Locked In:
Rate Hikes, Housing Markets, and Mobility

Aditya Aladangady, Jacob Krimmel, and Tess Scharlemann *

Click here for latest version

March 2025

Abstract

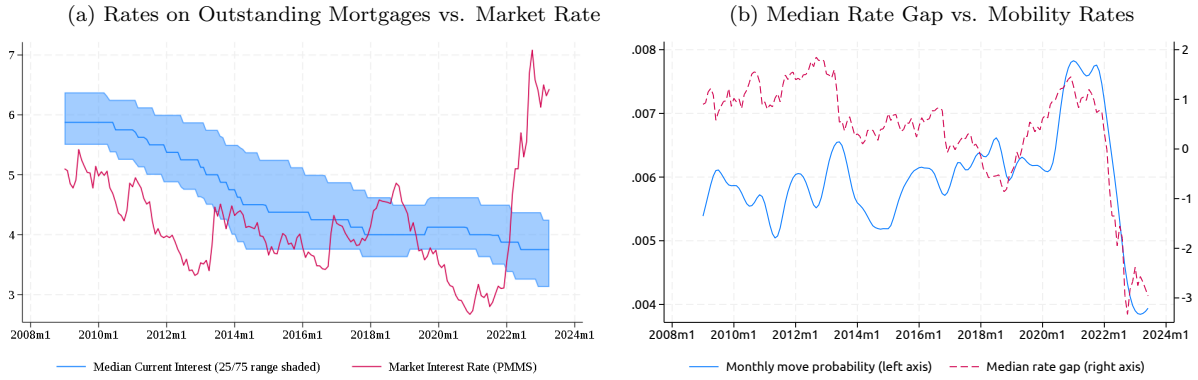
Rising interest rates in 2022 significantly increased moving costs for homeowners with low fixed-rate mortgages, leading to a sharp drop in mobility. After accounting for biases from selective refinancing, we find mortgage rate "lock in" – the decline in moves due to the rising gap between market rates and homeowners' fixed rates – explains 44% of the drop in mortgage borrower mobility from 2021 to 2022. This effect is primarily due to fewer local moves, which reduced home listings and pushed up prices. Consistent with a housing search model, we show that a lock-in shock tightens markets, driving up house prices – an effect that increases with a market's initial tightness. The model also implies the effect of lock-in grows non-linearly in shock size. We estimate the 2022 lock-in reduced time on market by 29% and increased house prices by 8%. However, these effects were entirely due to historically tight initial conditions. In a more balanced housing market, the same lock-in shock would have had little to no impact on prices or tightness.

*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.

1 Introduction

Mortgage interest rates rose by over 300 basis points in 2022, representing the largest increase in recent history (Figure 1, left panel, red line). Higher interest rates increase financing costs for all prospective home buyers, but they introduce an additional friction to moving for existing homeowners in the US, many of whom continued to enjoy low rates on existing long-term, prepayable, fixed rate mortgages (left panel, blue range). Economists have long noted the structure of fixed-rate mortgages creates incentives for borrowers to base mobility decisions (in part) on their "rate gap" – the difference between a borrower's current interest rates and market rates – which capture the savings (or costs, if negative) of financing a new mortgage of comparable size after a move (Quigley, 1987, 2002; Berger et al., 2021). Indeed, mobility fell sharply in 2022 just as rate hikes pushed market rates well above most homeowners' existing rates, lowering rate gaps into negative ranges (Figure 1, right panel).

Figure 1: Mortgage Rates and Residential Mobility



The left panel of 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. The right panel 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 Credit Risk Insight Servicing Data (CRISM), Freddie Mac Primary Mortgage Market Survey (PMMS)

The recent rate cycle has renewed interest in this topic and prompted widespread specu-

lation of 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] Importantly, any 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). However, the potential for such misallocation depends on whether lock-in curtails moves across labor market areas. Second, lock-in may severely disrupt local housing markets. While lower residential mobility would undoubtedly reduce home sales, lock-in's effect on house prices is theoretically ambiguous since a move forgone due to mortgage lock-in reduces both supply and demand in the market. From labor to housing markets, measuring lock-in's broader impacts is crucial for understanding how monetary policy passes through to the economy following sharp rates hikes.

This paper examines the link between mobility and interest rates over the most recent rate cycle and explores the broader economic consequences, using a novel dataset of mortgages, credit records, and deeds. Relative to the recent and contemporaneous literature, we offer three main contributions. First, we refine the now-common methodology used to estimate the relationship between mobility and interest rate gaps (see e.g. Fonseca and Liu (2023); Liebersohn and Rothstein (2023); Batzer et al. (2024)). We argue that the existing method overestimates the relationship between lock-in and mobility because it does not account for the role of households' private information about moving plans. Specifically, likely movers may forgo refinancing opportunities, leading them to be differentially selected towards higher rate gaps. We show that including a proxy for foregone refinancing opportunities addresses this selection bias, flattening the relationship between borrowers' rate gaps and their move probability. Even so, we find lock-in explains about 44% of the very large decline in mobility between 2021 and 2022. We also find that cross-metropolitan area moves are largely unaffected by lock-in, ruling out large effects on labor misallocation. Instead,

[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

most of the decline in mobility reflects a slowdown in local moves (or "churn"), leading us to focus on housing market effects.

Our second main contribution is to estimate the effects of the recent increases in lock-in on housing market outcomes, including prices, listings, time-on-market, rents, and building permits. We construct a metropolitan area-level (Core Based Statistical Area, or CBSA) measure of exposure to lock-in, estimated as the decline in local homeowner mobility driven by the rate increases.[2] We find that metro areas with more exposure to lock-in saw significantly higher house price growth, accompanied by substantial market tightening relative to areas with lower exposure to lock-in. More lock-in exposed areas also saw higher growth in construction permits and rents.

The reason prices rise in response to lock-in is not immediately clear, as lock-in-induced declines in local moves affect both supply and demand of homes for sale. As our third main contribution, we build a canonical housing search model to explain this apparent puzzle. The key intuition of the model is that when markets are tight - i.e. when potential buyers exceed sellers - a decline in local churn causes an equal reduction in buyers and sellers, further tightening the market- i.e., increases the ratio of buyers to sellers. A tighter market (1) allows home sellers to demand a higher price as they meet increasing numbers of prospective buyers, and (2) makes buyers–who now face more competition for a given home–more willing to pay a premium to stop searching. A key feature of the model is that the effect of lock-in on market tightness and house prices depends crucially on initial conditions: markets that are initially tighter will see larger price responses. We validate this fact empirically, showing not only that markets appear to tighten in response to lock-in–driving price responses–but also that markets that were tightest prior to rate hikes saw the largest responses to the lock-in shock.

Finally, we estimate aggregate effects of lock-in on prices. The fact that the marginal

[2]The identifying variation in this measure arises from differences in the initial distribution of rates combined with differences in how market rates evolved for the local mix of borrowers with different characteristics (i.e. time-in house, origination time, credit scores, loan-to-value ratios, etc.).

effect of lock-in on market tightness rises with the level of tightness implies a non-linearity in the aggregate effect on both tightness and prices with respect to the shock size. This strong convexity complicates calculating the aggregate house price effect of lock-in during the 2022 rate-increase. While our CBSA-level results provide insight into the effect of lock-in on prices, they exploit cross-CBSA variation which is quite small compared to the aggregate shock. As a result, estimates from cross-CBSA regressions – which reflect differences between CBSAs that were more or less exposed to lock-in – recover marginal effects of lock-in that are much higher than the average effect over the full shock.

Variation in initial CBSA-level tightness offers traction. We use this variation to estimate how the marginal effect of lock-in varies with the relevant state variable, tightness. Using these results, we are able to show that the aggregate lock-in driven reduction in moves led to a 29 percent increase in market tightness (as proxied by fewer days on market) and an 8 percent increase in home values.

To conclude, we show that unique housing market conditions before rate hikes nearly fully explain lock-in's price effects. The housing market was historically tight in 2021 (Anenberg and Ringo, 2021), with prices rising at 20% annually and homes being sold in record time. To assess how the 2022 lock-in shock to mobility would have played out under more typical conditions, we re-estimate its impact using each metro's 2019 level of market tightness as the initial condition. Strikingly, we find price effects would have been minimal — some cities would have seen more slack housing markets and price declines, while others would have experienced only minimal tightening and modest price growth. This suggests the large effects of the recent lock-in on the housing market arose from a perfect storm: rate hikes reduced mobility at a time when the market was already extremely tight, which further tightened markets and pushed prices up, partially offsetting the direct dampening effects of higher mortgage rates.

1.1 Related literature

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.[3] 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.

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. Anenberg and Ringo (2022) point out historical house prices are predominantly driven by demand shocks, which generate a positive correlation between sales volume and prices in housing search models similar to ours (see Han and Strange (2015) for further review of literature). We explore why, in the lock-in context, variations in the supply shock was important for price growth. 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. This result also offers insight on the relationship between interest rates, market tightness, and house prices in other contexts.

