

Mortgage Lock-In and Home Sales Volume Dynamics*

Joshua Abel[†]

June 9, 2025

Abstract

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 persistent and 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 amid recent fluctuations in interest rates and the impacts of proposed policies to stabilize sales volume. The model predicts that a tax credit for home-sellers in our current low-sales equilibrium 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. The model is also used to explore implications of mortgage lock-in for socially desirable mortgage design.

*Dayanara Diaz Vargas, Vedant Mundhra, Thomas Sargent, and Tarun Anbarasu provided excellent research assistance; funding was generously granted by the David A. Radivonyk Fund for Student Research in Economics at Northeastern University. I am grateful for helpful conversations with Andrew Abel, Santiago Caicedo, Jonah Coste, Edward Glaeser, Adam Guren, Justin Katz, Kyle Mangum, Robert Triest, and Paul Willen, as well as seminar participants in the Federal Housing Finance Agency’s Research Department. I also thank Luke Watson for his discussion at the 2025 AREUEA National Conference in the Monetary Policy and Public Finance session.

[†]College of the Holy Cross. E-mail: jabel@holycross.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 press¹ 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 new loans with high interest rates; 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 the consequences for their impact on the dynamics of home sales.

In the model, a household experiences persistent shocks to its idiosyncratic match quality to its current house. This match quality is a catch-all to encode every factor other than the interest rate that affects 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 this match quality and the nominal mortgage interest rate, which is driven by shocks to inflation and the real interest rate, the household decides 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.* \(2024\)](#) – who model 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, generating lock-in.

A calibration of the model does a good job of replicating recent movements in aggregate

¹See, 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.

²Other 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.

³Frictions 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 the institutional features of mortgages that they are not portable or assumable.

⁴It 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.

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. Under the interest rate path implied by the Federal Open Market Committee’s Survey of Economic Projections, it is likely that the tax credit would increase economic efficiency, as many households are eschewing efficient moves to avoid a transfer to their lenders in the form of higher mortgage payments. 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 serve primarily as a transfer from taxpayers to homeowners and their lenders.

Mortgage lock-in has attracted substantial attention from academic economists. This 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 assumptions to gain tractability.

One 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. [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

⁵For 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](https://www.whitehouse.gov) since March 7, 2024.

thickness, particularly for younger households in lower-income areas, while at the same time increasing prices. [Aladangady *et al.* \(2024\)](#) also model search-and-matching in the housing market and argue that lock-in increases market tightness and consequently prices. All 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 the COVID episode was so different than 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. So while that analysis does have dynamics, it seeks to explain home prices while taking sales volume as a given rather than modeling it. [Katz and Minton \(2024\)](#) also study dynamics, and while they allow for moves to be endogenous, they make the simplifying assumption that unobserved mismatch between households and their homes has no persistence, which offers them tractability. They then focus on home prices and show how lock-in can increase net housing demand by preventing own-to-rent moves and find that that friction can explain why home prices did not fall in response to rising interest rates, as a standard asset pricing model would predict.

The present paper therefore is unique in contributing to this literature an analysis of the *dynamics of home sales*. Unlike [Fonseca *et al.* \(2024\)](#), [Gerardi *et al.* \(2024\)](#), and [Aladangady *et al.* \(2024\)](#), there is aggregate uncertainty and the analysis is not confined to steady states. Unlike [Amromin and Eberly \(2024\)](#) and [Katz and Minton \(2024\)](#), the decision to move is endogenous and depends on a persistent component of mismatch between households and homes, which I will argue below is important for matching moments of the data.⁶ 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 price of a house is on both sides of the budget constraint.⁷ In this way, the paper is a complement to the existing literature because, while

⁶These papers also do not consider the option to refinance, which empirical studies show is an important dimension of household behavior, as discussed below.

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

it is silent on home prices, it 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 recent work 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 moving 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.⁸ They also show that homeowners who face larger disincentives are less likely to relocate in response to labor market incentives. [Katz and Minton \(2024\)](#) focus on the rise in interest rates of roughly 70bp at the end of November 2016 and find that homeowners who bought their homes after this rise were about 25% more likely to sell their homes in the next four years compared to those who bought prior the interest rate increase. [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%. [Batzner 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.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 reduced form empirical relationship between a household’s interest rate incentive and its likelihood of a move depends on the recent path of interest rates and

⁸In part motivated by this evidence, the model of the present paper incorporates the option to refinance, an empirically relevant factor that the previous structural literature discussed above abstracts away from.

⁹These 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 \(1995\)](#).

home sales. In particular, the model shows 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. This literature largely advocates for the adoption of an alternative mortgage design, along the lines of what [Abel \(2025\)](#) calls a “Self-Refinancing Mortgage” (SRM), which automatically refinances when interest rates decline. That strand of the literature focuses on benefits of the SRM relative to the FRM when interest rates are *falling*, but the lock-in phenomenon points

¹⁰The 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.

to a potential problem with the SRM: when interest rates *rise* after a period of decline, the SRM would seemingly worsen the lock-in effect because it would have allowed all borrowers to refinance when rates were low. However, the model shows that this intuition is incomplete. Critically, borrowers with an SRM would be more willing to move and accept a high interest rate than would a borrower with a standard FRM, since they know the high rate will be adjusted downward immediately and costlessly when interest rates again fall. In fact, the calibrated model finds that sales for 2022-2024 would have been *higher* if borrowers had SRMs due to that mechanism, a surprising result that points to a novel benefit of this sort of mortgage relative to the FRM. More generally, SRMs reduce sensitivity to interest rates for movers and so can reduce home sales volume volatility.

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 analyzes how homeowner behavior would differ under the alternative SRM 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 endogenous 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 related 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.

¹¹As 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.

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

2.1 Macroeconomic Variables

Consider a homeowner with 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 it is risk-neutral¹³ and that the mortgage is non-amortizing (interest-only).

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 $m_t = r_t + \pi_t$.^{14,15,16}

Letting m_{i0} be the market interest rate when household i last moved or refinanced, 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 household’s mortgage if it sells or refinances.

¹³Assuming 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 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 (2025) do quantitative analyses of household refinancing decisions and also maintain the assumption of risk-neutrality.

¹⁴This 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.

¹⁵Due 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 nominal interest rate is a pure Random Walk. Nonetheless, having the real interest rate follow a Random Walk greatly simplifies the model by allowing us to simply track the gap between the borrower’s interest rate and the current market rate (discussed below) without needing to keep track of both separately. This reduces the dimension of the state space which helps with both computational tractability and ease of interpretation. Adding even a simple feature of realism, such as a lower bound on m_t , would require an increase in the dimension of the state space and thus reduce tractability. As discussed in Section 3, the model is able to match critical moments of the data with this Random Walk assumption for r_t , and adding a lower bound on m_t would likely lead to nothing more than a mild improvement in its ability to match data on refinancing around the COVID pandemic.

¹⁶Quantitatively, inflation plays little role in model beyond driving the nominal mortgage interest rate: household behavior in the two states of inflation are very similar, suggesting a richer inflation process is not necessary.

2.2 Household Problem

To produce lock-in effects, the model must allow the household to have an endogenous 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 the flow of additional housing amenity the household would get if it moved to a new house. Normalize $x_{it} = 0$ in the immediate aftermath of a move and assume x_{it} follows a Random Walk with shocks drawn from $N(0, \sigma_x^2)$.¹⁷

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 0.¹⁸
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 housing 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 .¹⁹ 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 change in her mortgage interest rate of y_{t-} . The flow of housing utility changes by χ , which I set to

¹⁷So $x_{it} = x_{i,t-1} + \epsilon_{it}$, where $\epsilon_{it} \sim N(0, \sigma_x^2)$. The persistence of this state variable is important for a number of the paper’s results. In Section 3, I show that such persistence is supported by the data. If no persistence is assumed, so $x_{it} = \epsilon_{it}$, the model cannot recreate households’ non-monotonic empirical hazard of sale.

¹⁸As in much of the related literature, I abstract from the process of searching for a home. As shown in Ngai and Sheedy (2020), the main determinant of sales volume is the rate at which households list their homes for sale, with the speed of listed homes selling playing only a minor role. Gerardi *et al.* (2024) use a search model to impact of mortgage lock-in in a model with market thickness effects.

¹⁹I 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.

