Mortgage Lock-In and Home Sales Volume Dynamics*

Joshua Abel†

March 20, 2025

Abstract

The recent increase in interest rates has substantially reduced home-selling. This paper considers the dynamic problem of a household with the dual options to move and to refinance in a model where interest rates and idiosyncratic housing match quality are both stochastic. The model admits the possibility of "lock-in," whereby households reduce their likelihood of sale in the face of high mortgage interest rates. A calibrated version of the model is used to analyze the dynamics of home sales moving forward and the impacts of proposed policies. While the impact depends on the future conduct of monetary policy, the model predicts that a tax credit for home-sellers would likely increase economic efficiency, but it would require roughly $40 of government expenditure for every $1 of social surplus created, since many moves would occur even in the absence of the subsidy. Therefore, such a subsidy is in large part simply a transfer to homeowners and their lenders. This paper contributes to the literature as the first analysis to consider the impact of lock-in on the dynamics of home sales outside of a steady state.

Note: I have no financial interest in this research to disclose.

*Dayanara Diaz Vargas, Thomas Sargent, Vedant Mundhra, and Tarun Anbarasu provided excellent research assistance. I am grateful for helpful conversations with Andrew Abel, Santiago Caicedo, Jonah Coste, Edward Glaeser, Adam Guren, Justin Katz, Robert Triest, and Paul Willen, as well as seminar participants in the Federal Housing Finance Agency's Research Department.

†Northeastern University. E-mail: j.abel@northeastern.edu

1 Introduction

Sales volume of single-family homes in the United States fell by roughly one third from 2021 to 2023 as the Federal Reserve raised interest rates beginning in 2022 to combat high inflation. Both the popular press[1] and academic researchers have attributed this to "mortgage lock-in:"[2] moving would require homeowners with the standard fixed-rate mortgage (FRM) to prepay existing loans with low interest rates contracted in the preceding decades and take out a new loan with a high interest rate; this provides a disincentive that can explain a sharp decline in home sales.[3] This paper presents a quantitative yet parsimonious dynamic model of household behavior that generates such lock-in effects and allows for a novel analysis of how the issue has evolved to date and, more importantly, how monetary and fiscal policies can impact the future evolution of home sales.

In the model, the mortgage interest rate is driven by shocks to both inflation and the real interest rate. Critically, households also experience persistent shocks to their idiosyncratic match quality with their current house. This match quality is a catch-all to encode all factors other than the interest rate that affect a homeowner's willingness-to-pay to move, such as proximity to a new job opportunity, the composition of room types in a home, or the mixture of local amenities, to name a few. Based on the evolution of these stochastic variables, households decide when to exercise their two options: 1) refinance their mortgage; 2) move. This can be viewed as an extension of the framework developed by Agarwal *et al.* (2013) and Berger *et al.* (2024) – used for modeling refinancing behavior when moves are assumed to be exogenous – to allow for an endogenous margin of moves.[4] This margin is affected by rising interest rates, which raises the effective cost of moving borne by the household and therefore can deter moves (i.e. lock-in).

[1] See, e.g., "The Home Buyer's Quandry: Nobody's Selling," by Nicole Friedman in *The Wall Street Journal* on May 10, 2023; or "A 30-Year Trap: The Problem With America's Weird Mortgages," by Ben Casselman in *The New York Times* on November 19, 2023; or "Here's what upgrading to a nice home could cost you, and why it's locking up the market" by Diana Olick on CNBC.com on April 2, 2024.

[2] Other terms such as "mortgage lock" or "house lock" are variously used, as well. Such terms are also used to describe other frictions that can decrease sales activity, such as negative equity or various tax incentives. This paper will use "lock-in" to refer to the disincentive to sell that is generated when a homeowner has a mortgage with an interest rate below what she could secure on a newly-originated mortgage.

[3] Frictions in the rental market are also necessary to create this disincentive. If there were no effort or search costs associated with renting a home, a homeowner with a good interest rate could move to a new home and rent, while simultaneously renting out her initial home. In that way, she could avoid surrendering the favorable interest rate while still relocating. Furthermore, it depends on institutional the institutional features of mortgages that they are not portable or assumable.

[4] It is not an extension strictly speaking, as those models are solved in continuous time, and this paper performs its analysis in discrete time. Note that Agarwal *et al.* (2013) analyze the refinancing problem assuming perfect attention, while Berger *et al.* (2024) allow for households to be partially inattentive. The model in this paper allows for inattention.

A calibration of the model does a good job of replicating recent movements in aggregate home sales and refinances, and it is also able to match microdata on sale and prepayment hazard rates. This calibrated model is used to study the evolution of home sales moving forward under different policy environments. Specifically, I analyze how home sales volume is impacted by the path of future interest rates, and I also evaluate the much-discussed fiscal policy response of subsidizing home sellers.[5] The model makes clear that monetary and fiscal policies interact, and the subsidy cannot be analyzed separately from the path of interest rates. If interest rates fall sharply, the lock-in problem will disappear, and in fact there will be a significant boom in sales as households that delayed moves due to high interest rates will seize their opportunity.[6] In this case, the tax credit will over-stimulate sales – distorting households' view of the cost of moving – and reduce economic efficiency. If, on the other hand, interest rates stay high and the lock-in problem persists, the tax credit will increase efficiency by stimulating moves by households whose fundamentals (i.e. mismatch) suggest that they should move but who are not doing so because it would entail a transfer to lenders in the form of a higher interest rate. The paper shows that under the interest rate path implied by the Federal Open Market Committee's Survey of Economic Projections, the second outcome is more likely and the tax credit would increase economic efficiency. However, this comes with the critical caveat that, because most moves would occur even in the absence of the credit, it would take over $40 of government spending to create each $1 of social surplus. As a result, the subsidy would likely serve as little more than a transfer from taxpayers to homeowners and their lenders.

The issue has attracted substantial attention from academic economists. The paper is most closely related to a growing literature of structural housing market models used to evaluate the consequences of lock-in and evaluate policies to address it. Modeling this phenomenon is challenging. The decision to move depends on the household's housing mismatch and interest rate and can also depend on home prices. Home prices, in turn, can depend on the volume and characteristics of movers. This means that an individual's decision to move depends on the future evolution of the joint distribution of interest rate gaps and mismatch (via prices), which itself is determined by the moving decisions of individuals. The computational complexity of unwinding that high-dimensional interdependence has not been solved, so instead papers in the literature (including this one) have made different types of simplifying

[5]For example, President Joe Biden proposed such a tax credit during his 2024 State of the Union Address. See, "FACT SHEET: President Biden Announces Plan to Lower Housing Costs for Working Families," available from WhiteHouse.gov since March 7, 2024.

[6]The idea of home sales being delayed has been explored in a life-cycle context by Attanasio *et al.* (2012), who model how the level and uncertainty of income early in life can affect the timing of households' home purchase decisions.

assumptions to gain tractability.

The most common approach has been to ignore the dynamics of the problem and solve for steady states. The analysis then consists of comparisons across steady states ("comparative statics"). For example, Fonseca *et al.* (2024) develop an equilibrium model with a housing ladder and many housing markets, where prices of different tiers of housing are determined endogenously. They find that a tax credit, like that proposed by President Biden, targeted at the sellers of "starter homes" would in fact mostly accrue to households at the top of the housing ladder because the incidence of the credit would largely be felt as an increase in the price of their homes. In a similar vein, Katz and Minton (2024) shows how lock-in can affect net housing demand and uses a quantitative model to show how that friction can explain why home prices did not fall in response to rising interest rates as a standard asset pricing model would predict. Gerardi *et al.* (2024) develop a life-cycle model with search-and-matching in the housing market and find that lock-in meaningfully reduces welfare due to reduced market thickness, particularly for younger households in lower-income areas, while at the same time increasing prices. All three of these papers present interesting results on home prices by abstracting away from transition dynamics and focusing on steady states, so households do not have to forecast any aggregate objects, greatly reducing the complexity of the problem.

Amromin and Eberly (2024) navigate the tradeoff in a different way. They study the dynamics of prices, arguing that the dynamics following the COVID shock were so different than those from the Great Financial Crisis in part due to the lock-in effect from the recent rise in interest rates. In order to resolve the tractability problem outlined above, these authors assume moves occur at an exogenous and time-varying rate to mechanically match data on sales, which simplifies the analysis because household therefore not do not have to forecast the future evolution of moves and prices – after all, their own decision to move is out of their hands. So while that analysis does have dynamics, it seeks to explain home prices while taking sales volume (i.e. mobility) as a given rather than modeling it.

The present paper therefore is unique in contributing to this literature an analysis of the *dynamics of home sales*. Unlike Fonseca *et al.* (2024), Katz and Minton (2024), and Gerardi *et al.* (2024), there is aggregate uncertainty and the analysis is not confined to steady states. Unlike Amromin and Eberly (2024), the decision to move is endogenous. The tradeoff in this paper that allows for tractability is that – unlike those – I do not endeavor to explain home prices. Formally, this paper navigates the challenge by assuming that all buyers are sellers and vice versa (and there is only one housing tier), which makes the price of housing irrelevant to their decision, since the home's value is on both sides of the budget constraint.[7]

[7]Notably, Anenberg and Bayer (2020) study the dynamics of the joint buying-and-selling problem in a

In this way, the paper is a complement to the existing literature because, while it offers no insights about home prices, it is the first that can offer a quantitative answer to a question of clear policy relevance like, "How will the Fed's future conduct of monetary policy affect the evolution of home sales?" This requires a framework in which aggregate variables (like the interest rate) change and in which households' decision to move is endogenous.

The paper is also quite connected to the recent papers documenting the empirical relevance of lock-in. Fonseca and Liu (2024) focus on household mobility from 2010-2018 and show that facing an increase in the mortgage interest rate of an additional 1 percentage point (pp) from a move reduces the probability of such a move by 0.68pp, or 9%, on average. Importantly, they show that while the prospect of a higher interest rate deters mobility, the ability to get a lower interest rate does not spur mobility, as borrowers can refinance their mortgage to get that low rate without relocating. They also show that homeowners who face larger disincentives are less likely to relocate in response to labor market incentives. Liebersohn and Rothstein (2025) also show that mobility is hindered when households would have to give up a low interest rate to relocate. In particular, they focus on the current episode of sharply rising interest rates, 2022-2023, and show that a 1pp increase in the disincentive reduces mobility by roughly 15%. Batzer *et al.* (2024) also study the 2022-2023 period, but they study the effect on home sales, rather than mobility. Consistent with the other papers, they find that a 1pp increase in the disincentive reduces the probability that a home will be sold by roughly 18%. These effects were strongest among wealthier borrowers and those with more-expensive homes.[8,9] While the present paper is not an empirical study, its model provides some insights for contextualizing the empirical results described above. Specifically, due to the dynamic nature of the lock-in issue, the model shows that the empirical reduced form relationship between a household's interest rate incentive and its likelihood of a move depends on the recent path of interest rates and home sales. In particular, the model shows

search model and provide evidence that such housing market participants are the major driver of volatility in sales volume, and so this is not an unreasonable approach to take, but it is certainly an abstraction and one that comes at the cost of analyzing prices.

[8]Batzer *et al.* (2024) also find that the degree of lock-in in an MSA is associated with higher home price appreciation; Fonseca and Liu (2024) find a similar result on list prices in the earlier period they study.

