Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Locked In: Rate Hikes, Housing Markets, and Mobility

Aditya Aladangady, Jacob Krimmel, and Tess Scharlemann

2024-088

Please cite this paper as:
Aladangady, Aditya, Jacob Krimmel, and Tess Scharlemann (2024). "Locked In: Rate Hikes, Housing Markets, and Mobility," Finance and Economics Discussion Series 2024-088. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.088.

Locked In:
Rate Hikes, Housing Markets, and Mobility

Aditya Aladangady, Jacob Krimmel, and Tess Scharlemann *

October 2024

Abstract

Rising interest rates in 2022 introduced large moving costs for homeowners with low, fixed-rate mortgages. Using a novel dataset linking mortgage loans, consumer credit profiles, and property sales, we examine the effects of rate hikes on household mobility and the broader economic impacts of the resulting mortgage rate lock-in. As market rates rise relative to those on borrowers' existing loans, likelihood of moving falls with the highest elasticity among borrowers just "in the money." Our results suggest about 44% of the decline in moves among mortgage holders between 2021 and 2022 may be attributable to the widening gap between borrower's existing and market rates. We find limited scope for labor misallocation due to lock-in, as moves across labor market areas are rather unaffected. Instead, lock-in primarily reduces within-metro churn and moves up the housing ladder, leading to fewer real estate listings and greater house price growth. We explain lock-in-driven price increases through a housing search model: in a seller's market, reduced churn raises market tightness, driving up prices. Consistent with such a model, we show measures of market tightness increase in response to lock-in, with the most significant effects in markets that were already tight.

*Aladangady: aditya.aladangady@frb.gov, Krimmel: jake.krimmel@frb.gov, Scharlemann: tess.scharlemann@frb.gov. We thank Elliot Anenberg, Cynthia Doniger, John Krainer, Raven Molloy, Dan Ringo, Brigitte Roth-Tran, Ariel Sun, and participants of the 2023 Consumer Finance Round Robin, FRB Applied Micro Brownbag, Office of Financial Research seminar, 2024 UEA Copenhagen, 2024 UEA North America, 2024 SED Barcelona, and FRBSF/UCLA Conference on Real Estate, Financial Markets and Monetary Policy for their helpful feedback. Latest version available at: https://shorturl.at/gGvVE

1 Introduction

Mortgage interest rates rose by over 300 basis points in 2022, representing the largest increase in recent history (Figure 1, red line). This surge followed a decade of stable, low rates, which fell further during the pandemic, prompting many homeowners to refinance. Since most U.S. mortgages are fixed-rate, long term contracts, the 2022 hikes pushed many homeowners' existing mortgage rate below the market rate or 'out of the money' (Figure 1, blue range), making refinancing or purchasing a new home more expensive for the vast majority of mortgage-holders.

Figure 1: Rates on Outstanding Mortgages versus Market Purchase Origination Rate

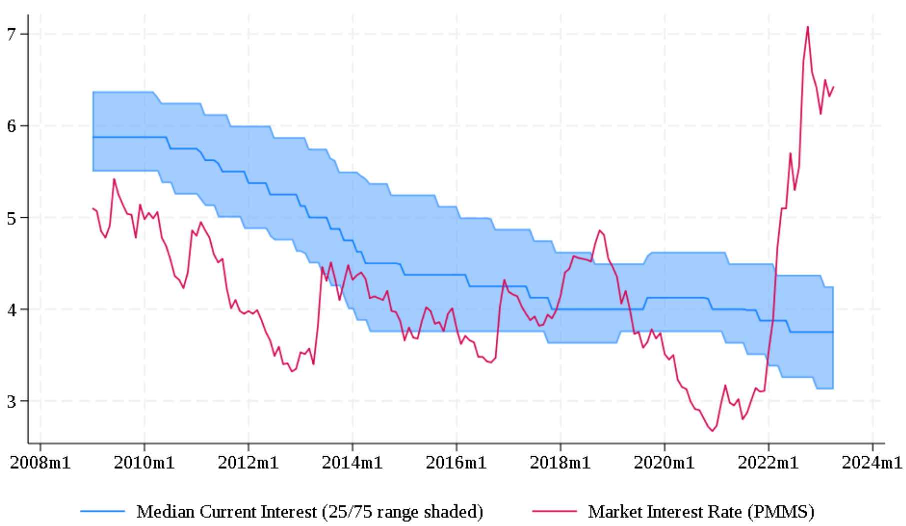

Figure 1 plots median and inter-quartile range of rates on outstanding mortgages (blue) compared with market rate for originating a purchase mortgage (red). Rate increases in 2022 have pushed current mortgage borrowers significantly out of the money for refinancing or moving.
Source: ICE, McDash and Equifax Credit Risk Insight Servicing Data (CRISM), Freddie Mac Primary Mortgage Market Survey (PMMS)

Economists have long noted the structure of long-term, pre-payable, fixed-rate mortgages (FRM) creates incentives for borrowers to base refinancing and mobility decisions (in part) on their "rate gap" – the difference between a borrower's current interest rates and market rates (Quigley, 1987, 2002; Berger et al., 2021). The rate gap provides a measure of a how much borrowers can save on financing costs when they refinance or purchase a home with a new mortgage. As market rates rise, borrowers' rate gaps decrease and may even turn negative, increasing the cost of financing a new mortgage of comparable size and discouraging homeowners from moving.

The recent rate cycle has renewed interest in this topic, as few prior episodes have pushed a large mass of borrowers into negative rate gaps. Accordingly, policymakers, market analysts, and the popular press alike have speculated about a 'lock-in effect,' whereby 'out-of-the-money' homeowners would be reluctant to sell their properties and move, thereby losing their low fixed-rate mortgage.[1] Indeed Figure 2 below provides prima facie evidence of lock-in: household mobility (in blue) declines sharply just as the borrower rate gaps (in red) plummet to new lows.

Figure 2: Median Rate Gap and Mobility Rates

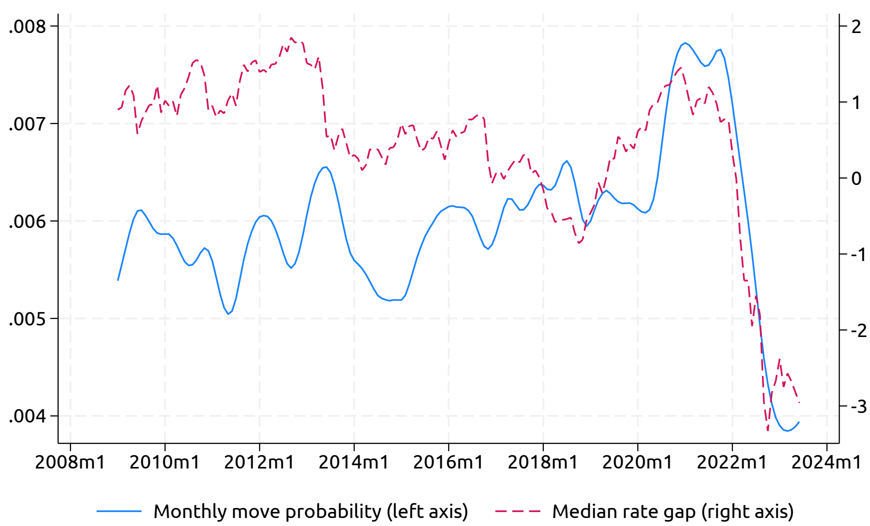

Figure 2 plots median rate gap–the difference between a borrower's current interest rate and prevailing purchase mortgage rates for this type of borrower–along with the seasonally-adjusted monthly probability of a mortgage borrower moving. As rate gaps fell following interest rate increases in 2022, mobility rates also declined sharply.
Source: ICE/McDash and Equifax CRISM, Freddie Mac PMMS

[1]See for example: https://www.cnbc.com/2023/08/01/why-many-homeowners-feel-trapped-by-low-rate-mortgages.html and https://finance.yahoo.com/news/fight-erupts-u-housing-market-212900803.html

The lock-in effect may have broader economic consequences as well. First, such a sharp reduction in mobility could cause a spatial mismatch between labor demand and supply if workers are unable to move to better job matches or highly-productive areas (Quigley, 1987; Hsieh and Moretti, 2019). Second, fewer moves and home sales may impact local housing markets – though lock-in's effect on prices is theoretically ambiguous, as a decline in residential mobility reduces both supply and demand for homes. Assessing lock-in's broader impacts is crucial for understanding how monetary policy passes through to the economy following sharp rate hikes.

This paper examines the link between mobility and interest rates over the most recent rate cycle and explores the potential broader economic consequences of the lock-in effect. We first estimate the causal relationship between rate gaps and mobility–including at low and negative rate gaps–and examine the implications of reduced mobility on housing markets and labor markets. To do so, we construct a novel data set combining mortgage servicing data from Equifax Credit Risk Insight Servicing and ICE/McDash (CRISM) with CoreLogic property deed transfers. The resulting loan-level panel tracks loan terms, property information, and homeowner mobility for 2.5 million mortgages from 2009 to 2023. Importantly, our link to property deeds data – along with household zip codes from credit records – allows us to better distinguish prepayments due to moves from those due to refinances. CRISM also provides detailed loan characteristics, which are crucial for constructing accurate borrower rate gaps over time. The data therefore allow us to compare move probabilities across homeowners for whom prevailing market rates impose different financial incentives.

In addition to controlling for a rich set of borrower and loan characteristics, our empirical approach addresses two key sources of bias in estimating the relationship between rate gaps in mobility. First, as Fonseca and Liu (2023) show, more financially sophisticated borrowers move more frequently but also originate loans at systematically lower interest rates, reducing their rate gaps and thereby biasing the estimated impact of rate gaps on mobility downward. Following Fonseca and Liu (2023) and other recent work on lock-in, we replace the actual origination rate with a market rate at the time of origination to remove potentially endogenous variation borrower-specific origination rates. Second, borrowers planning to move forgo refinancing opportunities when rates drop, leading likely movers to be selected into high rate gaps thereby biasing the estimated impact of rate gaps on mobility upward. To address this selection bias, we introduce a new control that captures foregone refinancing opportunities over the history of the loan, absorbing latent variation in mobility that may lead likely movers to higher rate gaps.

We find the probability of moving declines non-linearly as borrower rate gaps fall. Consistent with prior literature, mobility rises over slightly positive rate gaps but flattens at higher gaps (Berger et al., 2021; Fonseca and Liu, 2023). Notably, our results utilize data through 2024, allowing us to identify the mobility response even in the negative rate gap range. For such 'out-of-the-money' borrowers, we find a flatter relationship between rate gaps and mobility. While the shape of our mobility results aligns with recent studies on lock-in, we estimate a smaller effect size. This difference arises from the inclusion of the new control for foregone refinancing opportunities at high rate gaps, which addresses the positive selection bias. After correcting for this positive selection bias, we estimate that changes in rate gaps explain about 44% of the drop in mobility in 2022.

Given the impacts on residential mobility, we then explore the broader economic implications of lock-in. One hypothesized concern is that lock-in may reduce workers' ability to move to productive jobs. In practice, however, we find limited scope for labor misallocation due to rate-driven lock-in. Specifically, rate gaps mainly influence short-distance moves within the same metro area or within a 30-mile radius. These moves continue to provide access to the same labor market areas, and survey data suggests such short-distance moves are primarily motivated by consumption-related reasons such as changes in household structure or climbing the housing ladder.[2] Conversely, long-distance moves, typically linked to job

[2]In addition to survey-based evidence, using a sub-sample of our data linking movers to their subsequent mortgage, we find that the average increase in home value following a within-metro move is relatively large (around 20%), suggesting households utilize savings from lower rates to purchase more expensive housing. Meanwhile, long-distance moves appear less motivated purely by households moving up the housing ladder,

changes in survey data, show no significant response to rate gaps, even when borrowers face negative mortgage rate gaps. In fact, despite fluctuations in rate gaps and interest rates in recent years, the overall rate of long-distance moves has fallen only modestly, with aggregate changes in residential mobility driven mainly by short-distance moves.

Taken together, our results show lock-in likely has little impact on job-related moves. Instead, it primarily reduces within-metro housing churn–and moves up the housing ladder, in particular. Our findings imply a shift in the composition of move types over time and across different rate gaps. Namely, as interest rates rise and rate gaps fall, non-discretionary moves become more prevalent. This shift helps explain why the rate gap-mobility curve flattens over the negative range: rate gaps are inframarginal for people moving for non-discretionary reasons, such as job changes.

We then ask how local housing markets respond to the sudden drop in mobility induced by rate hikes in 2022. We construct a Core Based Statistical Area (CBSA) level measure of ex ante exposure to lock-in. Specifically, we compute how the 2022 rate hikes shifted the distribution of rate gaps for mortgages in each CBSA. Using our baseline results, we then estimate the implied change in mobility due to lock-in. Areas more exposed to this interest rate-driven decline in mobility experienced larger drops in home listings and higher house price growth. A one standard deviation increase in our exposure measure was associated with roughly 2% greater price growth over the next year, with a larger impact in bigger metro areas.

Lower churn, due to the drop in within-CBSA moves, has ambiguous theoretical effects on local prices since both the supply of houses for sale and the demand for existing homes decline. We develop a housing search model to understand the mechanisms behind lock-in's effect on house prices. Our model shows that the price impact of reduced churn depends heavily on market tightness. If buyers outnumber sellers prior to the shock, a decrease in within-market moves increases market tightness, leading to higher prices.[3] Consistent with the model, our

and the change in the house value is smaller - about a 10% increase on average.

[3]As a concrete example, in a market with 1000 buyers and 800 sellers, a decline in within-market moves

data show markets more exposed to lock-in become hotter along various dimensions: days on market decrease, list prices rise, and price cuts are less frequent. Additionally, the model predicts that ex-ante tighter markets (those with below-median days on market) should show stronger price responses due to reduced churn, a finding also confirmed by our data.

This paper is closely related to a growing literature on the link between mortgages and mobility. Early work by Quigley (1987, 2002) documented the relationship between rate differentials and mobility using survey data during rate hikes in the early 1980s and the 1990s.[5] A wave of more recent research builds on this older literature using various administrative data sources over the post-pandemic interest rate cycle. Fonseca and Liu (2023) utilizes credit bureau data to estimate a strong, positive relationship between rate gaps and mobility, which flattens at high rate gaps, consistent with a model incorporating simultaneous mobility and refinancing motives. In related work, Batzer et al. (2024) and Liebersohn and Rothstein (2023) find similar relationships between mobility and rate gaps, with results suggesting lock-in may be an important driver of depressed mobility since 2022.

Our work contributes to the existing literature in several ways. First, we build on on existing work to improve the measurement of the rate gap mobility relationship. In particular, we show that the estimated mobility response may be biased upwards, as borrowers who anticipate moving choose to forgo refinancing opportunities, leaving them to be differentially selected at higher rate gaps.[6] To address this concern, we introduce a new control–the maximum rate gap experienced through 3 months ago–which accounts for the history of refinancing opportunities borrowers have forgone. While this flattens the elasticity of moves,

of 400 raises the buyer-to-seller ratio from 1000/800 = 1.25 to 600/400 = 1.5. Higher tightness raises the value of search for sellers and lowers it for buyers, raising prices in equilibrium.[4] Empirically, this effect may be amplified by the fact that the marginal mover does so for non-discretionary reasons and is likely also less price-sensitive.

[5]Several papers (Ferreira et al. (2010), Schulhofer-Wohl (2011), Ferreira et al. (2011)) study whether negative equity induced by falling house prices in the 2008 recession also hampered mobility. Our work focuses instead on the impact of interest rate differentials, which are more prominent in the current environment.

[6]This bias is distinct from biases associated with variation in origination rates driven by cross-sectional heterogeneity or market timing, which Fonseca and Liu (2023) tackle by introducing alternate rate gaps and origination cohort dummies. In the appendix, we show how biases arising from endogenous origination rates or timing, as well as those arising from selective attrition influence estimates.

rate gaps still significantly impact mobility in a non-linear manner, similar to the shape found in prior studies.