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:

$$\begin{aligned}
V(x, y, \pi) = & \\
(1 - p_M) \cdot p_R \cdot p_A \cdot & \left(\underbrace{\max \left\{ \frac{1}{1 + \rho} E_{x, y, \pi} [V(x_+, y_+, \pi_+)], -q_\pi \cdot y - C^R + \frac{1}{1 + \rho} E_{x, 0, \pi} [V(x_+, y_+, \pi_+)], \frac{1 + \rho}{\rho} \cdot x - q_\pi \cdot y - C^M + \frac{1}{1 + \rho} E_{0, 0, \pi} [V(x_+, y_+, \pi_+)] \right\}}_{\text{Value of no action}} \right. \\
& \underbrace{\left. \frac{1}{1 + \rho} E_{x, y, \pi} [V(x_+, y_+, \pi_+)], -q_\pi \cdot y - C^R + \frac{1}{1 + \rho} E_{x, 0, \pi} [V(x_+, y_+, \pi_+)] \right\}}_{\text{Value of refinancing}} \left. + \frac{1 + \rho}{\rho} \cdot x - q_\pi \cdot y - C^M + \frac{1}{1 + \rho} E_{0, 0, \pi} [V(x_+, y_+, \pi_+)] \right) + \\
& \underbrace{\hspace{10em}}_{\text{Value of considering both a move and refi}} \\
(1 - p_M) \cdot p_R \cdot (1 - p_A) \cdot & \left(\underbrace{\max \left\{ \frac{1}{1 + \rho} E_{x, y, \pi} [V(x_+, y_+, \pi_+)], -q_\pi \cdot y - C^R + \frac{1}{1 + \rho} E_{x, 0, \pi} [V(x_+, y_+, \pi_+)] \right\}}_{\text{Value of no action}} \right. \\
& \left. \frac{1}{1 + \rho} E_{x, y, \pi} [V(x_+, y_+, \pi_+)], -q_\pi \cdot y - C^R + \frac{1}{1 + \rho} E_{x, 0, \pi} [V(x_+, y_+, \pi_+)] \right\}}_{\text{Value of refinancing}} \left. + \frac{1}{1 + \rho} E_{x, y, \pi} [V(x_+, y_+, \pi_+)] \right) + \\
& \underbrace{\hspace{10em}}_{\text{Value of considering a refi but not a move}} \\
(1 - p_M) \cdot (1 - p_R) \cdot & \frac{1}{1 + \rho} E_{x, y, \pi} [V(x_+, y_+, \pi_+)] + \\
& \underbrace{\hspace{10em}}_{\text{Value of considering neither move nor refi}} \\
p_M \cdot \left(x - q_\pi \cdot y - C^M + \frac{1}{1 + \rho} E_{0, 0, \pi} [V(x_+, y_+, \pi_+)] \right), & \\
& \underbrace{\hspace{10em}}_{\text{Value of exogenous move}}
\end{aligned} \tag{1}$$

where $q_\pi \equiv 1 + E[\sum_{t=1}^{\infty} (\frac{1}{(1+\rho)^t \cdot \Pi_{\tau=1}^t (1+\pi_\tau)}) | \pi_0 = \pi]$. 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 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, $\frac{1+\rho}{\rho} \cdot x$.²⁰ 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 of Equation 1 have the same structure but account for different combinations of exogenous moving and attention shocks.

²⁰Housing 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 .

2.3 Household Solution

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 is useful to consider as a benchmark the behavior of a homeowner with an adjustable-rate mortgage (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 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 $x < \tilde{x}$, the household is on the margin of refinancing rather than moving because the housing utility gain from moving is not sufficiently large; when $x > \tilde{x}$, it is on the margin of moving because the gain is sufficiently large.

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 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 making an efficient move. 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

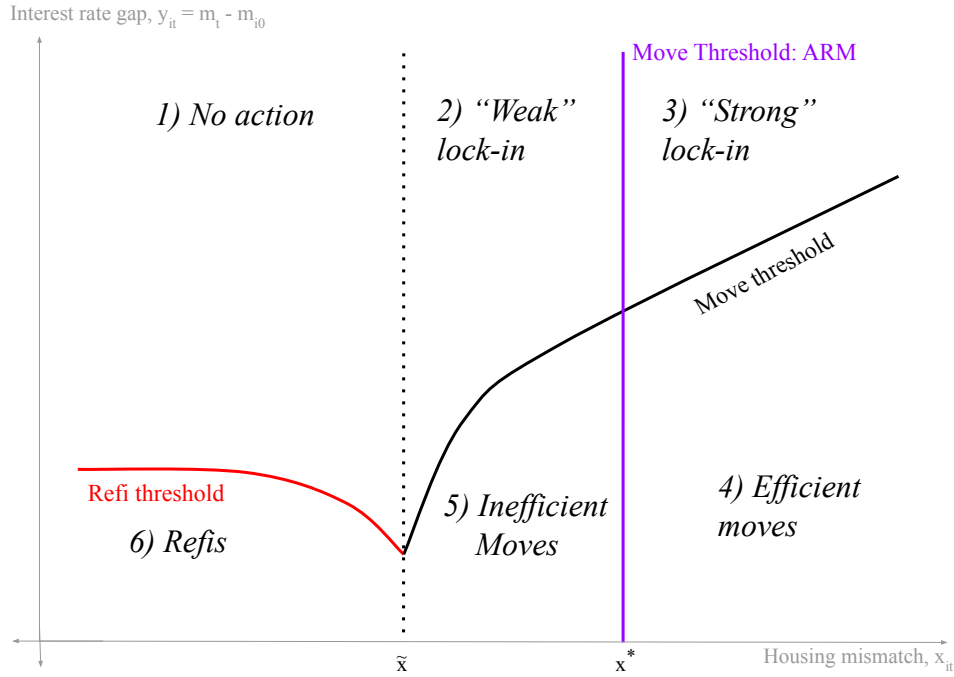


Figure 1: A household (that is paying attention) with a FRM moves if its state, (x, y) , is below the black curve and refinances if its state is below the red curve. A household with an ARM moves if its state is to the right of the purple line. The third state variable, inflation (π) is not shown as it does not qualitatively impact the results shown here.

“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.²¹

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

²¹In other words, if there were no transaction costs, the homeowner and lender could arrange for a side payment from the lender to the homeowner such that both would be better off if the homeowner moved. Another way of understanding the inefficiency is that the homeowner has sole authority over the decision to move but her decision affects the lender through the change in interest rates. Therefore, it is as if there is a positive externality from moving when $y_{it} > 0$, so moves are inefficiently low as borrowers ignore this social benefit. A final way of making this point is that in this model, the interest rate only serves a distributional role, since the mortgage balance is exogenous. Therefore, a social planner seeking to maximize efficiency would want moves to be based solely on housing fundamentals (x) and not interest rates (y). So if a homeowner chooses not to move because the interest rate is too high, that is inefficient.

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. Such moves are inefficient.²²

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) and Berger *et al.* (2024), 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.

To provide additional intuition for the model’s solution, Figure 2 explores how the model’s parameters impact the refinancing and moving rules.²³ In particular, it uses as a baseline the parameters shown in Table 1, which result from a calibration process described in Section 3, and then shows how varying parameters affects the rules around that baseline. Figure 2a gives the intuitive result that if housing mismatch, x , is more volatile, households should wait for a higher level of mismatch before deciding to move. This is a standard feature of options problems: the option is more valuable when the state variable is more volatile, which pushes against exercising the option.²⁴

²²The logic is simply the reverse of region 3. In region 5, the lender could pay the borrower *not* to move in order to keep the interest rate relatively high, since the borrower is not getting very much improvement in housing utility from the move. Or, with $y_{it} < 0$, we could think of moving as having a negative externality on the lender in the form of a lower payment, which does not concern the borrower and so we see too many moves.

²³The Appendix shows comparative statics additional with respect to additional parameters: ρ , p_A , p_M , and C^M .

²⁴This insight dates to at least Black and Scholes (1973) and Merton (1973).

The remaining panels of the figure demonstrate an important insight: when y is low, moving and refinancing serve as substitutes; and when y is high, they serve as complements. Intuitively, when y is high, the ability to refinance in the future makes homeowners more willing to move and accept a high interest rate, since they know they can subsequently refinance it. When y is low, refinancing immediately becomes a viable choice and so people will choose to do that rather than move unless their mismatch x is sufficiently high. That is why, in Figure 2b, the moving rule is steeper when the volatility of the (real) interest rate is higher. Higher σ_r increases future expected savings from refinancing and so makes moving more palatable at high values of y .²⁵ In contrast, when y is low and the refinancing option is competing with the option of moving, and higher σ_r makes it all the more appealing, lowering the likelihood of moving.²⁶

In the same vein, Figures 2c and 2d show that when refinancing is easier, the moving rule steepens. When refinancing is more costly, the borrower requires a higher value of x in order to move, particularly when y is large. Intuitively, since she is less likely to refinance due to its high cost,²⁷ she is more wary of taking on a high interest rate when moving because she will be stuck with it for a long time. Similarly, higher p_R makes the borrower more willing to move when y is large: because she pays more attention, she knows she will get more savings in the future whenever interest rates fall and she refinances, making her more willing to accept a high interest rate to move today.²⁸ These two comparative statics will be important in Section 5, which analyzes a mortgage design that refinances automatically when $y < 0$ (i.e. $C^R = 0$ and $p_R = 1$).

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

²⁵Another intuition for this is that when y is very high, the refinancing option is not very valuable in the short term because it is far out-of-the-money. Moving resets and y to zero and so infuses the option with value again. When σ_r is large, that option is more valuable, which means that when y is high, moving has greater value when σ_r is large.


²⁶The increase in σ_r causes the refinancing threshold to fall, as in standard analyses of refinancing in which moves are exogenous, as in Agarwal *et al.* (2013).

²⁷In fact, the “High C^R ” parameterization shown in Figure 2c has such high refinancing costs that it can be viewed as a version of the model with no refinancing option, like a FRM with extremely high prepayment penalties.


²⁸As in models with exogenous moves, the increase in p_R causes the refinancing threshold to fall (see Berger *et al.* (2024)) because more attention causes the option to refinance to be more valuable.




(a) Baseline model uses $\sigma_x = 0.00096$; Low σ_x model uses $\sigma_x = 0.00048$.



(b) Baseline model uses $\sigma_r = 0.022$; Low σ_r model uses $\sigma_r = 0.011$.



(c) Baseline model uses $C^R = 0.018$; High C^R model uses $C^R = 0.9$. The refinancing rule in the High C^R case is not shown because at such a high cost, refinancing is effectively removed as an option.



(d) Baseline model uses $p_R = 0.0372$; High p_R model uses $p_R = 0.5$. Additionally, p_A is adjusted downward from 0.96 to 0.14 so that the analysis isolates the impact of paying attention to the refinancing incentive more while not changing how often households consider moving.

