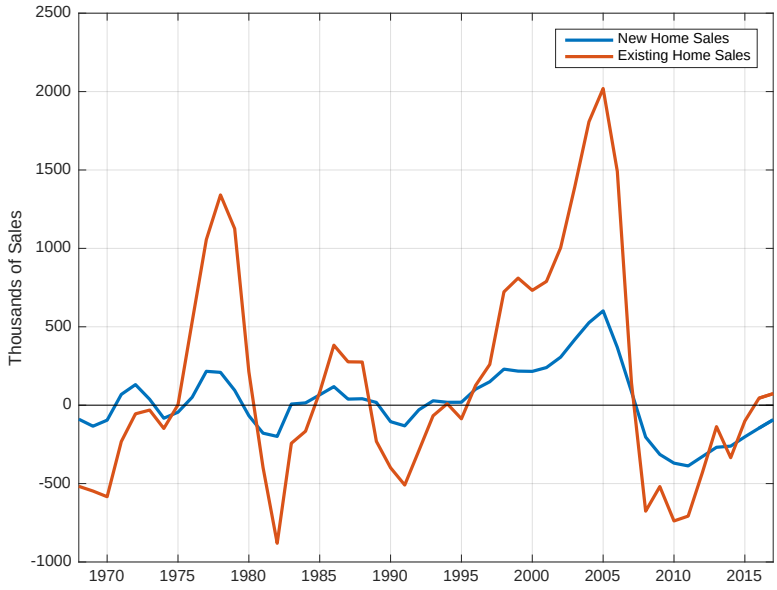

Figure 25: De-trended New and Existing home sales. Trend is calculated based on 1968-1994 data.

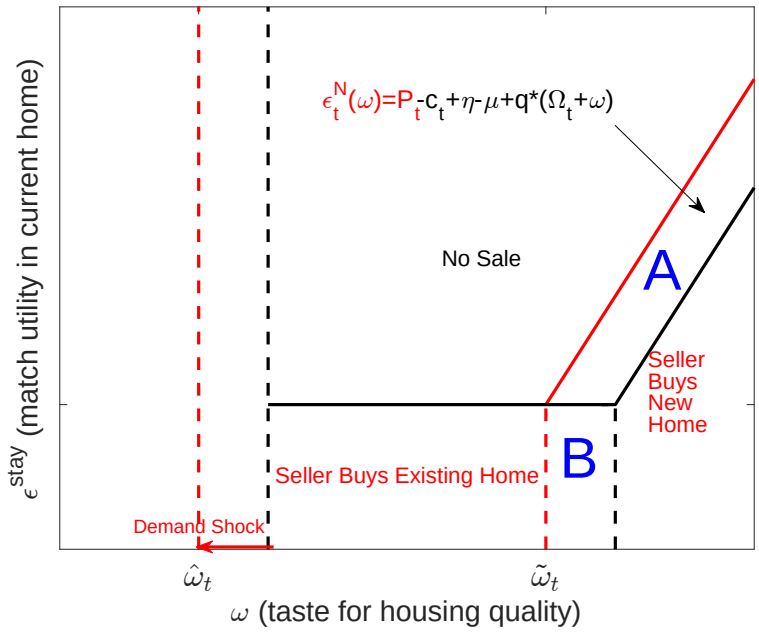

Figure 26: Graphical representation of the household's solution. The horizontal axis represents taste for housing quality and the vertical axis represents match utility from current home. The 4 cutoff rules are represented by $\hat{\omega}_t$, $\tilde{\omega}_t$, $\hat{\epsilon}_t^E$, and $\hat{\epsilon}_t^N$. The trapezoid labeled "A" represents people who move into New homes and sell Existing homes *who would not have moved at all in the absence of the boom*. The rectangle "B" represents people who who move into New homes and sell Existing homes, but they would have sold their Existing homes even in the absence of the boom, as they had quite bad match utilities to their current homes (ϵ^{stay}).