In addition to improving measurement of the rate gap-mobility relationship, our work contributes key insights into the mechanisms through which mortgage rate lock-in affects the economy at large. First, our paper provides concrete evidence that lock-in likely did not cause significant labor misallocation. Exploiting the richness and statistical power of our data, we explore the heterogeneous effects of rate gaps on various types of moves and show lock-in predominantly impacts churn within a labor market area. Lock-in has minimal effects on cross-labor-market moves, leaving households with access to similar jobs as they would in the absence of lock-in. This, along with survey-based evidence linking move motivations to move distances, provides the first robust finding in the literature against significant labor market disruptions due to rate-based lock-in.[7]

Second, while several papers in this literature document a link between lock-in and house prices, we are the first to provide direct evidence on the mechanisms driving this relationship. Specifically, the 2022 rate hikes occurred during a period of already-tight local housing markets, due to demographic trends and other factors (Anenberg and Ringo, 2021). We show both theoretically and empirically that in such a setting, lock-in reduces churn, which further tightens markets and drives up local home prices. To our knowledge, we are the first to propose and document this mechanism.[8] This result also offers insight on the relationship between interest rates, market tightness, and house prices in other contexts.

Finally, our work builds on the existing literature exploring the channels of monetary policy pass-through via mortgage markets. Berger et al. (2021) show that long-term, fixed-rate mortgage contracts induce path-dependence in monetary policy, as past rates influence

[7]While lock-in may impact moves across areas with different wage-growth differentials–as suggested by Fonseca and Liu (2023)–the potential for meaningful aggregate impacts on on labor mis-allocation appear limited.

[8]Our work is closely related to Mabille et al. (2024) and Gerardi et al. (2024) who study the effects of lock-in on moves up and down the housing ladder. These models study flows in and out of market segments, providing insights into the heterogeneous welfare impacts of lock-in related disruptions to mobility. Our work instead focuses on the implications of reduced mobility on CBSA markets as a whole, highlighting the state-dependent relationship between churn and prices when markets are tight or slack.

the distribution of rate gaps seen at any point in time. This induces asymmetry between rate cuts and rate hikes. While they focus on refinancing behavior, our work shows that similar arguments apply to mobility, as well. The recent episode also highlights a notable asymmetry in the channels of monetary policy between rate cuts and rate hikes. Our work sheds light on some of these channels: rate hikes do not pass through directly to households via financing costs, but may reduce mobility. Notably, rate hikes do not appear to significantly impact labor misallocation, but do hamper moves up the housing ladder. Lower local churn in a tight housing market raises house prices.

2 Data

2.1 Equifax CRISM and Corelogic Deeds Match

Our main analysis identifies mobility effects of lock-in using a novel merge of loan-level data from the Equifax Credit Risk Insight Servicing and ICE/McDash (CRISM, 2024) dataset and property deed records. The data, which run from 2009-2023, allow us to better classify loan terminations into moves and refinances than loan servicing data alone, facilitating our measure of moves under alternative rate counterfactuals. We supplement the loan-level data with real estate market and transactions data from Realtor.com and house price data from CoreLogic.[9]

The CRISM data link anonymized loan-level mortgage servicing information from Intercontinental Exchange-McDash (ICE McDash) with anonymized borrower credit records from Equifax allowing us to track loan and borrower characteristics over time. Importantly, the data allows us to observe the borrower's current interest rate and loan terms along with credit records and servicing information which allow us to create a measure of "rate gaps." The data also provide the property zip code from loan servicing records and the borrower's zip

[9]The Realtor.com data include monthly real estate listings and transactions data at the CBSA and zip code level. We also use CoreLogic repeat sales data and house price indices at the CBSA and zip code level.

code from credit bureau records. Because CRISM only provides a borrower's credit records for 6 months after their loan terminates, we merge on county deeds records for the property to determine whether a loan termination corresponds with a home sale. Information about sales from the deeds data improves our ability to classify loan terminations into refinances or moves. For some analyses, we use the household identifier in the CRISM data to match a subset of mortgagors who prepay to their subsequent loan on their next home.[10] We obtain the sale price of the old home from the deeds transfer data, while the value of the new home comes from the initial appraisal value in the CRISM data.

We describe specific definitions more in depth in the following sub-sections.

2.2 Defining Moves

Merging CRISM's loan and credit records with CoreLogic's property deeds records allows us to construct a comprehensive move definition. The CRISM data provide both the property's zip code from the loan data and the borrower's current zip code their credit record, identifying when a borrower changes location. However, we only observe the borrower's zip code in CRISM for 6 months following a loan termination, which makes it difficult to distinguish some moves from mortgage refinances. Specifically, because credit records may be slow to update addresses, the fact that only 6 months of credit records are observed following a loan termination means many terminations may appear to be refinances when they truly are associated with a move. Additionally, within-zipcode moves will be misclassified as refinances, understating mobility, particularly in short-distance moves.[11]

To address this issue of false-negatives, we augment the CRISM data by matching it to CoreLogic data on county mortgage lien records. In short, if a loan termination coincides with a property sale, the borrower has necessarily moved away. Matching on the loan close

[10]To account for the low count of movers and the low match-rate between prepaid and subsequent loans in our main data set, we draw a much larger sample (30% of loans in the CRISM database) to allow for greater precision.

[11]In a subset of our data where the post-move mortgage also appears in CRISM, we find about 2.4 percent of moves that we identify using property sales are within-zipcode moves.

date, zip code, and origination balances allows us to link about half of active loans during the 2009 to 2023 period to properties in the CoreLogic data base. Using the property identifier associated with this match, we link county deed transfers on the property prior to the loan origination date to determine a purchase date and price for the home associated with the loan. If a subsequent arms-length deed transfer is observed on the same property, we use it to determine the date and price of the home sale associated with the loan termination, and flag the termination as a move.[12]

We define a move as: (1) a persistent (at least 6 month) change in the borrower's zip code away from the property zip code; or (2) a sale of property recorded in county deeds records (at least 90 days after the loan origination date) and that coincides with a loan termination. Moves are classified as sales when the loan is paid in full within a 6 month window of the move or a deed transfer. Situations where moves are not associated with a prepayment or deed transfer are classified as non-sale moves and likely reflect situations where the property is either vacant or rented, but the borrower is living away from the property after having lived there for a time.[13]

In our matched sample, about 20 percent of loans have neither terminated nor moved zip codes as of the end of our sample period. Of the remaining loans, about 70 percent of terminations appear to be refinances, ending with no deed transfer or change in zip code. Another 18.25 percent of terminations are moves that we are able to infer from zip code changes in the CRISM data, though the bulk of these also have a deed transfer recorded. Notably, another 11.3 percent of loans terminate with no change of zipcode recorded within the 6 months of credit data that we observe, but appear to be moves based on a property sale coinciding with the prepayment event. Our ability to link CRISM with property sales

[12]We ensure the deed transfer *following* the purchase occurs alongside a loan termination to ensure it is a proper sale. A small fraction of deed transfers appear to be recorded even when a loan continues to be serviced for several months afterwards, likely because these are non-sale transfers or recording errors. In this event, we consider subsequent transfer records, up to 3 records following the purchase. If the deed transfer is recorded more than one year following the loan termination, the loan is considered a refinance or payoff without a move, as the subsequent loan is potentially matched to this sale if the refinanced loan is also in the CRISM data.

[13]Such moves without sale represent $\approx 13\%$ of all moves.

almost doubles the unconditional move rate from an annualized hazard of 4.2 percent to an annualized hazard of 6.4 percent–closely in line with move rates among households with mortgages in the CPS.

We subset our main sample to mortgages with 30-year fixed rate loans where the owner is occupying the home during the first year of the mortgage (owner's zip code matches property zip code within first few months of origination). We apply some sample restrictions to remove outliers that may be due to reporting errors, loan modifications, or severe delinquencies. Specifically, we require loans to have current and origination LTV's below 150 percent and borrowers to have credit scores (Equifax CRISM's Risk Scores) between 500 and 850. We drop a very small number of loans which show a change in the reported interest rate despite being 30 year fixed rate loans, likely reflecting modifications or reporting errors. Finally, we drop loans with origination rates more than 1.5 percentage points above or below our estimated borrower-specific origination rate, which we define in the next section in Equation (2.3).

2.3 Defining Rate Gaps

The CRISM data provide detailed information about the loan itself, including the mortgage rate and terms. To compute the rate gap, for every loan-month observation, we impute a borrower-specific market rate that would be offered to the mortgagor based on new originations to observationally similar borrowers, similar to the approach used by Berger et al. (2021). Observing borrower-specific current and market rates in this manner improves our measure of rate gaps relative to market rates quoted by survey-based measures like Freddie Mac's Primary Mortgage Market Survey (PMMS), which only reflect rates faced by prime borrowers with relatively low LTV ratios.

To predict the rate that would be offered to a mortgagor, we use a sub-sample of new originations from the CRISM data to estimate mortgage rates conditional on borrower and loan characteristics at each origination date. We allow origination rates to vary with a polynomial of LTV and FICO at origination (from ICE, McDash) including their interaction, loan type, occupancy status, all interacted with time dummies to allow for time-variation in pricing that differs across borrower types:

$$rate_{it}^{mkt} = \theta_t X_i + \nu_i$$

We then use time-varying estimates of $\hat{\theta}_t$ from this regression to impute the likely market rates each existing borrower in our main sample would face if moving and originating a new purchase mortgage at each point in time.

With a measure of borrower-specific market rates in hand, we define the gap between the (fixed) origination rate the borrower pays and the time-varying market rate as the "rate gap". We begin by defining the "actual rate gap" for a borrower as:

$$\begin{aligned} RateGap_{i,t}^{actual} &= rate_i^{orig} - \hat{rate}_{i,t}^{mkt} \\ &= rate_i^{orig} - \hat{\theta}_t X_{i,t} \end{aligned} \tag{1}$$

where borrower i has a loan originated at time $o(i)$ at rate $rate_i^{orig}$ and faces a market rate of $rate_{i,t}^{mkt}$ at time t based on their borrower characteristics $X_{i,t}$. We refer to $RateGap_{i,t}^{actual}$ defined here in Equation 1 as the "actual" rate gap, as it is best approximation to the rate gap the borrower faces when they decide to move or refinance.

As we will discuss in the next subsection, we consider alternate versions of the rate gap which may exclude potentially endogenous variation in rate gaps. Most importantly, in our main regressions, we consider an alternate definition which utilizes an imputed origination rate based on borrower characteristics at origination time $\hat{rate}_{i,o(i)}^{mkt}$ in lieu of their actual origination rate $rate_i^{orig}$. Relative to the "actual rate gap," this version removes idiosyncratic variation in borrower i's origination rate (ν_i in Equation 2.3), which may be correlated with mobility decisions, as shown by Fonseca and Liu (2023).[14]

[14]In addition to this alternate rate gap definition using imputed origination and market rates, in the appendix we also consider measures utilizing just the PMMS rate (in lieu of the imputed rates) to ensure our

Before turning to identification of mobility responses, we first explore the distribution of rate gaps, and how they evolve in response to changes in market rates over time. Figure 3 shows changes in the distribution of rate gaps between December 2021, when the prevailing market interest rate was 3.11 percent, and December 2022, when rates were over 6.4 percent on average.

Figure 3: Evolution of the Rate Gap Distribution, Dec. 2021 - Dec. 2022

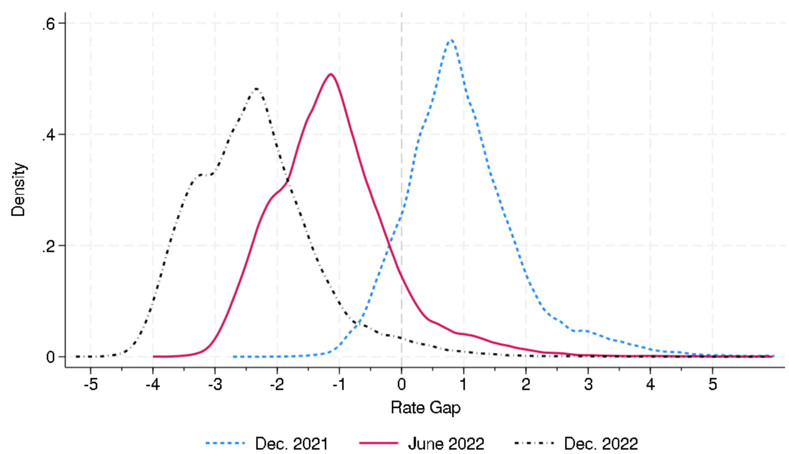

As we see in Figure 3, in late 2021, prior to the rate hike cycle, most borrowers had modestly positive rate gaps (i.e. stood to reduce their mortgage rate by about 1 percent from moving or refinancing), while few faced negative rate gaps (i.e. an increase in rates from moving or refinancing). Because market rates were already so low, refinancing and new home purchase activity was very high. This made the pre-hike rate distribution both unusually low and unusually narrow. As rates rose in 2022, the entire distribution of rate gaps shifted to the left as increasing numbers of borrowers were pushed out-of-the-money and faced negative rate gaps. As we noted in Figure 2, from late 2022 through 2023, the

results are not driven by our imputation procedure. Our main results include controls for LTV and credit score at both origination and observation to control for calendar-time invariant effects of borrower-type on mobility that may be independent of rate gaps. (Credit score at origination is from ICE McDash, and the credit score at observation is measured using the Equifax Risk Score.) Estimates in Appendix Figure A.1 utilize a PMMS-based definition of rate gap $RateGap_{it}^{mkt-mkt} = rate_o^{PMMS}(i) - rate_t^{PMMS}$ that removes all borrower-type variation from rate gaps entirely. These results are statistically indistinguishable from our baseline results using imputed market rate gaps.

fraction of borrowers facing rate increases from a move or refinance was at the highest in decades.

Within our data, we are able to link a subset of home sellers to the mortgage used to purchase the subsequent home. In Appendix Figure C.1 we show that we are able to accurately predict the change in the borrower's mortgage rate using our rate gap measure, and that the accuracy of our prediction does not vary by estimated rate gap.

3 Effects of Rate Gaps on Mobility

Our goal is to recover the relationship between rate gaps and mobility, allowing us to quantify residential mobility rates in counterfactual rate environments. Before turning to our empirical approach, we consider a conceptual discrete choice model of moving. Consider a household with a latent benefit of moving, which includes potential capital gains differences and labor market or amenity improvements, net of moving costs. Meanwhile, the option to refinance provides a financial benefit or cost due to the change in mortgage payments associated with the move. Fixing the mortgage size, the present value of the difference in mortgage payments would be proportional to the loan size and rate gap defined in the prior section. Of course, the household may simultaneously choose to increase or decrease their home or loan size in response to this gap. For example, a household may stand to save on monthly payments due to the change in financing costs, providing an income effect that may, in part, be put towards additional housing financed by a larger loan. Conversely, a household may choose lower-cost housing or a rental if the financial cost of moving is high, but the latent benefit of the move is large, for example, due to a job opportunity. We treat this change in house size and mortgage balance as a choice of the household, potentially in response to the rate gap, and utilize the rate gap as the relevant measure of financial incentives for mobility.

3.1 Empirical Approach

We estimate how responsive mobility is to rate gaps using the following empirical model:

$$move_{it} = \sum \beta_k \mathbf{1}(g_{k-1} < RateGap_{it} < g_k) + \gamma Z_{it} + \varepsilon_{it} \tag{2}$$

The coefficients β_k capture the move probabilities at various ranges of rate gaps, conditional on other observable characteristics Z_{it}. Variation across the β_k coefficients captures the potentially non-linear relationship between mobility and rate gaps where households may respond differently if their rate gap is very high (positive), nearly in the money (near-zero), or severely out of the money (negative). The length of our sample (2009-2023) allows us many different rate regimes over which to draw out non-linearities in the relationship between move probabilities and rate gaps. In particular, the inclusion of the recent rate hike episode allows us to understand how rate gaps affect mobility at very low and negative rate gaps for the vast majority of mortgage holders.[15]

The key identifying assumption to interpret β_k as a structural relationship between rate gaps and mobility behavior is that rate gaps for a borrower at a given point in time are not correlated with other factors influencing mobility decisions via the error term ε_{it}. Our data provides a rich set of loan and borrower characteristics, which allow us to isolate plausibly exogenous variation in rate gaps. Throughout all our specifications, we control for a vector of loan-time factors Z_{it} including flexible controls for borrower age and time since the household purchased the home to capture the time and age varying hazard of moving.[16] Polynomials in current and origination LTV's and credit scores,[17] and flags for purchase mortgage status and loan terms account for potential heterogeneity in loan and credit characteristics that are

[15] For context on how different historical rate increases can have very different effects on the rate gap distribution, consider the following: rate increases in 2013 pushed an 30% of borrowers out of the money over the course of 8 months; the 2022 hikes pushed 80% of borrowers out of the money in 7 months.