[9]These papers use large, administrative datasets to show the empirical importance of this friction in recent years. This builds on the work of Quigley (1987) and Ferreira *et al.* (2010), which demonstrated the relevance of the incentive using survey data. On a related topic, Ferreira (2010) and Imrohoroglu *et al.* (2018) showed that property taxes can affect mobility by creating financial disincentives against moving, in particular by looking at caps on taxes paid by incumbent owners in California. Best and Kleven (2018) look at a related policy in the United Kingdom. A number of papers, such as Ferreira *et al.* (2010), Chen (2001), Bernstein and Struyven (2021), and Genesove and Mayer (1997) have also documented another friction that mortgages can cause that reduce mobility, which is low/negative equity: if a homeowner has little or no equity in the home, selling it will not recover a sufficient amount to fund a down payment on a new, comparable home, which prevents some borrowers from moving. This conceptual point was popularized by Stein (1995).

that the probability of a move is higher after a lock-in episode than before it, conditional on a household's interest rate incentive. In other words, the reduced form results in the papers cited above are context-dependent and a new relationship – with higher sale probabilities – is likely to emerge in the coming years.[10] These dynamic effects are reminiscent of Berger *et al.* (2021), who make the point that the ability of monetary policy to stimulate spending through the refinancing channel depends critically on the past path of interest rates, as that determines the extent to which households will be able to lower their payments through a refinance.

The paper also contributes to the literature on refinancing. The model's solution shows that as a household's match quality with its current home deteriorates, it should require a larger decrease in the market interest rate before exercising the option to refinance. This is because the flow of benefits from refinancing is terminated when a household moves, and since a household with poor match quality should expect to move relatively soon, it should not waste effort refinancing unless the flow of benefits from doing so is very strong. This adds subtlety to interpreting empirical analyses that find many households wait too long to refinance (e.g. Agarwal *et al.* (2017) and Andersen *et al.* (2020)), as it demonstrates that a household's optimal refinancing threshold cannot strictly be known when the econometrician cannot observe its housing match quality. More specifically, it can help justify why borrowers with, for example, lower income are less likely to refinance, as they are also more likely to move.

Finally, this paper adds to the literature on mortgage design, which largely documents undesirable features of the dominant FRM design in the US. Due to heterogeneous financial sophistication and refinancing likelihoods, FRMs extract more heavily from unsophisticated borrowers, as pointed out in Gabaix and Laibson (2006) and quantitatively modeled by Zhang (2024) and Berger *et al.* (2024). This also has macroprudential consequences, as highlighted by Eberly and Krishnamurthy (2014), Campbell *et al.* (2021), and Guren *et al.* (2021), because the failure to refinance reduces the ability of monetary policy to deliver payment relief to households in periods when the marginal propensity to consume is high. While that strand of the literature focuses on undesirable features of the FRM when interest rates are *falling*, lock-in is a problematic consequence of that design when interest rates *rise*. Using the calibrated model, this paper provides an estimate of the aggregate increase in welfare that borrowers could expect from an adjustable-rate mortgage (ARM) relative to

[10]The dynamic framework also shows that even if interest rates were to never fall again, the lock-in effect would dissipate over time. Intuitively, households will be spurred to move eventually, even despite high interest rates; as each household moves, it gives up its low mortgage rate and no longer experiences lock-in, therefore returning to "normal" moving behavior.

the FRM. In particular, an ARM would increase welfare by 12% of their mortgage balance, with 2% coming from reduced refinancing costs and the rest from removing the distortion that FRMs can cause for the decision to move. This estimate is quite large, but notably the model does not consider risk aversion. As a result, because the model does not capture the oft-cited benefit of FRMs – that borrowers know that their payments are capped at a certain level – one way to interpret this estimate is as a lower bound on the value of that benefit.

The paper proceeds as follows. Section 2 describes the model of household behavior and its solution. Section 3 describes the model's calibration and its behavior through early 2025. Section 4 provides a forward-looking evaluation of monetary and fiscal policy affecting lock-in. Section 5 considers the implications of the paper's ideas for mortgage design, and Section 6 concludes.

2 Model

This section presents a parsimonious but quantitatively-relevant model of homeowner behavior with lock-in effects. The core is essentially a discrete-time version of the model presented in Agarwal *et al.* (2013), who found a closed-form solution for optimal refinancing behavior assuming moves are exogenous.[11] While they make a number of simplifications to allow the closed-form solution to be tractable, they show that their findings align closely with those in more-complex settings that are solved numerically.[12] The most important change I make to their setting is to allow for an active moving margin, as a household's match quality to its home follows a stochastic process and it can choose to move when it deems that another house would be a sufficiently better match. In contrast to the literature on refinancing, then, households in the model in this paper have two options (refinance, move), which interact. While this and some other wrinkles make a closed-form solution intractable, the model's solution is straightforward enough to provide clear intuition for how households behave and how they will respond to different policy environments.

2.1 Macroeconomic Variables

Consider a house whose owner has a FRM of size M, which is exogenous and without loss of generality is normalized to \$1. The homeowner has a discount rate of ρ and, for tractability, assume the borrower is risk-neutral[13] and that it is a non-amortizing (interest-only) mortgage.

[11] As described below, the model incorporates the generalization presented in Berger *et al.* (2024), which solves for refinancing behavior when households exhibit inattention to interest rates.

[12] They do a detailed comparison of their results to those of the numerical model in Chen and Ling (1989).

[13] Assuming risk-neutrality simplifies the model and its solution considerably, as it means that households need only worry about the difference between their current interest rate and the market rate. A risk-averse

There are two macroeconomic state variables. The real mortgage interest rate, r_t, follows a random walk with shocks drawn from $N(0, \sigma_r^2)$. Inflation follows a Markov process with two states: High and Low. In the high state, inflation is π_H, and the probability of remaining in the high state the following period is p_π^H; in the low state, inflation is π_L, and the probability of remaining in the low state the following period is p_π^L. The nominal mortgage interest rate is given by $m_t = r_t + \pi_t$.[14,15]

Letting m_{i0} be the market interest rate when house i was last purchased, define $y_{it} \equiv m_t - m_{i0}$ to be the gap between the current market rate and the homeowner's contract rate. This y_{it} represents the change in the interest rate on the home's mortgage if the current owner sells or refinances.

2.2 Household Problem

To produce lock-in effects, the model must allow the household to have an active margin of moves, which will be impacted by y_{it}. To that end, assume that match quality between a house and its owner follows a stochastic process. x_{it} will capture the borrower's mismatch to her current home at time t. In particular, it represents the willingness-to-pay to sell the house and move to one with a better idiosyncratic match. Normalize $x_{it} = 0$ in the immediate aftermath of a move and assume x_{it} follows a random walk with shocks drawn from $N(0, \sigma_x^2)$.

In each period, the borrower must take exactly one out of three actions:

1. Refinance at cost C^R: This will result in a change in her mortgage interest rate of y_{it^-} (the interest rate gap prior to the refinance) and reset the interest rate gap to 0.

2. Move at cost C^M: This, too, will result in a change in her mortgage interest rate of y_{it^-} and a change in the flow of housing utility of x_{it^-}. Both the interest rate gap and the housing mismatch value reset to 0.

household's choices would depend not just on the difference but also on the levels, adding a state variable. Agarwal *et al.* (2013) assume risk-neutrality for this reason, and as discussed above, they find that their results are very similar to models with risk-averse agents. Recent papers by Berger *et al.* (2024) and Abel (2024) do quantitative analyses of household refinancing decisions and also maintain the assumption of risk-neutrality.

[14]This model does not solve for equilibrium interest rates but rather takes them as given. Solving for equilibrium interest rates would require lenders to forecast refinancing and selling behavior, which as discussed throughout this paper depends on the evolution of the joint distribution of households' interest rates and housing mismatch. Given the computational difficulty of that problem, I assume this simple exogenous process for mortgage interest rates.

[15]Due to the mean reversion in the inflation process, the nominal mortgage interest rate also exhibits a degree of mean reversion, in contrast to Agarwal *et al.* (2013), in which the interest rate is a pure random walk.

3. Neither: This will result in no change in her interest rate or housing mismatch.

For quantitative relevance, households in the model exhibit inattention, similar to Berger *et al.* (2024), and exogenous moves that are unrelated to house mismatch, as in Agarwal *et al.* (2013). In particular, in any given period, the probability that the household is aware of the market interest rate is p_R. Whenever she is paying attention, she considers refinancing; but within those attention events, she only considers moving with probability p_A.[16] Finally, with probability p_M, a household receives an "exogenous" moving shock and will move, regardless of x, y, and whether it is "paying attention." If one of these exogenous moving shocks occurs, the household pays the moving cost C^M and experiences a permanent change in her mortgage interest rate of y_{t^-}. The flow of housing utility changes by χ, which I set to 0. Both the interest rate gap and the housing mismatch value (y and x) reset to 0. The exogenous moving and attention shocks are assumed to be stationary and uncorrelated.

The homeowner maximizes the present discounted value of its flow of housing amenity less its costs; costs include mortgage payments as well as the costs of refinancing and moving. Given this setup, the household's value function is characterized by:

$$
\begin{aligned}
&V(x,y,\pi) = \\
&(1-p_M)\cdot p_R \cdot p_A \cdot \underbrace{\Big(\max\Big\{\underbrace{\frac{1}{1+\rho}E_{x,y,\pi}[V(x_+,y_+,\pi_+)]}_{\text{Value of no action}},\ \underbrace{-q_\pi\cdot y - C^R + \frac{1}{1+\rho}E_{x,0,\pi}[V(x_+,y_+,\pi_+)]}_{\text{Value of refinancing}},\ \underbrace{\frac{1+\rho}{\rho}\cdot x - q_\pi\cdot y - C^M + \frac{1}{1+\rho}E_{0,0,\pi}[V(x_+,y_+,\pi_+)]}_{\text{Value of moving}}\Big\}\Big)}_{\text{Value of considering both a move and refi}} + \\
&(1-p_M)\cdot p_R \cdot (1-p_A)\cdot \underbrace{\Big(\max\Big\{\underbrace{\frac{1}{1+\rho}E_{x,y,\pi}[V(x_+,y_+,\pi_+)]}_{\text{Value of no action}},\ \underbrace{-q_\pi\cdot y - C^R + \frac{1}{1+\rho}E_{x,0,\pi}[V(x_+,y_+,\pi_+)]}_{\text{Value of refinancing}}\Big\}\Big)}_{\text{Value of considering a refi but not a move}} + \\
&(1-p_M)\cdot(1-p_R)\cdot \underbrace{\frac{1}{1+\rho}E_{x,y,\pi}[V(x_+,y_+,\pi_+)]}_{\text{Value of considering neither move nor refi}} + \\
&p_M\cdot \underbrace{\Big(-q_\pi\cdot y - C^M + \frac{1}{1+\rho}E_{0,0,\pi}[V(x_+,y_+,\pi_+)]\Big)}_{\text{Value of exogenous move}},
\end{aligned}
\tag{1}
$$

where $q_\pi \equiv 1 + E[\sum_{t=1}^{\infty}\left(\frac{1}{(1+\rho)^t\cdot\Pi_{\tau=1}^{t}(1+\pi_\tau)}\right)|\pi_0=\pi]$. Row 1 of Equation 1 considers the case where the household does not face an exogenous moving shock (probability $1-p_M$) and pays attention to both the refinancing and moving incentives (probability $p_R\cdot p_A$). In that case, she can either take no action, she can refinance, or she can move. The flow values of refinancing and moving both include the present discounted value of the reduction in

[16] I do not allow the household to consider moving without also considering refinancing. Such a situation would allow for the logically strange decision to move simply to get a lower interest rate. A household who is aware of interest rates enough to do such a thing should clearly just refinance, which is less costly anyway.

mortgage payments, $-q_\pi \cdot y$. Following the insight of Agarwal *et al.* (2013), this can be treated as a perpetuity because the possibility that payments will change subsequently is built into the continuation value. For the same reason, the change in housing amenity from moving is treated like a perpetuity, $\frac{1+\rho}{\rho} \cdot x$.[17] The continuation values of the three different options differ based on the state values that are brought into the next period: if no action is taken, housing mismatch and the interest rate gap remain at x and y (subject to the next period's shocks); refinancing causes the interest rate gap to reset to 0, but housing mismatch remains at x; moving causes both housing mismatch and the interest rate gap to reset to 0. The remaining rows have the same structure but account for different combinations of exogenous moving and attention shocks.