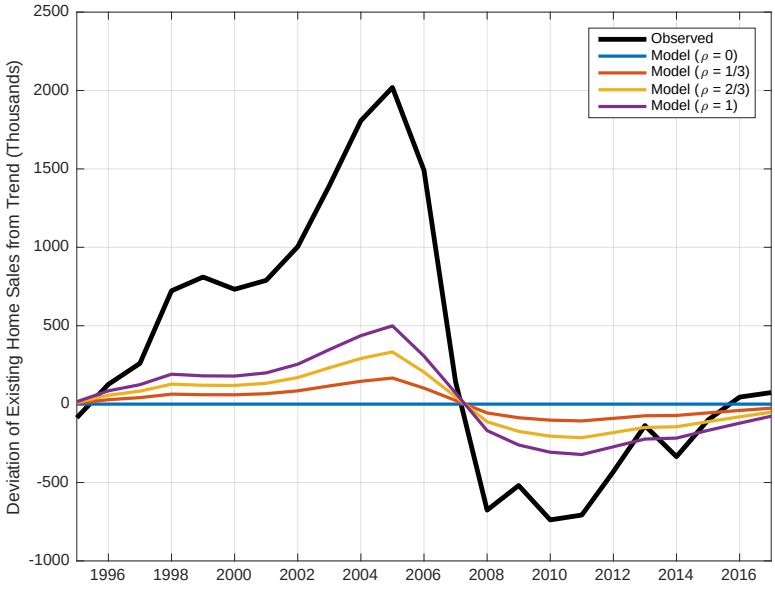

Figure 27: Time series of Existing home sales in the model, given the time series of New home sales in the data. This is computed from Equation 29, with $\sigma = 0.60$. Results are shown for different levels of ρ. The actual movement in Existing home sales is shown as well.

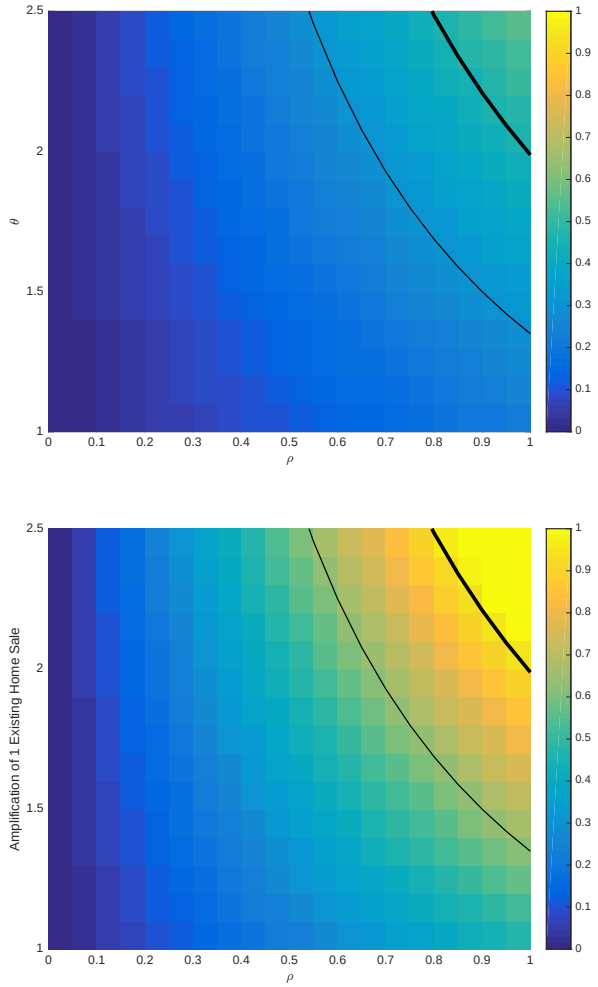

Figure 28: Share of variation in Existing home sales in the boom (top panel) and bust (bottom panel) that can be explained by the mechanism, as a function of 2 unknown sufficient statistics: the share of marginal buyers of New homes selling Existing homes that *also* would not have moved if the market had remained on the CCP (ρ); and the reduced form sale multiplier (θ). $\rho \in [0, 1]$, so I show its whole support. θ is unbounded; I simply show $\theta \in [1, 2]$. This is based on Equation 30, using the observed values of $\frac{\hat{NC}_t}{\hat{ES}_t}$ as 0.3 in the boom and 0.7 in the bust, and setting $\sigma = 0.60$, as discussed in the text.