[16] Time since home purchase provides a measure of time-varying hazard of moving, and is inferred from the most recent property deed transfer dated prior to the loan origination date. This may have occurred long before the loan origination if this loan is a refinance, and is therefore distinct from the loan origination cohort.

[17] Measured using Ice McDash's origination FICO and Equifax's Risk Score, respectively.

correlated with mobility. We include year, month, and CBSA level fixed effects to account for aggregate variation and seasonality in mobility and differences in mobility across metro areas.

Similar to Fonseca and Liu (2023), we also include monthly origination cohort fixed effects to account for different baseline move probabilities for groups of borrowers who originate at different points in time. Including cohort fixed effects also ensures we compare borrowers who originate in the same month and therefore face similar rate paths and macroeconomic conditions over time.

As discussed, the key identifying assumption is that rate gaps are not correlated with other factors influencing mobility decisions via the error term ε_{it}. Even conditional on the controls and fixed effects introduced thus far, there remain two main biases which work in opposite directions. First, the rate gap-mobility relationship will be biased downward (i.e. slope too flat) if borrowers with greater unobserved move probabilities are able to select into lower rate gaps. As shown by Fonseca and Liu (2023), for instance, borrowers who are more likely to move ex-ante are able to obtain lower origination rates, thereby systematically lowering their rate gaps.[18] Second, the rate gap-mobility relationship will be biased upward (i.e. slope too steep) if borrowers with greater unobserved move probabilities select into higher rate gaps. For example, borrowers who expect to move soon are unlikely to refinance into a lower interest rate, as they would not remain in their home long enough to offset the fixed refinancing costs. As a result, households with higher rate gaps are more likely to be those with a pre-existing intention to move, independent of the rate gap itself.

We first address the downward bias stemming from borrowers with higher move propen-

[18] Borrowers who obtain lower rates given the prevailing rate environment have different move probabilities, either because they paid down points or because they shopped differently. Theoretically, the sign of this bias is ambiguous. For example, Stanton and Wallace (1998) construct a model of mortgage demand where those who plan to stay in a home longer are more likely to purchase points to lower their current rate. Such behavior would result in households with higher move hazards selecting into higher ranges of rate gaps, biasing up the relationship between rate gaps and mobility. While this bias may be present, Fonseca and Liu (2023) show that in practice, instrumenting for the borrower's rate steepens the relationship between rate gap and move probability, as financially sophisticated borrowers may get lower rates and also move more frequently.

sities originating at lower interest rates, all else equal. In the context of of Equation 2.3, the bias is driven by the fact that idiosyncratic variation, ν_i, in origination rates is negatively correlated with mobility, and is present in the "actual" rate gap defined by Equation 1. We therefore follow the approach that has become common in the literature of utilizing the market rate faced by the borrower at origination time instead of the actual origination rate, thereby removing ν_i from the rate gap definition (Fonseca and Liu, 2023; Batzer et al., 2024; Liebersohn and Rothstein, 2023).

We define imputed rate gap as:

$$\begin{aligned} RateGap_{i,t} &= \hat{rate}^{mkt}_{i,o(i)} - \hat{rate}^{mkt}_{i,t} \\ &= \hat{\theta}_{o(i)} X_{i,o(i)} - \hat{\theta}_t X_{i,t} \end{aligned} \tag{3}$$

Under this definition, borrower i originating a loan at time $o(i)$ with characteristics $X_{i,o(i)}$ would have faced a market rate at origination of $\hat{rate}^{mkt}_{i,o(i)} = \hat{\theta}_{o(i)} X_{i,o(i)}$, which may be different from the borrower's actual rate at origination. The imputed rate gap is defined as the difference between this predicted origination rate and the market rate the borrower faces as time t, given by $rate^{mkt}_{i,t}$ as in the earlier definition of the "actual rate gap" defined in Equation 1. However, relative to the "actual rate gap," this version removes idiosyncratic variation in borrower i's origination rate (ν_i in Equation 2.3), which may be correlated with mobility decisions. At time t, this borrower now faces a market rate of $\hat{rate}^{mkt}_{i,t} = \hat{\theta}_t X_{i,t}$ if they were to originate a new purchase mortgage at similar terms to their existing loan, and $RateGap_{i,t}$ reflects the differential financing costs associated with that origination.[19]

We next address the upward bias arising from the fact that borrowers expecting to move in the future would not choose to refinance, resulting in differential selection of movers into higher rate gaps. Importantly, our inclusion of origination cohort fixed effects–a common ap-

[19] As we note above, our main regression includes controls for $X_{i,o(i)}$ and $X_{i,t}$, removing calendar-time-invariant differences in mobility across borrower type. We also include alternate specifications in the appendix that utilizes the PMMS rate at $o(i)$ and t to fully remove variation due to borrower-type from our rate gap definition. As shown in Appendix Figure A.1, results are statistically indistinguishable from our baseline results.

proach in the literature–does not address this bias. In a declining interest rate environment, households planning to stay are more likely to refinance, while those planning to move drift further into the money. This results in differential attrition within a cohort, with refinancers exiting and soon-to-be movers remaining at higher rate gaps. This endogeneity biases our estimated mobility slope upward.[20] Therefore, in addition to loan-origination cohort dummies, we must also include a variable that reflects a time-varying measure of forgone refinancing opportunities.

To address this concern, we introduce a new control variable which tracks the maximum rate gap available to the borrower between origination and three months ago. This variable provides a measure of the best refinancing opportunity the borrower has faced and forgone. By taking the maximum up to three months prior to the observation, we allow the household three months to move in response to a drop in rates (which pushes up their rate gap). After three months, if the household has not moved or refinanced, the control includes the higher rate gap, and only marginal effects from even higher (or lower) rate gaps would influence the move elasticity estimate.[21]

Our main source of variation comes from short-run time-series variation in the prevailing mortgage rate and in the rate available to a specific borrower based on evolution in the pricing environment for their credit characteristics. Our coefficient estimates reflect the average change in move rates as a borrower moves between rate gaps due to the changing rate and credit environment. Even after correcting for the two biases above, identification

[20] As an example, consider a cohort of borrowers who originate in January 2020, prior to pandemic-related rate cuts. Within this cohort, consider a household who knew they would need to move for unrelated reasons in the next 18 months. Many loans from this cohort refinanced by January 2021 as rates had come down significantly, but the household we are considering would forgo this opportunity, and subsequently move in June 2021 with a very high rate gap. The average mobility rate for loans in this cohort would reflect the average move hazard of entire cohort, and not differentially reflect move rates for loans that survived longer even amid attractive refinancing opportunities.

[21] In Appendix Figure A.2, we consider alternate definitions for this variable using a lag of only one month or six months. The six-month lag version is quite close to results from a specification that does not include a maximum rate gap control at all, and suggests a substantially steeper relationship between rate gaps and mobility than our baseline three-month version. As such, a six-month versions may not sufficiently address the bias. A one-month lag reduces the estimated slope further than our baseline, albeit only slightly. However, this version provides households almost no time to actually search for a home or close on a new loan, potentially over-controlling for mobility.

relies on the assumption that changes in market interest rates are largely unforeseen by borrowers. This assumption is consistent with Larsen and Martinez (2024), who show that mortgage rates are highly correlated with forecast errors, even for professional forecasters.

3.2 Results

We first estimate how rate gaps affect homeowners' monthly move probabilities to quantify lock-in's impact on mobility. Later, we use these estimates to assess broader economic impacts of the post-2021 lock-in effect.

Our estimate of Equation (2) reveals a positive but nonlinear relationship between rate gaps and the monthly probability of moving. Figure 4 illustrates this relationship, showing how move probabilities are generally increasing in rate gaps.[22] Importantly, the contour is stable across multiple definitions of rate gaps and specifications, demonstrating the robustness of the finding. Homeowners with very low or negative rate gaps (that is, those holding a mortgage with an interest rate far lower than current market rates) show a reduced propensity to move. This generally suggests an increase in market rates – all else equal – creates a stronger lock-in effect, discouraging homeowners from selling their current home and taking on new mortgages at higher rates. Graphically, as market rates increase, current mortgage holders move to the left, down the move probability slope.

There are, however, two important features to note on the nonlinearities in move probabilities. First, the relationship between rate gaps and mobility flattens at around 2%. This is consistent with results found by Fonseca and Liu (2023) and Berger et al. (2021), who note that many homeowners with such large rate gaps may find it more attractive to refinance rather than move. Second, the slope flattens out for negative rate gaps. For instance, note the relationship between rate gaps and mobility is strongest between rate gaps of 0 and 2, but weaker between 0 and -2.[23]

[22] Appendix Figure A.1 and A.2 a similar shape is revealed using several alternate specifications and alternate definitions for our rate gap history control.

[23] Importantly, this nonlinearity (i.e. the shallower slope for 'out of the money' borrowers is robust across alternate specifications. See Appendix A for details.

Figure 4: Mobility by Rate Gap, 2009-2023

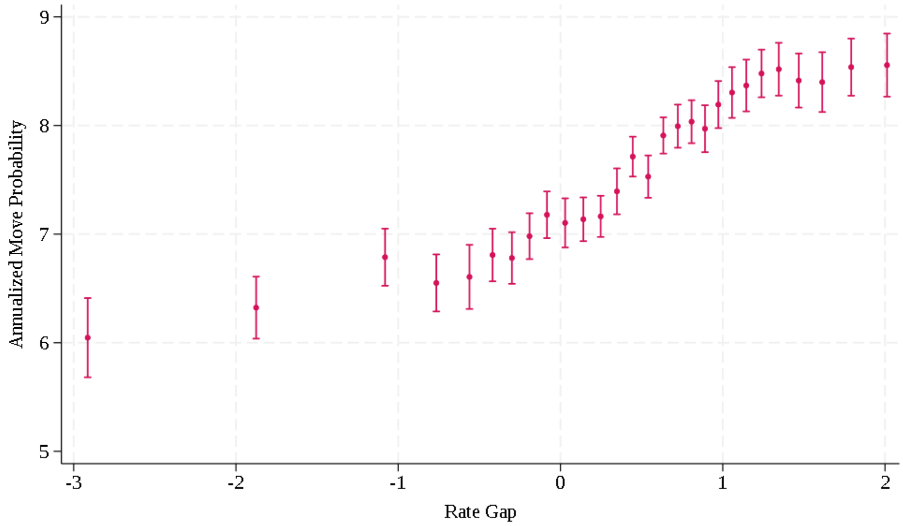

Figure 4 plots estimated monthly probability of moving conditional on borrower's rate gap using estimated regression coefficients from Equation (2) evaluated at mean of control variables. Vertical bars mark 95 percent confidence bands using standard errors clustered by origination-month cohort. Alternate specifications are shown in Appendix Figure A.1 and A.2 in Appendix A.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, CoreLogic deeds data

3.3 Counterfactual Analysis: Estimating "Missing Moves"

Rate hikes between 2022 and 2023 have pushed the distribution of rate gaps further down, lowering implied move probabilities. Our estimates allow us to quantify the effect of this shift on overall moves. To do so, we define a counterfactual rate gap $Rat\tilde{e}Gap_{it}$ as of month t by utilizing prevailing market rates on new originations one year ahead in month $t+12$:

$$\begin{aligned} Rat\tilde{e}Gap_{it} &= r\hat{a}te^{mkt}_{i,o(i)} - r\hat{a}te^{mkt}_{i,t+12} \\ &= r\hat{a}te^{mkt}_{i,o(i)} - \hat{\theta}_{t+12} * X_{i,t} \end{aligned} \tag{4}$$

The counterfactual distribution of rate gaps is similar to the actual movement in rate gaps we showed in Figure 3 earlier, although the counterfactual distribution removes the compositional effects driven by refinancing activity, new originaitons, and moves that occur over the year. The counterfactual instead would give us the distribution of rate gaps on the stock of loans active in period t using market rates in $t+12$. In effect, this tells us the rate gap today's borrowers would face if they were to not move or refinance in one year's time. Given estimates relating rate gaps to move probabilities, we can translate the shift in rate gaps into changes in expected moves for each borrower, holding borrower characteristics constant at time t levels. Aggregating across borrowers provides a measure of the change in moves driven by the change in rates over the t to $t+12$ period.

$$RateDrivenMoves_t = \sum_i \hat{B}(Rat\tilde{e}Gap_{it}) - \hat{B}(RateGap_{it}) \tag{5}$$

where $\hat{B}(\dots)$ is our estimate of the potentially nonlinear relationship between rate gaps and move probabilities.[24]

Figure 5 below shows the results of the simulation exercise at the national level. For each year in our data, the figure first shows the actual number of moves observed nationwide in the gray bars. The red bars represent the implied change in the number of moves for each year's counterfactual rate gap. Taking 2022 as an example, the gray bar shows that about 2.9 million moves occurred, down from over 4.5 million in 2021. The red bar in 2022 suggests around 719 thousand more moves would have occurred in 2022 had borrowers faced the same rate gap as they did before the 2022's mortgage rate hikes. In other words, about 44% of the drop in moves between 2021 and 2022 can be explained by shifts in the rate gap distribution and the resulting lock-in effect.[25]

Looking over the long run, the figure also shows year to year changes in interest rates and their ensuing effect on rate gaps had very impact on mobility in almost all other years. Certainly prior to the pandemic, adding back rate-implied changes in moves does very little

[24] In a linear model where $move_{it} = \beta_0 RateGap_{it} + \gamma Z_{it} + \varepsilon_{it}$, our estimate of moves driven by changes in rates would simplify to $RateDrivenMoves_t = \sum_i \hat{\beta}(\hat{\theta}_t - \hat{\theta}_{t+12})X_i$. In our setting, we allow for the effect of rate gaps to vary non-linearly as captured by β_k, and our estimate of rate-driven changes in moves also accounts for the mass of loans at various points in the initial rate gap distribution.

[25] A potential concern in this calculation is that the estimated non-linearity at extremely high rate gaps may be noisy, resulting in large effects for a small number of borrowers in the tails. To ensure this is not the case, we confirm that enforcing a flat relationship above a rate gap of +2.0 only raises the fraction of the mobility drop explained by lock-in by 1.7 percentage points.

to change the overall mobility picture. Interestingly though, 2020's red bar is quite negative, hinting that far fewer moves would have occurred in 2020 without large pandemic-related rate cuts.

Figure 5: Simulating Missing Moves Due to Rate Changes

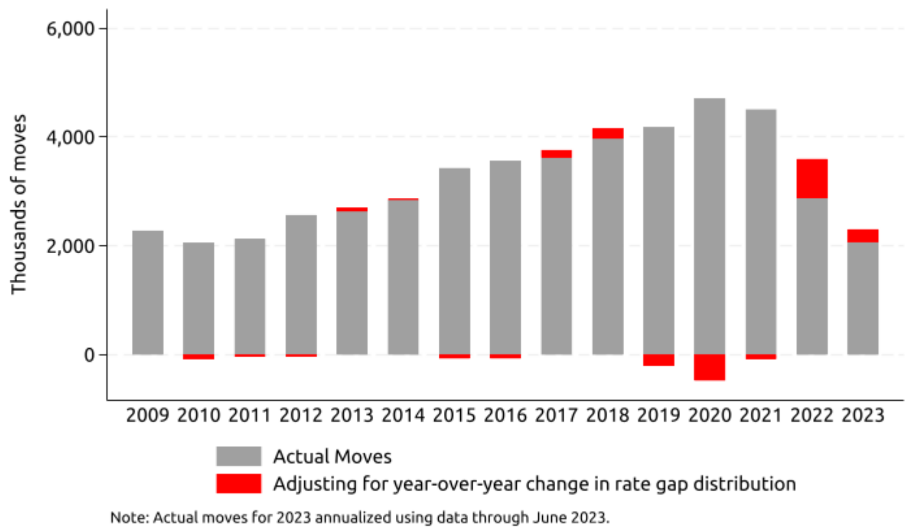

Figure 5 plots actual moves in each calendar year (gray bars). Changes in the rate gap distribution from year to year, coupled with the non-linear relationship between rate gaps and moves in Figure 4 provides an estimate of the contribution of rate changes to the year-to-year change in mobility (shown in red bars). For example, about 44% of the decline in moves from 2021 to 2022 can be explained by the sharp increase in rates over the period.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, CoreLogic deeds data, Freddie Mac PMMS

4 Evaluating Lock-In's Broader Economic Impacts

The previous section showed that mortgage rate lock-in primarily reduces residential mobility, as higher interest rates increase the financial burden of moving for existing homeowners with low fixed rates. However, understanding the full impact of the lock-in effect requires considering broader spillovers and externalities that come from this decline in mobility.