Our work is also closely related to Mabille et al. (2024) and Gerardi et al. (2024) who study

[3]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.

the effects of lock-in on moves up and down the housing ladder over the life-cycle. While our work focuses on understanding the impact of lock-in on metropolitan area housing markets as a whole, these models study flows in and out of market segments, providing insights into the heterogeneous welfare impacts of lock-in related disruptions to mobility.

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 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. Given the large durables expenditures driven by home purchases (Benmelech et al. (2022)), our research suggests that mobility could be an important driver of the path-dependent monetary policy documented in many of these studies.[4]. 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, well as rents, hence putting upward pressure on housing inflation.

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 which we link to property deed records. The data, which run from 2009-2023, allow us to

[4]Anenberg et al. (2023) rule out a meaningful path-dependent effect on consumption from the cash-out refinance pathway

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, house price data from CoreLogic, rents from Zillow, and single-family permits data from the Census Building Permits Survey (BPS).

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 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.[5] 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. The link provides a measure of the net purchase of housing assets, as well as a means to validate our imputed market rates measures.

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

[5]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.

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.[6]

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 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.[7]

We define a move as: (1) a persistent (at least 6 month) change in the borrower's zip code

[6]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.

[7]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.

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.[8]

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 significantly increases the unconditional mean 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,

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

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.

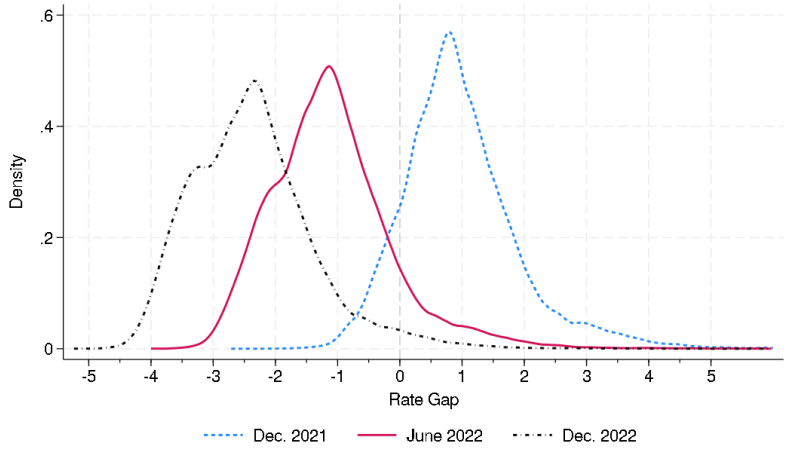
As we will discuss in the next subsection, we consider alternate versions of the rate gap which may exclude potentially endogenous variation in rate gaps. Most importantly, in our main regressions, we consider an alternate definition which utilizes an imputed origination rate based on borrower characteristics 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).[9]

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 2 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.

As we see in Figure 2, 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

[9]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 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.

Figure 2: Evolution of the Rate Gap Distribution, Dec. 2021 - Dec. 2022



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 1, from late 2022 through 2023, the fraction of borrowers facing rate increases from a move or refinance was at the highest in decades.

Our rate gap definition also appears to align with the actual difference in rates observed for movers in our data. 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 E.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.[10]

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.[11] Polynomials in current and origination LTV's and credit scores,[12] and flags for purchase mortgage status and loan terms account for potential heterogeneity in loan and credit characteristics that are 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.

[10]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.

[11]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.

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

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.[13] 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 propensities 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 $X_{i,o(i)}$

[13] 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.

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.[14]

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 approach 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.[15] 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

[14] 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.

[15] 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.

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.[16]

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 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 3 illustrates this relationship, showing how move probabilities are generally increasing in rate gaps.[17] All else equal, homeowners

[16]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.

[17]Appendix A discusses alternate specifications and provides details on how various features of our specification affect identified responses. Alternate specifications consistently yield a similar non-linear shape,

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. As market rates increase, current mortgage holders move to lower or negative rate gaps (left on the x-axis),, discouraging homeowners from selling their current home and taking on new mortgages at higher rates.

Figure 3: Mobility by Rate Gap, 2009-2023

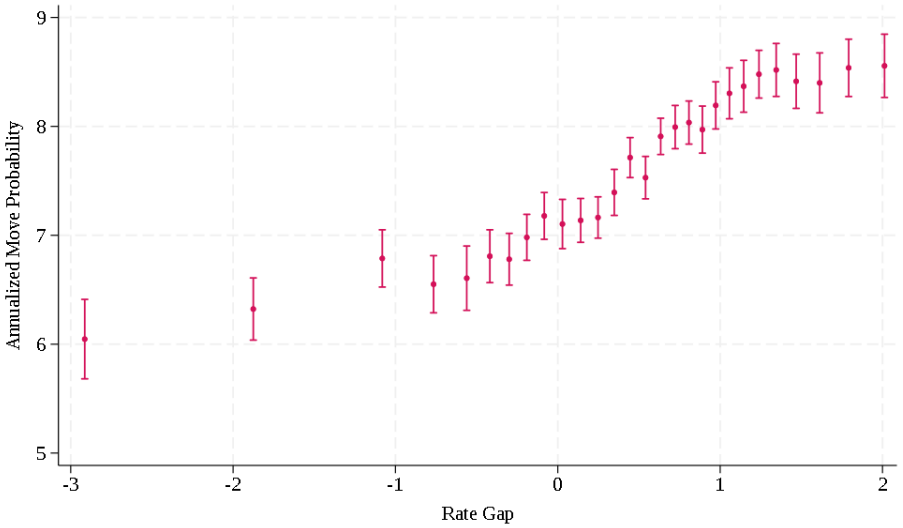


Figure 3 plots estimated monthly probability of moving conditional on borrower's rate gap using estimated regression coefficients from Equation (2) evaluated at mean of control variables. Vertical bars mark 95 percent confidence bands using standard errors clustered by origination-month cohort. Alternate 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

There are, however, two important features to note about 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.

albeit with different slopes reflecting how they address the two competing biases discussed in Section 3.1. Results appear robust to the use of alternate rate gap measures, as well.

3.3 Counterfactual Analysis: Estimating "Missing Moves"

As shown in Figure 1, rate hikes between 2022 and 2023 pushed the distribution of rate gaps down sharply just as mobility among mortgage borrowers fell. Of course the time series correlation may be driven by a number of common factors, such as the direct effect of rate hikes or pay-back for increased moves during the pandemic. The identified relationship between mobility and rate gaps in Figure 3 allows us to quantify the role that lock-in played, over and above direct effects of rates or aggregate trends in mobility.

To quantify the role of lock-in during this episode, we define a counterfactual rate gap $\tilde{RateGap}_{it}$ as of month t by utilizing prevailing market rates on new originations one year ahead, in month $t + 12$:[18]

$$\begin{aligned} \tilde{RateGap}_{it} &= \hat{rate}^{mkt}_{i,o(i)} - \hat{rate}^{mkt}_{i,t+12} \\ &= \hat{rate}^{mkt}_{i,o(i)} - \hat{\theta}_{t+12} * X_{i,t} \end{aligned} \tag{4}$$

Focusing on $t = Dec2021$, the measure provides a counterfactual distribution of rate gaps on the stock of loans active in December 2021 before rate hikes using market rates a year later in December 2022 after average market rates had risen about 4 percentage points. Using 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 total moves driven by lock-in caused by the change in rates over the t to $t + 12$ period.

$$LockInDrivenMoves_t = \sum_i \hat{B}(\tilde{RateGap}_{it}) - \hat{B}(RateGap_{it}) \tag{5}$$

where $\hat{B}(\dots)$ is our estimate of the nonlinear relationship between rate gaps and move

[18]The counterfactual distribution of rate gaps is similar to the actual movement in rate gaps we showed in Figure 2 earlier, although the counterfactual distribution holds the sample of loans constant, thereby removing compositional effects driven by refinancing activity, new originations, and moves that occur over the year.

probabilities. The measure can be seen as the number of additional moves we would have seen had the rate gap distribution remained fixed. Notably, the estimate is purged of trends in mobility, compositional changes in borrower characteristics, and direct aggregate effects rates or other macroeconomic conditions. In Section 4, we utilize the same measure scaled by the number of borrowers to compute the change in move probabilities, $LockInDrivenMoveProb_{m,t}$, for a given metropolitan area.

Appendix Figure E.2 shows the results of the simulation exercise at the national level over all years in our sample. Unsurprisingly, rate gaps contribute little to the change in mobility in most years prior to the pandemic, and we focus here on the sharp rate hike episode in 2022. In our data, about 2.9 million moves (among mortgage borrowers) occurred in 2022, down from over 4.5 million in 2021. Our estimate of $RateDrivenMoves_t$ 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.[19][20]

3.4 Evaluating Lock-In's Broader Economic Impacts

The economic implications of a sudden drop in mobility depend on the nature and purpose of moves forgone by "locked-in" homeowners, as well as any potential externalities imposed by those moves. As a start, understanding the types of moves most 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

[19] In Appendix A, we show how the share of the mobility drop explained by lock-in varies depending on specification used. Notably, our estimates imply the two competing biases discussed in 3.1 are quantitatively large, contributing 50-60 percentage points to the variation explained by lock-in.