Appendix

A.1 Empirical Results for Seattle-Tacoma-Bellevue MSA

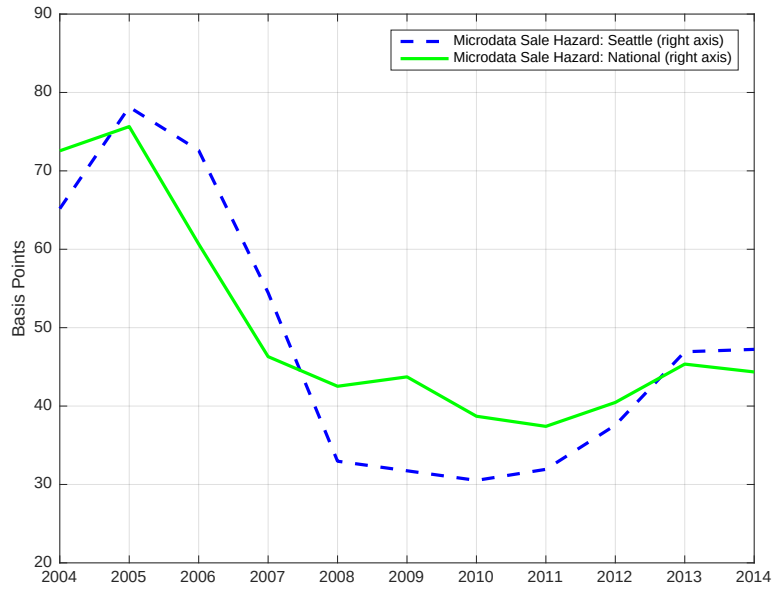

Figure A-1: Comparison of sale hazards in the Seattle and national samples.

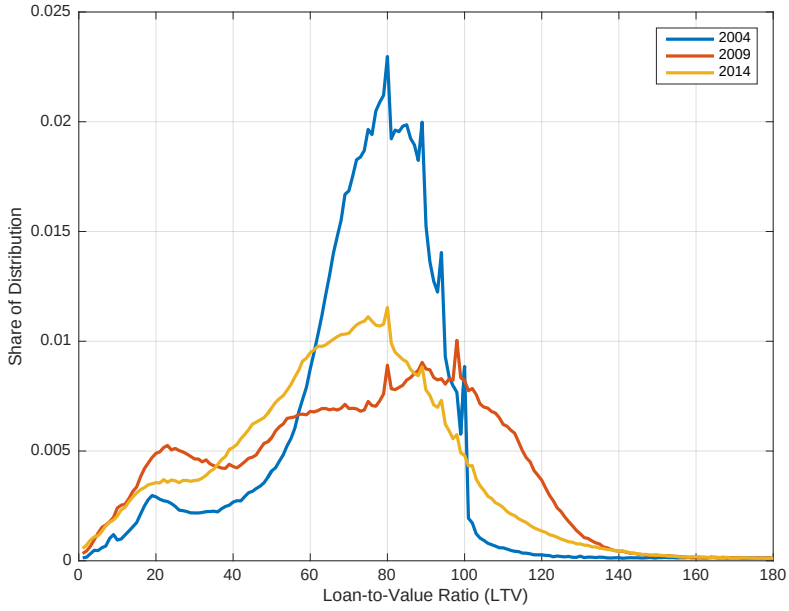

Figure A-2: LTV distribution in Seattle sample for selected years.

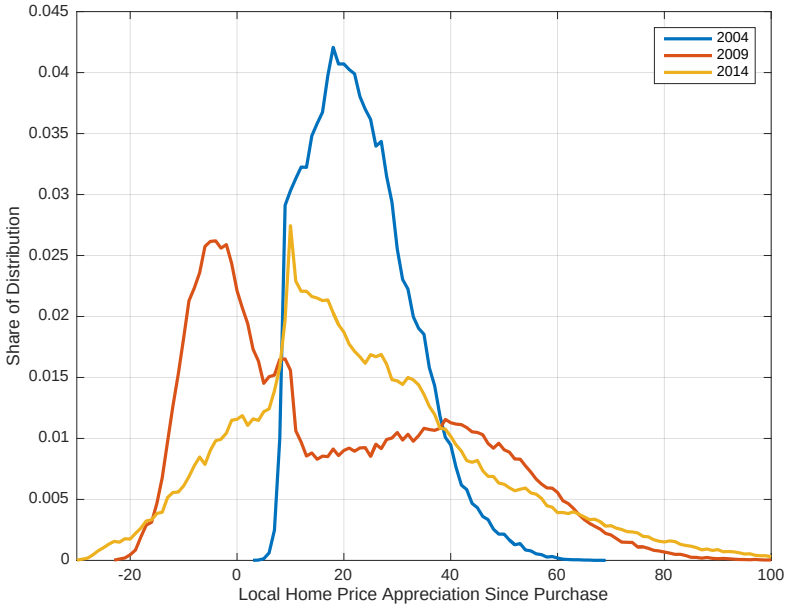

Figure A-3: LTV distribution in Seattle sample for selected years.