The chief concern is that the economic impacts of lock-in extend beyond short-run changes in homeowner behavior and also influence the labor market, house prices, and within-housing market dynamics. As such, the following subsections consider two broader impacts of rate hike-induced lock-in: (1) potential inefficiencies and mismatch in the labor market; (2) potential effects on market tightness and house prices across metro areas.

4.1 Effect of Lock-In on Labor Market Mismatch

The economic implications of a sudden drop in mobility depend on the nature and purpose of moves forgone by “locked-in” homeowners. As a result, understanding the types of moves most and least responsive to changes in rate gaps is crucial. For example, if the lock-in effect were to prevent workers from relocating for new job opportunities, it could lead to labor market inefficiencies by creating a spatial mismatch between labor demand and supply (Fonseca and Liu, 2023). Conversely, if lock-in mainly reduces *within-metro* mobility, the potential for labor market inefficiencies is more limited. Therefore, to uncover the potential labor market effects of lock-in, we first identify the types of moves most related to job changes and then estimate how responsive these moves are to changes in rate gaps.

Using several other datasets, we first document key stylized facts about mobility, focusing on the distinction between job-related and non-job-related moves. Figure 6 displays the distribution of moves by distance and metro area (top panel) and reasons for the move and distance (bottom panel). Overall, the Figure reveals two distinct types of moves: (1) shorter distance, within-metro moves made for family or house/neighborhood match reasons; and (2) much longer distance, across-metro moves made for employment-related reasons. This distinction informs the manner in which we map move distances and origin-destination observed in our data to a move’s likely labor market motivations.

The top panel of Figure 6 shows the distribution of move distances among borrowers with mortgages in the FRBNY/Equifax Consumer Credit Panel (hereafter, CCP), a data set comprised of a 5 percent sample of credit bureau records in the US (of New York and Equifax, 2024). The takeaways are threefold: first, over half of moves are short distances (within 30 miles of the origin) and are typically within the same CBSA, likely providing access to the same jobs as the origin zip code; second, there are few intermediate-distance moves; third, there is a long tail of moves over 200 miles to different CBSAs, likely associated with job changes.

While we cannot observe reasons for moves in the CCP, the Current Population Survey Annual Social and Economic Supplement (CPS ASEC) asks respondents who moved over the prior year why they did so. This question allows us to verify the hypothesis that job-driven moves are more likely to be at longer-distances. The lower panel collapses these detailed reasons for moving codes into broad categories, restricting to households with mortgages. Short-distance moves (within the same county) are overwhelmingly driven by either family reasons (change in household structure) or motivated by desire to improve housing or neighborhood. We therefore dub these shorter-distance moves "consumption-driven" moves, as they appear motivated by preferences for housing services rather than changes in employment. Conversely, "job-driven" moves–particularly for new jobs–are typically across longer distances (to different states).

These two types of moves are driven by very different household motivations, and understanding how borrower rate gaps affect the likelihood of each type of move is crucial to understanding the broader economic impacts of the lock-in driven decline in mobility. On the one hand, lower (more negative) rate gaps could inhibit "consumption-driven" moves. In particular, formerly in-the-money borrowers would have saved on monthly mortgage payments from a refinance. The income-effect of these savings may have driven spending on various non-housing goods and services (Berger et al., 2021), but also may provide the ability to purchase more housing services by moving. As rate gaps fall into low or negative regions, the financial benefit of the refinance diminishes, reducing this income effect. In such a case, rising rates may lock in homeowners by reducing their ability to climb the housing ladder, but impacts on labor markets may be more limited.

On the other hand, negative rate gaps could inhibit "job-driven" moves: workers who are locked-in to their current homes may find it unattractive to move to job opportunities

Figure 6: Heterogeneity in Move Distances and Reasons

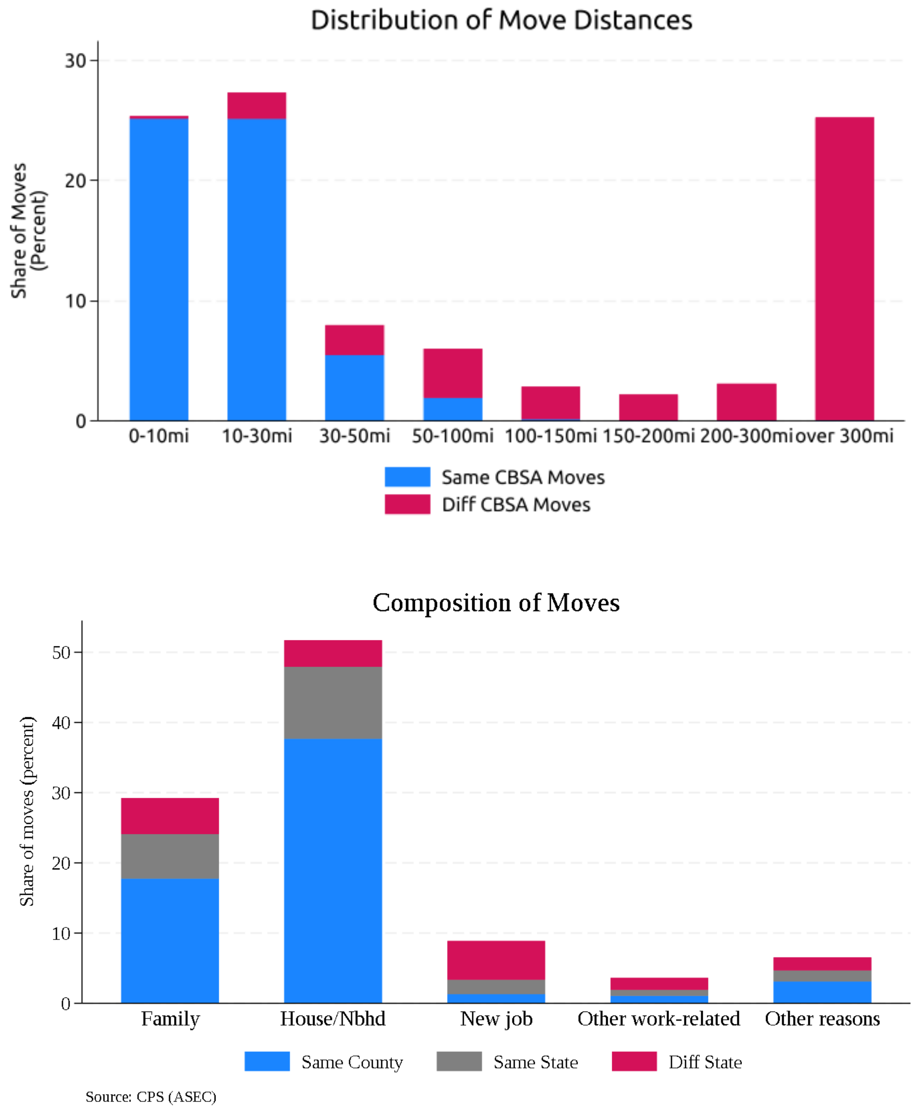

Figure 6 shows heterogeneity in move distances and reasons. The top panel shows the distribution of move distances among households with mortgages. About 1/2 of moves are within 30 miles and 2/3 are within the same CBSA, likely providing access to the same jobs as before. Few moves occur to intermediate distances, and most longer distance moves are to different CBSAs. The bottom panel shows most job-related moves tend to be across state lines, while most short-distance moves are motivated by desire to improving housing or neighborhood quality/size or due to family structure reasons.

Sources: Top panel – FRBNY/Equifax Consumer Credit Panel, sample of borrowers with mortgage prior to move. Bottom panel – CPS (ASEC), 2010-2023, sample of homeowners with mortgage in other labor market areas unless they offer sufficient wage gains to compensate for rate gaps.[26] A reduction in "job-driven" moves would lead to labor misallocation, with potential impacts on aggregate productivity and structural unemployment.

Before estimating the elasticity of each move type with respect to rate gaps, we document aggregate trends in these types of moves using the CCP data. Restricting our sample again to borrowers with mortgages prior to moving, we see in the top panel of Figure 7 that the decline in mobility in recent years is driven by falling within-CBSA moves (blue line) coinciding with the timing of rate hikes (vertical dashed line). Meanwhile, the rate of across-CBSA moves fell only modestly (red line) and did so only several quarters after rate hikes began. Overall, Figure 6 provides suggestive evidence that the decline in mobility driven by lock-in effects reflects fewer consumption-related moves up the housing ladder, with more limited effects on labor mobility.

Of course, trends in recent data are only suggestive evidence, as many other factors – such as the surge of moves in 2020 – may drive aggregate move rates independently of the rate hikes in 2022. Given this suggestive evidence on recent trends in the data, we return to our baseline model (as in Section 3.2) to estimate the relationship between each type of move and borrower rate gap.[27]

Our results also suggest limited scope for large effects on labor misallocation. Figure 8 shows rate gaps primarily affect moves by changing the prevalence of within-CBSA moves

[26] This is closely related to the similar LTV-driven lock-in effect described by many authors following the 2008 housing market crash (Ferreira et al., 2010; Schulhofer-Wohl, 2011; Ferreira et al., 2011; Sterk, 2015; de Francisco and Powell, 2020). In that setting, negative or insufficient equity prevents moves following a significant decline in home values. The constraint is driven more by available equity and required down-payments on the subsequent house rather than changes in monthly servicing costs. In more recent work focused on rate-driven lock-in, Fonseca and Liu (2023) show that low and negative rate gaps prevent moves to nearby locations (50-150 mile rings from origin) with higher wage growth.

[27] Mobility results splitting moves by the destination of the move are subset to moves where the final destination is observed in the data. This excludes prepayment events that have no zip-code change in the CRISM data, but are inferred to be sales based solely from property deed transfers (ie, home sale occurs around a prepayment, but no zip-code change recorded in credit records within 6 months). Using a subset of moves which we can link to the subsequent loan (about 20 percent of the sample), we find that moves resulting in a missing destination by our main definition are predominantly (70 percent) within-metro and within-30-mile moves, and results for missing-destination moves appear quite consistent with an appropriately weighted average of results from observed-destination moves to near and far locations.

Figure 7: Time Series of Moves by Type

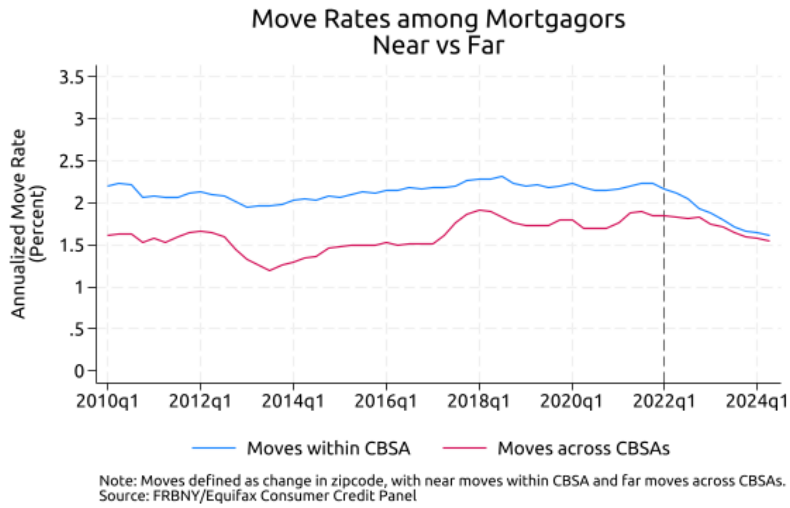

Figure 7 shows moves (zip code changes) in credit records for borrowers that had a mortgage prior to the move. The figure splits moves into those within CBSA (blue) and across CBSAs (red), and shows aggregate declines in moves over 2022 were driven primarily by short-distance (within CBSA) moves, which tend to be driven by consumption-related reasons. Across-CBSA moves remained relatively flat.
Sources: FRBNY/Equifax Consumer Credit Panel

(left panel, blue) or at short distances (right panel, blue), while cross-CBSA moves (left panel, red) and longer distance moves (right panel, red) are almost unaffected by rate gaps. Therefore, consistent with suggestive aggregate trends, recent changes in rates appear to primarily reduce local moves up the housing ladder, but do not affect the jobs these households can access. At the same time, productive, job-related moves–which tend to be over longer distances–appear little affected by changes in rate gaps.

Our results provide two broad implications. First, we find limited scope for labor misallocation driven by lock-in. Overall, rate gaps appear to affect mobility, but do not distort the jobs which households can access. Rising rates do lock in households, preventing within-CBSA moves, but these households remain within commuting distance to the jobs they could have accessed had they moved. Moves across labor market areas appear little affected by such a lock-in effect. This may reflect the fact that wage differentials due to job changes are typically large enough to make rate gaps inframarginal, or that wages can adjust to

Figure 8: Moves Within and Across Labor Market Areas versus Rate Gap

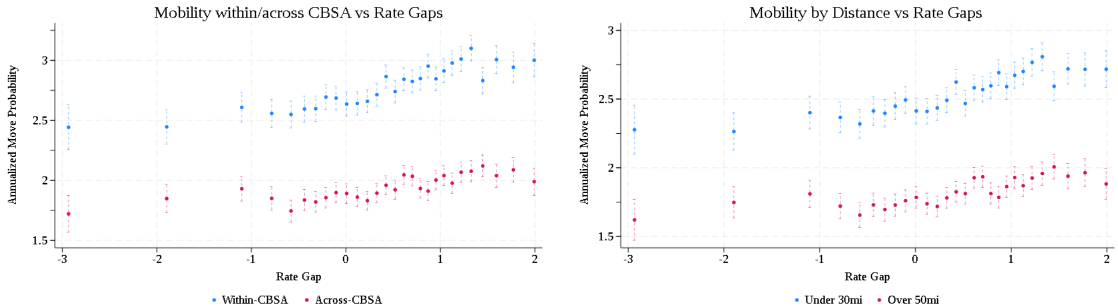

The left panel of Figure 8 plots estimated probability of moves either within the same CBSA (blue) or to a different CBSA (red) using coefficient estimates from regression equation (2) evaluated at the mean of controls using the imputed market rate to create rate gaps. The right panel similarly plots estimated probability of moves under 30 miles from the origin zip code (blue) to the probability of moving more than 50 miles away (red). Moves between 30-50 miles are excluded, but are quite uncommon. Dashed bars are 95 percent confidence bands. Changes in rate gaps appear to affect same-CBSA moves (left panel) and moves within 30 miles of the origin zip code (right panel), each of which are likely more discretionary. Moves across CBSAs or more than 50 miles away–which are more likely to reflect changes in jobs–are almost unaffected by rate gaps, suggesting effects on labor allocations may be minimal.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS

compensate for rate gaps when moves are sufficiently productive.

Second, lower (and negative) rate gaps imply a change in the marginal mover to one who is more insensitive to rate gaps. Typically, over half of moves are within a 30-mile radius, and about 2/3 of moves are within the same CBSA. The bulk of these are consumption-related moves, typically to higher-priced homes or neighborhoods. As rates rise and rate gaps fall, these types of moves decline sharply, and longer-distance, job-related moves comprise a larger share of overall moves. Put differently, the marginal mover in an environment with high rate gaps is one who is more likely to be moving for consumption-related reasons, whereas the marginal mover in a rising rate/low rate gap environment is more likely to be moving across cities and changing jobs. Such movers also appear less sensitive to rate gaps, as shown in Figure 8. Indeed, this is consistent with the flattening of the move rate slope at low and negative rate gaps in Figure 4.

Finally, we look at how the rate gap affects housing decision on the intensive margin - i.e.

how much house to buy (conditional on moving). Using the sample of loans matched within CRISM described in Section 2, we run a regression of the form in Equation 2, with the log change in home value between the sold and purchased homes as the dependent variable.[28] Appendix Figure C.3 shows the results of this regression. There are two key takeaways. First, the magnitude of the price upgrade is gently increasing in the rate gap, but most of the effect of the rate gap on housing demand operates through the extensive margin. Second, within-MSA moves are associated with much larger housing upgrades (measured by price increases) than between-MSA moves. On average, within-MSA movers' new home is about 20% more expensive than their old home, whereas between-MSA movers' new home is about 10% more expensive than their old home. This result is consistent with our broader finding that within-MSA moves seem primarily motivated by consumption considerations, with moves generally involving material upgrades (either in space or neighborhood quality).