[20] 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.

mainly reduces *within-metro* mobility, the potential for labor market inefficiencies is more limited, but the potential for housing-market effects may be greater.

In order to explore these possibilities, we estimate mobility responses to moves within and across metro areas. Two stylized facts help motivate our approach to this question.[21] First, data from the Federal Reserve Bank of New York/Equifax Consumer Credit Panel, or CCP, (a 5% random sample of anonymized credit bureau records) shows that among households with mortgages, half of moves are within 30 miles and 2/3 are within the same CBSA, suggesting households maintain access to the same labor market following most moves. Second, Current Population Survey data show that, among mortgage borrowers, reported reasons for long- and short-distance moves differ considerably: long-distance moves are more likely to be job-related, whereas moves within a labor market are more likely to be driven by housing or consumption-related factors (access to schools, home size, proximity to family, etc.).

Appendix B shows that long- and short-distance 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. Borrowers who are in-the-money stand to save on monthly mortgage payments from a refinance. The income-effect of these savings may drive 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. As we show later in the paper, reduced mobility thins the market, and in the presence of search frictions can create an externality on other buyers and sellers. However, dampening households' moves up the housing ladder within a labor market area has limited impact on labor allocations.

[21] Specific results supporting these facts are provided in Appendix B.

On the other hand, negative rate gaps could inhibit "job-driven" moves by increasing the compensation households require to move to job opportunities in other labor market areas.[22] A reduction in "job-driven" moves would lead to labor misallocation, with potential impacts on aggregate productivity and structural unemployment.

To understand the effect of lock-in on long-distance (job-driven) and short-distance (consumption-driven) moves, we return to our baseline model (as in Section 3.2) to estimate the relationship between each type of move and borrower rate gap. [23] Results are shown in Figure 4.

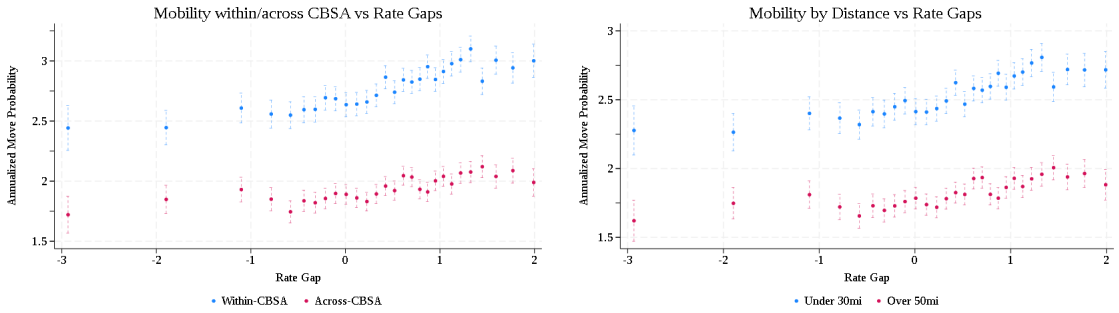
Our results suggest limited scope for large effects on aggregate labor misallocation. Figure 4 shows rate gaps primarily affect moves by changing the prevalence of within-CBSA moves (left panel, blue) or at short distances (right panel, blue), while cross-CBSA moves (left panel, red) and longer distance moves (right panel, red) are less affected by rate gaps. 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 compensate for rate gaps when moves are sufficiently productive. Results from Figure 4 is also consistent with evidence in Appendix B that long-distance moves have been roughly stable since the 2022 interest rate increases, even as local moves have declined materially.[24]

[22]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. In that setting, the constraint is driven 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.

[23]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 moves, and comprise a similar distribution of move distances to ones where we observe the destination directly. In addition, results for missing-destination moves appear quite consistent with an appropriately weighted average of results from observed-destination moves to near and far locations.

[24]A secondary implication of this finding is that lower (and negative) rate gaps result in the marginal mover 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,

Figure 4: Moves Within and Across Labor Market Areas versus Rate Gap



The left panel of Figure 4 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

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 across moves within CRISM described in Section 2, we run a regression similar to the specification in Equation 2, with the log change in home value between the sold and purchased homes as the dependent variable.[25] Appendix Figure E.4 shows the results of this regression. There are two key takeaways. First, 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

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 4. Indeed, this is consistent with the flattening of the move rate slope at low and negative rate gaps in Figure 3.

[25]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

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). Second, 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 decision to move.

4 Effect of Lock-In on Housing Markets

As described in the prior section, the sharp increase in mortgage rates in 2022 led to a sudden reduction in local mobility, which likely has impacts on metro-area housing markets. Owners who elected not to move within a local housing market market lowered not only demand for housing they would have bought, but also the supply of homes available for sale. Indeed, in 2022 following rate hikes, the aggregate number of home sales fell sharply and inventory dried up, though the impact of this reduction in sales on prices is more uncertain.

We begin by estimating the response of house prices to a reduction in local churn by comparing prices across CBSAs that differed in their exposure to lock-in before developing a housing search model that helps explain the mechanisms through which churn impacts prices. The model provides a framework to understand both heterogeneity in price responses across housing markets as well as a means to quantify aggregate impacts of lock-in on house prices.

4.1 House Prices and Housing Market Activity

To examine the lock-in effect on local housing markets, we first construct a CBSA-level measure of "lock-in driven move probability" similar to the aggregate lock-in driven move calculation described 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.[26] To the extent that our estimates recover a structural relationship between

[26]Unlike the aggregate measure which counts the total *number* of moves held back by lock-in, the CBSA-level exposure is a decline in the *average probability* of moving, hence normalized by the loan count in the

rate gaps and mobility, the resulting rate-gap driven decline in moves provides a sufficient statistic for each metro area's differential exposure to lock-in.

Our analysis reveals widespread exposure to lock-in, as nearly all metros experience a significant lock-in-driven decline in moves. Though the aggregate shock is much larger than cross-metro variation does exist.[27] Variation across metro areas reflects differences in how rate gaps for local borrowers were shifted across the mobility curve in Figure 3. 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, as mortgage rate increases may be less pivotal i mobility decisions for borrowers already at lower points in the rate gap distribution.[28]

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 5).[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 6. We find a modest positive effect of exposure to lock-in on CBSA-level house price growth. In our preferred specification, a one standard deviation increase in

area.

[27]See Appendix Figure E.5 for map of lock-in exposure. While we establish trends in housing market outcomes were similar across metro areas with higher and lower exposure to lock-in, areas with larger predicted declines in churn are somewhat smaller and less expensive.

[28]To see this, note the slope is much steeper from 1.5 to -0.5 than from -0.5 to -2.5 such that a 200 basis point increase in rates would lower mobility more among those with slightly positive rate gaps ex ante.

[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 5: Effect of Lock-In on CBSA-level Active Listings

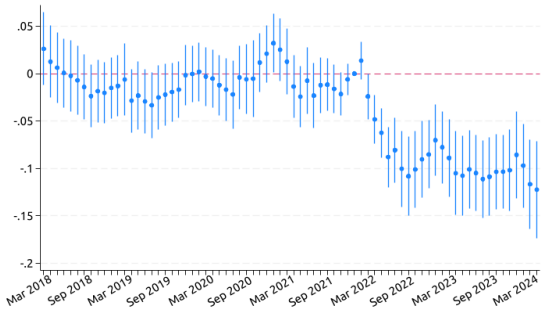


Figure 5 plots coefficients from two-way fixed effect difference-in-difference style regression. The figure shows the effect of a one standard deviation increase in CBSA exposure to lock-in (as measured by CBSA average 'missing move' rate) on the natural log of monthly active real estate listings in 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

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 4), 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 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 lock-in related prices increases are driven by the tightening of already-tight housing markets.