2.3 Household Behavior

For a set of parameters, Equation 1 is solved through value function iteration. Figure 1 below shows the decision rules that characterize the solution. Before understanding the solution of a homeowner with a FRM, it useful to consider as a benchmark the behavior of a homeowner with an ARM. Such a homeowner's interest rate tracks market interest rates regardless of whether it moves and gets a new mortgage. As a result, the change in the market interest rate since mortgage origination, y_{it}, is irrelevant for decisions and the household moves based entirely on housing fundamentals. As shown by the vertical purple line in Figure 1, an ARM household moves if and only if $x_{it} > x^*$ – in other words, when the increase in housing utility from moving is enough to justify the cost (taking into account option value).

Turning now to the full FRM model, the red and black curves show the refinancing and moving thresholds, respectively. The black curve shows, as a function of mismatch to the current house (x), the largest interest rate increase (y) the household is willing to accept as part of a move. The red curve shows the cutoff for the interest rate gap below which the household prefers to refinance as opposed to taking no action. The solution to the household's problem is to refinance when her state (x,y) is below the red curve and to move when it is below the black curve. The household is indifferent between moving and refinancing when $x = \tilde{x}$. When $x < \tilde{x}$, refinancing is preferred to moving because the housing utility gain from moving is not sufficiently large; when $x > \tilde{x}$, moving is preferred to refinancing because it is.

The black curve slopes up because the larger its housing mismatch is, the more of an interest rate increase the household is willing to accept during a move. The red line slopes down because as x increases, the household recognizes that the expected time until its next move

[17] Housing mismatch is measured in real terms and so can simply be discounted by the rate of time preference, ρ; since mortgage payments are nominal, discounting them requires considering future inflation, as encoded in q_π.

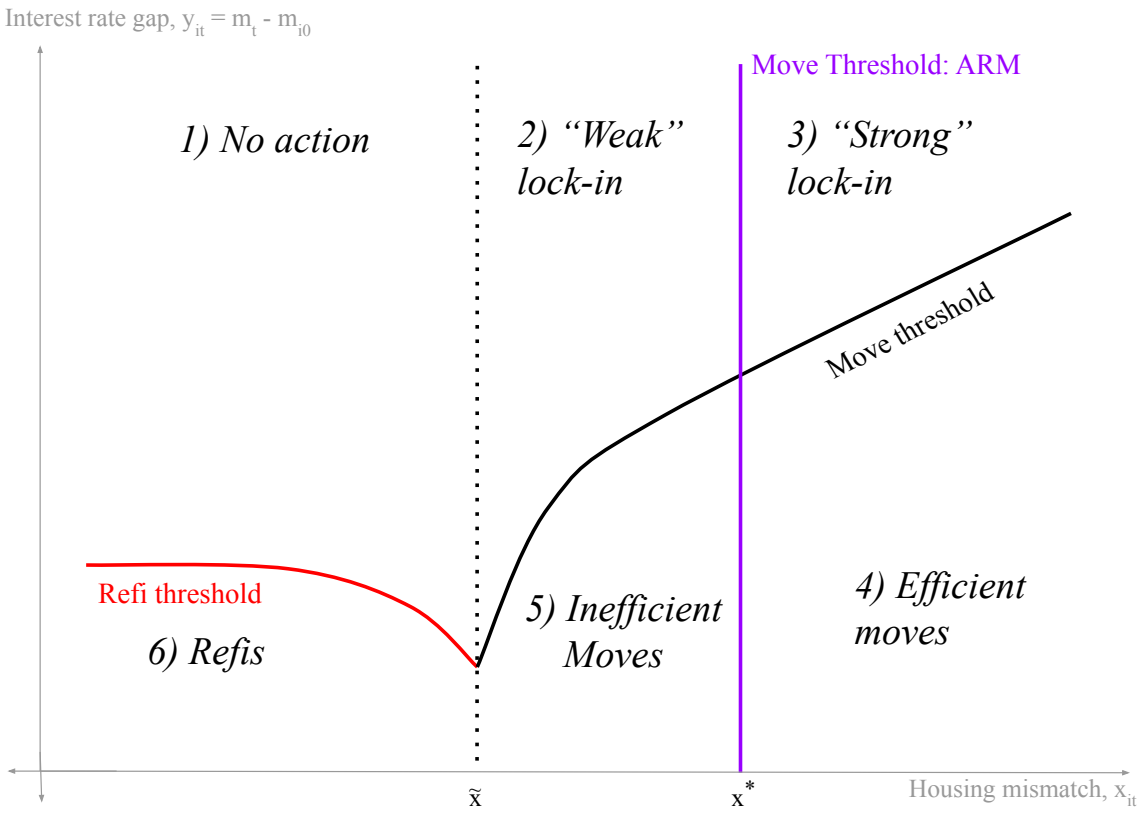


Figure 1: A household (that is paying attention) with a FRM moves if its state, (x, y), is below the black curve and refinances if its state is below the red curve. A household with an ARM moves if its state is to the right of the purple line. In regions 1 and 2, the household takes no action, regardless of the mortgage type. In region 3, the household moves if it has an ARM but takes no action if it has a FRM. In region 4, the household moves regardless of the mortgage type. In region 5, the household moves if it has a FRM but takes no action if it has an ARM. In region 6, a household with a FRM refinances. The third state variable, inflation (π) is not shown as it does not qualitatively impact the results shown here.

decreases, as it is getting nearer to the black curve. This decreases the benefit of refinancing, since the new refinanced interest rate will be surrendered upon moving anyway, and therefore the household requires a larger decline in the interest rate to justify the cost of refinancing.

The various curves in Figure 1 split the state space into 6 regions. Region 3 is the most relevant given the subject of this paper: this is where lock-in is operative. In particular, this is the region of the state space where high current mortgage rates prevent a homeowner from moving. In particular, because $x_{it} > x^*$, a household with an ARM would move because doing so would increase its housing utility sufficiently. However, a household with a FRM chooses not to move because doing so would result in an increase in its mortgage rate (y_{it}) that would

overwhelm the improved housing utility. In other words, while the "housing fundamentals" justify moving, a move would also entail a transfer from the household to lenders in the form of higher mortgage payments, and this prevents the household from moving. This is inefficient because the housing fundamentals favor the move but the household stays in their house to avoid making a transfer to the lender.

Interestingly, region 5 shows that the opposite of lock-in can occur. In this region, housing fundamentals are not strong enough to induce a move on their own ($x_{it} < x^*$), but a household with a FRM still chooses to move, as the decrease in mortgage payments is large enough to justify the move (because y_{it} is low). The household does not want to wait for x_{it} to increase, because doing so risks losing the favorable interest rate, as y_{it} could rise. And importantly, the household does not want to refinance to take advantage of the low interest rate because it knows it is likely to move in the near future; moving "early" allows it to avoid incurring the cost of refinancing and then the cost of moving shortly thereafter.

Region 4 is the other region (in addition to 5) in which FRM households move. This is where the housing fundamentals of moving are strong enough to justify a move. Notably, moves occur even when interest rates are high, but the threshold level of x_{it} needed to justify such a move is correspondingly high, as discussed above, and so moves are less likely at high levels of y_{it}.

Region 6 is the part of the state space in which a FRM household wants to refinance. As discussed above, the interest rate gap required to induce a refinance falls as x_{it} increases, because the household requires a larger decline in its monthly mortgage payments when it knows that the expected time until its next move is shorter. When x_{it} is very low and no endogenous move is likely to happen soon, the refinancing rule essentially collapses back to the model of Agarwal *et al.* (2013), in which moves are exogenous.

Finally, in regions 1 and 2, no action is taken, regardless of whether the borrower has a FRM or ARM. I refer to region 2 as "'weak' lock-in" because high interest rates are preventing moves in a causal sense. In other words, a household in region 2 could be incentivized to move if the interest rate (and thus, y_{it}) were to fall sufficiently. However, those moves would be inefficient, as discussed above, in contrast to the "'Strong' lock-in" in region 3, where efficient moves are being prevented by high interest rates.

3 Lock-In Dynamics Through Early 2025

This section describes how the model of Section 2 is calibrated and what it implies about the dynamics of home sales in recent years. The section begins with a description of the relevant

data, followed by the calibration procedure and a discussion of how the model fits the data during the sample period. The section concludes with an assessment of the magnitude of the lock-in phenomenon as of March 2025.

3.1 Data

For the macroeconomic variables, the analysis relies on monthly data on mortgage interest rates[18] and inflation[19] from April 1991 through March 2025. The inflation measure is lagged by 12 months, as the lagged inflation series has a much stronger correlation with nominal mortgage rates and so seems to be a better measure of expected inflation that markets used to set mortgage rates.[20] The analysis begins in April 1991, capturing the Great Moderation, the Great Financial Crisis, COVID, and the post-COVID inflation. The choice of start date has very little impact on the analysis, as the large driver of behavior in recent years is the recent path of mortgage interest rates, with earlier data merely being relevant for setting "initial conditions" before the interest rate increase; the impact of developments on these initial conditions is quite minimal for earlier years. These series are shown in Figure 2.

Home sales activity is measured using the quarterly sales rate as reported in Batzer *et al.* (2024) from 2005Q1-2024Q1.[21] Using microdata on mortgages guaranteed by Fannie Mae and Freddie Mac (the government-sponsored enterprises, or GSEs), these authors linked mortgage terminations to deeds records to identify which are sales as opposed to refinances or other mortgage terminations. Their series displays a very strong correspondence with aggregate home sales data as reported by the National Association of Realtors (NAR), as shown in Figure 3a. The data from Batzer *et al.* (2024) is a preferable measure for this calibration for two reasons. First, it reports the rate of sales (i.e. share of homes sold), which is a more natural analog to the model than is the aggregate number reported by the NAR. More importantly, as described in the next paragraph, the present paper uses microdata from the GSEs to infer the time series of refinancing activity during this time period. This makes it appealing to have a sales measure based on mortgages guaranteed by the GSEs, as the Batzer *et al.* (2024) measure provides.

Public, loan-level origination and performance data from Fannie Mae's "Single-Family Fixed

[18] In particular, the average interest rate on a 30-year FRM as reported by Freddie Mac's Primary Mortgage Market Survey is used. This is available from the Federal Reserve Economic Database (FRED) under the name "MORTGAGE30US."

[19] In particular, the average Consumer Price Index year-over-year inflation for urban consumers, less food and energy, as reported by the Bureau of Labor Statistics is used. This is available from FRED under the name "CPILFESL."

[20] In particular, $\text{corr}(m_t, \pi_t) = 0.35$ and $\text{corr}(m_t, \pi_{t-12}) = 0.57$

[21] I am grateful to the authors of Batzer *et al.* (2024) for making auxiliary materials from their analysis available. These can be found at https://www.fhfa.gov/research/papers/wp2403.