Figure 2: Decision rules under different model parameters. The Baseline model uses the parameters shown in Table 1.

the lock-in phenomenon as of March 2025.

3.1 Data

For the macroeconomic variables, the analysis relies on monthly data on mortgage interest rates²⁹ and inflation³⁰ from April 1991 through March 2025 as shown in Figure 3. 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.³¹ 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, and the impact of decades-old developments on these initial conditions is quite minimal.

Home sales activity is measured using the quarterly sales rate as reported in *Batzer et al.* (2024) from 2005Q1-2024Q1.³² Using microdata on mortgages guaranteed by Fannie Mae and Freddie Mac (the government-sponsored enterprises, or GSEs), these authors linked mortgage prepayments to deeds records to identify which are sales. 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 4a.

Public, loan-level origination and performance data from Fannie Mae’s “Single-Family Fixed Rate Mortgage” dataset and Freddie Mac’s “Standard” loan-level mortgage performance dataset are used to construct a time series of refinancing activity.³³ The loan-level data is used to construct a quarterly time series of mortgage prepayments. From there, the quarterly refinancing rate (i.e. share of households that refinance) is calculated as the share that prepaid their mortgage (from the microdata) minus the share that sold their home (from

²⁹In 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.”

³⁰In 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.”

³¹In particular, $\text{corr}(m_t, \pi_t) = 0.35$ and $\text{corr}(m_t, \pi_{t-12}) = 0.57$

³²I 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>.

³³These 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.

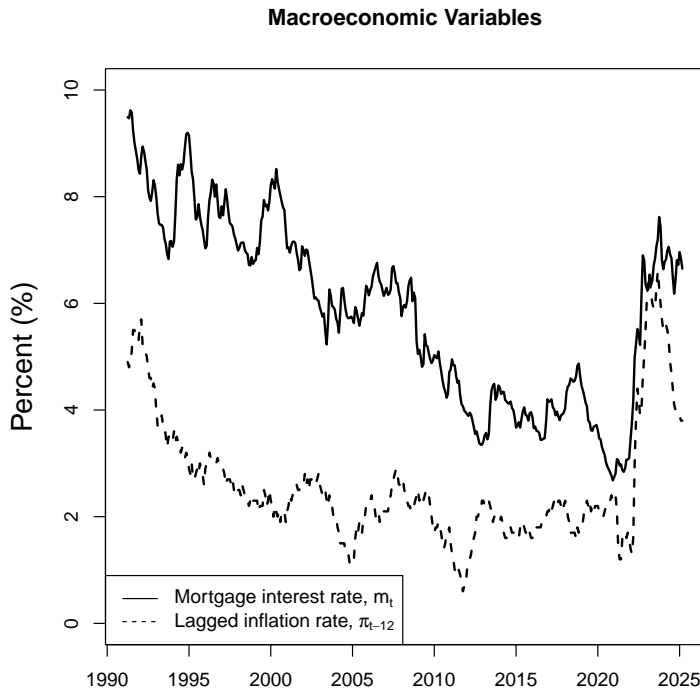


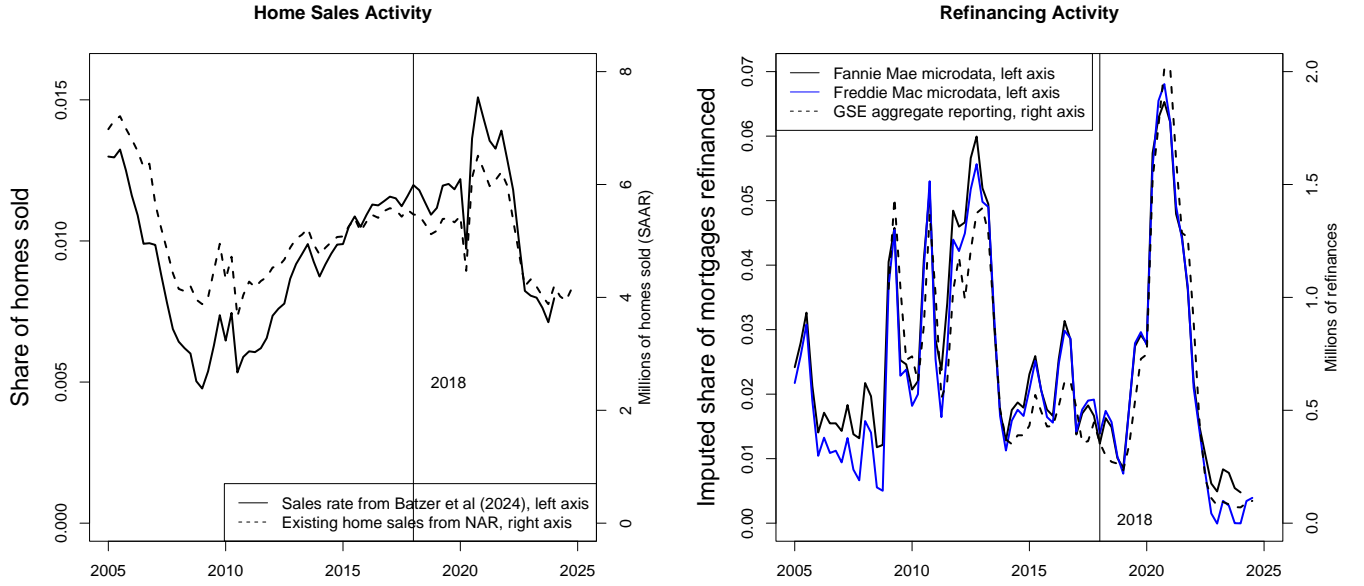
Figure 3: Nominal mortgage interest rate and inflation rate. The inflation rate shown here is lagged by 12 months.

Batzer *et al.* (2024)). This makes the strong assumption that all mortgage prepayments result either from sales or refinances, which is not true. Nonetheless, Figure 4b 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).³⁴

Finally, the American Housing Survey (AHS) is used 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.³⁵

³⁴In 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 2009Q1 through 2019Q1; and Foreclosure Prevention, Refinance, and FPM Reports, which took the place of the Refinance Reports after March 2019. They are all available at <https://www.fhfa.gov/reports?topic=fannie-mae-and-freddie-mac-reports>.

³⁵As 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 in year t . Note that because the survey is only conducted every other year, hazard rates for durations less than two years cannot



(a) Home sales series.

(b) Refinancing series.

Figure 4: Sales and refinancing data. Figure 4a 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 4b 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.”

3.2 Calibrated Model

The model is parameterized at the monthly frequency in order to capture richness in interest rate variation across households.³⁶ 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 March 2025. 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:

$$dm_t = d\pi_t + dr_t. \quad (2)$$

be calculated.

³⁶Results, however, will be shown at the quarterly frequency to smooth out high frequency variation in outcomes.

Parameter	Meaning	Value	Source
ρ	Discount rate (annualized)	0.05	Agarwal <i>et al.</i> (2013)
C^R	Refinancing cost	0.018	Agarwal <i>et al.</i> (2013) (for \$250k mortgage)
C^M	Moving cost	0.09	Fonseca and Liu (2024)
π_L	Low inflation rate (annualized)	0.022	Empirical averages in the US from April 1991 - March 2025, using 4% as the cutoff value between low and high inflation
π_H	High inflation rate (annualized)	0.052	
p_π^L	Probability of remaining in low-inflation state	0.994	
p_π^H	Probability of remaining in high-inflation state	0.941	
σ_r	Real mortgage rate volatility (annualized)	0.022	Matches nominal mortgage rate volatility from April 1991 - March 2025 (0.017)
If x is Random Walk: $x_t = x_{t-1} + \epsilon_t$, $\epsilon_t \sim N(0, \sigma_x^2)$			
σ_x	Housing mismatch volatility	0.00096	Match moments of refinancing and sales activity from Q1 2018 - Q1 2024, as described in body of paper
p_M	Probability of exogenous move	0.00135	
p_R	Probability of being attentive to interest rate	0.0372	
p_A	Probability of considering a move, if attentive to interest rate	0.96	
----- If x is i.i.d.: $x_t = \epsilon_t$, $\epsilon_t \sim N(0, \sigma_x^2)$ -----			
σ_x	Housing mismatch volatility	0.00125	
p_M	Probability of exogenous move	0.00285	
p_R	Probability of being attentive to interest rate	0.038	
p_A	Probability of considering a move, if attentive to interest rate	0.82	

Table 1: Calibrated parameters. The top three panels show parameters that are taken from the literature or calibrated to match moments of macroeconomic data. The bottom panel shows four parameters jointly calibrated to match aggregate sales and refinancing data from 2018Q1-2024Q1. These four parameters are calibrated under two alternative specifications of the model: one in which a household’s housing mismatch follows a Random Walk, meaning shocks are completely persistent; and one in which a household’s housing mismatch has no persistence and so is i.i.d.

The four parameters in the final panel of Table 1 (σ_x , p_m , p_R , and p_A) are calibrated via simulation. To demonstrate the importance of persistence in housing mismatch, x , these four parameters are calibrated once assuming x follows a Random Walk and once assuming it has no persistence – meaning x is independent and identically distributed (i.i.d.) over time – as indicated in the table.

For each quadruple in a grid of those four parameters, the following steps are performed for 100,000 households in the calibration:

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 396 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 3 from April 1991 to March 2024 (396 months).