4.2 Effect of Lock-In on Housing Markets

Analyzing the lock-in effect at the CBSA (or metropolitan area) level is valuable for several reasons. First, a key macroeconomic question is whether lock-in suppressed housing market activity, constraining the supply of homes for sale and thereby increasing house prices. Exploiting variation across local housing markets can help answer such a question. Second, since housing markets are inherently local, it is important to understand whether there was differential exposure to lock-in – and therefore different effects – across metros.

To examine the lock-in effect by CBSA, we conduct a similar "rate-driven moves" calculation to that in Section 3.3. Specifically, we use equation 5 to aggregate the decline in move probability induced by changes in rates from December 2021 to December 2022 for each CBSA. To the extent that our estimates recover a structural relationship between rate gaps and mobility, the resulting rate-driven decline in moves provides a sufficient statistic for each metro area's differential exposure to lock-in.

[28] The sales price of the old house is determined using deeds records. The purchase value of the new home is estimated using the initial appraisal recorded in the servicing records

Our analysis reveals widespread exposure to lock-in, as nearly all metros experience a significant rate-driven decline in moves. Though there is not a lot of variation across metros, smaller and less expensive CBSAs tended to have greater exposure to lock-in. Perhaps counterintuitively, those most 'exposed' to mortgage rate lock-in (whether individuals or geographic areas) had positive rather than negative rate gaps prior to the rate hikes. Though absolute move probabilities are declining in rate gaps (as shown in Figure 4), a 200 basis point increase in the market mortgage rate differentially reduces the likelihood of moving to a much greater degree for those with an initial rate gap of 1.5 compared to those with an initial rate gap of -0.5. (To see this, note the slope is much steeper from 1.5 to -0.5 than from -0.5 to -2.5.) In other words, market mortgage rate increases are less likely to be pivotal in the mobility decisions for borrowers already at lower points in the rate gap distribution ex ante.

4.2.1 House Prices and Housing Market Activity

With the exposure measure in hand, we first look at the effect of lock-in on real estate listings. More locked-in CBSAs saw differentially larger declines in active real estate listings (Figure 9).[29] A one standard deviation increase in a CBSAs exposure to lock-in was associated with a 10 percent decline in active listings in the 18 months following the initial rate hikes.[30] Intuitively, since listings and moves go hand in hand, this finding is consistent with the decline in mobility we show earlier in Section 3.2.[31]

Next, we analyze the effects of a lock-in induced decline in mobility on CBSA level prices as shown in Figure 10. We find a modest positive effect of exposure to lock-in on CBSA-

[29]Per Realtor.com: "The active listing count tracks the number of for sale properties on the market, excluding pending listings where a pending status is available. This is a snapshot measure of how many active listings can be expected on any given day of the specified month."

[30]On average, our exposure measure suggests a little over 1 percent decline in move probabilities, but exposure across CBSA's is quite correlated. The cross-sectional standard deviation of exposure across CBSA's roughly corresponds to a 0.1 percentage point difference in move probabilities.

[31]Importantly, these two findings come from different data sets. Moves are measured in the CRISM-Deeds match; MLS listings data are from Realtor.com. It is therefore reassuring to see corroboration across the two data sources.

Figure 9: Effect of Lock-In on CBSA-level Active Listings

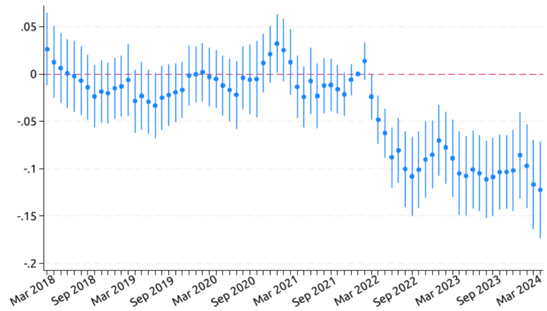

Figure 9 plots coefficients from two-way fixed effect difference-in-difference style regression. The figure shows the effect of a one standard deviation increase in CBSA exposure to lock-in (as measured by CBSA average 'missing move' rate) on the natural log of monthly active real estate listings in that CBSA, according to Realtor.com data. The omitted month is December 2021, just prior to the Federal Reserve's rate hikes. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.
Source: Authors' calculations using ICE/McDash and Equifax CRISM, PMMS, Realtor.com MLS data

level house price growth. In our preferred specification, a one standard deviation increase in exposure increases year-over-year prices by about 2 percent (left panel), with larger effects observed in more populous CBSAs (right panel). This finding suggests lock-in affects metro area housing markets by reducing mobility and tightening inventory, which in turn marginally raises prices. Importantly, since rate gaps seem to be orthogonal to cross-CBSA mobility decisions (as shown in Figure 8), it is unlikely that internal migration plays a pivotal role in either the inventory or house price effects here.

Overall, the lock-in effect has important implications for local housing markets. By reducing within-CBSA churn, exposure to lock-in tightens inventory and marginally increases house price growth. Although high interest rates cooled economic activity in local housing markets, rate hikes also caused lock-in, which put upward pressure on prices.

4.2.2 Lock-In Through the Lens of a Housing Search Model

It is not obvious that prices would rise due to lock-in, as reduced churn lowers both housing demand and the supply of homes for sale. Using a housing search and matching model similar to Han and Strange (2015), Genesove and Han (2012), and Novy-Marx (2009), we show that

Figure 10: Effect of Exposure to Lock-In on CBSA House Price Growth

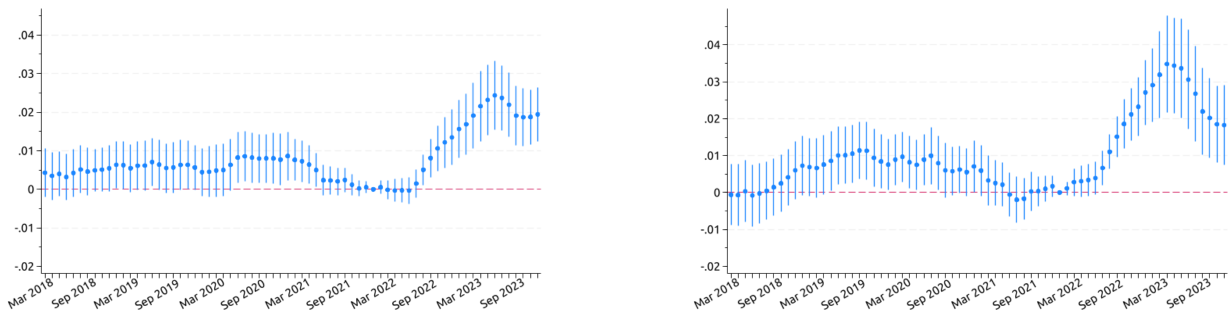

Figure 10 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses log year-over-year house price changes at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. Effect sizes are for a one standard deviation increase in CBSA exposure to lock-in (as measured by CBSA average 'missing move' rate). The left panel includes all 699 CBSAs (out of 767 in our data) for which we have at least 500 loans in the Dec. 2021 cross-section of the Equifax CRISM dataset. The right panel is restricted to the 300 largest CBSAs in our sample. Magnitudes are greater when limiting the sample to the largest metro areas. For both panels, the omitted month is December 2021, just prior to the Federal Reserve's rate hikes. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.
Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, CoreLogic HPI

lock-in related prices increases are driven by the tightening of already-tight housing markets. Even when both buyers and sellers decrease equally in such markets, the buyer-to-seller ratio rises. We then validate this with data, showing that various measures of market tightness increase more sharply in lock-in-exposed metros, confirming the model's prediction of the effects of reduced churn. The model also yields some testable hypotheses on heterogeneity in price growth, which we confirm in the data. Specifically, we show ex ante tighter markets (those that favored sellers more prior to lock-in) saw larger price effects.

Model Setup: A constant-returns matching function $m(B, S)$ maps the number of buyers B and number of sellers S into the number of meetings between them.[32] Because the function is constant returns, it is useful to define market tightness or the buyer-to-seller ratio $\theta = \frac{B}{S}$ and write the probability of a seller (or buyer) matching to a trading partner as

[32] As is standard in the literature, we assume the function $m(B, S)$ is increasing in both arguments, concave, and constant-returns (homogeneous of degree 1). We also assume no matches occur when either buyers or sellers fall to zero $m(0, S) = m(B, 0) = 0$. The monotonicity assumption implies that higher numbers of buyers or sellers increase the likelihood of matches occurring. The CRS assumption allows us to simplify the matching rates to be a functions of the buyer-to-seller ratio, or market tightness.

$q(\theta)$ (or $h(\theta)$) given by:

$$q(\theta) \equiv \frac{m(B,S)}{S} = m(\frac{1}{\theta}, 1)$$
$$h(\theta) \equiv \frac{m(B,S)}{B} = m(1, \theta)$$

Under the standard assumptions about the matching function, $q(\theta)$ is increasing in θ such that a higher proportion of buyers to sellers increases the likelihood of a seller meeting a buyer. Similarly, the function $h(\theta)$ is declining in θ, as a higher proportion of buyers to sellers decreases the likelihood of a buyer meeting a seller.

A buyer values a given home at $X \sim 1 - G(X)$ where $G(X)$ is the complement of the CDF of X. The distribution of X is known to both buyers and sellers, and the realization of X is observed by both parties upon matching. We define V_S and V_B as the (endogenous) values of continuing to search for sellers and buyers respectively. Transactions occur when $X \geq V_S + V_B \equiv y$, where the value to the buyer X exceeds the combined continuation value for both parties. This occurs with probability $G(y)$, and the surplus $X - y$ is split between the buyer and seller via Nash Bargaining, where β reflects the bargaining power of the seller and $(1 - \beta)$ reflects the bargaining power of the buyer.[33] The price for a match with positive surplus is then given by $P = V_S + \beta(X - y) = (1 - \beta)V_S - \beta V_B + \beta X$ such that the seller receives a surplus of $P - V_S = \beta(X - y)$ and the buyer receives surplus of $X - P - V_B = (1 - \beta)(X - y)$.

This allows us to now write the values of continuing to search for both the seller and buyer.

$$rV_S = -c_S + q(\theta)\beta G(y)(E[X|X \geq y] - y) \tag{6a}$$

$$rV_B = -c_B + h(\theta)(1 - \beta)G(y)(E[X|X \geq y] - y) \tag{6b}$$

[33] Alternate forms of bargaining considered in the literature such as Rubinstein bargaining result in the surplus share for the seller $\beta(\theta)$ being an increasing function of θ. As we will show, such a scenario would amplify the effect of lock-in as rising θ increases prices not only via compensating buyers and sellers for changes in their value of search, but also because the match surplus to the seller increases.

The value of continued search valued at discount rate r. A seller listing their home incurs an exogenous flow cost c_S. With a probability of $q(\theta)$ the seller meets a matched buyer. Conditional on this meeting, $G(y) = P(X - y \geq 0)$ is the likelihood that this match has positive surplus and will result in a trade, and $E[X|X \geq y] - y$ is the expected surplus conditional on a positive surplus. The seller keeps a fraction β of this surplus. Similarly, the buyer's value reflects a flow cost of search c_B, $h(\theta)$ is the probability of meeting a seller, and the buyer retains $(1 - \beta)$ share of the expected surplus.

The system of equations (6) and depend on three endogenous variables V_S, V_B and θ. Typical approaches to closing the model assume some entry conditions for either buyers, sellers, or both, in order to pin down outside option values for buyers or sellers. We assume a general form for entry such that buyers and sellers enter as linear functions of the value of search.[34]

$$\dot{B} = \alpha_B + \gamma_B V_B$$

$$\dot{S} = \alpha_S + \gamma_S V_S$$

Equilibrium: For values and prices to be constant along a balanced growth path equilibrium, the proportion of buyers and sellers (θ) must also be constant. Setting $\frac{d}{dt}\frac{B}{S} = 0$ using the entry conditions for buyers and sellers yields.

$$\theta^* = \frac{\alpha_B + \gamma_B V_B}{\alpha_S + \gamma_S V_S} \tag{7}$$

We show in Appendix B that there is a unique equilibrium $\theta, V_S, V_B, P(X)$ that satisfies Equation (7), along with Equations (6) and the pricing function $P(X) = V_S + \beta(X - V_S - V_B)$.

Lock-in Driven Missing Moves Shock: We now consider a "lock-in" type shock to the steady state equilibrium such that the number of movers falls, reducing both the

[34] This general form embeds several possible sub-models. For example, we can rewrite the linear functions as $\dot{B} = \alpha^{inbuyer} + \alpha^{movers} + \gamma_B V_B$ and $\dot{S} = \alpha^{outseller} + \alpha^{movers} + \gamma_S V_S$. A flow of "movers" move within the market, increasing the number of both buyers and sellers. In addition, "in-buyers" and "out-sellers" are the net in-flow and out-flow of buyers and sellers from this market, reflecting cross-market moves and moves in/out of owner-occupying within the market. In this setting, a reduction in movers α^{movers} would result in an equivalent level shock to both α_B and α_S in our baseline case, and this is the shock we consider.

number of entering buyers and sellers by an equal amount: $d\alpha_S = d\alpha_B = -d\ell$. This shock approximates the effect of reduced within-market moves we see in the data and allows us to understand mechanisms driving prices.

Appendix B provides a more complete mathematical treatment of the comparative statics, but much of the intuition behind the mechanism can be understood by differentiating the steady-state value of market tightness θ given by (7) with respect to the lock-in shock.

$$\frac{d\theta}{d\ell} = \frac{\theta - 1}{\alpha_S + \gamma_S V_S} + \frac{\gamma_B \frac{dV_B}{d\ell} - \theta \gamma_S \frac{dV_S}{d\ell}}{\alpha_S + \gamma_S V_S} \tag{8}$$

The first term reflects the direct effect of the shock on θ, holding fixed the values of search (and therefore also P).[35] Notably, the sign of this effect depends crucially on whether the market is a buyer's market ($\theta < 1$) or a seller's market ($\theta > 1$). The second term captures indirect effects of equilibrating forces as entry of buyers and sellers respond to the respective changes in search values. In a setting where $\theta > 1$ prior to the shock, the shock raises θ, raising the value of search for sellers and lowering it for buyers. This draws in new sellers, but deters buyers from entering the market, attenuating the increase in θ.[36]

In reality, demographic forces and other factors likely result in potential buyers exceeding sellers, causing $\theta > 1$. In such a setting, the lock-in shock that reduces both buyers and sellers by similar amounts *raises* θ further away from 1.[37] Our results also suggest lock-in caused a shift in the composition of movers away toward those moving for more non-discretionary reasons. While our simple model does not provide endogenous reasons for moving, it is

[35] This term maps closely to a simpler model version of the model where $\gamma_B = \gamma_S = 0$ described in B. In that setting, θ is pinned down by entry flows alone because entry does not respond to changes in search value. This results in no equilibrating forces wherein rising θ draws in more sellers and deters buyers.

[36] The endogenous response captured by the second term cannot exceed the direct effect of the shock in the first term, since this would result in a lower θ which would not result in the response in flows described here. Therefore, the sign of the net effect must be driven by the direct effect (first term), with equilibrating forces (second term) attenuating the magnitude. For example, if $\theta > 1$, the first term suggests tightness θ would rise, lowering the value of search to buyers $\frac{dV_b}{d\theta} < 0$ and raising the value for sellers $\frac{dV_s}{d\theta} > 0$. As the flow of buyers and sellers respond, θ falls, as described by the second term in 8. However, if this response leaves θ lower than the initial value, the signs of these derivatives would reverse, reversing the sign of the second term.

[37] This can be seen concretely by considering a market with 1000 buyers and 800 sellers, such that $\theta = 1000/800 = 1.25$. A reduction in both buyers and sellers of 400 results in $\theta = 600/400 = 1.5$.

quite likely that the marginal mover after the shock is less sensitive to financial incentives when making the decision to move, including considerations of the value of search. This may suggest the marginal mover also has relatively lower values of γ_B and γ_S, reducing the attenuation provided by the second term. Regardless, the net effect of the lock-in shock is a tighter housing market.

Increased market tightness raises the value of search for sellers V_S and lowers values for buyers V_B as seen in Equation (6). Prices for a home that a buyer values at X are given by $P(X) = V_S + \beta(X - V_S - V_B) = \beta X + (1-\beta)V_S - \beta V_B$. Sellers must now be compensated for a higher value of search, raising prices. Buyers have a lower value of search, and are willing to pay higher prices. Both these effects result in a higher price for homes.