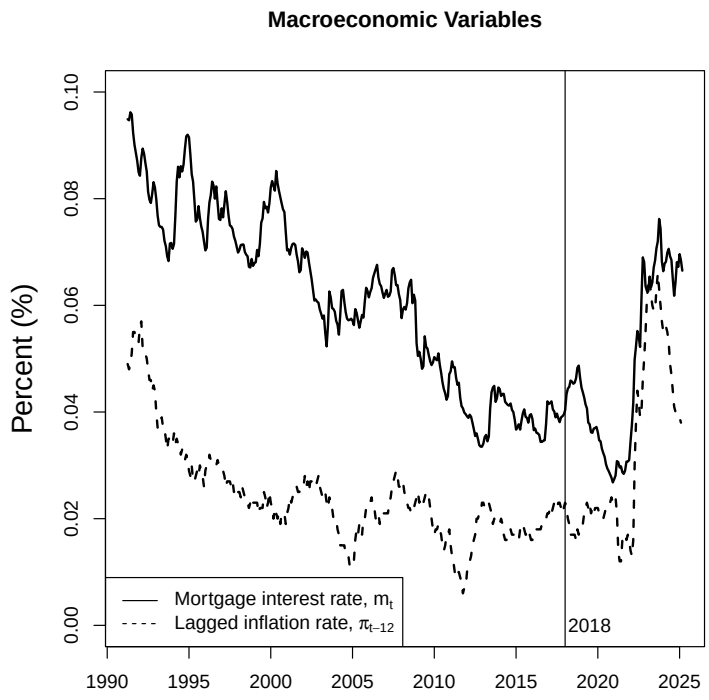
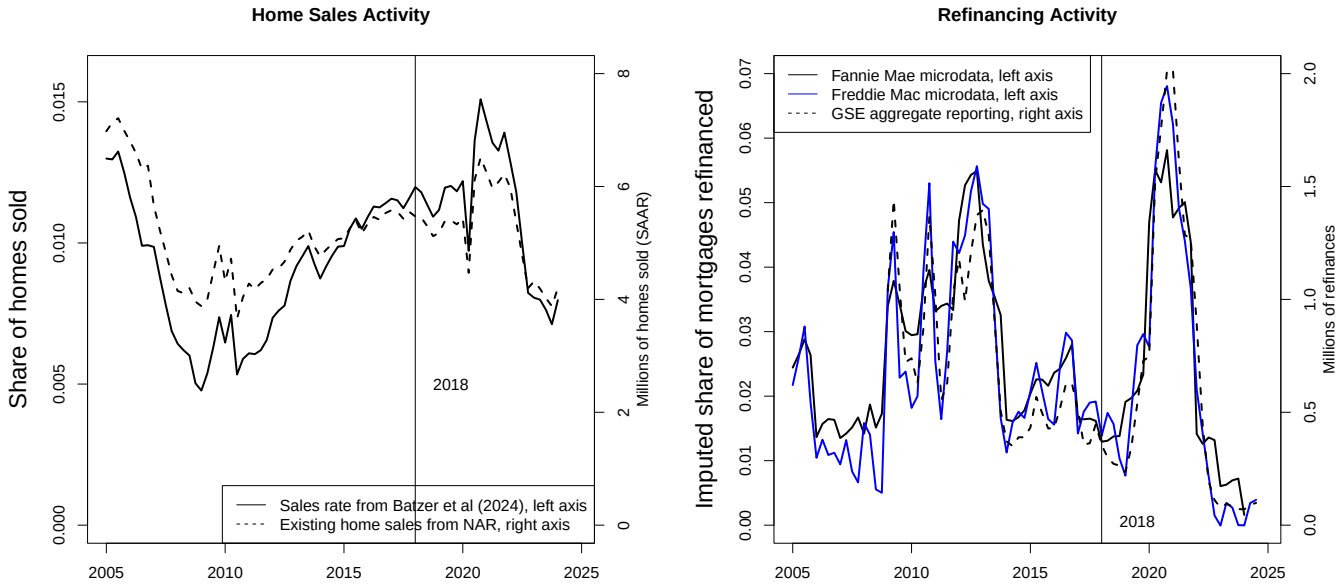


Figure 2: Nominal mortgage interest rate and inflation rate. The inflation rate shown here is lagged by 12 months.

Rate Mortgage" dataset and Freddie Mac's "Standard" loan-level mortgage performance dataset are used to construct a time series of refinancing activity.[22] The loan-level data is used to construct a quarterly time series of mortgage terminations. From there, the quarterly refinancing rate (i.e. share of households that refinance) is calculated as the share that terminated their mortgage (from the microdata) minus the share that sold their home (from Batzer *et al.* (2024)). This makes the strong assumption that all mortgage terminations result either from sales or refinances, which is not true. Nonetheless, Figure 3b shows that the series of refinancing rates constructed in this way has a strong correspondence with aggregate refinancing activity of GSE mortgages as reported by the Federal Housing Finance Agency (FHFA).[23] Furthermore, as described in Section 3.2, the calibration will depend not on the

[22]These datasets track 30-year FRM GSE mortgages that are fully-amortizing and fully-documented. There are additional restrictions that remove high-risk mortgages, such as Fannie Mae removing loans with origination loan-to-value ratios greater than 97%. While some of these sample selection criteria are more restrictive than is ideal for this paper, note that the data is used here to measure refinancing activity, rather than mortgage default, making the loss of higher-risk loans less troubling.

[23]In particular, two sets of reports produced by FHFA can be used to construct a time series of refinance volume of GSE mortgages: Refinance Reports, which provide the relevant data from Q1 2009 through Q1 2019; and Foreclosure Prevention, Refinance, and FPM Reports, which took the place of the



(a) Home sales series.

(b) Refinancing series.

Figure 3: Sales and refinancing data. Figure 3a is a validation of similarity between the sales rate calculated in Batzer *et al.* (2024) and aggregate data reported by the National Association of Realtors. Figure 3b is a validation of similarity between the refinancing rates calculated by combining microdata from Fannie Mae and Freddie Mac with sales data from Batzer *et al.* (2024) and aggregate refinancing data of GSE mortgages reported by the FHFA. "SAAR" stands for "seasonally-adjusted annual rate."

level of refinancing but rather on a second moment. As a result, to the extent that this approach ignores relatively acyclical sources of mortgage terminations, such as households completing their amortization schedules, it will not affect the results of the analysis.

Finally, the American Housing Survey (AHS) is used to to evaluate the model's prediction of households' baseline hazard function. Following Ngai and Sheedy (2020), the hazard function is calculated for each cohort (i.e. year of prior purchase) as the decline in the number of homes from that cohort in successive surveys. The baseline hazard function is then a weighted average across cohorts.[24]

Refinance Reports after March 2019. They are all available at https://www.fhfa.gov/reports?topic=fannie-mae-and-freddie-mac-reports.

[24] As the AHS is conducted every two years, the hazard rate for cohort c in year t is calculated as $1 - \sqrt{N_{ct}/N_{c,t-1}}$, where N_{ct} is the number of homes in cohort c that remain in the survey from year t. Note that because the survey is only conducted every other year, hazard rates for durations less than two years cannot be calculated.

Parameter	Meaning	Value	Source
ρ	Discount rate (annualized)	0.05	Agarwal *et al.* (2013)
C^R	Refinancing cost	0.018	Agarwal *et al.* (2013) (for \$250k mortgage)
C^M	Moving cost	0.09	Fonseca and Liu (2024)
π_L	Low inflation rate (annualized)	0.022	
π_H	High inflation rate (annualized)	0.053	Empirical averages in the US from April 1991 - December 2023, using 4% as the cutoff value between low and high inflation
p_π^L	Probability of remaining in low-inflation state	0.997	
p_π^H	Probability of remaining in high-inflation state	0.975	
σ_r	Real mortgage rate volatility (annualized)	0.024	Matches nominal mortgage rate volatility from April 1991 - December 2023 (0.017)
σ_x	Housing mismatch volatility	0.0010	
p_M	Probability of exogenous move	0.0013	Match moments of refinancing and sales activity from 2018-2023, as described in body of paper
p_R	Probability of being attentive to interest rate	0.038	
p_A	Probability of considering a move, if attentive to interest rate	0.80	

Table 1: Calibrated parameters

3.2 Calibrated Model

Table 1 shows the calibrated values of the model's parameters that are used in the analysis in subsequent sections. The first panel shows parameters whose values are set based on other papers in the literature. The second panel lists the inflation-related parameters, which are calibrated based on empirical analogs from April 1991 through December 2023. Given those values, σ_r is calibrated to match the empirical volatility of the nominal interest rate, assuming the nominal interest rate follows the process described in Section 2.1:

$$dm_t = d\pi_t + dr_t. \tag{2}$$

The four parameters in the final panel of Table 1 (σ_x, p_m, p_R, and p_A) are calibrated via simulation.[25] Specifically, for each quadruple in a grid of those four parameters, the following

[25]While other papers in the literature, such as Agarwal *et al.* (2013), have calibrated values for the arrival of exogenous moving shocks, those are not appropriate for this model. In those models, where all moves are exogenous, those shocks need to account for the entire hazard of moving, whereas in the model in this paper, where there is also an endogenous source of moves, the arrival of exogenous moving shocks must be lower. Therefore, p_M must be recalibrated in this paper. Similarly, the model in Berger *et al.* (2024) allows for inattention to interest rates, and their calibration finds that the average borrower pays attention to interest rates in 2.6% of months. Again, that parameter value is not appropriate for the model in the present paper because, as discussed above, the refinancing decision is affected by the possibility of endogenous moves. Reassuringly, though, the calibrated value of p_R in Table 1, 3.8%, is fairly close to what is found by Berger

steps are performed for 100,000 households:

1. Start in time $t = 0$ with $x_{it} = y_{it} = 0$;
2. Simulate 1,000 months of behavior with shocks to x drawn from $N(0, \sigma_x^2)$, assuming y remains constant at 0 and inflation is in its low state;
3. Store x_{1000}, the household's housing mismatch at the end of the 1,000 months;
4. Simulate forward for 393 additional months, where shocks to x continue to be drawn from $N(0, \sigma_x^2)$ and the nominal mortgage interest rate and (lagged) inflation follow the paths shown in Figure 2 from April 1991 to December 2023 (393 months).

The parameters shown in Table 1 are the ones that best matched the moments of sales and refinancing activity from 2018-2023 shown in Table 2.

The rationale for this procedure is to find a set of parameters that is able to mimic aggregate measures of housing market activity in recent years. The moments being matched draw from the period starting in 2018 because the model is not designed to capture some important dynamics of the housing market for the decade or so prior, which was heavily influenced by a series of destabilizing developments, including: notable relaxation and then tightening of lending standards, large gyrations in expectations and realizations of home price growth, a large rise and decline in the use of housing as method of speculation among investors (i.e. non-owner occupants), the Great Recession, and the long recovery and stabilization of the market from all of these factors.[26] As shown in Figure 3a, home sales activity seemed to stabilize after these fluctuations in 2018, and so that is where I begin evaluating the model. Nonetheless, I cannot simply start the simulations in 2018 because, while the distribution of y (the household's interest rate gap) could be observed, the analysis requires the joint distribution y with x (housing mismatch), and x is not observable. The 1,000 months of simulation prior to introducing observed macroeconomic data approximates the ergodic distribution of x in a steady state where interest rates are constant. Then, by introducing a long pre-period of observed data (April 1991 through December 2017), the procedure

et al. (2024). It makes sense that the model with endogenous moves finds a greater rate of attention to interest rates because it is able to justify some refinancing inaction by borrowers by the possibility that their housing mismatch might be close to triggering a move, lessening the degree of inattention that needs to be assumed.