Figure 6: Effect of Exposure to Lock-In on CBSA House Price Growth

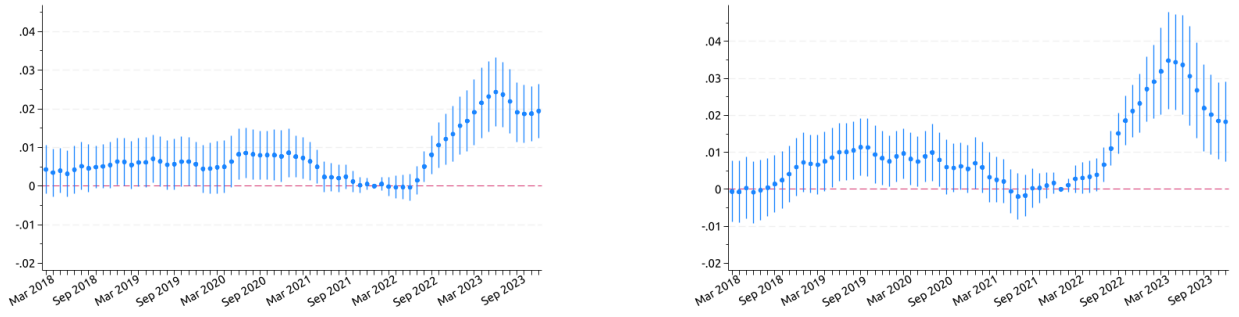


Figure 6 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses log year-over-year house price changes at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. Effect sizes are for a one standard deviation 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

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, which we confirm in the data. Specifically, we show lock-in driven declines in churn not only led to increases in home prices, but also in rents and construction as new buyers and sellers responded to the tightening market. In addition, we show responses depended on intial conditions with ex ante tighter markets (those that favored sellers more prior to lock-in) saw larger price effects. This heterogeneity also causes large shocks to have disproportionately large impacts as lock-in driven tightening in markets raises the marginal effect of additional tightening. We exploit heterogeneity in initial tightness across cities to inform the aggregate effects of the lock-in shock.

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 $q(\theta)$ (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 $X - P - V_B = (1 - \beta)(X - y)$.

[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.

[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.

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

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

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

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 $\frac{d}{dt}\frac{B}{S} = 0$

[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.

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 C 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 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 C 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). 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.[35]

[35]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

In reality, demographic forces and pandemic-era demand likely resulted in potential buyers exceeding sellers in 2021, causing $\theta > 1$ prior to the 2022 rate hikes. In such a setting, the lock-in shock that reduces both buyers and sellers by similar amounts *raises* θ further away from 1.[36][37]

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.

The model therefore provides some testable predictions in the data. First, the model suggests increased tightness is a key mechanism by which lock-in raises house prices. If the decline in churn did indeed cause the market to tighten, thereby raising prices, we should see empirical measures of market tightness rise in markets that were more exposed to the lock-in shock. Second, all the arguments here hinge crucially on whether the market favored buyers or sellers prior to the rate hikes, and the extent to which this was true. Put differently, the model provides a clear testable prediction that responses to lock-in depend on the ex ante 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). The same is true of prices: prices rise more sharply in markets that were tighter prior to the shock. In the next section, we test these predictions in the data.

The fact that the effect of lock-in depends on the initial level of tightness also shows that

second term.

[36] 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 $\theta = 600/400 = 1.5$.

[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 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.

large shocks to churn–such as the one seen in 2022–can have a non-linear effect. As churn falls, the market tightens, increasing the marginal effect of additional declines in churn. As a result, a large shock will have a disproportionately larger impact than a smaller shock. We return to this point when considering aggregate effects in Section 4.4.

Before turning to empirical tests, 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 in recent history. 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.

4.3 Lock-In and Housing Market Tightness

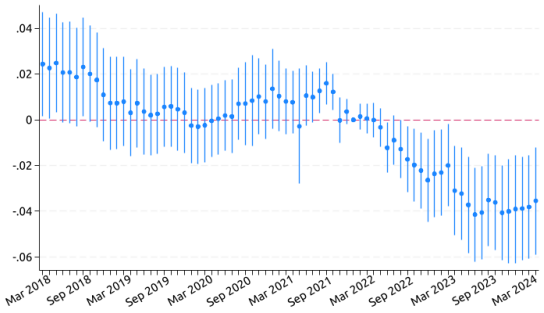
The search and matching framework yields some testable implications in the data. First, the model suggests a key mechanism driving the link between churn and prices is the fact that reduced churn in an already-tight (seller's) market, further increases tightness. While we cannot measure the number of 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.

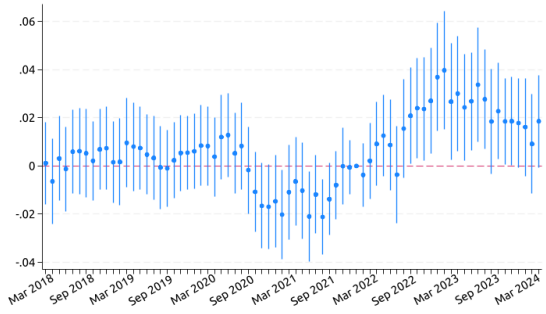
Figure 7 shows how each of these measures responds to lock-in, utilizing a specification similar to the one used earlier. Looking through seasonality in sales, metro areas that were more exposed to lock-in saw properties stay on market for 5 percent shorter times than prior to the shock (top panel). 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 7: Effect of Lock-In Exposure on Housing Market Tightness

(a) Log Median Days on Market



(b) List Prices, YoY change



(c) Log Price Reductions

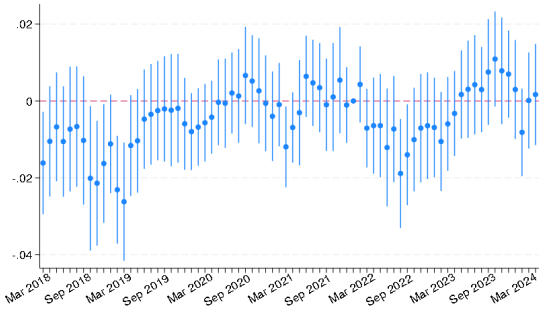


Figure 7 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses various CBSA-level measures of market tightness on CBSA-level exposure to lock-in. The top panel (a) plots the effect of lock-in on the log of the 12-month moving average of median days 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 8, 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).

4.4 Aggregate Price Effects of Lock-in

How much did lock-in affect house prices altogether? A simple approach would be to re-scale event study results by the magnitude of the overall shock. For example, Figure 6 shows that prices rose by 2.5 percent per standard deviation of cross-CBSA exposure to lock-in. However, the cross-CBSA variation in lock-in exposure is only about 1/10th as large as the aggregate shock, and a simple rescaling would imply lock-in drove a rather implausible 25 percent increase in house prices between 2021 and 2022.

The reason for this implausible effect size is twofold. First, the effect of lock-in driven declines in local churn are non-linear in the size of the shock. Equation 7 shows the marginal effect of lock-in on tightness rises with the level of tightness. Just as CBSA's with initially tighter markets saw larger effects, a given housing market may see increasingly large effects as tightness rises. Consider the aggregate lock-in shock of 2022 as the sum of smaller shocks. The nonlinearity we see in initial conditions implies that the aggregate price-impact of the

Figure 8: Lock-in Driven House Price Growth is Stronger in Tighter Markets

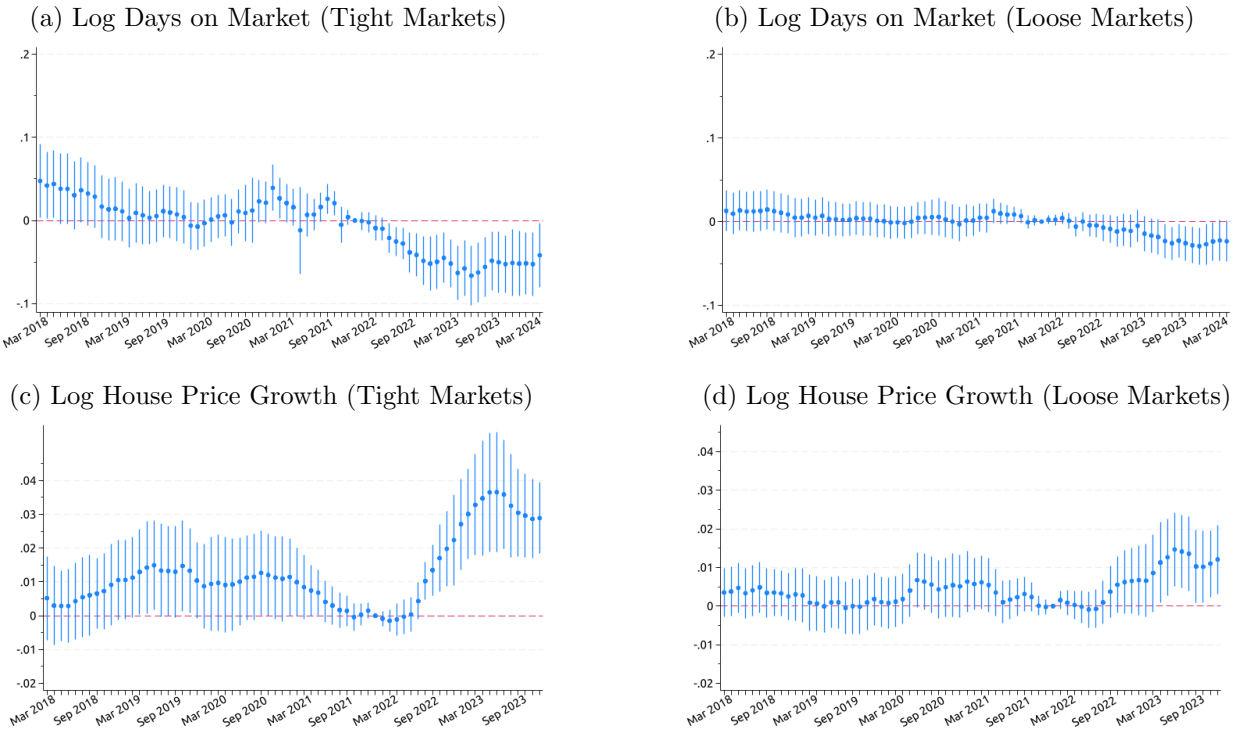


Figure 8 plots coefficients from a two-way fixed effect difference-in-difference style specification. Panels (a) and (b) regress log of the 12-month moving average of median days on market (according to Realtor.com) on CBSA-level exposure to lock-in. Panels (c) and (d) do the same for the log year-over-year 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

first n basis point increase in mortgage rates was much smaller than the price-impact of the final n basis point-increase. This is because each incremental rate increase tightens the housing market, pushing up the marginal tightening effect of the subsequent n basis-point rate increase.