While we do not explicitly model rental markets, the model offers some insights into the impacts of lock-in on rents. In particular, renters may be deterred from transitioning to home ownership if lock-in tightens the housing market and the value of searching for a home falls. The lower inflow of new buyers is captured in a reduced-form way via γ_B, and the second term of 8 describes how this endogenous response attenuates the impact on market tightness. While prices may not react as sharply, the fact the reduced outflow from rentals would raise demand for rental housing, thereby raising rents.[38]

Notably, all the arguments here hinge crucially on whether the market favored buyers or sellers prior to the rate hikes. The model also provides a clear testable prediction about heterogeneity in how a market responds depending on the level of tightness. In particular, Equation (8) shows market tightness rises more in response to a given lock-in shock when the initial market is already tighter (ie, the response $\frac{d\theta}{d\ell}$ is larger when θ is initially larger).

[38]Another channel by which lock-in could affect rental markets is by raising the transition of current owners to rentals. For example, current owners may move despite a low or negative rate gap by substituting from owner-occupied homes to rentals in order to avoid higher interest costs. The resulting increased inflow to rental markets would raise rents. We find little quantitative evidence for this channel, however. Appendix Figure C.2 shows the decline in mobility among mortgagors is driven by a drop in moves to new owned homes. There is essentially no change in moves to likely rentals (where the borrower does not have a mortgage at the destination), suggesting no significant substitution toward rental housing. This is consistent with the fact that lock-in primarily impacts local moves up the housing ladder, and few rental alternatives are likely to be available for such a move.

The same is true of prices: prices must rise more sharply in markets that were tighter prior to the shock. In the next section, we test these predictions in the data.

Before turning to that test, we also note that our model is very similar to ones that have been utilized to explain the *positive* relationship between volume and prices in the housing market (Genesove and Han, 2012). This relationship can be explained via shocks to the distribution of buyer valuations of homes $G(X)$, which generate an increase prices along with an increase in the overall surplus, drawing in both buyers and sellers. Indeed, Anenberg and Ringo (2024) show that historical variation in sales is driven predominantly by demand shocks affecting the flow of buyers. Obviously, the price-volume correlation has reversed in 2022 as moves declined sharply, reducing home sales even as prices continued to rise. The discrepancy is easily explained by the fact that the source of shocks are quite different from typical situations. As we show, the model can produce negative price-volume correlations when the shock impacts the inflow of both buyers and sellers in a hot market.

4.2.3 Lock-In and Housing Market Tightness

The search and matching framework yields some testable implications in the data. First, the model suggests a key mechanism driving the link between churn and prices is the fact that reduced churn in an already-tight (seller's) market, further increases tightness. While we cannot measure the number of buyers in a market to compute θ directly, we do observe a number of outcomes that are reflective of market tightness. Specifically, we utilize measures of list prices, price cuts, and seller's time on market from Realtor.com. A higher ratio of buyers to sellers would lower a seller's time on the market–a relatively direct measure of the seller's matching hazard and likelihood of sale given by $q(\theta)G(V_B + V_S)$ in the model. In addition, higher tightness would put upward pressure on the seller's value of search, pushing up asking list prices and reducing seller's willingness to cut asking prices.

Figure 11 shows how each of these measures responds to lock-in, utilizing a specification similar to the one used earlier. Looking through seasonality in sales, metro areas that were more exposed to lock-in saw properties stay on market for 5 percent shorter times than prior to the shock (top panel). This reduced time on market occurred even as listing prices rose (bottom left panel). It was also not the case that sellers reduced prices in order to get their homes off the market (bottom right panel), and the reduced churn led to higher overall sale prices as shown previously. Taken together, the results show markets that were more exposed to lock-in appear to have tightened, pushing up prices as described in the model.

Figure 11: Effect of Lock-In Exposure on Housing Market Tightness

(a) Log Median Days on Market

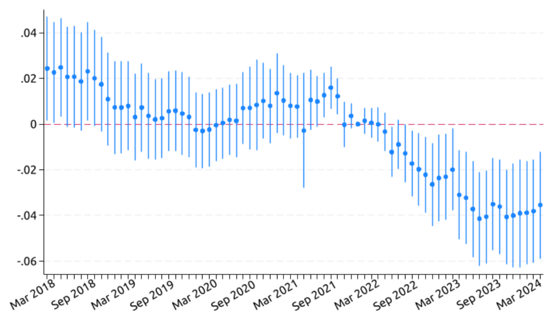

(b) List Prices, YoY change

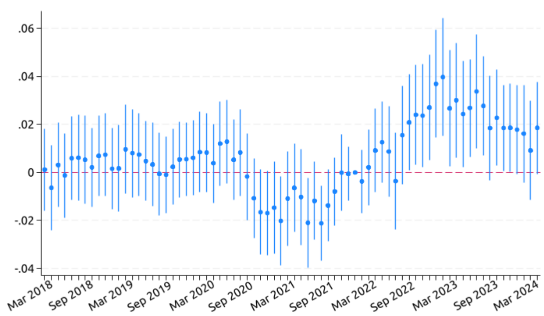

(c) Log Price Reductions

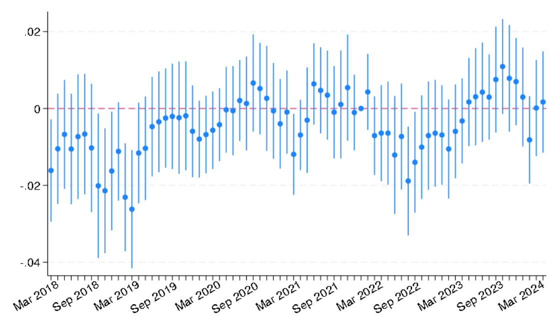

Figure 11 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses various CBSA-level measures of market tightness on CBSA-level exposure to lock-in. The top panel (a) plots the effect of lock-in on the log of the 12-month moving average of median days on market (according to Realtor.com data). The bottom left panel (b) shows the 12 month change in list prices and the bottom right panel (c) shows the log count of listings with price reductions. For all panels, the omitted month is December 2021, just prior to the Federal Reserve's rate hikes. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.
Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, Realtor.com

The model yields an additional prediction that markets that were initially tighter would have a sharper response to reduced churn from lock-in. To test this prediction, we sort metro areas based on the average time on market for homes in the Fourth Quarter (October-

December) of 2021, prior to rate hikes. Markets with below-median time on market are tighter initially, and should display larger responses in both post-2021 tightness and prices due to reduced churn. As seen in Figure 12, splitting the sample of metro areas based on initial tightness reveals considerable heterogeneity. In line with the model's predictions, the left panels show a one standard deviation increase in exposure to lock-in decreased time on market (top left) and increased prices (bottom left) more sharply in initially tight markets compared to markets that were less tight to begin with (right panels). Notably, lock-in appears to raise prices in both types of areas, suggesting even relatively looser markets still favor sellers initially ($\theta_0 > 1$). CBSAs see a similar decline in active listings of about 10% regardless of pre-hike tightness (not shown).

5 Conclusion

The 30-year fixed-rate mortgage regime causes the current distribution of mortgage rates on outstanding loans to affect both refinancing and mobility incentives. Our results show that the "rate gap"–the difference between the current rate a borrower is paying and the market rate they would receive on new financing–is an important indicator of mobility. These results highlight the asymmetry between the effects of rate hikes and rate cuts introduced by the presence of fixed rate mortgages. In particular, while rate cuts are typically met with refinancing activity, passing through interest rates to households via lower mortgage payments, rate hikes may pass through different channels. Our results highlight mobility as an important mechanism for this pass-through.

The so-called 'lock-in effect' resulting from large rate hikes has a pronounced impact on overall residential mobility. Households are much less likely to move and sell homes when available mortgage rates are not sufficiently below rates on their existing loans. Our results suggest lock-in may account for nearly half of the drop in mobility since rates went up in 2022.

Figure 12: Lock-in Driven House Price Growth is Stronger in Tighter Markets

(a) Log Days on Market (Tight Markets)

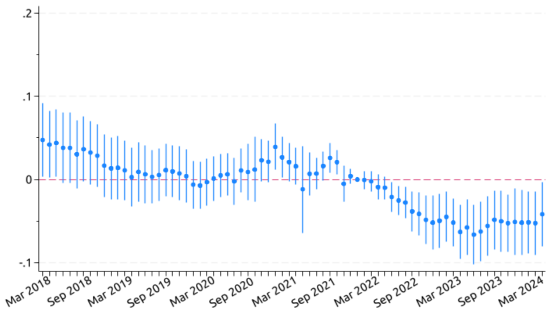

(b) Log Days on Market (Loose Markets)

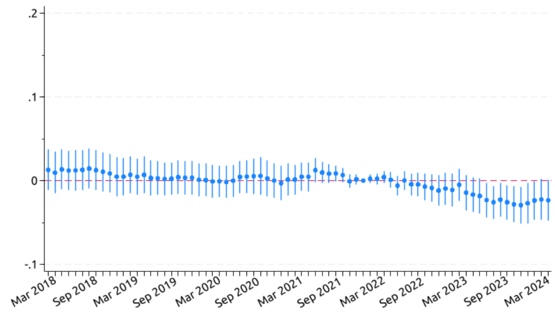

(c) Log House Price Growth (Tight Markets)

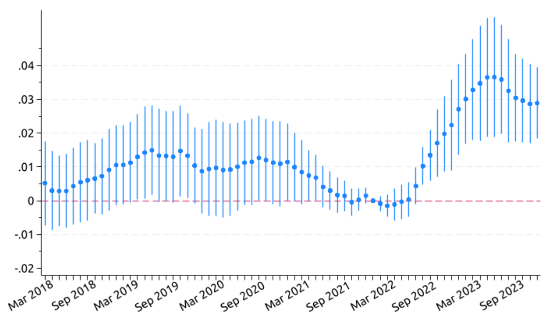

(d) Log House Price Growth (Loose Markets)

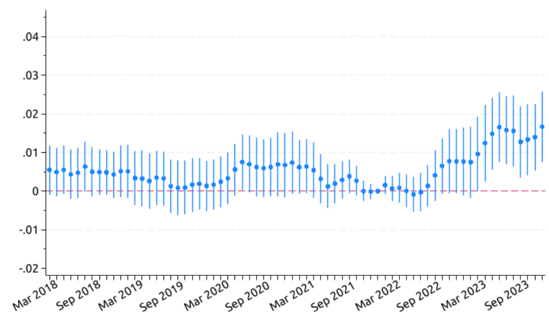

Figure 12 plots coefficients from a two-way fixed effect difference-in-difference style specification. Panels (a) and (b) regress log of the 12-month moving average of median days on market (according to Realtor.com) on CBSA-level exposure to lock-in. Panels (c) and (d) do the same for the log year-over-year house price changes at the CBSA level (according to CoreLogic's CBSA House Price Index). In both cases, the set of left panels (Panels (a) and (c)) include CBSAs with above median market-tightness prior to the rate hikes where market tightness is measured as the CBSAs median days on market in 2021Q4. The right panels (Panels (b) and (d)) include CBSAs with below median market-tightness. Magnitudes of both time on market effects and price effects are two to three times larger for markets above median tightness. In all panels, the omitted month is December 2021, just prior to the Federal Reserve's rate hikes. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.
Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, CoreLogic HPI, Realtor.com

We also assess the potential for broader economic impacts resulting from this drop in residential mobility. Contrary to Fonseca and Liu (2023), we find limited scope for lock-in to affect labor markets. Specifically, our analysis shows that changes in rate-gaps primarily affect relatively short-distance moves within metros, which are likely to be more discretionary and motivated by location choice rather than jobs. Rate gaps have limited impact on cross-CBSA moves, leaving workers largely able to access jobs they would have had access to absent lock-in. Consequently, it is unlikely that lock-in disrupts labor market efficiency or introduces spatial mismatches between jobs and workers.

We find that lock-in significantly impacted local churn. As such, over the recent rate hike period, lock-in contributed to tighter local housing markets, including fewer active real estate listings, higher list prices, and lower time on market. We show this empirical pattern reflects the macroeconomic environment of the time, in which housing markets were already tight: inventories were historically low and house prices were rising by 20% per year. As a result of these initial conditions, CBSAs with greater exposure to lock-in experienced more rapid home price growth, with the effects concentrated in the tightest markets. Though higher mortgage rates curtailed housing market activity, the resulting lock-in put positive pressure on prices.

References

Anenberg, E. and Ringo, D. (2021). Housing market tightness during covid-19: Increased demand or reduced supply? *FEDS Notes https://doi.org/10.17016/2380-7172.2942*.

Anenberg, E. and Ringo, D. (2022). The propagation of demand shocks through housing markets. *American Economic Journal: Macroeconomics*, 14(3):481–507.

Anenberg, E. and Ringo, D. (2024). Volatility in home sales and prices: Supply or demand? *Journal of Urban Economics*, 139:103610.

Batzer, R., Coste, J., Doerner, W., and Seiler, M. (2024). The lock-in effect of rising mortgage rates. Technical report, Federal Housing Finance Agency.

Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2021). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review*, 111(9):2829–2878.

CRISM (2024). Equifax credit risk insight servicing and ice, mcdash data.

de Francisco, Eva, G.-C. J. and Powell, T. (2020). Stuck at home? the drag of homeownership on earnings after job separation. Technical report, FEDS Note.

Ferreira, F., Gyourko, J., and Tracy, J. (2010). Housing busts and household mobility. *Journal of Urban Economics*, 68(1):34–45.

Ferreira, F., Gyourko, J., and Tracy, J. (2011). Housing busts and household mobility: An update. Working Paper 17405, National Bureau of Economic Research.

Fonseca, J. and Liu, L. (2023). Lock-in, mobility, and labor reallocations. Technical report, Jacobs Levy Equity Management Center for Quantitative Financial Research Paper.

Genesove, D. and Han, L. (2012). Search and matching in the housing market. *Journal of Urban Economics*, 72(1):31–45.

Gerardi, K., Qian, F., and Zhang, D. (2024). Mortgage lock-in, lifecycle migration, and the welfare effects of housing market liquidity. *Lifecycle Migration, and the Welfare Effects of Housing Market Liquidity (July 28, 2024)*.

Han, L. and Strange, W. C. (2015). Chapter 13 - the microstructure of housing markets: Search, bargaining, and brokerage. In Duranton, G., Henderson, J. V., and Strange, W. C., editors, *Handbook of Regional and Urban Economics*, volume 5 of *Handbook of Regional and Urban Economics*, pages 813–886. Elsevier.

Hsieh, C.-T. and Moretti, E. (2019). Housing constraints and spatial misallocation. *American economic journal: macroeconomics*, 11(2):1–39.

Larsen, W. and Martinez, A. (2024). House prices, debt burdens, and the heterogeneous effects of mortgage rate shocks. Technical report, George Washington University Working Paper.

Liebersohn, J. and Rothstein, J. (2023). Household mobility and mortgage rate lock. Technical report, SSRN.

Mabille, P., Liu, L., and Fonseca, J. (2024). Unlocking mortgage lock-in: Evidence from a spatial housing ladder model. *Available at SSRN 4874654*.

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

of New York, F. R. B. and Equifax (2024). Federal reserve bank of new york/equifax consumer credit panel.

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

Quigley, J. M. (2002). Homeowner mobility and mortgage interest rates: New evidence from the 1990s. *Real Estate Economics*, 30(3):345–364.

Schulhofer-Wohl, S. (2011). Negative equity does not reduce homeowners' mobility. Working Paper 16701, National Bureau of Economic Research.

Stanton, R. and Wallace, N. (1998). Mortgage choice: What's the point? *Real estate economics*, 26(2):173–205.

Sterk, V. (2015). Home equity, mobility, and macroeconomic fluctuations. *Journal of Monetary Economics*, 74:16–32.

A Mobility Regressions: Alternate Specifications

In this section, we compare alternate specifications for our baseline regression given by Equation 2 to better understand potential biases and sensitivity of our results. As we note in Section 3.1, in addition to addressing heterogeneity in mobility and rate gaps through including common loan-time controls such as borrower and loan characteristics, our baseline model attempts to address two broad sources of bias. The first is bias arising from cross-sectional variation in origination rates which may be correlated with mobility. The second is bias arising from differential selection of likely-movers into higher rate gaps as they forego refinancing opportunities.