A.2 Model of House Lock and Nominal Loss Aversion

In this appendix, I present a simple model of housing transactions that highlights the role of mortgage debt and nominal losses in determining transactions volume. The model is an extended version of the one found in Stein (1995). Households have some degree of mismatch to their current homes, generating an incentive to move. Importantly, the cost of moving must be financed out of the equity from the household's initial home, meaning that households with high LTVs will be forced to down-size if they choose to move.[49] As a result, high LTVs can deter home sales. In addition, I allow for the possibility of default, so that borrowers with *very high* LTVs may in fact have heightened incentive to move, defaulting on their cumbersome mortgage and allowing a lender to sell their home. Finally, I will assume that households have some reference price against which they evaluate the sale price of their initial home. If we interpret the reference price as the price at which they *bought* the home initially, this allows for nominal loss aversion to play a role in determining sales volume.

Now, consider a household that owns a house of size 1. The household paid price p_0 for it and currently has a mortgage balance of L. Assume that the household lives for one more period, and the key decision it will have to make is whether to move and, if it does so, how large its new home will be. The timing of the model is as follows:

1. The household learns its mismatch to its current home, θ, and its income level, Y, which it will receive at the end of the period;
2. The household decides whether to move and how much housing, h, to purchase if it does;
3. If it moves, the household sells the current house at price p, pays off the mortgage, and then buys the new home at price $p \cdot h$;

[49] Note that in contrast to the housing ladder model presented in the main text, in which homes are of discrete qualities, here I assume that housing quality (or size) is a continuous variable, as in Stein (1995).

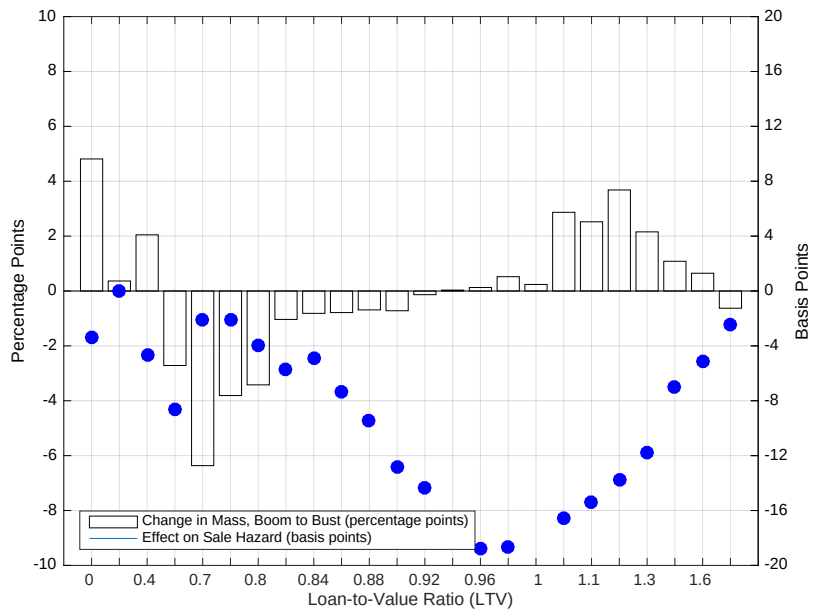

Figure A-4: The relationship between the probability of sale and LTV, estimated from the Seattle-Tacoma-Bellevue sample, is shown in the blue dots. The white bars show the change from 2004-6 to 2007-11 in the percent of borrowers in each bin in.

4. The household receives income, consumes (net of any remaining debt), and lives in its final house.

Utility is given by:

$$U = f + v(h) + I_{MOVE} \cdot \big(\theta + \alpha_G (p - p_0)^+ - \alpha_L (p - p_0)^-\big). \tag{31}$$

The household gets utility from consuming the numeraire, food (f), and living in a larger house (h), where $v(\cdot)$ is a concave function. I assume the above quasilinear form for utility to abstract from wealth effects. θ is a reduced-form measure of the household's mismatch to its current home. Therefore, if the household moves ($I_{MOVE} = 1$), it will gain θ. The final terms parameterize the idea of nominal loss aversion, where it is typically assumed that $\alpha_L > \alpha_G \geq 0$. This captures the possibility that the initial purchase price, which rational models predict will not affect agents' decisions, will affect the household's valuation of a transaction. $(p - p_0)^+$ is the positive piece of the difference between the current price and the purchase price (0 if negative), while $(p - p_0)^-$ is the negative piece (0 if positive).