The parameters shown in Table 1 are the ones that best matched the moments of sales and refinancing activity from 2018Q1-2024Q1 shown in Table 2 under the alternative assumptions of a Random Walk or an i.i.d. process for x . One interesting result to note is that the model requires less inattention to interest rates than in the model of Berger *et al.* (2024), who model refinancing when moves are exogenous. Their calibration finds that the average borrower pays attention to interest rates in 2.6% of months, whereas my model with endogenous moves sets the calibrated value of p_R to 3.7%. Intuitively, the model with endogenous moves is able to justify some refinancing inaction by the possibility that borrowers’ housing mismatch might be close to triggering a move, lessening the degree of inattention that needs to be assumed. Having said that, $p_R = 3.7\%$ still connotes a great deal of inattention, as they found.³⁷

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.³⁸ As shown in Figure 4a, 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

³⁷While 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. Agarwal *et al.* (2013)’s calibration implies a roughly 10% likelihood of moving in a given year, where as this model with endogenous moves only has exogenous moves arrive in roughly 1.5% of years, as most moves occur through the endogenous margin.

³⁸The developments include a 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. 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 - x is a Random Walk	Model - x is i.i.d.
Average sales rate, 2018Q1 - 2020Q1	0.012	0.012	0.011
Average sales rate, 2020Q3 - 2022Q1	0.014	0.013	0.015
Average sales rate, 2022Q2 - 2024Q1	0.009	0.008	0.009
(Average refinance rate, 2020Q3 - 2022Q1)- (Average refinance rate, 2018Q1 - 2020Q1)	0.030	0.044	0.047

Table 2: Moments of data targeted in calibration

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 estimates what the joint distribution of y and x looked like prior to the evaluation period of 2018Q1-2024Q1. I cut off in 2024Q1 because that is when the sales series from [Batzner *et al.* \(2024\)](#) stops.

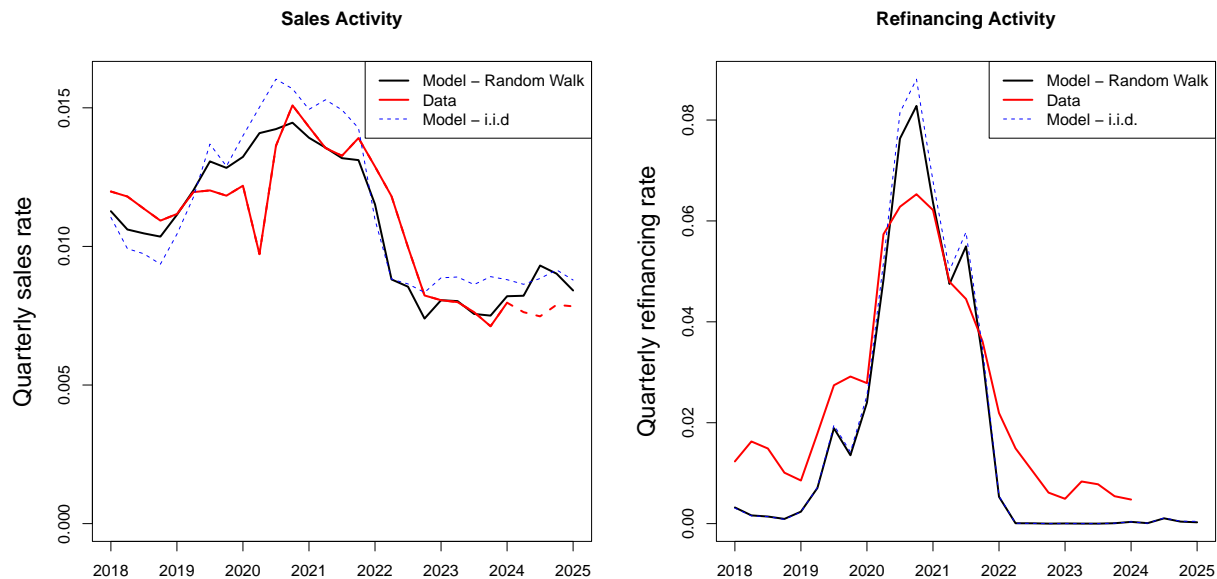
As shown in [Table 2](#), the Random Walk version with persistent shocks is capable of matching the main patterns of home sales activity from 2018Q1-2024Q1.³⁹ In particular, it is able to essentially match the roughly 20% increase in sales that occurred after the immediate onset of the COVID-19 pandemic, when interest rates declined steeply. 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 5a](#) 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.^{40,41} Note that the i.i.d. version also does a fairly good job at matching this data, though it struggles to get recent sales activity to be substantially lower than sales around 2018.

[Figure 5b](#) 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 some

³⁹While 100,000 households were used during the calibration process, I increased that number to 300,000 to perform all analyses with the model after calibration.

⁴⁰It is this time shift that causes the Random Walk model to appear to undershoot the data in row 3 of [Table 2](#). [Figure 5a](#) 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.

⁴¹The 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, 2020Q2 is not used in the calculation of any moments used in the model’s calibration, as shown in [Table 2](#).



(a) Home sales series.

(b) Refinancing series.

Figure 5: Comparison between data and model of sales and refinancing activity from 2018Q1-2025Q1. The [Batzer *et al.* \(2024\)](#) sales data are only available through 2024Q1; this series is extended through 2025Q1 by applying the growth rate of sales as observed in the NAR series. 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 through 2024Q1.

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. 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 behavior, there are also interesting differences. As shown in [Figure 3](#), 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.⁴² Then, as mortgage rates continued to fall during the pandemic, refinancing in the model surges, making up for any prior lag. When 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. Note that here, too, the Random Walk and i.i.d. versions of the model perform comparably.

As another check on the performance of the model, Figure 6 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 between the mortgage rate in June 2024 (6.92%) and the mortgage rate at the time of loan origination. The black and blue curves represent the analogous objects from Random Walk and i.i.d. versions of the model, respectively. Both correspond 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 7 shows that the Random Walk version of the model is successful at replicating patterns in household-level microdata, even though these patterns were not targeted in the calibration. Figure 7a 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 households and the sale hazard rises correspondingly. The relationship reverses eventually 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.⁴³ Critically, the i.i.d. version of the model

⁴²This is the moment of the data alluded to in Section 2, where imposing a lower bound on m_t would likely allow the model to fit better. As discussed there, doing so would come at the expense of significantly increasing computational complexity.

⁴³Ngai and Sheedy (2020) present a model that matches the empirical baseline sale hazard in a different

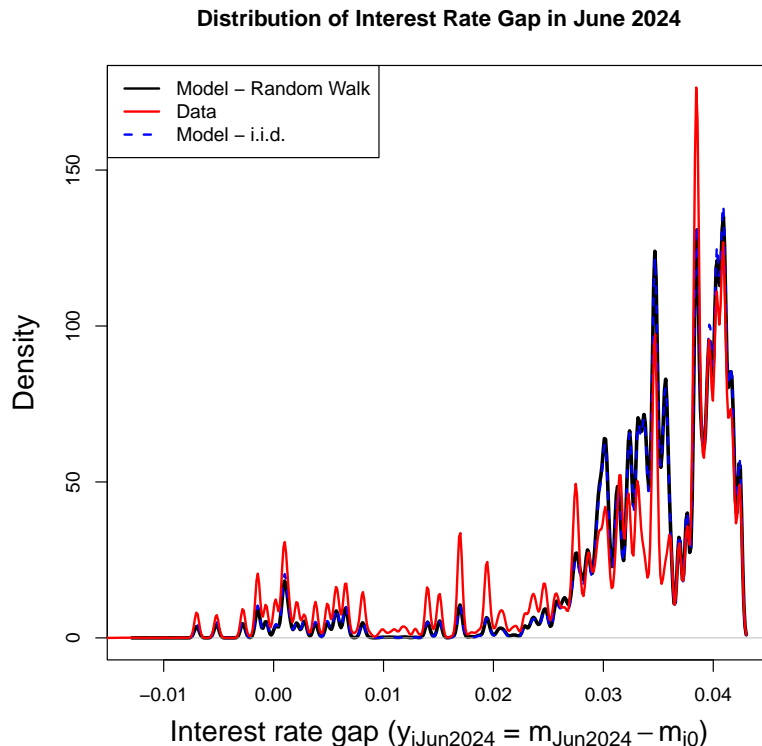
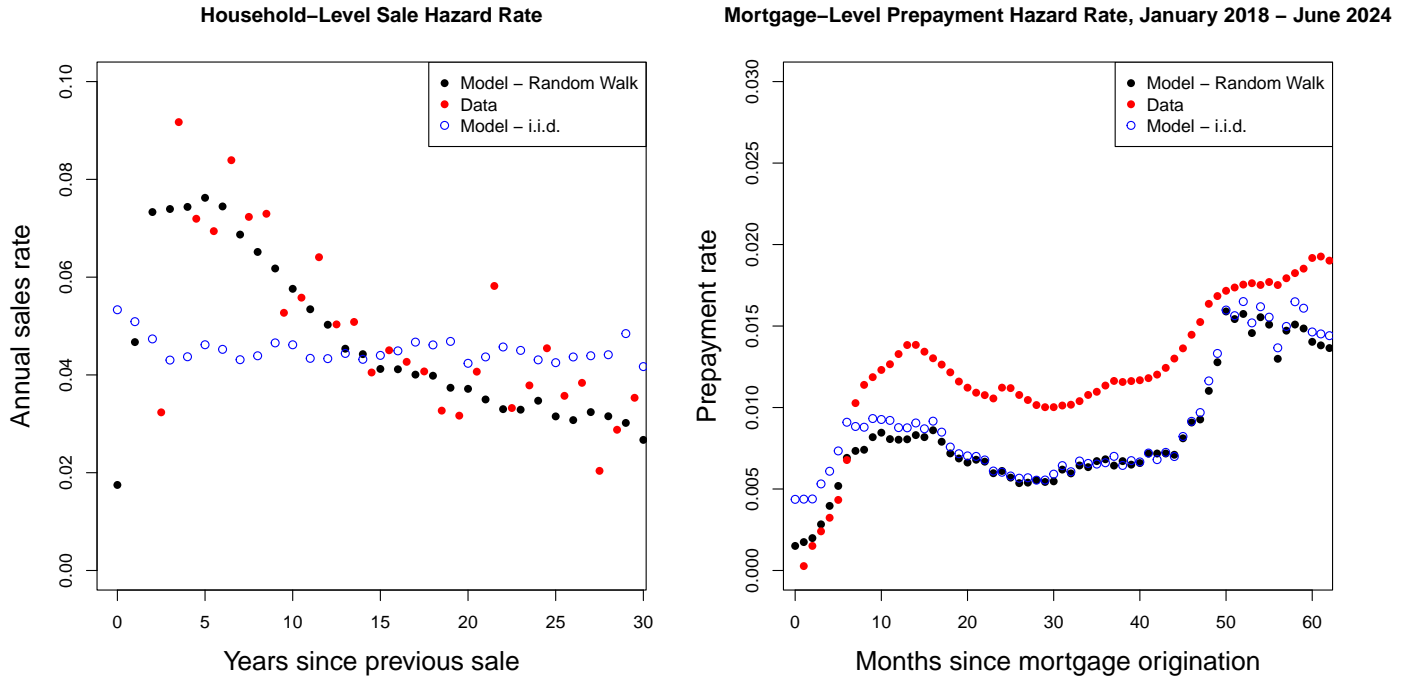