Figure A.1 shows our baseline specification (black) in comparison to a number of alternate specifications in an effort to understand how various components of our empirical model address each bias.

The specification shown in red is the simplest, and utilizes the "actual rate gap" (defined in Equation 1), which includes cross-sectional variation in origination rates. It also excludes our maximum rate gap control. As such, it is subject to both downward bias due to more sophisticated borrowers shopping for lower origination rates and also having higher move hazards, as well as the upward bias from survivors who forgo refinancing opportunity having both high rate gaps and high move hazards.

The blue dots show a specification which replaces the "actual rate gap" with the imputed rate gap defined in Equation 3.[39] By doing so, the specification removes all cross-sectional variation in origination rates–shutting down the first bias–resulting in a steeper relationship between mobility and rate gaps. In fact, the slope is nearly three times as high, and suggests lock-in may explain as much as 94 percent of the decline in mobility following the 2022 rate hike episode. Even so, the overall shape of the relationship appears similar, with a flatter slope at high and low elasticity ranges, and cross-sectional variation in "rate-driven moves"

[39]This specification is very similar to the one used by Fonseca and Liu (2023) and Batzer et al. (2024), and delivers an elasticity similar to those papers.

Figure A.1: Alternate specifications – Mobility by Rate Gap, 2009-2023

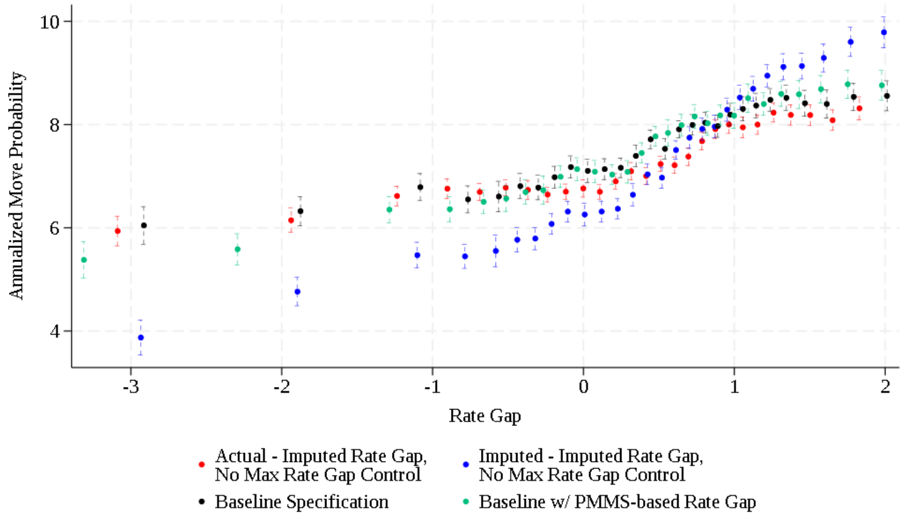

Figure A.1 plots estimated monthly probability of moving conditional on borrower's rate gap using estimated regression coefficients from Equation (2) evaluated at mean of control variables. The plot shows several alternate specifications compared to our baseline specification in black. Red dots are estimated using rate gaps defined as the difference between actual origination rates and imputed market rates at the time of observation, and do not include controls for forgone refinancing opportunities. Blue dots also exclude controls for forgone refinancing opportunities, but utilizes rate gaps constructed using imputed borrower rates at both origination and current time. Green dots repeat our baseline specification using rate gaps constructed using PMMS prime rates for all borrowers at both origination time and current time. In all specifications, dashed lines mark 95 percent confidence bands using standard errors clustered at origination cohort.
Source: Authors' calculations using ICE/McDash and Equifax CRISM, CoreLogic deeds data, Freddie Mac

across CBSA's is quite similar.

The black dots show our preferred specification, using the same model as the blue dots, but adding in our maximum rate gap control. In conjunction with origination cohort fixed effects–which were already included in the earlier models discussed–the maximum rate gap variable controls for forgone refinancing opportunities that may signal an intent to move, helping address the second bias. Indeed, the slope flattens considerably when this control is included, with the estimates suggesting about half of the drop in mobility in 2022 reflects lock-in.

Finally, we include a specification which utilizes PMMS prime market rates at both origination and observation periods to construct rate gaps, shown in green.[40] This removes all cross-sectional variation in rates when constructing rate gaps, utilizing only time series variation. This version yields results that are statistically indistinguishable from our baseline (black).

Next, we consider alternate definitions for our maximum rate gap control. The purpose of this control is to address the potential that rate gaps today may reflect the fact that borrowers who plan to move soon may find it sub-optimal to refinance even when rates are low. The control measures forgone refinancing opportunities in the past by tracking the maximum rate gap that the borrower faced through 3 months prior to observation. Our choice of 3 months is driven by a desire to balance two concerns. On the one hand, we want to use a window that includes recent enough observations that it reflects refinancing opportunities the borrower has faced. On the other hand, we want to allow the borrower enough time to react to a rate gap before we assume the borrower has truly forgone the opportunity.

Figure A.2 compares our baseline three month definition (black) with alternatives defined as the maximum rate gap through 1 month ago (green) and 6 months ago (red). We also include a version with no maximum rate gap control at all (blue, identical to blue dots in A.1 above).

Results using a 6 month offset (red) appear only slightly flatter than the version that excludes rate gap history controls (blue), but quite a bit steeper than the baseline version using a 3 month offset (black). Moving from the baseline 3 month to the 1 month version (green) further flattens the slope, although only slightly.

We opt for the three month version as a baseline as it provides a sufficient history to account for refinancing opportunities in the past without constraining households to move immediately in response to a high rate gap. Refinancing waves appear quickly after rate cuts, and a household who does not refinance within three months after a rate cut may be signaling

[40] Specifically, a borrower i who originated a loan in period $o(i)$, faces a PMMS-based rate gap $RateGap^{mkt-mkt}_{i,t} = rate^{PMMS}_{i,o(i)} - rate^{PMMS}_{i,t}$ at time t

Figure A.2: Alternate Max Rate Gap Controls – Mobility by Rate Gap, 2009-2023

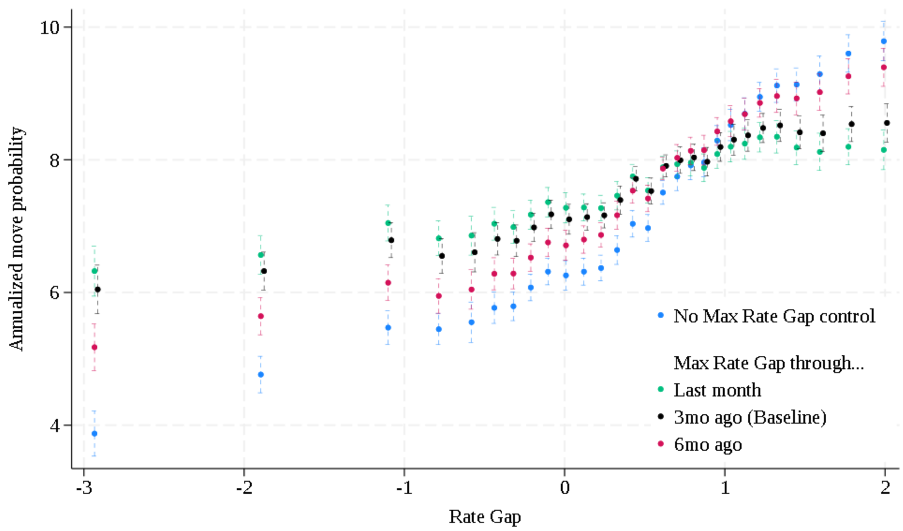

Figure A.2 plots estimated monthly probability of moving conditional on borrower's rate gap using estimated regression coefficients from Equation (2) evaluated at mean of control variables. The plot shows several alternate definitions of our maximum rate gap control compared to our baseline specification in black. Blue dots show estimates that exclude any control for rate gap history (same as blue dots in A.1). Green, black, and red dots include a control for the maximum rate gap observed between origination and 1 month ago, 3 months ago (baseline), and 6 months ago, respectively.
Source: Authors' calculations using ICE/McDash and Equifax CRISM, CoreLogic deeds data

either inattention or a desire to move in the near future. Shortening the time frame beyond three months provides households little time to actually find a home and move in response to a high rate gap. However, Figure A.2 shows the shape of the move probability relationship is broadly similar across all these definitions, and results in similar cross-sectional variation across CBSA's in rate-driven moves. As such, results in the remainder of the paper are robust to small alterations to the window used in defining this variable.

B Model Appendix

B.1 Simple Version: Entry of buyers and sellers at constant rates

In this section, we consider a simpler version of our baseline model to build basic intuition. In particular, we assume $\gamma_S = \gamma_B = 0$ such that buyers and sellers enter at constant flow rates $\dot{B} = \alpha_B$ and $\dot{S} = \alpha_S$. Along a balanced growth path equilibrium, it must be the case that market tightness θ is constant:

$$\begin{aligned} \dot{\theta} &= \frac{d}{dt}\frac{B}{S} \\ &= \frac{\dot{B}S - \dot{S}B}{S^2} \\ &= 0 \\ \theta^* &= \frac{\alpha_B}{\alpha_S} \end{aligned} \tag{9}$$

Given any initial level of market tightness, the level of buyers and sellers in the long run is determined by the inflow of each: $\theta_t = \frac{B_0 + t\alpha_B}{S_0 + t\alpha_S} \rightarrow \frac{\alpha_B}{\alpha_S} = \theta^*$. Therefore, this equilibrium is stable.

Values for buyers and sellers are given as in (6), and equilibrium prices are given as $P(X) = V_S^* + \beta(X - y^*) = (1-\beta)V_S^* - \beta V_B^* + \beta X$ for homes with match quality $X > y^*$.

Now consider a "lock-in" shock which exogenously reduces (within-market) moves. Considering that $\alpha_B = \alpha_B^{netbuyers} + \alpha^{movers}$ and $\alpha_S = \alpha_S^{netsellers} + \alpha^{movers}$, such a shock would reduce both α_B and α_S by a similar amount: $d\alpha_B = d\alpha_S = -d\ell$.

Differentiating the equilibrium level of market tightness shows the impact on the market depends crucially on whether the market favors buyers or sellers:

$$\begin{aligned} \frac{d\theta}{d\ell} &= \frac{d}{d\ell}\left[\frac{\alpha_B}{\alpha_S}\right] \\ &= \frac{-\alpha_S + \alpha_B}{\alpha_S^2} \\ &= \frac{\theta - 1}{\alpha_S} \end{aligned} \tag{10}$$

If the number of buyers exceeded the number of sellers when the shock occurred, $\theta > 1$, then θ will rise further away from 1. The value for sellers from the increased market tightness rises, and the value for buyers falls (from Equations (6)). Equilibrium prices will therefore rise since $P(X) = (1-\beta)V_S^* - \beta V_B^* + \beta X$.

B.2 Equilibrium under linear entry

On a balanced growth path equilibrium, the proportion of buyers and sellers (θ) is constant. This implies...

$$\begin{aligned}\dot{\theta} &= \frac{d}{dt}\frac{B}{S} \\ &= \frac{\dot{B}S - \dot{S}B}{S^2} \\ &= 0 \\ (\alpha_S + \gamma_S V_S)\theta &= \alpha_B + \gamma_B V_B\end{aligned}$$

Letting $y = V_S + V_B$, and solving for values as functions of y and θ yields...

$$V_S = \frac{\alpha_B - \theta\alpha_S}{\gamma_B + \theta\gamma_S} + \frac{\gamma_B}{\gamma_B + \theta\gamma_S} y \tag{11a}$$

$$V_B = \frac{\theta\alpha_S - \alpha_B}{\gamma_B + \theta\gamma_S} + \frac{\theta\gamma_S}{\gamma_B + \theta\gamma_S} y \tag{11b}$$

Plugging these back into value functions (6) provides...

$$r\frac{\alpha_B - \theta\alpha_S}{\gamma_B + \theta\gamma_S} + r\frac{\gamma_B}{\gamma_B + \theta\gamma_S} y = -c_S + q(\theta)\beta G(y)(E[X|X \geq y] - y) \tag{12a}$$

$$r\frac{\theta\alpha_S - \alpha_B}{\gamma_B + \theta\gamma_S} + r\frac{\theta\gamma_S}{\gamma_B + \theta\gamma_S} y = -c_B + h(\theta)(1-\beta)G(y)(E[X|X \geq y] - y) \tag{12b}$$

Differentiating these shows that (6a) is upward sloping in (θ, y) space and (6b) is downward sloping, suggesting any intersection satisfying this system yields a unique equilibrium.

Existence follows from assumptions about the shape of the matching function at limits of B and S.

$$
\begin{aligned}
S: \frac{\partial y}{\partial \theta} &= \frac{R + q'(\theta)\beta G(y) E[X - y | X \geq y]}{r\frac{\gamma_B}{\gamma_B + \theta\gamma_S} + q(\theta)\beta G(y)} \\
&= \frac{[+] + [+]}{[+] + [+]} > 0
\end{aligned}
$$

$$
\begin{aligned}
B: \frac{\partial y}{\partial \theta} &= \frac{-R + h'(\theta)(1-\beta) G(y) E[X - y | X \geq y]}{r\frac{\theta\gamma_S}{\gamma_B + \theta\gamma_S} + h(\theta)(1-\beta) G(y)} \\
&= \frac{[-] + [-]}{[+] + [+]} < 0
\end{aligned}
$$

where...

$$
R = r\frac{\alpha_S(\gamma_B + \theta\gamma_S) + \gamma_S(\alpha_B - \theta\alpha_S) + \gamma_B\gamma_S y}{(\gamma_B + \theta\gamma_S)^2} > 0
$$

B.3 Comparative Statics following Lock-in Shock

Differentiating the value functions (6) with respect to this shock yields...

$$
\begin{aligned}
S : r\frac{\theta - 1}{\gamma_B + \theta\gamma_S} =& \frac{\partial\theta}{\partial\ell}\Bigg[R + q'(\theta)\beta G(y)E[X - y|X \geq y]\Bigg] \\
&+ \frac{\partial y}{\partial\ell}\Bigg[-r\frac{\gamma_B}{\gamma_B + \theta\gamma_S} - q(\theta)\beta G(y)\Bigg] \\
=& \frac{\partial\theta}{\partial\ell}\Bigg[[+] + [+]\Bigg] + \frac{\partial y}{\partial\ell}\Bigg[[-] + [-]\Bigg]
\end{aligned}
$$

$$
\begin{aligned}
B : -r\frac{\theta - 1}{\gamma_B + \theta\gamma_S} =& \frac{\partial\theta}{\partial\ell}\Bigg[-R + h'(\theta)(1-\beta)G(y)E[X - y|X \geq y]\Bigg] \\
&+ \frac{\partial y}{\partial\ell}\Bigg[-r\frac{\theta\gamma_S}{\gamma_B + \theta\gamma_S} - h(\theta)(1-\beta)G(y)\Bigg] \\
=& \frac{\partial\theta}{\partial\ell}\Bigg[[-] + [-]\Bigg] + \frac{\partial y}{\partial\ell}\Bigg[[-] + [-]\Bigg]
\end{aligned}
$$

where...

$$
R = r\frac{\alpha_S(\gamma_B + \theta\gamma_S) + \gamma_S(\alpha_B - \theta\alpha_S) + \gamma_B\gamma_S y}{(\gamma_B + \theta\gamma_S)^2} > 0
$$

We can clearly sign each component above and write them in matrix form below where parameters z, x_i are all positive.

$$
\begin{aligned}
\begin{bmatrix} +z \\ -z \end{bmatrix} &= \begin{bmatrix} +x_1 & -x_2 \\ -x_3 & -x_4 \end{bmatrix} \begin{pmatrix} \frac{d\theta}{d\ell} \\ \frac{dy}{d\ell} \end{pmatrix} \\
\begin{pmatrix} \frac{d\theta}{d\ell} \\ \frac{dy}{d\ell} \end{pmatrix} &= \frac{1}{-x_1 * x_4 - x_2 * x_3} \begin{bmatrix} -x_4 & +x_2 \\ +x_3 & +x_1 \end{bmatrix} \begin{bmatrix} +z \\ -z \end{bmatrix} \\
&= \begin{bmatrix} [+] & [-] \\ [-] & [-] \end{bmatrix} \begin{bmatrix} [+] \\ [-] \end{bmatrix} \\
&= \begin{bmatrix} [+][+] + [-][-] \\ [-][+] + [-][-] \end{bmatrix} \\
&= \begin{bmatrix} [+] \\ [?] \end{bmatrix}
\end{aligned} \tag{13}
$$

The result above allows us to sign effects on equilibrium quantities θ and $y = V_S + V_B$. As argued in the main text, the effect on market tightness θ is clearly determined by the sign of $\theta - 1$. Under the assumption that $\theta > 1$ prior to the shock, θ must rise. The sign of $y = V_S + V_B$ is ambiguous, as it reflects both an increase in V_S and a decline in V_B in response to the change in tightness. However, the impact on prices can be clearly determined from $P(X) = \beta X + (1 - \beta)V_S - \beta V_B$. Increasing V_S and decreasing V_B in response to rising tightness θ must result in higher prices for a given home.