In a model with no equity constraints and no default option, the single constraint on the household would be:

$$Y + p - L = f + ph. \tag{32}$$

This says that net worth (income plus the value of its house, less debt) is equal to spending. However, as in Stein (1995), I will add the key friction that a household must pay for the move with money raised from selling the initial home.[50] This adds an additional constraint:

[50]For simplicity of notation, I will assume that the new home is purchased in cash - with no mortgage. One could

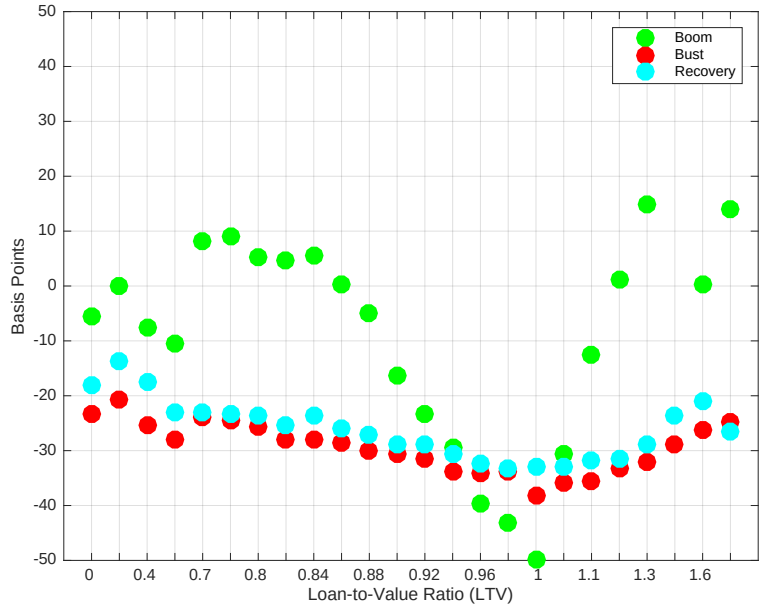

Figure A-5: The time-varying relationship between the probability of sale and LTV in the Seattle-Tacoma-Bellevue MSA. The "boom" (2004-6) is shown in green dots, the "bust" (2007-11) is shown in red dots, and the "recovery" (2012-4) is shown in blue dots. The estimates come from a regression with all covariates described in the caption of Figure 7 and MSA fixed effects).

$$h \leq \frac{p - L}{p}. \tag{33}$$

This generates the possibility of "house lock": a sufficiently indebted household (L high) may be constrained by its current level of home equity and be unable to purchase its desired house.

To allow for mortgage default, assume the household has the choice to live in a home of size 0 and not pay back its debt. Therefore, the constraint 32 is replaced by 32':

$$Y + p - L \cdot (1 - I_{Default}) = f + ph, \tag{32'}$$

along with the requirement that defaulting households live in a home of size 0:

$$h = 0 \text{ if } I_{Default} = 1. \tag{34}$$

To solve the model, begin by defining h^* to be the level of housing a household would choose if unconstrained by their home equity, conditional on moving.[51] It is characterized by equality of the marginal rate of substitution between housing and food and the ratio of their prices:

allow the households to buy a new home while only putting down some fraction of its value, but this has no effect on the results.

[51] This concept is independent of loss aversion because, unlike mortgage debt, the reference point only affects the decision of whether to move (as we will see), not the size of the home conditional on moving.

	(1)	(2)	(3)
LTV	0.9	1.3	1.1
	(0.3)***	(0.4)***	(0.4)***
(LTV-0.8)$\cdot I_{\{LTV>0.8\}}$	-17.0	-14.6	-14.3
	(0.5)***	(0.5)***	(0.5)***
HPA	40.0	34.5	35.2
	(1.0)***	(1.0)***	(1.1)***
HPA$\cdot I_{\{HPA>0\}}$	-29.5	-21.8	-21.1
	(1.2)***	(1.3)***	(1.3)***
LTV + (LTV-0.8)$\cdot I_{\{LTV>0.8\}}$	-16.0	-13.3	-13.2
	(0.3)***	(0.3)***	(0.3)***
HPA + HPA$\cdot I_{\{HPA>0\}}$	10.5	12.7	14.1
	(0.4)***	(0.5)***	(0.5)***
Controls		✓	✓
Year FEs			✓