Second, the cross-CBSA event study results provide a local average treatment effect with respect to the cross-CBSA variation in lock-in exposure. Because all CBSA's experienced a large decline in mobility, with some variation across CBSAs, the effect identified is the marginal effect of additional lock-in exposure from a CBSA with a somewhat large decline in

mobility to one with a very large decline in mobility. Figure 9 provides a stylized picture considering a simple difference-in-difference between two cities experiencing large, but slightly different sized shocks. The diff-in-diff recovers the slope of the true relationship between the low-shock (blue) and high-shock (orange) city. Because marginal effects rise with tightness and shock size, the estimated slope is much larger than the average slope across the full range of the shock, resulting in implausibly large effects when scaling (purple dashed line).

Figure 9: Stylized Example: Local Estimate to Large Lock-in Shock with Convex Aggregate Response

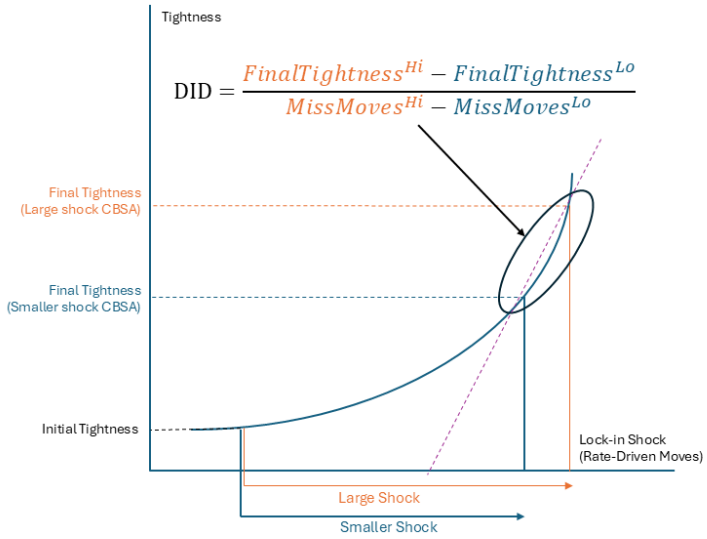


Figure 9 shows a stylized depiction of our baseline estimate. In the example, two cities begin at the same initial tightness and experience a large shock to churn, though one city has a larger shock (blue) and the other has a smaller shock (orange). Cross-CBSA estimates reduce to a simple difference-in-difference in the two-city case, and recover a LATE for the slope shown in purple reflecting variation in the circled range. Due to convexity in the aggregate effect (ie, marginal effects rising in tightness), the LATE recovered by the DiD (purple) is much steeper than the average slope over the range of the shock.

A key feature is that the reason for the curvature in the marginal effect reflects the changing level of tightness as a result of the shock. As shown in Figure 8, cities varied in initial levels of tightness, providing us with a way to estimate how the marginal effect evolves with tightness. Specifically, we use the log of a CBSA's median days on market as a proxy

for θ, and consider a linear approximation of Equation 7.[38]

$$\Delta lnDOM_m = (\kappa_0 + \kappa_1 lnDOM_m^{init}) * \ell_m + \gamma lnDOM_m^{init} + u_{mt} \tag{9}$$

The change in tightness in CBSA m from December 2021 to December 2022 ($\Delta lnDOM_m$) depends on the size of lock-in exposure ℓ_m as a linear function of initial tightness ($\kappa_0 + \kappa_1 lnDOM_m^{init}$).[39]

Notably, even though the regression still utilizes cross-CBSA variation in ℓ_m, the *change* in the marginal effect driven by the level of tightness is captured up to a linear approximation. As such, the regression provides globally valid estimates $\hat{\kappa}_0, \hat{\kappa}_1$ which tell us how the marginal effect changes across the range of $lnDOM$ in the data. We can then estimate aggregate effects on tightness that account for the changing marginal effect by solving the differential equation given by a linear approximation of Equation 7:

$$lnDOM_m(\ell_m) - lnDOM_m(0) = \left(lnDOM_m^{init} + \frac{\kappa_0}{\kappa_1} \right) \left(e^{\kappa_1 \ell} - 1 \right) \tag{10}$$

We translate the aggregate effect on tightness into prices by estimating an elasticity of prices with respect to tightness. As we have argued, the reason prices change in response to mobility is due to rising tightness. Therefore, the lock-in driven reduction in mobility provides a valid instrument to estimate impact of changes in tightness (log days on market) on price growth over this period, which we find to be about -.277.[40]

Results suggest lock-in driven reductions in mobility had an overall impact of reducing days on market by about 29 percent on average across CBSAs, with the interquartile range

[38] See Appendix D for additional details and intermediate results.

[39] The regression is nearly identical to the event study, but isolates the 12 month horizon akin to a local projection where the effect varies linearly in $lnDOM_m^{init}$. We include a direct effect of initial tightness (and a constant) to account for potential trends in $lnDOM$ that may have varied across cities, akin to separate time fixed effects for tight and slack markets in the event study. As we show in the appendix, cities that had experienced sharper market tightening during the pandemic saw a faster unwinding as tightness between 2021 and 2022. This mean-reversion in tightness is accounted for by the constant and direct effects of $lnDOM$.

[40] See Appendix D additional details.

across cities ranging from 16 to 40 percent declines. This tightening led to an 8 percent increase in prices on average across the country, with an interquartile range running from 4.5 to 11 percent.[41]

A challenge for external validity is that the housing market in 2022 was historically tight (Anenberg and Ringo, 2021), and our model suggests that initial housing market conditions are extremely important. Before the interest rate increases, house prices were rising at about 20% per year, and new home listings were being sold in record time (average of 45 days on market across CBSAs in 2021q4 compared with 65 days on market in fourth quarters of years prior to the pandemic).

In addition to understanding the effect of lock-in during the 2022 episode, our approach allows us to construct a counterfactual estimate of how a similar sized decline in local churn would have impacted markets under alternate distributions of initial tightness. As a concrete example, we consider how the 2022 lock-in shock would have impacted housing markets had it occurred under market tightness conditions in 2019. For context, average days on market across CBSAs in the fourth quarter of of 2021 was about 45.8 days. By comparison, homes were on market an average of 65 days in the fourth quarter of 2019 was about 65.0 days, comparable to the same quarter in the prior several years. Repeating our aggregate effects exercise using the distribution of days on market in 2019 across cities suggests a 2022-sized decline in churn would have had a rather negligible effect on home prices on average, with cities at the 25th and 75th percentiles seeing a roughy 3 percent decline or increase in prices respectively. This suggests, based on our estimates, markets were somewhat balanced between buyers and sellers in the years leading up to the pandemic.[42]

[41] All aggregate effect statistics are weighted by owner-occupied units across CBSAs.

[42] As described in Appendix D, $-\kappa_0/\kappa_1$ provides a threshold for log days on market that determines whether impacts of lock-in raise or lower prices, correspond to the model's threshold $\theta = 1$. In our baseline estimates, this corresponds to approximately 64 days on market, very close to the average days on market in 2019q4. Note that this occurs despite our estimate not utilizing any information about market tightness prior to 2021q4.

4.5 Spillovers to Rents and Construction

While we have not explicitly modeled rental markets above, the model offers some insights into the impacts of lock-in on rents. We extend the model to formally include a rental and construction sector in Appendix C.3, but discuss intuition for the implications on rents and construction here.

In particular, construction may respond to tighter housing markets, impacting the supply of housing and thereby prices. As the value of search for sellers rises, builders may have an incentive to provide new supply until marginal costs of new construction match the value of search. The resulting increase in sellers in the market reduces tightness, mitigating the impact of lock-in on prices.