Figure 6: Distribution of $y_{i,Jun2024}$, the interest rate gap in June 2024, from the data and the model. The “Data” series is constructed from Fannie Mae’s loan-level data. Freddie Mac’s loan-level data is not used for this because it reports the quarter of mortgage origination, rather than the month, 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.

is not capable of reproducing this pattern: without persistence in housing mismatch x , the sale hazard is flat. This is the key moment of the data that points to the importance of persistence in x and therefore the Random Walk assumption.

Figure 7b shows the hazard of mortgage prepayment at the monthly level for households

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.



(a) Sale Hazard.

(b) Prepayment hazard.

Figure 7: 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 monthly level but does not distinguish between different types of prepayments, such as sales and refinances.

in the Fannie Mae microdata.⁴⁴ 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 5b.

⁴⁴Recall that the Fannie Mae microdata does not distinguish among sales, refinances, and other types of prepayments. That is why Figure 7b evaluates the model based on the combined outcome and why Figure 7a required the AHS to evaluate the hazard of sale.

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 8 is useful for doing so. The two panels show heatmaps of where in the x - y state space households in the model were at two points in time: February 2020 (prior to COVID; left panel) and March 2025 (most recent data; right panel). Laid on top of this heat map are the moving and refinancing rules from the calibrated (Random Walk) model, analogous to what was discussed for Figure 1 and shown in Figure 2. In February 2020, when mortgage interest rates were near historic lows, only 0.2% households were in the strong lock-in region of the state space (region 3 of Figure 1), deterred from moving despite housing fundamentals that made moving efficient (i.e. $x_{it} > x^*$). In fact, only 0.9% of households were in regions 2 or 3, where a sufficiently-sized decline in the interest rate could have induced a move, efficient or not. In contrast, by March 2025, when interest rates had risen sharply, 5.9% 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 7.0%. In other words, while the interest rate gap disincentivized 0.9% of households from moving (if and when they received the appropriate attention shocks) in February 2020, that number had spiked to 7.0% in March 2025.

Table 3 provides a more detailed breakdown. It shows that in March 2025, 11.2% 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 50% (or 5.6% 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 February 2020. Given that the quarterly moving rate is roughly 1%, the large shift of households upward in Figure 8 and rightward in Table 3 represents the removal of many potential participants from the active housing market, generating the steep decline in home sales beginning in 2022 displayed in Figure 3.⁴⁵

⁴⁵As a back-of-the-envelope calculation, consider that households in the model consider moving in $0.0372 \cdot 0.9675 = 3.6\%$ of months. Given that the share of households with $x_{it} > \tilde{x}$ and $y_{it} < y^*(x_{it})$ fell from 8.5% to 4.2% between February 2020 and March 2025, this implies a reduction in the monthly sales probability of $3.6\% \cdot (8.5\% - 4.2\%) = 0.16\text{pp}$, or 0.46pp at the quarterly level. That is very similar to the 0.5pp decline in sales from before and after the interest rate spike, as shown in Table 2.

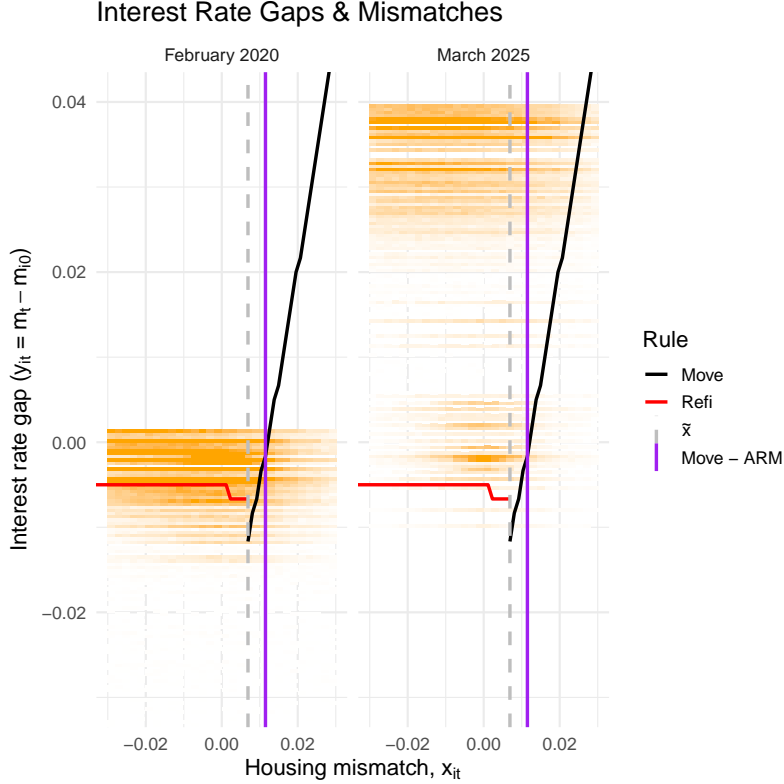


Figure 8: Heatmap of household level values of housing mismatch, x_{it} , and interest rate gap, y_{it} , from model in February 2020 and March 2025. Solid lines show households’ cutoff rules for moving and refinancing, as described in Section 2.3 and Figure 1. The dashed line at $x_{it} = \tilde{x}$ separates FRM households who would refinance if y_{it} were sufficiently low from those who would move.

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	9.4%	8.5%	0.9%	0%	0%	0%	0%
March 2025	11.2%	4.2%	1.5%	1.6%	1.8%	1.7%	0.5%

Table 3: Distribution of decline in mortgage interest rate required to incentivize a move for households in the model. 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

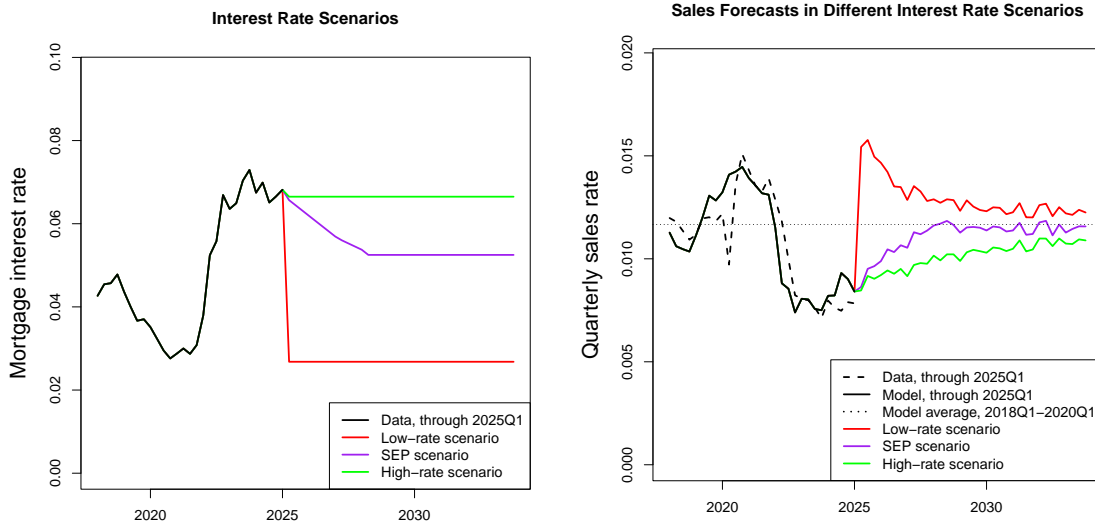
This section first considers the evolution of home sales under different paths of interest rates, i.e. monetary policies. It then considers the implementation of a subsidy (or tax credit) for home sellers and how such a fiscal policy interacts with the interest rate environment created by monetary policy.

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 9a. 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 Projections⁴⁶ and that changes in the Federal Funds Rate translate one-for-one into changes in the mortgage interest rate.