Table A-1: Regression results using MLS listings data in Seattle. Sample is non-REO homes listed for sale. "HPA" is home price appreciation since property was last purchased—my measure of nominal gains/losses. Due to reduced sample size (from only looking at listed properties), I use a parametric relationship (linear spline) to study the effects of LTV and nominal gains/losses on selling behavior. All specifications control for months since purchase, calendar month, and months since put on market (listed). Specification (2) further controls for variables listed in the footnote of Figure 7. The final columns adds indicators for each year. Note:$N = 405,702$.

$$v'(h^*) = p. \tag{35}$$

The moving household's choice can then be written as:

$$h = \max\left\{\min\left\{h^*, \frac{p-L}{p}\right\}, 0\right\}. \tag{36}$$

The inner minimum function incorporates the requirement that current home equity be sufficient to pay for the new house, while the outer maximum function incorporates the default option, which says that a household can walk away from its current home and cannot be forced to pay off its debt. In that case, it will set $h = 0$ and $f = Y$. As a result, utility from moving can be expressed as follows:

$$U_{\mathrm{M}} = \begin{cases} Y - L - p(h^* - 1) + v(h^*) + \theta + \alpha_G(p - p_0)^+ - \alpha_L(p - p_0)^- & \text{if } L \leq p(1 - h^*) \\ Y + v(\frac{p-L}{p}) + \theta + \alpha_G(p - p_0)^+ - \alpha_L(p - p_0)^- & \text{if } p(1 - h^*) < L \leq p \\ Y + v(0) + \theta + \alpha_G(p - p_0)^+ - \alpha_L(p - p_0)^- & \text{if } L > p. \end{cases} \tag{37}$$

Utility from not moving is simply:

$$U_{\mathrm{N}} = Y - L + v(1). \tag{38}$$

Define $\Delta(L, p, Y, \theta, p_0)$ to be the difference between utility when moving and utility when not moving. First, consider the effect of L on the household's incentive to move:

$$\frac{\partial \Delta}{\partial L} \begin{cases} = 0 \text{ if } LTV < (1 - h^*) \\ < 0 \text{ if } LTV \in \left[(1 - h^*), 1\right] \\ > 0 \text{ if } LTV > 1 \end{cases}, \tag{39}$$

where $LTV \equiv L/p$. Intuitively, when LTV is low and the equity constraint does not bind, tightening it with a little more debt does not affect the household's incentive to move: every additional dollar of debt is a dollar that is not consumed, regardless of whether the household moves or not.[52] With an intermediate level of debt, when the down-payment constraint binds, the household is being forced to under-consume housing. The more that constraint tightens (i.e. LTV increases), the further away from the optimum h^* a moving household will be, so the incentive to move decreases. Effectively, these households are being forced to down-size if they move, and so only those with relatively high mismatch realizations θ will find it worthwhile to do so. Finally, when the household's debt is so high that it is underwater ($LTV > 1$), it cannot afford to buy a house larger than $h = 0$ if it moves. Therefore, a marginal increase in L does not affect house size. Instead, every additional dollar of debt is a dollar that must be paid if the household stays in their current home but that the household can walk away from by defaulting and allowing the lender to sell the foreclosed home. As a result, sales become more likely when debt increases in this region—though they are distressed sales by lenders.

If the distribution of θ is similar across LTV levels, this simple model predicts the following:

1. The probability of a home selling is flat at low levels of LTV;
2. For some $L\tilde{T}V < 1$, the probability of a home selling is a decreasing function of LTV when $LTV \in [L\tilde{T}V, 1]$;
3. The probability of a home selling is an increasing function of LTV when $LTV > 1$.

Turning to nominal loss aversion, the effect of p_0 is quite straightforward. Note that unlike L, p_0 does not affect the household's decision of home choice, conditional on moving. It does, however, affect whether the household decides to move. In particular, if $p < p_0$, moving requires the household to realize a nominal loss. As this has a negative effect on utility, a higher θ is required to induce the household to transact, and so the probability of sale falls, as the household will require a higher mismatch (θ) to be induced to move. Therefore, we expect the probability of sale to be an increasing function of the difference between p and p_0. Note a subtlety here, which is that this may only apply to non-distressed sales, where the homeowner is actually receiving the proceeds of the sale. When the household defaults, it does not receive any payment for the home and may in fact have moved well before the sale occurs. As a result, I will focus on non-distressed sales in the empirical tests of this portion of the model.

[52] This is the one place where the assumption of quasilinear utility is substantive. With wealth effects, the additional dollar of consumption would be valued differently depending on whether the household moves, since the mixture of f and h depends on that decision. As a result, marginal increases in debt could affect the household's incentive to move.