[26] The literature on this episode of the housing market is far too large to cite adequately. Two papers that are particularly germane are Berger *et al.* (2020), who study the role of investors in driving the joint dynamics of price and sales volume in the housing market, and Kaplan *et al.* (2020), who quantify the importance of home price expectations and credit conditions in driving the housing cycle of the 2000s.

Moment	Data	Model
Average sales rate, Q1 2018 - Q1 2020	0.012	0.011
Average sales rate, Q3 2020 - Q1 2022	0.014	0.013
Average sales rate, Q2 2022 - Q4 2023	0.009	0.007
(Average refinance rate, Q3 2020 - Q1 2022)-		
(Average refinance rate, Q1 2018 - Q1 2020)	0.026	0.045

Table 2: Moments of data targeted in calibration

estimates what the joint distribution of y and x looked like prior to the evaluation period of 2018-2023.

As shown in Table 2, the model is capable of matching the main patterns of home sales activity from 2018-2023. In particular, it is able to essentially match the roughly 20% increase in sales that occurred after the immediate onset of the COVID-19 pandemic, which initiated a steep decline in interest rates as the Federal Reserve pursued aggressive expansionary policy. Then, following the onset of higher inflation, the Federal Reserve increased interest rates substantially, generating a roughly 35% decline in home sales, which the model is also able to match. Figure 4a shows finer detail of how the model compares to the data. One clear difference is that households in the model react more quickly to changes in interest rates, as the data lags behind by 1-2 quarters.[27,28]

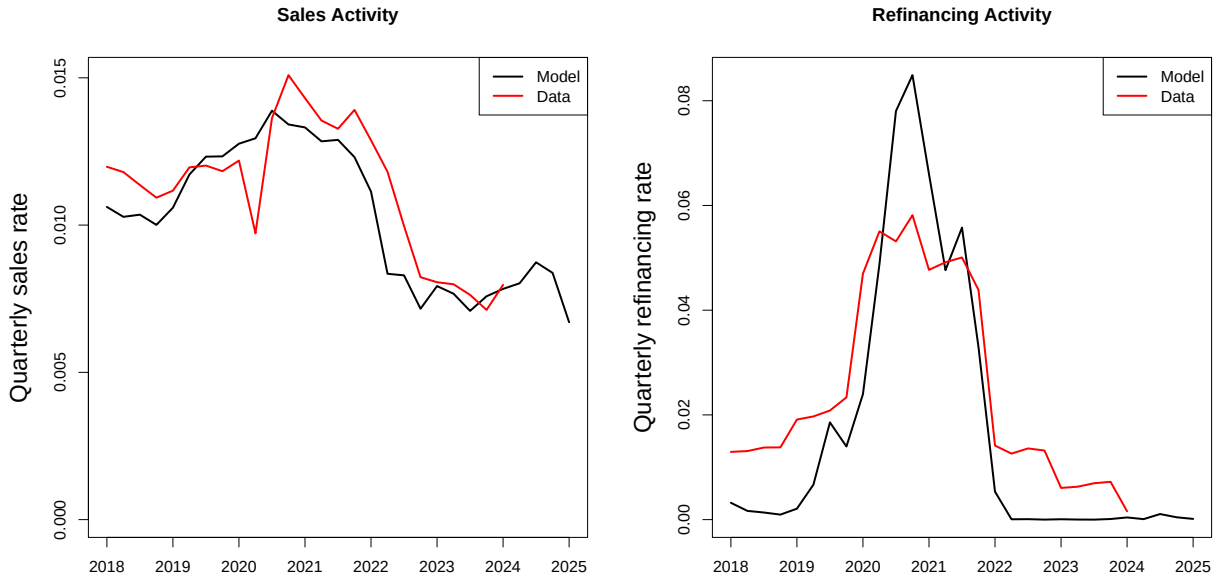
Figure 4b shows the model's performance with regards to refinancing behavior. In periods of high interest rates, the model predicts essentially zero refinancing, whereas there is always some refinancing in the data at all times. In the model, there is no reason to refinance when market interest rates have risen, but in reality some borrowers might refinance in such environments for reasons not considered by the model, such as equity withdrawal, term change, or an improved credit profile that could allow a household to lower its interest rate even if market rates have gone up since their mortgage origination.[29] As a result, it is not surprising that the model does not match the *level* of refinancing seen in the data. That is the rationale for using the *change* in refinancing activity that occurred after interest rates fell sharply during the pandemic, as shown in row 4 of Table 2.

While there is broad alignment between the data and the model with respect to refinancing

[27]It is this time shift that causes the model to appear to undershoot the data in row 3 of Table 2. Figure 4a shows that the model closely matches the level and decline of sales in the later period quite well, but the calculation of the average sales rate is somewhat skewed by the shifted timing of the decline.

[28]The model does not replicate the decline in sales in the first quarter of the pandemic, when stay-at-home restrictions were most stringent. Given the extremely unusual circumstances surrounding that event, Q2 2020 is not used in the calculation of any moments used in the model's calibration, as shown in Table 2.

[29]There is also the possibility that the data processing discussed in Section 3.1 misclassified some terminations as refinances, such as when a borrower fully amortizes its loan.



(a) Home sales series.

(b) Refinancing series.

Figure 4: Comparison between data and model of sales and refinancing activity from 2018-2025Q1. The "Data" series are in part on the sales series provided by Batzer *et al.* (2024) and so are only available through 2024Q1. The refinancing "Data" series is based on Fannie Mae data, though results are extremely similar with Freddie Mac data. As described in Table 2, these data series are targeted by the calibration.

behavior, there are also interesting differences. As shown in Figure 2, mortgage rates declined significantly during 2019, which caused an increase in refinancing before the onset of the pandemic. In this case, the model actually appears to lag the data. This may be because real households believed that interest rates could not go much lower than their levels in 2019 and so were willing to refinance at mortgage rates that households in the model deemed to be too high. Then, as mortgage rates continued to fall during the pandemic, refinancing in the model surges, making up for any prior lag. As mortgage rates rose in 2022, households in the model scaled back their refinancing more quickly than those in the data, replicating the lag structure seen in the sales series. This more detailed account of refinancing activity shows that while the model certainly does not perfectly replicate the data, it does a better job than row 4 of Table 2 seems to indicate, as much of the discrepancy is due to relatively subtle differences in timing of refinancing activity.

As another check on the performance of the model, Figure 5 shows loans that were active in the last month of available Fannie Mae microdata described in Section 3.1, June 2024. The figure shows in red the histogram of their interest rate gaps, $y_{iJun2024}$, as the difference

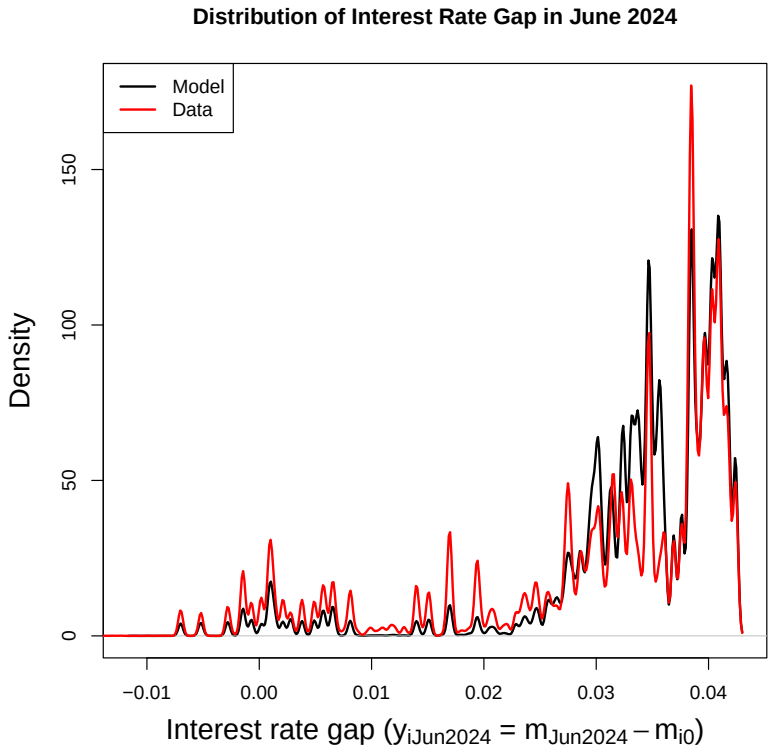


Figure 5: Distribution of $y_{iJun2024}$, the interest rate gap in June 2024, from the data and the model. The "Data" series is constructed from Fannie Mae's loan-level data. Freddie Mac's loan-level data is not used for this because it reports the quarter of mortgage origination, rather than the month, which leads to a coarser calculation of y. As the Fannie Mae data begins with mortgage originations in January 2000, the households from the model that are used in this figure are also those whose most recent origination was in January 2000 or later. The densities are estimated using a Gaussian kernel with bandwidth of 0.0002.

between the mortgage rate in June 2024 (6.92%) and the mortgage rate at the time of loan origination. The black curve represents the analogous object from the model, which corresponds very closely to the red empirical series. The main discrepancy between the two distributions is that the model produces more households with $y_{iJun2024}$ in the 0.030-0.035 range, owing to the differential timing of refinances around the time of the onset of the COVID-19 pandemic, as discussed above – those are households who get very low interest rates. Note that this distribution was not targeted by the calibration – only the four moments in Table 2 were targeted.

Finally, Figure 6 shows that the model is successful at replicating patterns in household-level microdata, even though these patterns were not targeted in the calibration. Figure 6a shows the relationship between probability of sale and the amount of time since the household

purchased the house – the baseline sale hazard. The figure compares the model to the empirical analogue computed from the AHS, following Ngai and Sheedy (2020). While the empirical sales hazard is quite noisy, it is closely aligned with the model's hazard, showing a peak around 5-7 years and then a steady decline thereafter. The model's interpretation of the data is that households are well-matched to their homes immediately after purchasing them and so are unlikely to move in the first few years. However, mismatch builds up for many household and the sale hazard rises correspondingly. The relationship reverses eventually (apparently after 5-7 years) because the households who remain in their homes that long are disproportionately well-matched and so unlikely to move. That selection effect gets stronger over time, driving down the sale hazard.[30] Figure 6b shows the hazard of mortgage prepayment/termination at the monthly level for households in the Fannie Mae microdata.[31] The *level* of the prepayment hazard in the model is too low – just as it generates aggregate refinancing shares that are too low, as discussed above – because the model does not consider refinances that are motivated by equity extraction, term change, or credit score improvement. But while the level is too low, the time pattern of the hazard in the model is a very close match to that of the data, just as was the case for aggregate refinancing shares in Figure 4b.