Figure 9b 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. 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. However, 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

⁴⁶Based 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.

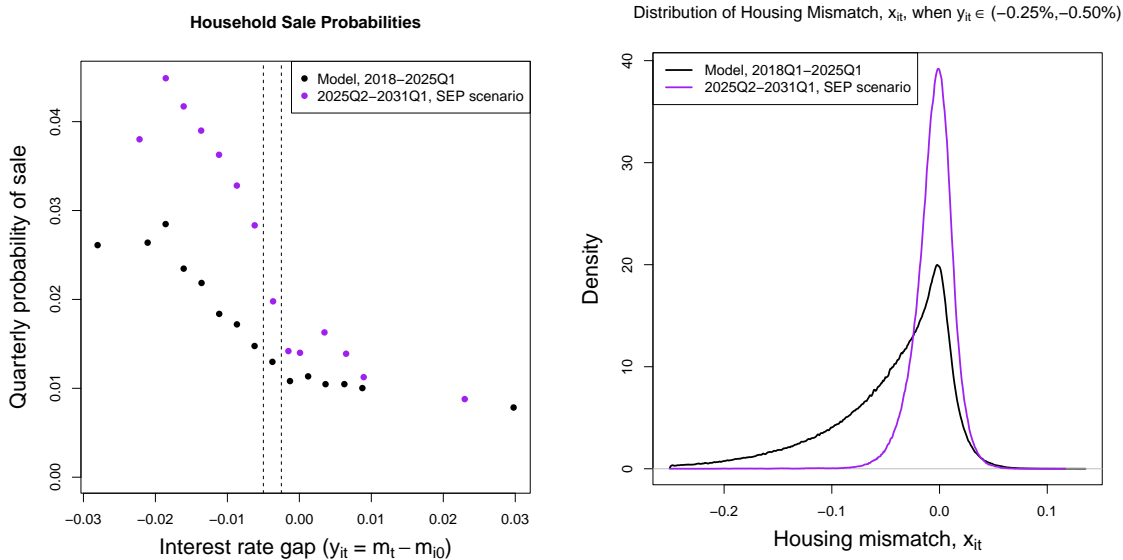
Figure 9: Assumed scenarios of interest rate paths and resulting series of home sales.

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 essentially splits the difference, with minimal overshooting.

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 10 examines it more closely. The black dots in Figure 10a are a binned scatterplot of sales against the interest rate gap, y , from 2018Q1-2025Q1, using data generated by the model. The downward-sloping relationship between y and the probability of sale, which reflects mortgage lock-in, is reminiscent of the empirical work of Fonseca and Liu (2024) and Batzer *et al.* (2024), as is the flat relationship when y is very negative and the refinancing option is a viable option.⁴⁷

The purple dots in the figure represent the same empirical object but evaluated from 2025Q2-2031Q1 under the SEP-based scenario. Importantly, while the general pattern is the same, the probability of sale is notably higher, conditional on y . Figure 10b 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

⁴⁷One inconsistency between the model and the empirical work cited in the main text is the decrease in the magnitude of the gradient as y becomes positive. This points to a shortcoming of the model, which is overly-reliant on “exogenous” moves for high values of y and therefore produces a flatter gradient in this region. Nonetheless, the sale hazard does continue to decline in that region, consistent with the evidence.



(a) Binned scatterplot of sale probability against interest rate gap, y_{it} , in different periods. (b) Distribution of housing mismatch, x_{it} , for households with $y_{it} \in (-0.25\%, -0.50\%)$, during different periods.

Figure 10: Comparison of household selling behavior in different periods. For the earlier period in Figure 10a (in black), the interest rate gap between -2.25% and 1% is partitioned into bins of width 0.25pp . Within each bin, the average level of y_{it} and the sale probability are plotted against each other. The leftmost dot uses all observations with $y_{it} < -2.25\%$ and the rightmost dot uses all observations with $y_{it} > 1\%$. The same is done for the later period (purple), except that its leftmost dot groups together all observations with $y_{it} < -2\%$, since that region is very sparsely populated during that time. The vertical dashed lines isolate the region of y that is analyzed in Figure 10b. Figure 10b 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 10b are estimated using a Gaussian kernel with bandwidth of 0.0002

rightward shift in the later period relative to the earlier one. 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 9b. It also highlights the benefit of using a structural model to understand future market dynamics, since the reduced-form empirical relationship is context-dependent

and liable 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.⁴⁸ The model can accommodate this policy by considering it as a perceived reduction in the cost of moving. In particular, 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 11 shows how the subsidy would affect home sales volume under the different interest rate scenarios described above. Figure 11a 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 11b, 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 11c, the impact is more muted as the degree of lock-in is so high that the tax credit of \$10k does little to stimulate sales, though there is a clear burst of sales before expiration. This again points to the importance of a quantitative structural model, as this impact is determined by the joint distribution of x and y at the time of the subsidy’s implementation; since x is not observed, this must be evaluated by simulating a model.

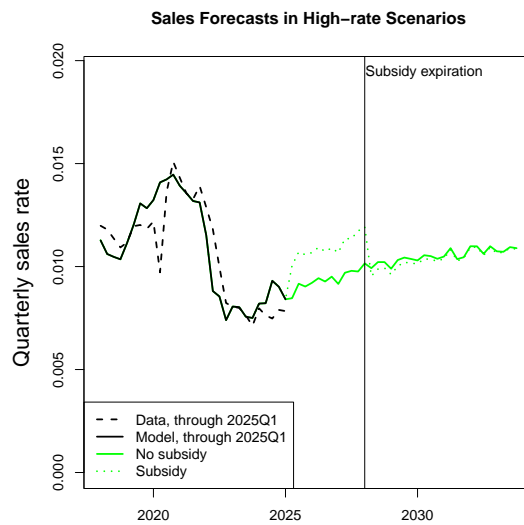
4.2.2 Efficiency and Welfare

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 endogenously move. The first term represents the present value 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

⁴⁸The 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 11: Model forecasts of home sales under different interest rate scenarios, with and without subsidy.

and lenders. C^M is the social cost of the move.

In the low-interest rate scenario, the subsidy lowers efficiency by about 0.008% 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 8 even without the subsidy – so the subsidy simply serves to over-stimulate sales activity as households are shielded from the true cost of moving. In contrast, the subsidy increases efficiency by roughly 0.021% 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, the subsidy would increase efficiency by about 0.017% of the size of the housing market. Given that the US housing stock is worth roughly \$50 trillion,⁴⁹ 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. It is beyond the scope of this paper to model how this expenditure would be financed, but doing so would require increased taxation or reduced government programs at some point in time. Considering the efficiency costs of raising tax revenue and the equity concerns associated with cutting government programs – given that homeowners tend to be higher in the income distribution – this policy would likely strike most as being undesirable. 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 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.

⁴⁹See, e.g., “Total Value of U.S. Homes on the Brink of \$50 Trillion: Redfin,” by Jeff Andrews in *HousingWire* on August 8, 2024.

5 Mortgage Design

A growing literature – including [Campbell \(2013\)](#), [Zhang \(2024\)](#), and [Berger *et al.* \(2024\)](#) – has pointed out that heterogeneous refinancing of FRMs leads to worse outcomes for borrowers with less financial sophistication. [Eberly and Krishnamurthy \(2014\)](#), [Campbell *et al.* \(2021\)](#), and [Guren *et al.* \(2021\)](#) point out that the failure of borrowers to refinance FRMs also has adverse macroprudential consequences. One clear and oft-discussed solution to both of these problems is a mortgage that refinances automatically and costlessly when interest rates fall – what we could call a “Self-Refinancing Mortgage,” or SRM. Advocates of mortgages like this have made clear the equity, efficiency, and macroprudential benefits of the SRM over the standard FRM when interest rates decline, as they often do during recessions, including the COVID recession. But what about when interest rates come back up? This literature has not considered the possibility of mortgage lock-in, and it seems plausible that in fact the SRM would have *adverse* consequences when interest rates rise. After all, if all borrowers got the record-low interest rates available during COVID, wouldn’t the lock-in problem have been more severe once interest rates rose, with home sales reduced even further?

To explore this, I rerun the model under the counterfactual assumption that borrowers have SRMs rather than FRMs.⁵⁰ Specifically, I keep all of the parameters in [Table 1](#) the same with the following exceptions:

- I set $C^R = 0$ to capture that refinancing a SRM is costless;
- I set $p_R = 100\%$, allowing the mortgage to be refinanced in any month in which $y_{it} < 0$;
- I set $p_A = 0.0372 \cdot 0.96 = 0.036$, so that $p_r \cdot p_A$ is unchanged in the counterfactual, meaning households consider moving as frequently as they did in the FRM baseline.