In addition, 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 a reduction in the inflow of new buyers mitigates the response in prices, the reduced outflow from rentals would raise demand for rental housing, thereby raising rents.[43]

Notably, both the construction and rent-to-own margins provide "relief valves" for the increase in market tightness in the owner-occupied market. In both cases, as lock-in reduces churn and tightens an already-tight market, endogenous increases in the entry of sellers or decrease in the entry of buyers dampens the impact of the initial shock. This is evident in the second term of Equation 8 where equilibrating forces in buyer and seller entry dampen the direct price effect from the first term. In the appendix, we endogenize buyer and seller

[43]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 E.3 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.

entry by modeling the rent-to-own and construction decisions.

Using rental data from Zillow and construction permits from the Census Building Permits Survey (BPS), we repeat our CBSA-level analysis of lock-in effects to test the predictions above. The left panel of Figure 10 shows year-over-year growth in rents rose about 2 to 3 percentage points higher in CBSAs more exposed to lock-in.[44] Decline in churn put upward pressure on rents as marginal renters faced higher prices in the owner-occupied market when considering transitioning to home-ownership. In the right panel, we see permitting for single-family homes also rose in areas more exposed to lock-in as low inventories and higher prices incentivized builders to increase construction.

Figure 10: Lock-in Spillovers to Rental and Construction Markets

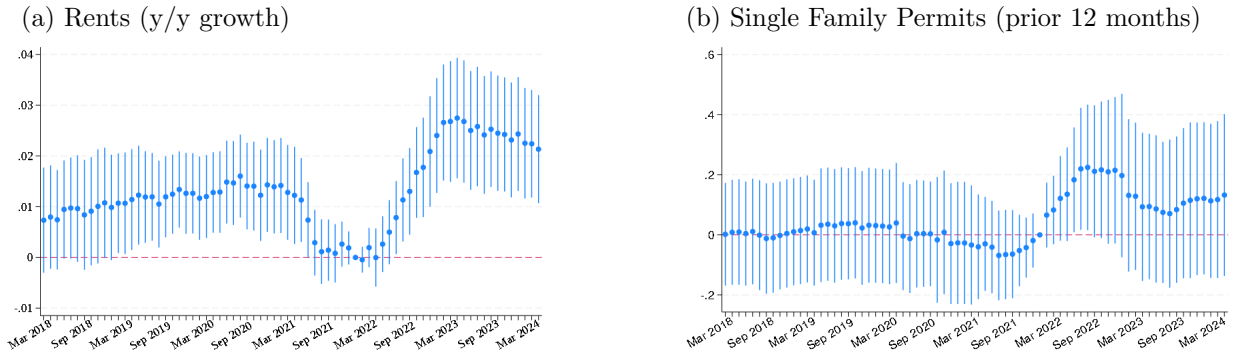


Figure 10 plots the effects of lock-in on rents (left panel) and single-family permits (right panel). Year-over-year growth in rents rose about 2 to 3 percentage points higher in areas with higher lock-in exposure as tighter owner-occupied housing markets disincentivized renters from transitioning to ownership. Rents figure shows 12 month changes in Zillow rental index for all rental units 467 CBSAs accounting for 93 percent of loans observed. Permits figure uses total single family permits from prior 12 months from Census's Building Permits Survey for 712 CBSAs accounting for 99 percent of loans observed. Permits series are normalized by historic annual permitting rate for CBSA to account for differences in permitting rates across places.
Source: Authors' calculations using ICE/McDash and Equifax CRISM, Corelogic Deeds, Freddie Mac PMMS, Zillow Rents, Census Building Permit Survey

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

[44]We plot the series for all rental units, not just single-family units since deterred new home buyers may place upward pressure on rents in all units. Results restricting to rents on single-family units are similar to those plotted here.

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, and quantify the implications of reduced mobility on both labor and housing markets.

Our work builds on existing literature to further refine the identification of mobility responses. Our approach not only addresses biases arising from highly mobile skilled workers receiving systematically lower borrowing rates, but we also quantify and address an additional bias arising from selective refinancing leading likely movers to be differentially selected into higher rate gap ranges. Addressing the two issues jointly suggests lock-in is an important driver of the recent decline in mobility, albeit somewhat less than prior estimates suggest.

Even so, we find sizeable impacts from lock-in driven decline in moves, particularly on housing markets. Using the 2022 rate hikes as a laboratory, we quantify spillovers from the resulting large shock to mobility. While lock-in appears to have quantitatively small impacts on cross-labor-market mobility–suggesting effects on labor misallocation may be limited–the decline in within-metro moves has sizeable impacts on housing markets as reduced local churn tightened housing markets, pushing up prices by 8 percent. The market tightening may go a long way to helping reconcile how house price growth remained positive despite the sharp increase in borrowing costs.

Importantly, the sizeable impact impact of lock-in on house prices during this episode reflects a confluence of factors. Our work highlights that the impact of reduced mobility on housing markets depends crucially on the tightness of the housing market prior to the shock. In addition, we show that effects of large shocks can be disproportionately large as incremental market tightening raises the marginal effect of further declines in churn. Rate

hikes in 2022–following a period of low rates–led to a large lock-in effect on mobility in a housing market that had tightened significantly over the pandemic.

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.

Anenberg, E., Scharlemann, T., and Van Straelen, E. (2023). Borrowing and spending in the money: Debt substitution and the cash-out refinance channel of monetary policy. *FEDS https://doi.org/10.17016/FEDS.2023.073*.

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

Benmelech, E., Guren, A., and Melzer, B. T. (2022). Making the House a Home: The Stimulative Effect of Home Purchases on Consumption and Investment. *The Review of Financial Studies*, 36(1):122–154.

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.

FRBNY and Equifax (2024). Federal reserve bank of new york/equifax consumer credit

panel.

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.

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.[45] 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 "lock-in-driven moves"

[45]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

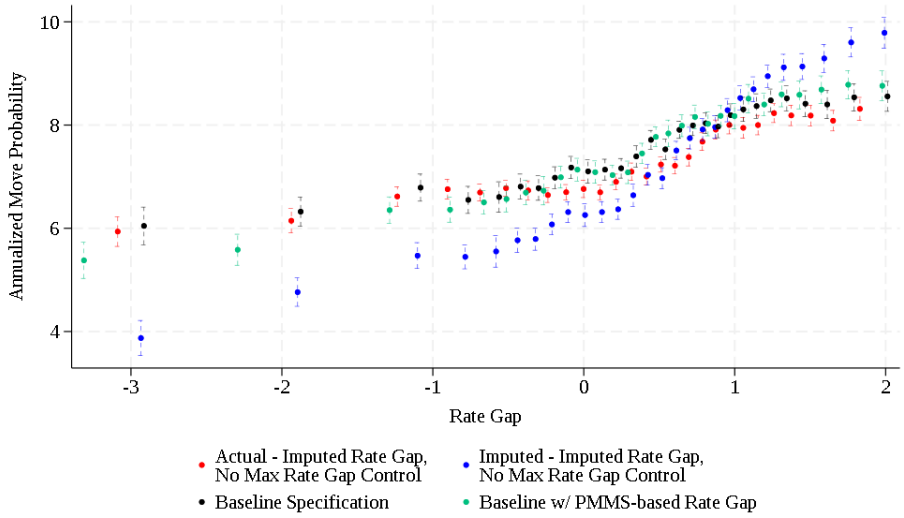


Figure A.1 plots estimated monthly probability of moving conditional on borrower's rate gap using estimated regression coefficients from Equation (2) evaluated at mean of control variables. The plot shows several alternate specifications compared to our baseline specification in black. Red dots are estimated 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.[46] 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

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

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

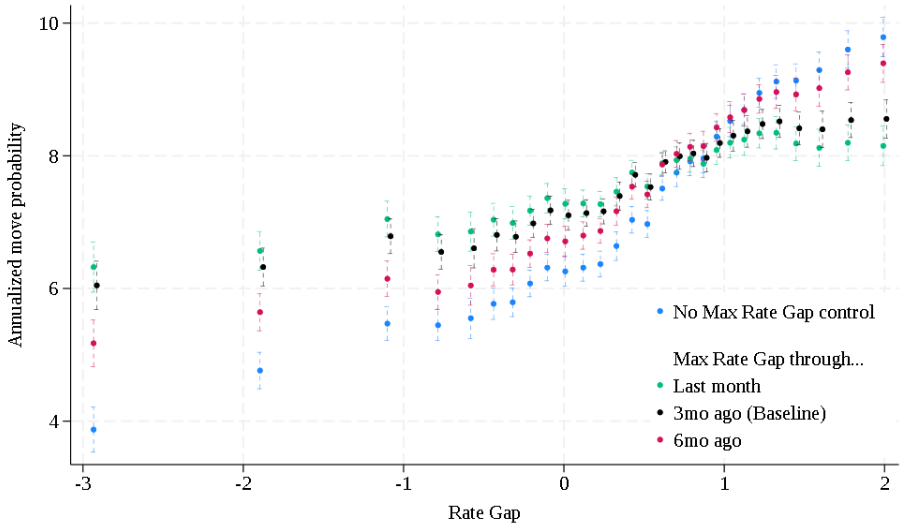


Figure A.2 plots estimated monthly probability of moving conditional on borrower's rate gap using estimated regression coefficients from Equation (2) evaluated at mean of control variables. The plot shows several alternate definitions of our maximum rate gap control compared to our baseline specification 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 lock-in-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 Move distance and move purpose

Using several datasets, we document key stylized facts about mobility, focusing on the distinction between job-related and non-job-related moves. Figure B.3 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 B.3 shows the distribution of move distances among borrowers with mortgages in the Federal Reserve Bank of New York/Equifax Consumer Credit Panel (hereafter, CCP), a data set comprised of a 5 percent sample of anonymized credit bureau records in the US (FRBNY 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.