3.3 Measuring the Current State of Lock-In

The calibrated model can be used to track the state of lock-in in a detailed way. Figure 7 is useful for doing so. The curves show the moving and refinancing rules from the calibrated model, analogous to what was discussed for Figure 1. Overlaid on those rules are individual model households' values of housing mismatch, x_{it}, and interest rate gap, y_{it}, in select months. In February 2020, when mortgage interest rates were near historic lows, only 0.4% households were in the strong lock-in region of the state space (3), deterred from moving despite housing fundamentals that made moving efficient (i.e. $x_{it} > x^*$). In fact, only 1.0% of households were in regions 2 or 3, where a sufficiently-sized decline in the interest rate could have

[30] Ngai and Sheedy (2020) present a model that matches the empirical baseline sale hazard in a different way. In their model, households become systematically more mismatched to their homes over time. This implies an upward-sloping sale hazard at the household level. They match the non-monotonic pattern in the data by assuming *ex-ante* heterogeneity in households' mismatch processes, so that households remaining in their homes more than 5-7 years are those whose mismatch builds more slowly and so are less likely to sell. In a sense, that explanation is similar to the selection explanation presented in the present paper. The main difference is that the present paper produces the non-monotonicity *at the household level* and does not assume any *ex-ante* heterogeneity. The key to matching the data is that mismatch can both rise and fall, so as it initially rises for some households, the sale hazard rises but leaves the remaining households with disproportionately low levels of mismatch, leading to a decline in the hazard.

[31] Recall that the Fannie Mae microdata does not distinguish among sales, refinances, and other types of prepayments. That is why Figure 6b evaluates the model based on the combined outcome and why Figure 6a required the AHS to evaluate the hazard of sale.



(a) Sale Hazard.

(b) Prepayment/termination hazard.

Figure 6: Panel a) shows the baseline sale hazard (i.e. annualized probability of sale conditional on time since purchase) from the data and the model. The "Data" series is constructed the AHS following the approach used by Ngai and Sheedy (2020). As the AHS is conducted every two years, the sale hazard at durations less than two years cannot be calculated. Panel b) shows the prepayment hazard, which is calculated based on the Fannie Mae microdata. The Fannie Mae microdata is available at the monthlyl level but does not distinguish between different types of prepayments, such as sales and refinances.

induced a move, efficient or not. In contrast, by March 2025, when interest rates had risen sharply, 6.8% of households were in region 3, prevented from efficient moves by their interest rate disincentive. When including region 2, which contains households who would make inefficient moves if the interest rate were lower, this number rises to 7.6%. In other words, while the interest rate gap disincentivized 1.0% of households from moving (if and when they received the appropriate attention shocks) in February 2020,that number had spiked to 7.6% in March 2025.

Table 3 provides a more detailed breakdown. It shows that in March 2025, 12.6% of households would consider moving if interest rates were sufficiently low (i.e. $x > \tilde{x}$). Of those households, 4% (or 0.5% of all households) would need a decline of more than 4pp in the interest rate before moving and 44% (or 5.5% of all households) would require a decline of more than 1pp. In contrast, no households required a decline of more than 1pp to incentivize

a move in Feburary 2020. Given that the quarterly moving rate is roughly 1%, the large shift of households upward in Figure 7 and rightward in Table 3 represents the removal of many potential participants from the active housing market, generating the steep decline in home sales in beginning in 2022 displayed in Figure 2.[32]

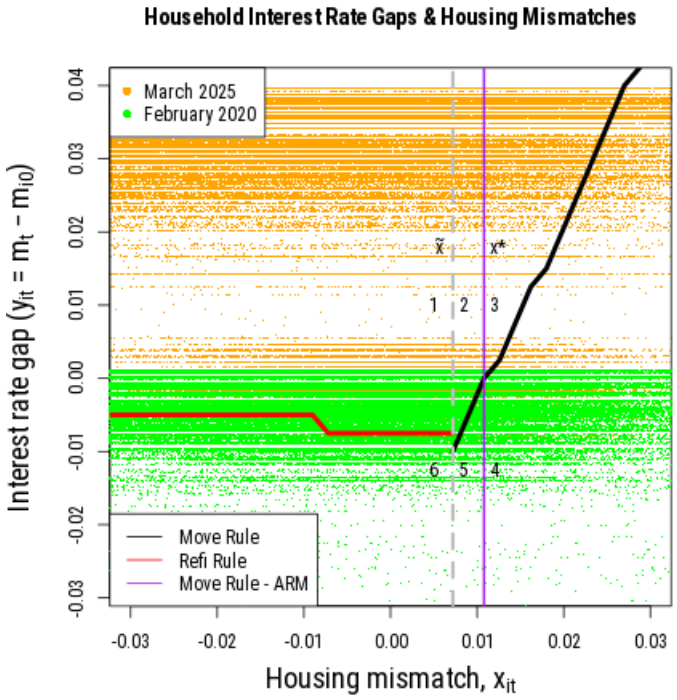


Figure 7: Household level values on housing mismatch, x_{it}, and interest rate gap, y_{it}, from model in February 2020 and March 2025. Solid lines show households' cutoff rules for moving and refinancing, as described in Section 2.3 and Figure 1. The dashed line at $x_{it} = \tilde{x}$ separates households who would refinance if y_{it} were sufficiently low from those who would move. In regions 1 and 2, the household takes no action, regardless of the mortgage type. In region 3, the household moves if it has an ARM but takes no action if it has a FRM (lock-in). In region 4, the household moves regardless of the mortgage type. In region 5, the household moves if it has a FRM but takes no action if it has an ARM. In region 6, a household with a FRM refinances. Taking action is conditional on paying attention.

[32] As a back-of-the-envelope calculation, consider that households in the model consider moving in $0.038 \cdot 0.8 = 3\%$ of months. If we take the 12.6% of households who are on the margin of moving rather than refinancing (i.e. $x_{it} > \tilde{x}$) in March 2025 but assume that 90.6% of them had $y_{it} < y^*(x_{it})$, as was true in February 2020 – rather than just 38.9%, as in March 2025 – then the quarterly sales hazard would have been 0.6pp higher. That accounts for essentially the entire decline in sales starting in 2022.

Month	$x_{it} > \tilde{x}$	$x_{it} > \tilde{x}$ and $y_{it} - y^*(x_{it}) \in$ $(-\infty, 0)$	$(0, 1\%)$	$(1, 2\%)$	$(2, 3\%)$	$(3, 4\%)$	$(4, 5\%)$
February 2020	10.6%	9.6%	1.0%	0%	0%	0%	0%
March 2025	12.6%	4.9%	1.6%	1.7%	2.0%	1.8%	0.5%

Table 3: Household distribution of decline in mortgage interest rate required to incentivize a move. Column 2 shows the share of households for whom moving is preferred to refinancing (i.e. a decline in interest rates could incentivize a move, rather than a refinance). Columns 3-8 then break that group down into bins of required interest rate declines. A negative value of $y_{it} - y^*(x_{it})$ indicates that a household is ready to move and will do so if it receives the appropriate attention shocks.

4 Policy

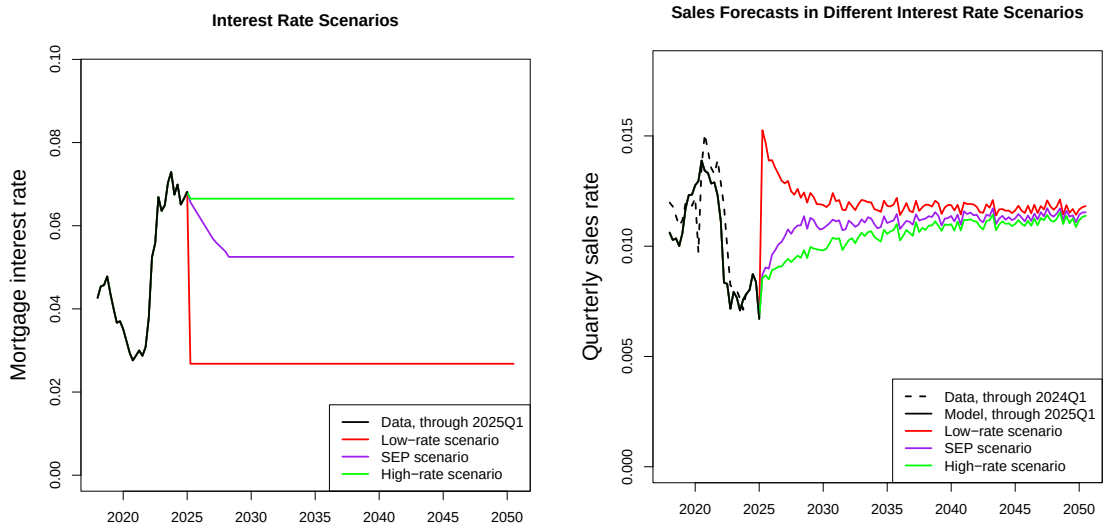
This section first considers the evolution of home sales under different paths of interest rates, i.e. monetary policies. It then considers the implementation of a subsidy (or tax credit) for home sellers and how such a fiscal policy interacts with the interest rate environment created by monetary policy.

4.1 Monetary Policy

The forward-looking analysis of policy begins with the model households' joint distribution of housing mismatches and interest rate gaps as of March 2025, $(x_{iMar2025}, y_{iMar2025})$. From there, household behavior is simulated under the three interest rate scenarios shown in Figure 8a. The "High-rate scenario" assumes that the mortgage interest rate remains at its March 2025 level (6.65%) for the foreseeable future. The "Low-rate scenario" assumes it immediately falls to its minimum observed level (2.68%, as in December 2020) and remains there for the foreseeable future. The "SEP scenario" assumes the path of the Federal Funds Rate follows the median projection in the Federal Open Market Committee's Survey of Economic Projections[33] and that changes in the Federal Funds Rate translate one-for-one into changes in the mortgage interest rate.

Figure 8b shows the level of home sales forecasted by the model under these different scenarios. Though the SEP-based projection is likely the most reasonable scenario, it is instructive to first consider the other two to better understand the dynamics of the model. Interestingly, even if interest rates never fall, as in the green "high-rate" scenario, home sales will recover.

[33]Based on the Survey of Economic Projections released on December 18, 2024, the median projection of the Federal Funds Rate falls by 0.5pp from 2024 to 2025, another 0.5pp from 2025 to 2026, another 0.3pp from 2026 to 2027, and then another 0.1pp as it settles into the long-run. See https://www.federalreserve.gov/monetarypolicy/files/fomcprojtabl20241218.pdf.



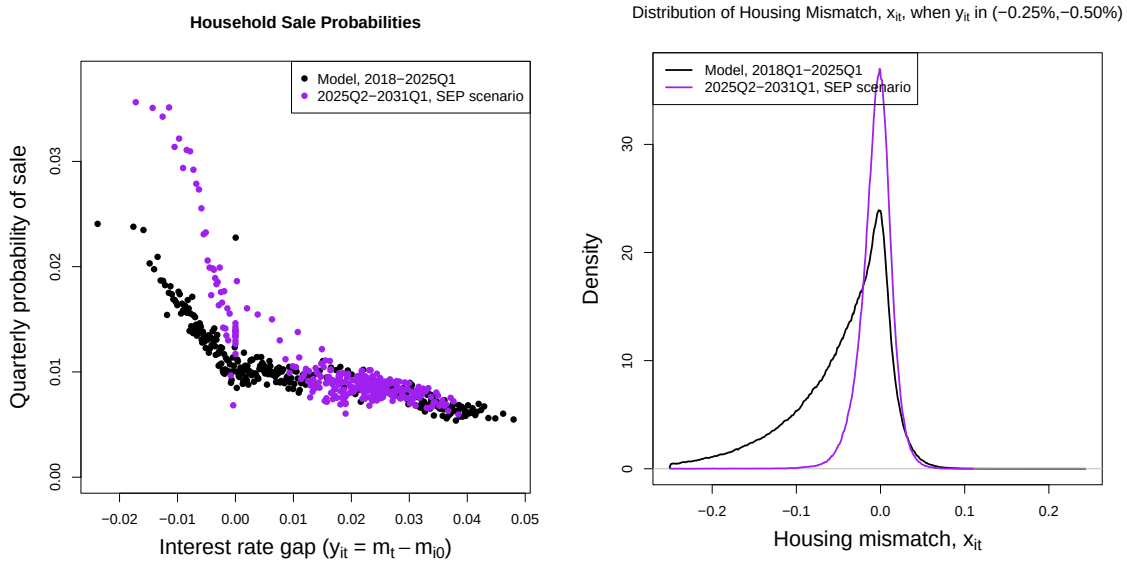
(a) Interest rate scenarios.