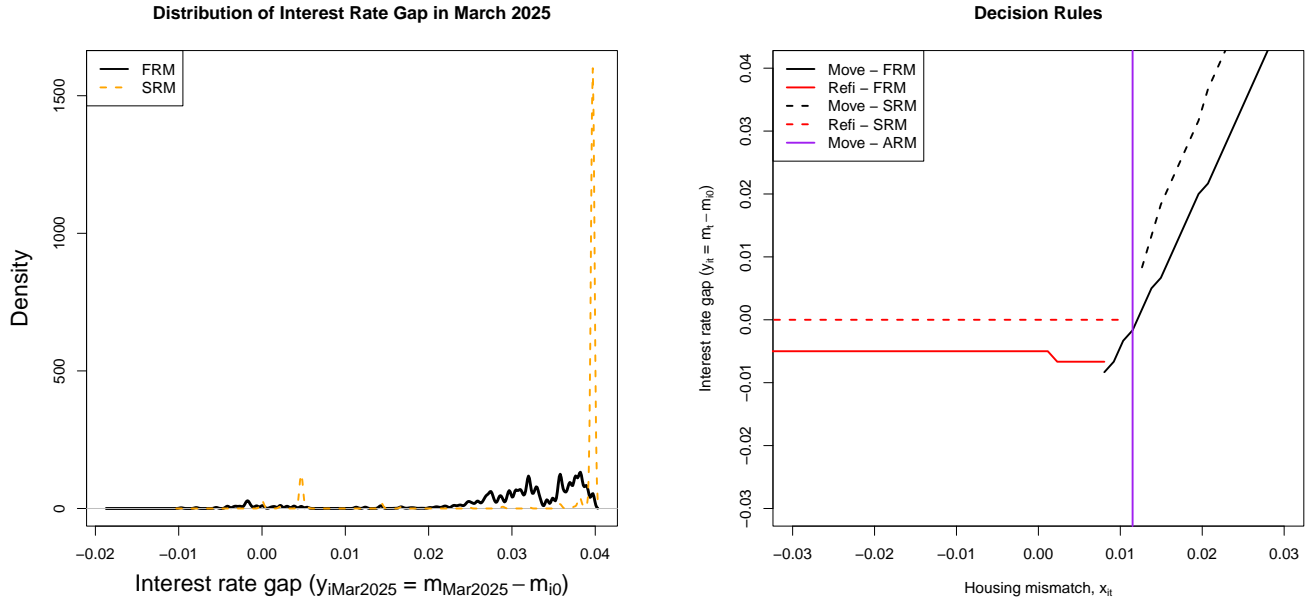
[Figure 12a](#) confirms the intuition from above: under the SRM counterfactual, in March 2025 the large majority of households would have very high interest rate gaps of roughly

⁵⁰SRMs would have higher interest rates than FRMs due to their higher prepayment risk, as pointed out by a number of papers. In this model, which has risk-neutral borrowers and no payment-to-income constraints, the level of the interest rate is immaterial – all that matters is how much the market interest rate has changed since the last purchase or refinance. The counterfactual SRM simulation assumes that movements in the SRM interest rate would have matched observed movements in the FRM interest rate. While this need not be the case, [Abel \(2025\)](#) presents a model where interest rates are determined in equilibrium and shows that, under a number of assumptions maintained in this paper – such as risk-neutrality, and interest-only mortgages – FRMs and SRMs would each have interest rates that are a constant wedge above lenders’ cost of capital. As a result, movements in both interest rates would be equal to movements in lenders’ cost of capital and so would be equal to each other.

4pp. This is because all borrowers would have automatically refinanced into the record-low interest rates in December 2020 (2.68%), and interest rates in March 2025 were 4pp higher (6.83%).⁵¹ In contrast, the large bulk of homeowners have interest rate gaps strictly between 3pp and 4pp with the FRM, with a sizable portion below 3pp, too.

However, Figure 12b shows there is a countervailing effect, which is that the black moving threshold pivots up/left for the SRM borrower: for a given level of mismatch x in the relevant region, households are willing to tolerate a higher increase in their interest rate y in order to move when they borrow with an SRM as opposed to a FRM. Intuitively, while the borrower does not like accepting a higher interest rate in either case, it is more of an impediment to moving for the FRM borrower who knows that she is likely to be stuck with that interest rate for awhile, given that she only refinances once interest rates have fallen sufficiently and she happens to be paying attention, and doing so is costly. The SRM borrower can move and accept the high interest rate with more comfort, knowing that she will start getting savings costlessly as soon as interest rates fall again. As discussed in Section 2, because moving and (future) refinancing are complements at high levels of y , a mortgage like the SRM that streamlines refinancing makes moves more likely when y is high. So while the SRM would have led to worse interest rate gaps in recent years, it would increase the likelihood of sale *conditional on an interest rate gap*. As a result of these two opposing forces, it is a quantitative question of how the SRM would affect sales volume.

⁵¹The only people with smaller interest rate gaps would be those who have bought homes since December 2020.



(a) Distribution of interest rate gap, y_{it} , in March 2025, evaluated with the baseline FRM model and the counterfactual SRM model. The densities are estimated using a Gaussian kernel with bandwidth of 0.0002.

(b) Solid lines show households' cutoff rules for moving and refinancing for the FRM and ARM, as described in Section 2.3 and Figure 1. Dashed lines show the analogous rules for the SRM.

Figure 12: Decision rules and interest rate gaps in March 2025 under the FRM and counterfactual SRM scenarios.

To investigate this, Figure 13 shows the paths of home sales under the three interest rate scenarios discussed in Section 4, under the FRM and SRM. Focusing first on the results up through the present (2025Q1), a few interesting results appear. The increase in sales starting in 2020 is far less pronounced with the SRM. As discussed in Section 2, low interest rates encourage FRM borrowers on the margin of moving to pull the trigger and do so, so they can lock in that low rate. This is better than refinancing to get the rate, since a move is likely going to happen soon anyway, and so moving right away allows the household to avoid duplicating the effort. This incentive is a lot weaker with the SRM because the household gets the low rate automatically and so need not worry about duplicating effort or missing the good interest rate.

More interestingly, after the rise in interest rates in 2022, sales would actually have been higher with the SRM than the FRM. As suggested above, this surprising result occurs because, even though the interest rate gaps would have been higher with the SRM, borrowers would be more willing to move with a high interest rate gap because the SRM would strengthen their ability to get savings subsequently. Indeed, in the relevant region of the

state space, SRM borrowers are willing to accept a higher interest rate increase of roughly 1pp compared to FRM borrowers (see Figure 12b). Because the interest rate gaps for FRM borrowers are mostly less than 1pp below those of SRM borrowers (see Figure 12a), that increased willingness to move translates into more sales in the SRM scenario.⁵²

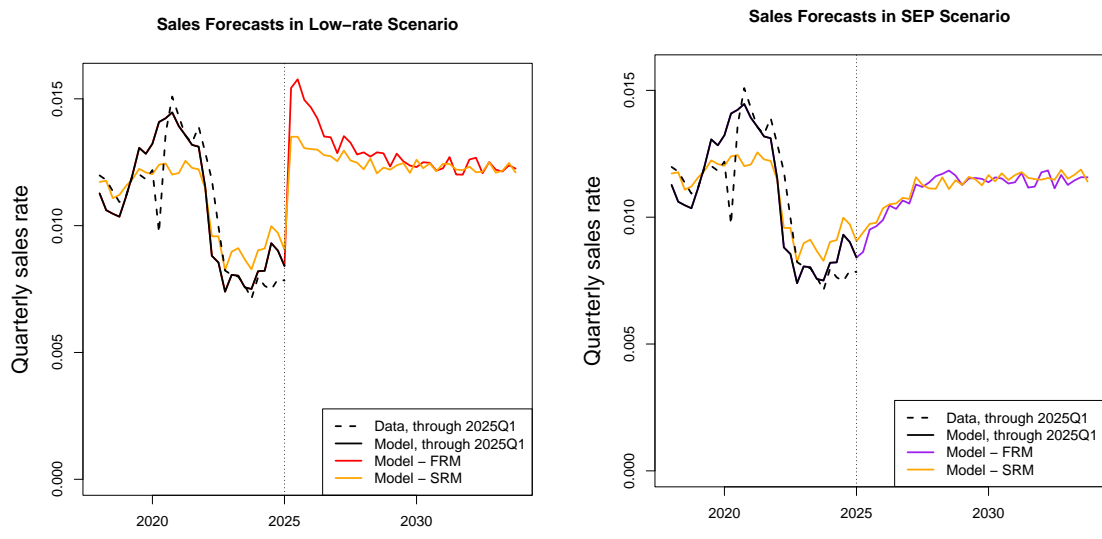
Turning to the forecasts, Figure 13a shows that if interest rates fall sharply, sales in the SRM counterfactual would increase but not as much as in the FRM baseline. Intuitively, less housing mismatch would be pent-up in the SRM scenario because sales did not fall as far from 2022-2024; also, as discussed above, SRM borrowers are not as enticed into moving by low interest rates as FRM borrowers. Figure 13c shows that if interest rates remain high, both series would recover slowly, with SRM sales remaining a bit higher for the reasons discussed above. In the SEP scenario shown in Figure 13b, where interest rates fall modestly over the next few years, FRM sales would catch up with SRM sales, with perhaps a little overshooting.

More generally, this analysis shows that the SRM – in addition to the other benefits already identified in the literature – can help to moderate home sales volume volatility stemming from movements in interest rates. Because the moving threshold steepens, as shown in Figure 12b, the SRM decreases the relevance of the interest rate gap in households’ moving decisions, inducing them to act more like ARM borrowers, whose moving decisions are based entirely on housing fundamentals. Because SRM borrowers would be less tempted to move by low interest rates and less deterred from moving by high interest rates – relative to FRM borrowers – home sales volatility would be reduced.


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. This paper is unique in focusing on how the options to move and refinance impact each other in a dynamic context, particularly when the incentive to move is persistent, which

⁵²This analysis is certainly not dispositive in showing that sales since 2022 would have been higher with an FRM. For instance, the model assumes that households’ attention parameters are time-invariant, and that parameter comes from empirical studies focused on periods of falling interest rates. In reality, it may be that after interest rates rise, households will increase their attention. Intuitively, it may be that there are people who are buying with FRMs with the intention of paying close attention to interest rates and refinancing aggressively when they fall, much as in the SRM intuition discussed in the main text. However, while it is certainly true that the evolutions of sales under these mortgage designs depends on details like that, the model shows that this benefit of the SRM is a real one, and potentially quite large.