C Appendix Figures

Figure C.1: Rate gap change vs predicted rate gap

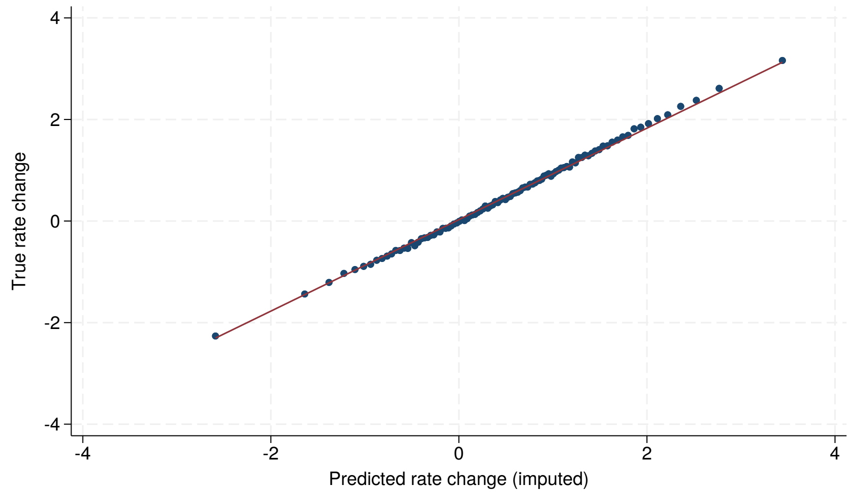

Figure C.1 plots the relationship between the rate gap we predict and the true change in rate among movers who both sell their home and subsequently purchase a new home using a mortgage. The graph is a binscatter with 100 bins, estimated on the subset of movers whose subsequent mortgage we could identify in the CRISM database.

Figure C.2: Move rates by tenure at destination

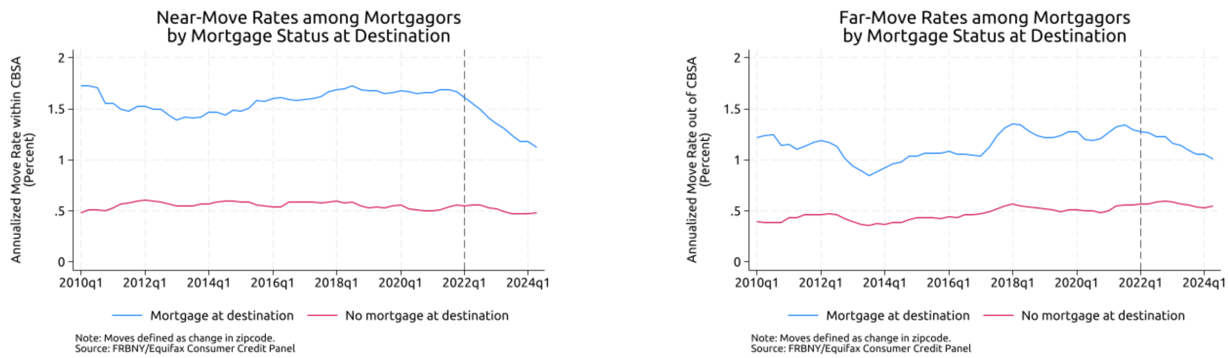

Figure C.2 plots mobility rates in the FRBNY/Equifax Consumer Credit Panel for borrowers with mortgages split by distance of destination and tenure at destination. The left panel shows annualized likelihood of moving within a CBSA split into moves where the borrower had a mortgage at the destination location (own-to-own moves) in blue and moves where the borrower did not have a mortgage at the destination (likely own-to-rent moves) in red. The decline in within-metro churn appears to be driven by own-to-own moves falling, while own-to-rent moves remain rather flat. The right panel repeats the same for moves out of a CBSA, with the small drop in cross-CBSA moves again driven by a decline in own-to-own moves.

Figure C.3: Effect of Rate Gap and Home Value Change for Within- and Between-CBSA Movers

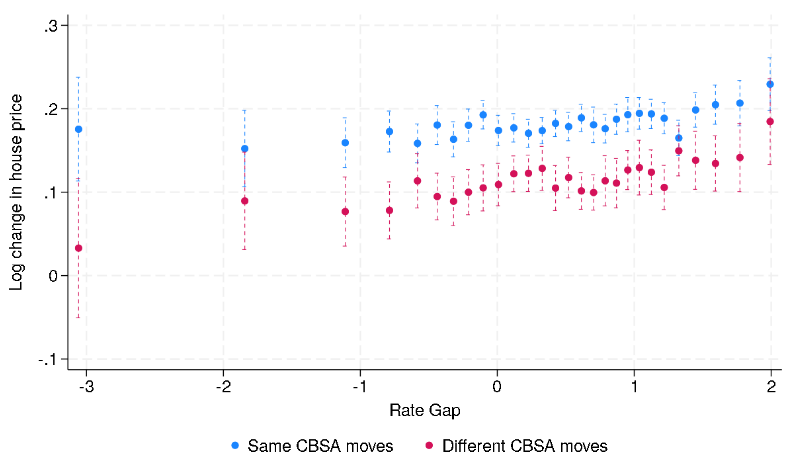

Figure C.3 plots the estimated log change in house value between movers' old and new houses. As in Figure 4, value changes are created using coefficients from Equation 2 evaluated at the mean of control variables. The rate gap is defined as the difference between the rate available to the borrower at origination and the imputed market rate for the borrower, with all the same controls as in Figure 4. Dashed lines mark 95 percent confidence bands.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, PMMS, and Corelogic Deeds

Figure C.4: Geographic Variation in Exposure to Lock-In

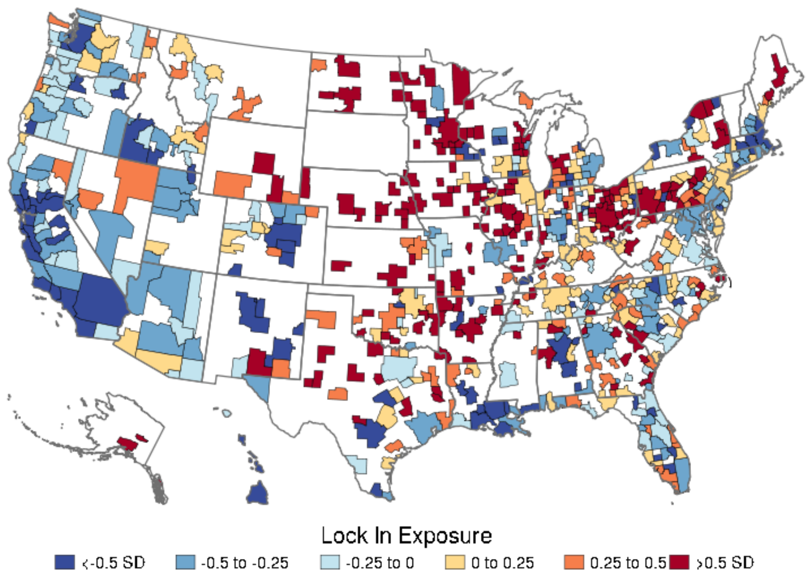

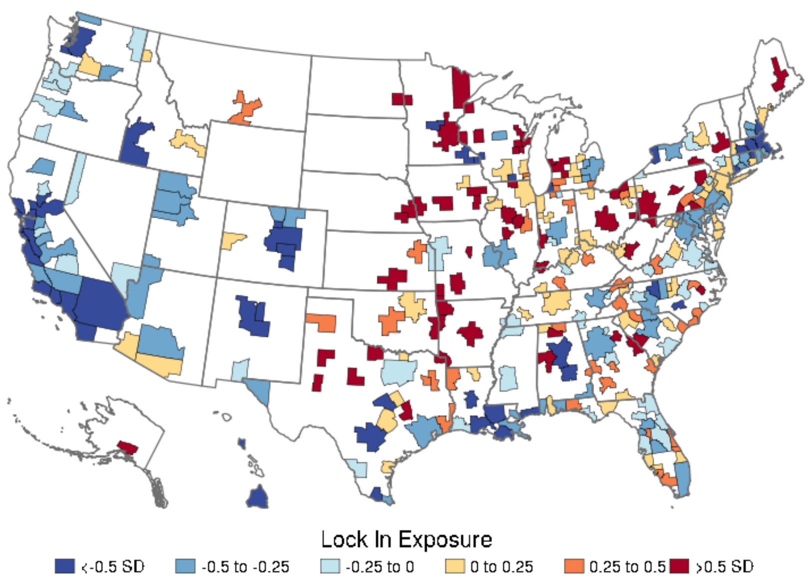

Figure C.4 plots estimated CBSA-level exposure to the Lock-In shock, as described in Section 4.2. Exposure is measured in standard deviations from the mean CBSA. The top map includes all metros in our sample; the bottom includes just the 300 most populated CBSAs. Exposure thresholds are consistent across the two maps.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, PMMS, and Corelogic Deeds

Figure C.5: Lock-in Driven House Price Growth is Strongest in Tighter, Larger Markets

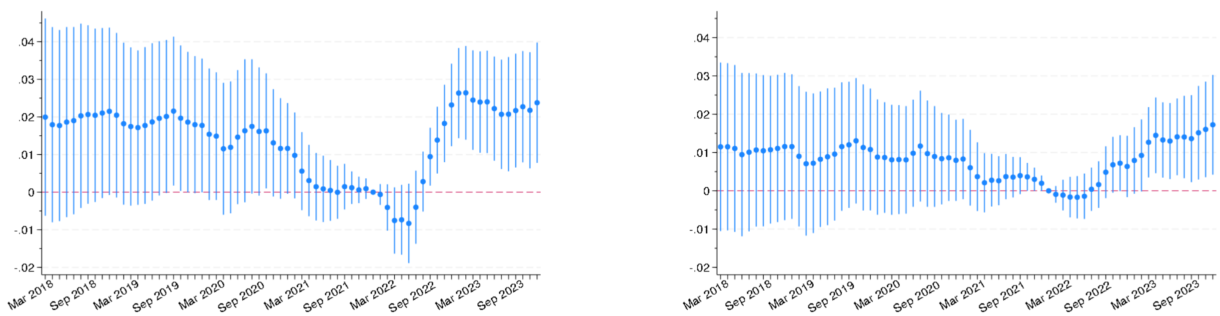

Figure C.5 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses log house prices at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. This figure includes only the top 300 most-populated CBSAs. The left panel includes CBSAs with above median market-tightness prior to the rate hikes. Market tightness is measured as the CBSAs median days on market in December 2021, according to data from Realtor.com. The right panel includes CBSAs with below median market-tightness, among the top 300 CBSAs. Magnitudes are larger for markets above median tightness. For both panels, the omitted month is December 2021, just prior to the Federal Reserve's rate hikes. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, CoreLogic HPI

Figure C.6: Lock-in Driven House Price Growth is Weakest in Looser, Smaller Markets

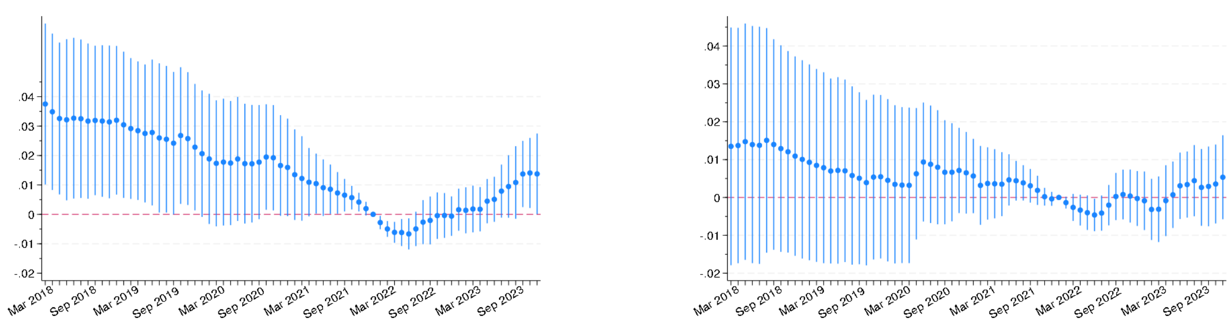

Figure C.6 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses log house prices at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. This figure includes only CBSAs outside the 300 most-populated. The left panel includes CBSAs with above median market-tightness prior to the rate hikes. Market tightness is measured as the CBSAs median days on market in December 2021, according to data from Realtor.com. The right panel includes CBSAs with below median market-tightness. Magnitudes are similar across smaller markets, regardless of tightness. For both panels, the omitted month is December 2021, just prior to the Federal Reserve's rate hikes. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, CoreLogic HPI

Figure C.7: Effect of Exposure to Lock-In on CBSA House Price Growth
Overall vs. Top300 CBSAs vs. All Others

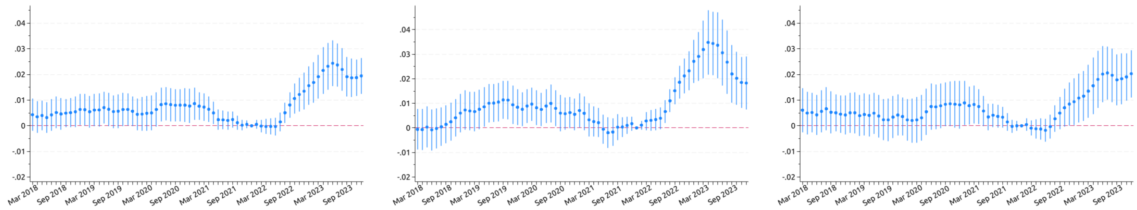

Figure C.7 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses ***log year-over-year house price changes*** at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. Effect sizes are for a one standard deviation increase in CBSA exposure to lock-in (as measured by CBSA average 'missing move' rate). The left panel includes all 699 CBSAs (out of 767 in our data) for which we have at least 500 loans in the Dec. 2021 cross-section of the Equifax CRISM dataset. The middle panel is restricted to the 300 largest CBSAs in our sample; the right panel contains all CBSAs not in the top 300. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, CoreLogic HPI

Figure C.8: Effect of Exposure to Lock-In on (log) CBSA House Prices
Overall vs. Top300 CBSAs vs. All Others

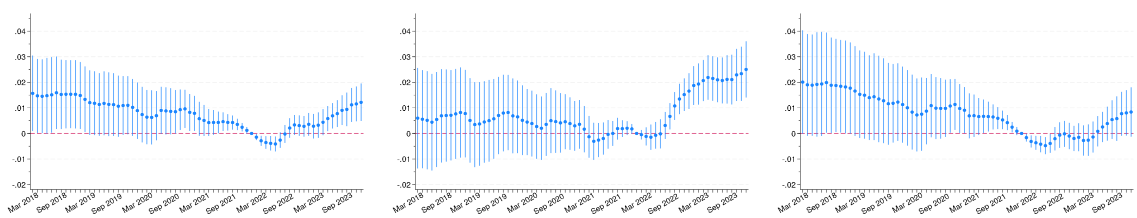

Figure C.8 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses ***log house prices*** at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. Effect sizes are for a one standard deviation increase in CBSA exposure to lock-in (as measured by CBSA average 'missing move' rate). The left panel includes all 699 CBSAs (out of 767 in our data) for which we have at least 500 loans in the Dec. 2021 cross-section of the Equifax CRISM dataset. The middle panel is restricted to the 300 largest CBSAs in our sample; the right panel contains all CBSAs not in the top 300. The coefficients can be interpreted as a percentage change from that baseline. Regressions include CBSA and month-year fixed effects. Standard errors are clustered at the CBSA level.

Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, CoreLogic HPI