(b) Home sale forecasts.

Figure 8: Assumed scenarios of interest rate paths and resulting series of home sales.

Even when the interest rate creates a strong disincentive against moving, the market does continue to grind out some sales, and as it does so, more and more homeowners give up their low interest rates and re-equilibrate with the market level of rates, weakening the lock-in effect, though the figure suggests this could take a long time. If, in contrast, interest rates immediately fall back to the low experienced in December 2020, the home sales series will not simply move to its long-run level because lock-in has been resolved. Rather, as shown in red, sales volume will surge past its long-run level for a number of years. This results from the release of pent-up desire to move that was building during the lock-in period starting in 2022. Intuitively, the stock of households whose housing fundamentals indicated they should move but did not do so because of the high interest rates built up while interest rates remained high. A sharp fall in interest rates would release them back into the market as active participants. Between these two extremes, the SEP-based scenario in purple predicts a the recovery in home sales will be far faster than in the high-rate scenario as interest rates fall but with minimal overshooting, as interest rates are not projected to fall very steeply.

The potential for overshooting of home sales if and when interest rates decline is an interesting feature of the dynamics of this model, and Figure 9 examines it more closely. The black dots in Figure 9a are a binned scatterplot of sales against the interest rate gap, y, from 2018Q1-2025Q1, using data generated by the model. The relationship is very similar to that shown in the empirical work of Fonseca and Liu (2024) and Batzer *et al.* (2024). When the interest



(a) Binned scatterplot of sale probability against interest rate gap, y_{it}, in different periods and scenarios.

(b) Distribution of housing mismatch, x_{it}, for households with $y_{it} \in (-0.25\%, -0.50\%)$, during different periods.

Figure 9: Comparison of household selling behavior in different periods. For each period in Figure 9a, the interest rate gap is partitioned based on y_{it} into 300 bins with the same number of observations. Within each bin, the average level of y_{it} and the sales rate are plotted against each other. Figure 9b shows the distributions of housing mismatch, x_{it}, in the different periods for households with $y_{it} \in (-0.25\%, -0.50\%)$. The densities in Figure 9b are estimated using a Gaussian kernel with bandwidth of 0.0002

rate gap is positive, the relationship is negative but weak, as moves are mostly driven by the exogenous source, but higher levels of y do discourage some of the endogenous moves that continue to happen. When y is negative and small in magnitude, the relationship is far stronger, as a larger share of moves are for endogenous housing mismatch reasons, and so increases in y discourage a larger share of moves, as households decide they will tolerate housing mismatch so as not to suffer an increase in their interest rate. Finally, the relationship flattens out again as y becomes very negative and the refinancing option becomes a viable alternative to moving.

The purple dots in the figure represent the same empirical object but evaluated in 2025Q2-2031Q1 under the SEP-based scenario. Importantly, while the general pattern is the same, there is an upward-shift in the probability of sale, particularly when $y < 0$. Figure 9b shows why this happens by focusing on borrowers with interest rate gaps within (-0.25%,-0.50%). Conditional on that range of interest rate gaps, the distribution of housing mismatch, x,

shows a pronounced rightward shift in the later period relative to the earlier. Intuitively, in the earlier period – when there had not been a recent, major episode of lock-in – the current stock of homeowners was relatively well-matched to their homes (low mismatch), and so their likelihood of moving was relatively low. However, during the lock-in episode beginning in 2022, many households remained in their homes despite fairly high mismatch, and so if and when interest rates do eventually fall, they will be eager to sell their homes and move. This demonstrates a change in the reduced form relationship between y and the likelihood of sale, which underlies why the model predicts home sales can overshoot during the market recovery, as shown in Figure 8b. It also highlights the benefit of using a structural model to understand future market dynamics, since the reduced-form empirical relationship is not stable and is likely to shift.

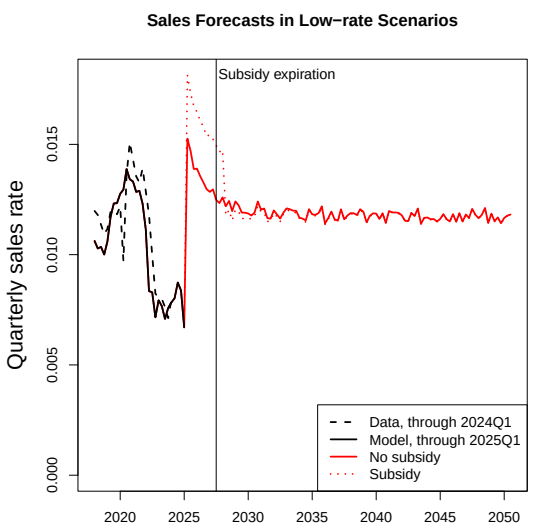
4.2 Home Sale Subsidy

In his 2024 State of the Union address, President Biden proposed a tax credit of $10,000 for home sellers as a way to counteract lock-in.[34] The model can accommodate this policy by considering it as a perceived reduction in the cost of moving. In particularly, because $10,000 is 4% of $250,000, the subsidy is modeled as a reduction in the cost to households of moving from 0.09 to 0.05. The subsidy is assumed to last three years, corresponding to the period until the mortgage interest rate reaches its long-run level in the SEP-based scenario. The temporary nature of the subsidy means that the households' problem is no longer stationary: starting in month 37 of the forecast, the value function reverts to what is characterized in Equation 1, but in each of the 36 months prior, a separate value function is needed to account for the approach of the subsidy's expiration. This is solved via backwards induction.

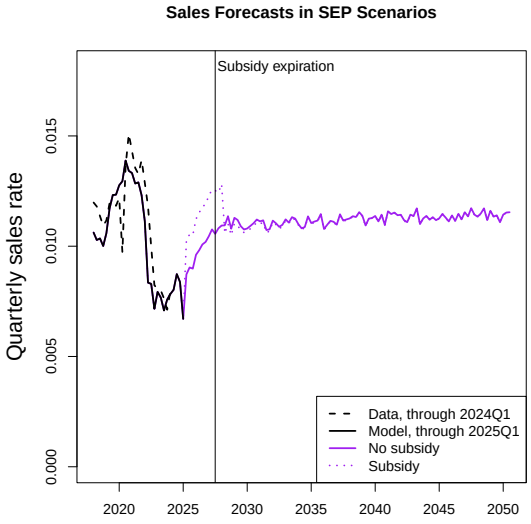
4.2.1 Sales Volume

Figure 10 shows how the subsidy would affect home sales volume under the different interest rate scenarios described above. Figure 10a shows that if interest rates were to fall quickly and essentially erase the lock-in incentive, then the subsidy would greatly magnify the overshooting of home sales up until the point of subsidy expiration. In the SEP-based scenario shown in Figure 10b, the subsidy would speed up and magnify the recovery of home sales, culminating in a burst of sales just before subsidy expiration. In the high-interest rate scenario shown in Figure 10c, the impact is more muted as the degree of lock-in is so high

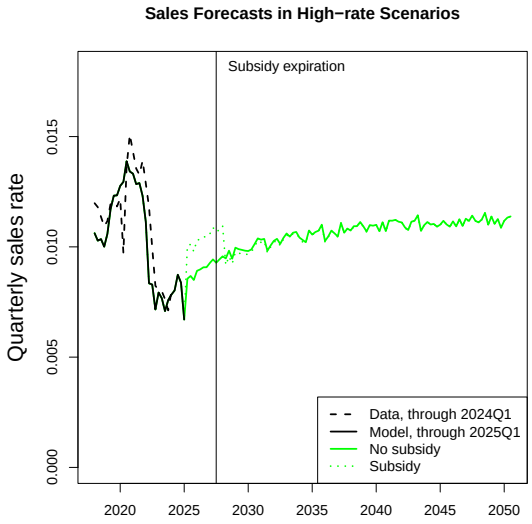
[34]The tax credit is in fact intended to be targeted at sellers of "starter homes." For an analysis of the impact of this subtlety, see Fonseca *et al.* (2024), who analyze the subsidy in a model with homes of different qualities.



(a) Forecasts in low-interest rate scenario.



(b) Forecasts in SEP scenario.



(c) Forecasts in high-interest rate scenario.

Figure 10: Model forecasts of home sales under different interest rate scenarios, with and without subsidy.

that the tax credit of $10k does little to stimulate sales, though there is a clear burst of sales before expiration.

4.2.2 Efficiency and Welfare

To assess the efficiency implications of the subsidy, I track the flow of $x_{it} \cdot \frac{1+\rho}{\rho} - C^M$ among movers who are motivated by housing mismatch. The first term represents the PDV of the increased flow of housing utility they receive, which is the social value of the move – any change in interest payments is ignored because that simply serves as a transfer between households and lenders. C^M is the social cost of the move.

In the low-interest rate scenario, the subsidy lowers efficiency by about 0.003% of the size of the value of the housing market. Intuitively, lock-in effects disappear in this scenario – households move out of region 3 as shown in Figures 1 and 7 even without the subsidy – so the subsidy simply serves to over-stimulate sales activity as households are shielded from its true cost. In contrast, the subsidy increases efficiency by roughly 0.024% of the size of the housing market in the high-interest rate scenario. In that scenario, many households are eschewing moving not because it does not justify the social cost, C^M, but because it cannot additionally justify the increase in the interest rate that moving would require. Not moving is inefficient, and the subsidy spurs some of these people to move, reducing the inefficiency. In the more likely and less extreme SEP-based scenario, it turns out that the subsidy would increase efficiency by about 0.021% of the size of the housing market. Given that the US housing stock is worth roughly $50 trillion,[35] this is a non-trivial effect on the order of about $10 billion.

While this is a point in favor of the subsidy, it is important to keep in mind that most moves that occur during the subsidy's existence would have happened even in the absence of the subsidy, and these people receive the $10,000 payment anyway. As a result, the model finds that it takes roughly $40 of government expenditure to generate each $1 of additional economic surplus mentioned above. As a result, this policy would largely be a transfer from the government to homeowners and their lenders, with a small increase in efficiency generated in the process. While this does not make the policy objectively "bad," many people and policymakers may deem this an undesirable transfer.[36] This complements the findings of Fonseca *et al.* (2024), who show that while targeting such a subsidy to the sellers of starter homes – as President Biden's proposal does – would have a modest impact on

[35] See, e.g., "Total Value of U.S. Homes on the Brink of $50 Trillion: Redfin," by Jeff Andrews in HousingWire on August 8, 2024.

[36] A full accounting of this policy would require modeling where the funds for the subsidy come from. To the extent that funding the program requires cuts to social programs, this would probably be a very regressive policy. To the extent that it comes from increased taxes on, say, high incomes, the policy would not be regressive but funding it may cause inefficiency that could erase the efficiency gain that it created in the housing market in the first place.

first-time home-buying (the primary goal), it would mostly benefit homeowners higher in the housing ladder. The analyses in both papers, then, suggest that while the subsidy can help to address lock-in, policymakers should be conscious of unintended consequences that may make it undesirable on distributional grounds.