We also document aggregate trends in long and short-distance 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 B.4 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 B.3 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.

Figure B.3: Heterogeneity in Move Distances and Reasons

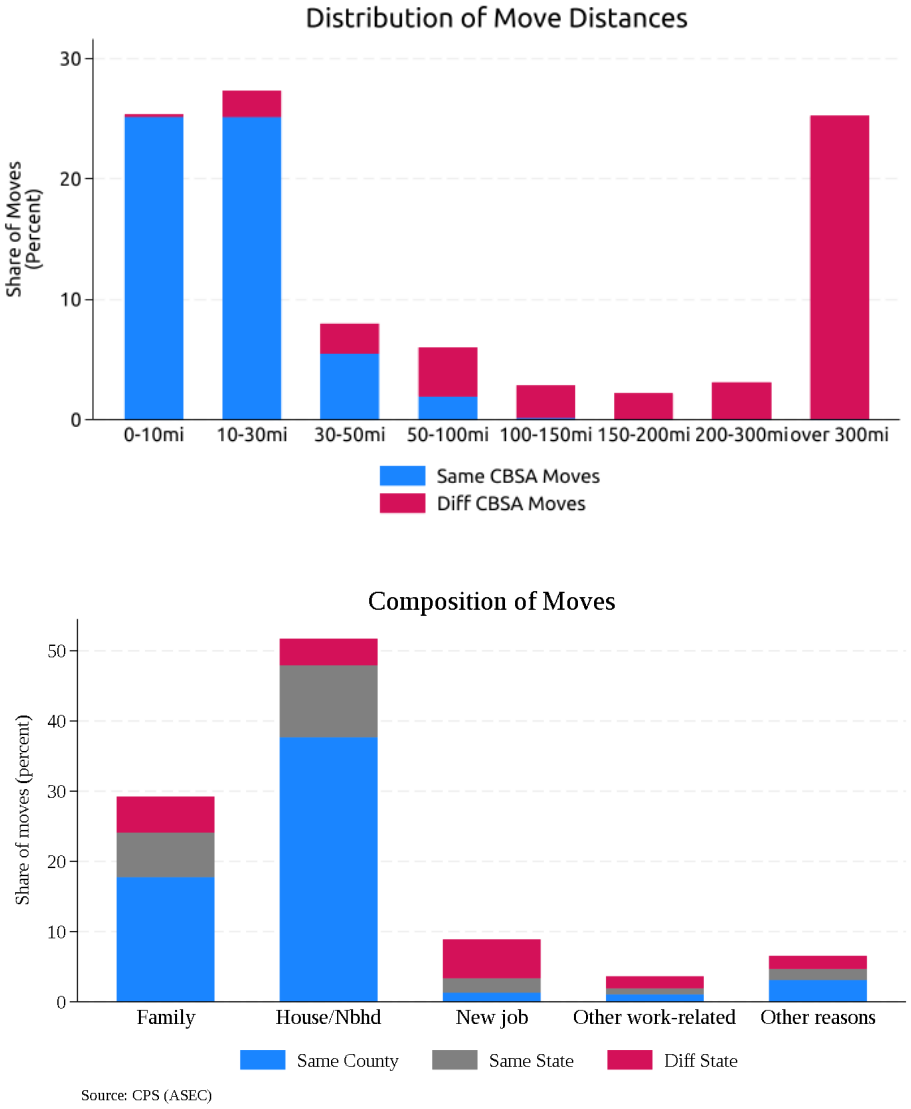


Figure B.3 shows heterogeneity in move distances and reasons. The top panel shows the distribution of move distances among households with mortgages. About 1/2 of moves are within 30 miles and 2/3 are within the same CBSA, likely providing access to the same jobs as before. Few moves occur to intermediate 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

Figure B.4: Time Series of Moves by Type

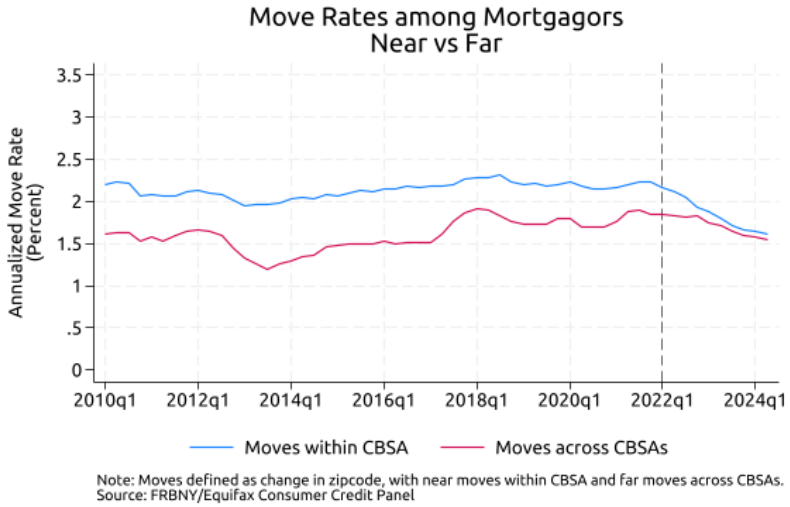


Figure B.4 shows moves (zip code changes) in credit records for borrowers that had a mortgage prior to the move. The figure splits moves into those within CBSA (blue) and across CBSAs (red), and shows aggregate declines in moves over 2022 were driven primarily by short-distance (within CBSA) moves, which tend to be driven by consumption-related reasons. Across-CBSA moves remained relatively flat.
Sources: FRBNY/Equifax Consumer Credit Panel

C Model Appendix

C.1 Steady-State Equilibrium

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

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

$$(\alpha_S + \gamma_S V_S)\theta = \alpha_B + \gamma_B V_B$$

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 the seller's value of search (6a) is upward sloping in (θ, y) space and the buyer's value of search (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...

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

C.2 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 x_i are all positive and z has the same sign as $\theta-1$. In the following, we assume

$\theta > 1$ such that z is 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 $\theta - 1$, which determines the sign of z above. 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.3 Incorporating rental and construction markets

Our baseline model does not explicitly consider renters or home builders, though both may interact with or influence the impacts of lock-in on the owned housing market. Broadly, our model assumes the entry of buyers and sellers is linear in the value of buying or selling, but does not account for the trade-offs that marginal buyers and sellers make when deciding whether to enter into search. On the buyer side, first-time home buyers likely weigh the

trade-off between home search (V_B) and the value of remaining a renter. At the margin, tighter markets may deter some first-time home buyers, who remain renters, pushing up rents. Similarly, home builders weigh the marginal cost of constructing a house with the value of searching for a buyer for that house (V_S), leading them to respond to low inventories and a tight housing market by building new homes. In this section, we consider extensions to our baseline model to account for these channels.

Renters

In the baseline model, we assume the entry of new buyers is linear in the buyers' value of search. In reality, the marginal *new* home buyer weighs the trade-off between searching for a home to buy V_B and remaining a renter V_R, which is a declining function of rent. This allows us to describe the flow of renters who become first time home buyers:

$$\begin{aligned} \dot{B} &= \alpha_B + \gamma_{oob} V_B + \gamma_{nb} \left(V_B - V_R \right) \\ \dot{R} &= \alpha_R - \gamma_{NB} \left(V_B - V_R \right) \\ V_R &= u_R - \rho_R \\ \rho_R &= \rho_0 R^{\epsilon} \end{aligned} \tag{14}$$

In the modified version of the entry conditions above, we see owner-occupier-buyers (oob) respond linearly to the buyer's value of search, as before, but the flow of new buyers (nb) is linear in the *difference* between the buyer's value of search and the value of remaining a renter.[47] In addition, the mass of renters evolves based on exogenous net household formation at α_R less the outflow of new home buyers $\gamma_{nb} \left(V_B - V_R \right)$. We assume the value of renting for a period is simply the utility of renting a home u_R less the flow rental cost ρ_R. To close the market, we assume the supply of rental housing is increasing in rents with a constant

[47] The extended model nests our baseline as a special case. Equation 14 can be rearranged such that $\dot{B} = (\alpha_B - \gamma_{nb} V_R) + (\gamma_{oob} + \gamma_{nb}) V_B$. Letting $\gamma_B = \gamma_{oob} + \gamma_{nb}$ maps the entry condition back to the baseline model with buyer entry being linear in V_B. The key difference is that the value of renting V_R appears in the constant inflow term $(\alpha_B - \gamma_B R$. If the value of renting is a constant – such as with free entry of new rental units – we revert to the baseline model.

elasticity $1/\epsilon$.