(a) Forecasts in low-interest rate scenario. (b) Forecasts in SEP scenario.



(c) Forecasts in high-interest rate scenario.

Figure 13: Model forecasts of home sales under different interest rate scenarios, with FRM and SRM.

allows for a novel analysis of the dynamics of home sales under different policy environments. The main policy result is that the oft-discussed policy remedy of subsidizing homesellers is probably undesirable: even though it is likely to increase economic efficiency (though this depends on how monetary policy evolves), there will be so many payments to people who would have moved anyway that 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.

The paper makes a novel and perhaps surprising contribution to the discussion of mortgage design. While there is broad consensus that a SRM would have numerous social benefits when interest rates fall, there has been no consideration of what its effects would be once interest rates rise again. Mortgage lock-in raises the specter that the SRM could in fact be harmful in that context, since more borrowers would be locked into low interest rates obtained when the market rate was down. The model in this paper shows, however, that while the distribution of interest rate gaps would have been “worse” entering 2022 if borrowers had SRMs instead of FRMs, people are substantially more likely to move, conditional on an interest rate gap, with an SRM. As a result, the SRM could have served to moderate the home sales decline, yet another – and perhaps unexpected – benefit of it relative to the FRM.

A clear drawback of the analysis of this paper is the absence of any discussion of 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 make simplifying assumptions that reduce their ability to study the dynamics of home sales. This paper has attempted to fill that gap.

Appendix

Calibration Procedure

This appendix provides additional detail on how the calibration of the parameters σ_x , p_R , p_A , and p_M was performed. As mentioned in Section 3.2, this was done via grid search over quadruples of those parameters. The process of running simulations and how that output was evaluated was discussed in the main text, so here I discuss the grid that was searched. I used an iterative process in which a relatively large portion of the parameter space was searched in a coarse manner, and then a smaller subsection around the best-performing quadruple was searched more finely. I did three iterations.

Parameter values for the first iteration were chosen in the following way:

- $p_A \in [0, 1]$ so I searched from 0.2 to 1.0, incrementing by 0.1;
- $p_R \in [0, 1]$, but Berger *et al.* (2024) found an average of 0.026, so I searched from 0.01 to 0.06, incrementing by 0.01;
- $p_M \in [0, 1]$, but the fact that the quarterly sales rate from 2022Q2 - 2024Q1 was 0.9% suggested a maximum value of 0.003, so I searched from 0 to 0.003, incrementing by 0.001;
- σ_x has no theoretical bound, and I began by searching from 0.002 to 0.02, measured at the annual level, incrementing by 0.002.

For each of the 2,160 quadruples available for those parameter values, I solved Equation 1 and then simulated behavior as described in Section 3.2. The combination that best matched the moments in Table 2 for the Random Walk version of the model was ($\sigma_x = 0.01$, $p_R = 0.04$, $p_A = 0.9$, $p_M = 0.002$). As all of these were on the interior intervals chosen, I did not expand the region explored and instead search over a finer region around that quadruple. In particular, the second iteration searched over quadruples with the following parameter values:

- For p_A , I searched from 0.8 to 1.0, incrementing by 0.025;
- For p_R , I searched from 0.03 to 0.05, incrementing by 0.004;
- For p_M , I searched from 0.001 to 0.003, incrementing by 0.002/3;
- For σ_x , I searched from 0.008 to 0.012, incrementing by 0.004/9.

The combination that best matched the moments in Table 2 was ($\sigma_x = 0.012$, $p_R = 0.038$, $p_A = 1$, $p_M = 0.0017$). For the final iteration, I searched over quadruples with the following parameter values:

- For p_A , I searched from 0.9 to 1.0, incrementing by 0.0125;
- For p_R , I searched from 0.034 to 0.042, incrementing by 0.0016;
- For p_M , I searched from 0.00125 to 0.002, incrementing by 0.00025;
- For σ_x , I searched from 0.01 to 0.0125, incrementing by 0.0005.

The combination that best matched the moments in Table 2 was $(\sigma_x = 0.0115, p_R = 0.0372, p_A = 0.9625, p_M = 0.00175)$. (A similar process was used to find the parameters when assuming x is i.i.d.)

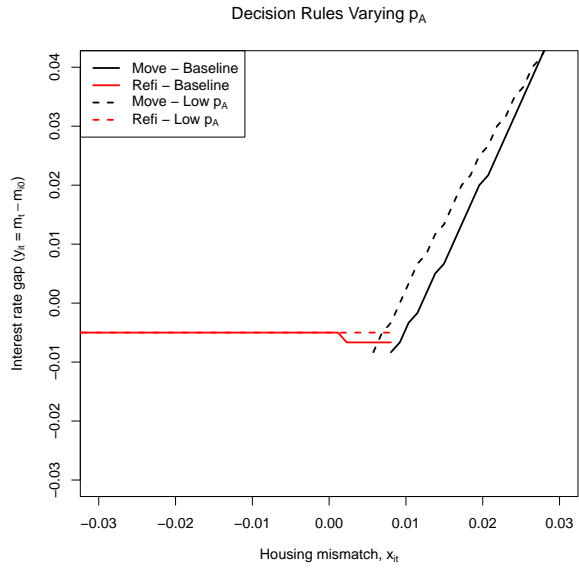

For all analysis, the grid for y was bounded by -0.1 and 0.1. The grid for x was bounded by $-35 \cdot \sigma_x$ and $35 \cdot \sigma_x$.⁵³ During the calibration process, the grids were kept relatively coarse to relieve some of the computational burden, with 10 positive gridpoints, 10 negative gridpoints, and one at zero, with the grid being finer closer to 0 and coarser closer to the bounds. For the calibrated model, I used finer grids for x and y , allowing each to have 60 positive and 60 negative gridpoints. Note that when the quadruple above – $(\sigma_x = 0.0115, p_R = 0.0372, p_A = 0.9625, p_M = 0.00175)$ – was simulated using the finer grid, the series of simulated sales was visibly upward-shifted by approximately 0.0004 at the monthly frequency. As a result, I used $p_M = 0.00135$ in the calibrated model, as reported in Table 1, rather than $p_M = 0.00175$ (recall that p_M is the monthly probability of an exogenous moving shock).

Additional Comparative Statics

Figure 14 supplements the analysis in Section 2 showing how parameter values affect households' decision rules.


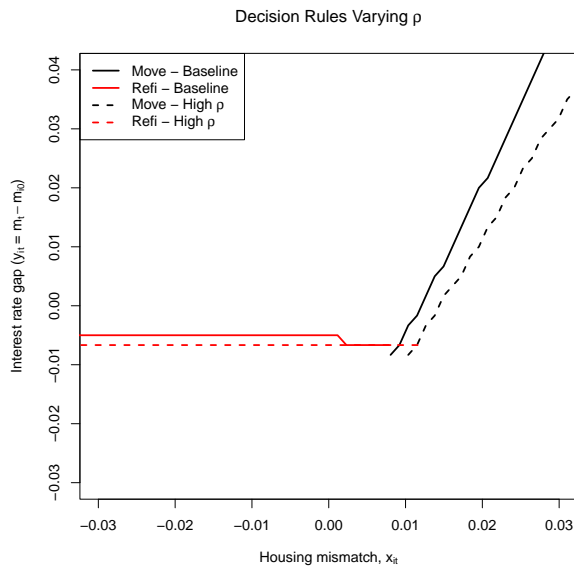
Figure 14a shows the unsurprising result that when moving is costlier, the household will require greater mismatch, x , in order to move. In a similar vein, Figure 14b shows that lower p_A induces the household to move at lower levels of mismatch: if she is less likely to be paying attention later, maintaining the option is less valuable and she is therefore more likely to move immediately. In Figure 14c, we see that with a higher discount rate, ρ , the household requires greater mismatch to justify a move, since moving has upfront costs and delayed benefits. Figure 14d shows that when the risk of an exogenous move is higher, the moving rule steepens. This is an artifact of the difficulty of interpreting exogenous moves in this model. In the relevant region, an increase in p_M causes households to worry that they may receive an exogenous shock that forces them to accept the high y but miss out on their high x , and so they are more eager to claim the high x by moving immediately.

⁵³For the i.i.d. version of the model, the bounds were $-3.5 \cdot \sigma_x$ and $3.5 \cdot \sigma_x$.



(a) Baseline model uses $C^M = 0.09$; High C^M model uses $C^M = 0.18$.

(b) Baseline model uses $p_A = 0.96$; Low p_A model uses $p_A = 0.48$.



(c) Baseline model uses $\rho = 0.05$; High ρ model uses $\rho = 0.1$.

(d) Baseline model uses $p_M = 0.00135$; High p_M model uses $p_M = 0.0054$.

Figure 14: Decision rules under different model parameters. The Baseline model uses the parameters shown in Table 1.

Bibliography

- ABEL, J. (2025). *Equilibrium Mortgage Design wit 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.
- ALADANGADY, A., KRIMMEL, J. and SCHARLEMANN, T. (2024). *Locked in: Rate Hikes, Housing Markets, and Mobility*. Working Paper.
- 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.
- 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.
- BLACK, F. and SCHOLES, M. (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy*, **81** (3), 637–654.
- CAMPBELL, J. Y. (2013). Mortgage market design. *Review of Finance*, **17** (1), 1–33.
- , 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**.
- MERTON, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, **4** (1), 141–183.
- 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.
- 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 Inefficiencies in the Mortgage Market*. Working Paper.