5 Mortgage Design

A large literature considers the relative merits of different mortgage designs. Among others, Campbell (2013), Zhang (2024), Berger *et al.* (2024), and Abel (2024) discuss the consequences of heterogeneous refinancing of FRMs and potential benefits of alternatives, such as ARMs or automatically-refinancing mortgages, which do not create worse outcomes for borrowers with less financial sophistication. Campbell and Cocco (2003) and Piskorski and Tchistyi (2010) show that ARMs can be particularly beneficial to borrowers who are more constrained and have riskier income streams. Eberly and Krishnamurthy (2014), Campbell *et al.* (2021), and Guren *et al.* (2021) consider the macroprudential impacts of different mortgage designs. This paper adds to this strand of the literature by studying a model in which moves are endogenous. As this paper has shown, this creates an additional burden of the FRM, which is that households do not maximize their flow of housing utility (less moving costs) because their decision to move is distorted by interest rates in a way that would not occur with an ARM.

Notably, this friction increases the prepayment risk faced by lenders, since moves are more likely when interest rates are low rather than high. To investigate how this affects mortgage pricing, I simulate 10,000 paths of the model's state variables, π, y, and x, and I keep track of mortgage payments, housing mismatch, moves, and refinances. I find that for a given initial interest rate, lenders receive an additional 52 basis points (bp) in mortgage payments from an ARM relative to an FRM. In other words, the premium they would need to charge to compensate for the two joint options to move and refinance is 52bp. This comes not only from ARMs' removal of refinancing, but also because ARMs make it so that households are not more likely to move when mortgage rates are low, relative to FRMs. To quantitatively assess the importance of the option value associated with moving, I rerun the simulations with the same average rate of moves (0.3% per month), but assuming that these moves are exogenous and so do not introduce any prepayment risk. In that version of the model, the premium that lenders would require is 45bp. So, the optionality of moves adds modestly to the prepayment risk that lenders face on FRMs, increasing the premium from 45bp to 52bp.

Turning to household outcomes, the simulations show that households see an increase in the expected PDV of their flow of housing utility (less moving costs) of 10.3% of the size of their

home value from having an ARM rather than a FRM. This accounts both for the fact that households would not experience lock-in, as in region 3 of Figures 1 and 7, as well as the fact that they would not move "too quickly" when in region 5. They would additionally save 2.2% of their home's value in foregone refinancing effort if they used an ARM rather than a FRM.[37] While these effects are large, it is worth remembering that households in this model are risk-neutral, and real households likely assign value to an FRM's guaranteed cap on mortgage payments. In that sense, one can think of the numbers above as the cost of choosing a FRM over an ARM; if borrowers do indeed prefer FRMs, then they must value the stability of payments above those costs.

6 Conclusion

Motivated by the recent rise in interest rates and decline in home sales – and empirical work by Fonseca and Liu (2024), Batzer *et al.* (2024), and Liebersohn and Rothstein (2025) linking the two – this paper has presented a dynamic model of household behavior in which interest rates are stochastic and households consider both an option to move and an option to refinance. On the empirical side, the dynamics of the model show that the reduced form relationship between a household's interest rate gap and its likelihood of moving is context-dependent, and in particular it depends on the past history of interest rates, much as in the study of refinancing in Berger *et al.* (2021). If interest rates have been high – and sales volume correspondingly low – in the recent past, then housing mismatch has had a chance to build up and the likelihood of moving will be higher than if recent interest rates had been low. This has important consequences for the dynamics of aggregate sales volume. If and when interest rates fall from their current levels, many of the sales that have been "lost" during the current lock-in episode will finally be realized, showing that they were in fact merely *delayed* as households waited out the high interest rates before pulling the trigger on a move. In other words, we can expect a boom in sales to follow a slowdown due to lock-in.

The model also allows for analysis of a home sales subsidy and, importantly, welfare analysis of such a policy. The effect and desirability of that policy depends critically on how monetary policy evolves. While the subsidy will boost home sales regardless of interest rates, the analysis showed that the implications for economic efficiency depend on the interest rate environment. The subsidy would likely be efficiency-enhancing (so long as interest rates do not decline rapidly), but the model estimates that it takes roughly $40 of government

[37] Underlying the calculation is the assumption that the refinancing costs borrowers bear are resource costs. To the extent one believes the costs are captured by the lending industry, then competition would lead borrowers in the FRM market to be compensated in the form of lower interest rates, meaning switching to an ARM would not save them the 2.2% mentioned above.

spending to generate $1 of additional social value. As a result, the subsidy's first-order effect is simply to transfer money to homeowners and their lenders, with increased efficiency being second-order. This is an important caveat for policymakers to consider.

A clear drawback of the analysis of the paper is the absence of any insight on home prices, which is an important and interesting dimension of the lock-in phenomenon. As discussed in Section 1, a number of other papers have explored the implications of lock-in for home prices, due to its impacts on net housing demand or market thickness. For tractability, those papers either assume that moves are exogenous (Amromin and Eberly (2024)) or they assume away dynamics and study comparative statics across steady states (Fonseca *et al.* (2024), Gerardi *et al.* (2024), Katz and Minton (2024)). The present paper therefore contributes to the literature and the discussion of policy by presenting a model with aggregate dynamics – notably, with regards to policy variables like interest rates and tax credits – in which moves are endogenous, allowing for analysis of how that critical housing market outcome is likely to be affected by policy, along with the corresponding welfare implications. Future work on a model like this with aggregate uncertainty and endogenous moves that also solves for equilibrium prices would be interesting and valuable.

Bibliography

Abel, J. (2024). *Equilibrium Mortgage Design with Heterogeneous Borrowers*. Working Paper.

Agarwal, S., Ben-David, I. and Yao, V. (2017). Systemic Mistakes in the Mortgage Market and Lack of Financial Sophistication. *Journal of Financial Economics*, **123** (1), 42–58.

—, Driscoll, J. C. and Laibson, D. I. (2013). Optimal Mortgage Refinancing: A Closed-Form Solution. *Journal of Money, Credit, and Banking*, **45** (4), 591–622.

Amromin, G. and Eberly, J. (2024). *Macro Shocks and Housing Markets*. Working Paper.

Andersen, S., Campbell, J. Y., Nielsen, K. N. and Ramadorai, T. (2020). Sources of inaction in household finance: Evidence from the danish mortgage market. *American Economic Review*, **110** (10), 3184–3230.

Anenberg, E. and Bayer, P. (2020). Endogenous Sources of Volatility in Housing Markets: The Joint Buyer-Seller Problem. *International Economic Review*, **61** (3), 1195–228.

Attanasio, O. P., Bottazzi, R., Low, H. W., Nesheim, L. and Wakefield, M.

(2012). Modelling the Demand for Housing Over the Life Cycle. *Review of Economic Dynamics*, **15** (1), 1–18.

Batzer, R., Coste, J., Doerner, W. and Seiler, M. (2024). *The Lock-In Effect of Rising Mortgage Rates*. FHFA Staff Working Paper Series.

Berger, D., Milbradt, K., Tourre, F. and Vavra, J. (2024). *Refinancing Frictions, Mortgage Pricing and Redistribution*. Working Paper.

—, Turner, N. and Zwick, E. (2020). Stimulating Housing Markets. *The Journal of Finance*, **75** (1), 277–321.

Berger, D. W., Milbradt, K., Tourre, F. and Vavra, J. (2021). Mortgage Prepayment and Path-Dependent Effects of Monetary Policy. *American Economic Review*, **111** (9), 2829–78.

Bernstein, A. and Struyven, D. (2021). Housing Lock: Dutch Evidence on the Impact of Negative Home Equity on Household MObility. *American Economic Journal: Economic Policy*, **14** (3), 1–32.

Best, M. and Kleven, H. (2018). Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the UK. *The Review of Economic Studies*, **85** (1), 157–93.

Campbell, J. Y. (2013). Mortgage market design. *Review of Finance*, **17** (1), 1–33.

— and Cocco, J. F. (2003). Household Risk Management and Optimal Mortgage Choice. *The Quarterly Journal of Economics*, **118** (4), 1449–1494.

—, Nuno, C. and Cocco, J. F. (2021). Structuring Mortgages for Macroeconomic Stability. *The Journal of Finance*, **76** (5), 2525–76.

Chen, A. H. and Ling, D. C. (1989). Monetary policy according to hank. *Real Estate Economics*, **17** (3), 278–99.

Chen, S. (2001). Spatial Lock-In: Do Falling House Prices Constrain Residential Mobility? *Journal of Urban Economics*, **49** (3), 567–86.

Eberly, J. and Krishnamurthy, A. (2014). Efficient Credit Policies in a Housing Debt Crisis. *Brookings Papers on Economic Activity,*, **Fall 2014**.

Ferreira, F. (2010). You can take it with you: Proposition 13 tax benefits, residential mobility, and willingness to pay for housing amenities. *Journal of Public Economics*, **94**, 661–673.

—, Gyourko, J. and Tracy, J. (2010). Housing busts and household mobility. *Journal of Urban Economics*, (68), 34–45.

Fonseca, J. and Liu, L. (2024). Mortgage Lock-In, Mobility, and Labor Reallocation. *Journal of Finance*.

—, — and Mabille, P. (2024). *Unlocking Mortgage Lock-In: Evidence From a Spatial Housing Ladder Model*. Working Paper.

Gabaix, X. and Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, **121** (2), 505–540.

Genesove, D. and Mayer, C. (1997). Equity and time to sale in the real estate market. *American Economic Review*, **87** (3), 255–269.

Gerardi, K., Qian, F. and Zhang, D. (2024). *Mortgage Lock-In, Life-Cycle Migration, and the Welfare Effects of Housing Marke Liquidity*. Working Paper.

Guren, A. M., Krishnamurthy, A. and McQuade, T. J. (2021). Mortgage Design in an Equilibrium Model of the Housing Market. *Journal of Finance*, **76** (1).

Imrohoroglu, A., Matoba, K. and Tuzel, S. (2018). Proposition 13: An equilibrium analysis. *American Economic Journal: Macroeconomics*, **10** (2), 24–51.

Kaplan, G., Mitman, K. and Violante, G. L. (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy*, **128** (9), 3285–345.

Katz, J. and Minton, R. (2024). *Rising Interest Rates, Mortgage Rate Lock, and House Price Fluctuations*. Working Paper.

Liebersohn, J. and Rothstein, J. (2025). Household mobility and mortgage rate lock. *Journal of Financial Economics*, **164**.

Ngai, L. R. and Sheedy, K. D. (2020). The Decision to Move House and Aggregate Housing-Market Dynamics. *Journal of the European Economic Association*, **18** (5), 2487–531.

Piskorski, T. and Tchistyi, A. (2010). Optimal Mortgage Design. *The Review of Financial Studies*, **23** (8), 3098–3140.

Quigley, J. (1987). Interest rate variations, mortgage prepayments and household mobility. *The Review of Economics and Statistics*, **69** (4), 636–643.

Stein, J. (1995). Prices and trading volume in the housing market: A model with down-payment effects. *The Quarterly Journal of Economics*, **110** (2), 379–406.

Zhang, D. (2024). *Closing Costs, Refinancing, and Inefficients in the Mortgage Market*. Working Paper.