We already showed that lock-in tightens the market, reducing V_B as buyers must compete over a thinner stock of available homes. The marginal new buyer is deterred from searching for homes, and thereby remains a renter. Holding rents fixed, the reduced outflow of renters ($\dot{R}$) increases the demand for rental housing (R). In a setting where rentals are supplied perfectly elastically ($\epsilon = 0$), we return to our baseline model setting: rental costs are exogenous at ρ_0 and rental markets fully absorb the decline in new home buyers by supplying additional rental housing. However, the story differs when rental supply is not perfectly elastic: higher demand for rental housing pushes up rents, diminishing the outside option of renting for the marginal buyer. Relative to our baseline setting, the impact of the lock-in shock on $\theta = B/S$ is diminished as higher rents push some new home buyers back to the purchase market. Therefore, rental markets provide a "relief valve" on the tightness in the home buying market: the decline in new home buyers reduces the impact of lock-in on tightness in the home buying market as well as its resulting effect on the price of owned homes. Instead, the lock-in shock passes through to rents. The extent of the pass-through to rents is determined by a combination of the supply elasticity of rental housing and the sensitivity of new home buyers to rents.

Home builders

In the baseline model, we also assume the entry of new sellers is linear in the sellers' value of search. In reality, new home construction may respond as lock-in reduces inventories, tightens the housing market, and raises the value of being a seller. In response, builders may find it profitable to build additional housing, at least if they can place the homes on market before they expect the tightness to subside.

We modify our model to include a construction sector with convex costs in investment in new housing. Specifically, the construction sector can create new homes I at a cost $C = \frac{\phi}{2} I^2$. Newly constructed homes I are placed on sale, adding to the net inflow of sellers

in the market.

Since V_S is declining in the number of sellers, builders will continue to build and put houses on market until their marginal cost of construction equals to the value of becoming a seller:

$$I = \frac{1}{\phi} V_S$$

The net inflow of sellers becomes the sum of both existing owners movers (as in the baseline model) and new investment:

$$\begin{aligned} \dot{S} &= (\alpha_S + \gamma_S V_S) + \frac{1}{\phi} V_S \\ &= \alpha_S + \left(\gamma_S + \frac{1}{\phi}\right) V_S \end{aligned}$$

Effectively, the inclusion of the construction sector impacts how sensitive the inflow of sellers is to a given change in V_S. In a setting where marginal costs of building (ϕ) are very high and supply of new homes is perfectly inelastic, changes in V_S have no impact on the behavior of builders and we return to the setting of our baseline model. If marginal costs are more moderate, the lock-in shock tightens the market, raising the value of search for sellers, increasing investment in new housing. This introduces additional homes for sale, reducing offsetting the increase in market tightness.

As with the inclusion of a more realistic rental sector, endogenizing home construction provides a "relief valve" for rising market tightness, as new home construction responds to higher values of search for sellers and introduces additional sellers to the market. In both cases, it is notable that these endogenous reactions may reduce the magnitude of house price responses to the lock-in shock due to equilibrating forces, but do not alter the qualitative result that (in an already-tight market,) lock-in increases market tightness and puts upward pressure on house prices. However, the model extensions provide clear mechanisms by which mortgage lock-in may have indirectly raise rents and new home construction.

D Details on Estimating Aggregate Effects

In this appendix, we provide additional details on our approach to estimating aggregate effects of lock-in. We begin by taking a linear approximation of Equation 7, using $lnDOM$ as a proxy for θ

$$\frac{dlnDOM}{d\ell} = \kappa_0 + \kappa_1 lnDOM \tag{15}$$

Noting that $lnDOM(0)$ is the initial (steady-state) level of tightness prior to the shock, the solution to this linear ordinary differential equation is given by the following formula for the impact of a lock-in shock of size ℓ:

$$lnDOM(\ell) - lnDOM^{init} = \left(lnDOM^{init} + \frac{\kappa_0}{\kappa_1}\right) * \left(e^{\kappa_1 \ell} - 1\right) \tag{16}$$

Note that the sign of the aggregate effect is determined by $\left(lnDOM^{init} + \frac{\kappa_0}{\kappa_1}\right)$, and that $\kappa_0/\kappa_1 < 0$. Therefore, $-\kappa_0/\kappa_1 < 0$ provides a threshold for (de-meaned) log days on market that corresponds to the threshold $\theta = 1$ such that any location tighter than this threshold sees a tightening due to lock-in, and any location slacker than this threshold sees a slackening.

In order to recover κ_0 and κ_1 to compute the aggregate effect on tightness in Equation 16, we estimate the empirical counterpart of the differential equation 15:

$$\Delta_h lnDOM_m = \kappa_0 \ell_m + \kappa_1 lnDOM_m^{init} * \ell_m + \gamma_1 lnDOM_m^{init} + \gamma_0 + u_m \tag{17}$$

Note that the $\Delta\ell$ is the exposure to lock-in for a CBSA given by the decline in average move probability for loans in the CBSA implied by the change in rate gaps $\Delta\ell = \frac{1}{N_m}\sum_{i \in m}\left(\hat{B}(RateGap_i) - \hat{B}(\tilde{RateGap_i})\right)$ as described in Section 3.3. We taking a difference in $lnDOM$ over a given horizon h (taken to be 12 months from December 2021 to December 2022) to align with the counterfactual change in rate gaps, and utilize $lnDOM_m^{init}$ as of December 2021 on the right hand side.

Results are given in Table 1 below:

Table 1: Estimates of Tightness State-variable Response

	(1) Growth in Days on Market (Dec2021 - Dec2022)
Lock-in shock	-235.0***
	(82.28)
Lock-in shock × Init Log Days on Market	640.7*
	(337.0)
Init Log Days on Market	-1.126***
	(0.306)
Constant	0.384***
	(0.0749)
Observations (CBSAs)	728

Regression weighted by owner-occupied housing units. Standard errors in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Because initial log days on market has been de-meaned, the direct effect of lock-in exposure can be interpreted at the effect on the average location. As expected, we find the direct effect (corresponding to κ_0) to be negative as lock-in tightened the market, reducing days on market. In addition, the interaction of lock-in with initial tightness (corresponding to κ_1) is positive, suggesting locations with initially slacker markets (higher days on market) had saw smaller lock-in driven declines in days on market. We also note that the constant term shows a 38 percent increase in days on market over this period accounting for a renormalization of pandemic-era tightening. This renormalization was also more pronounced for locations that had tightened more sharply, which is accounted for by the direct effect of initial days on market.

With estimates of κ_0 and κ_1 in hand, we have all terms necessary to compute the aggregate effect on days on market from Equation 16 for each CBSA, with the aggregate national effect (28.8 percent decline in days on market) given by the mean, weighted by owner-occupied unit counts. Cross-CBSA variation in effect size is driven almost entirely by variation in initial tightness, with the (weighted) interquartile range of estimates across CBSA's ranging

from a 16.4 percent decline in days on market to a 39.5 percent decline. We also note that in December 2021, nearly all CBSA's fell on the "tight" side of the point estimate for the threshold $-\kappa_0/\kappa_1$, which corresponds to roughly 64 days on market.

We translate the aggregate effect on market tightness for each CBSA into the implied effect on prices by estimating the elasticity of price changes with respect to changes in days on market driven by the lock-in shock. Specifically, we consider a regression as follows:

$$\Delta_h lnHPI_m = \beta \Delta_h lnDOM_m + v_m \tag{18}$$

As argued in the main body of the paper, the lock-in exposure ℓ_m provides a valid instrument when estimating β since a pure reduction in local churn impacts prices via its effect on tightness. Consistent with results that lock-in significantly reduced days on market, results of the IV regression given below show the instrument is strong ($F = 19.22$) and the resulting elasticity $\hat{beta}$ is -0.277 and significant.

Table 2: Estimates of Elasticity of House Prices to Tightness

	(1) Growth in HPI (Dec2021 - Dec2022)
Growth in Days on Market	-0.277***
	(0.0768)
Constant	0.118***
	(0.0134)
First-Stage F-stat	19.22
Observations (CBSAs)	728

Regression weighted by owner-occupied housing units. Standard errors in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Using the elasticity to translate aggregate effects on tightness into prices, we find the weighted-average price effect across CBSA's is 8 percent, with the cross-CBSA variation in effect sizes ranging from 4.5 percent at the 25th percentile to 11.0 percent at the 75th percentile.

E Additional Appendix Figures

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

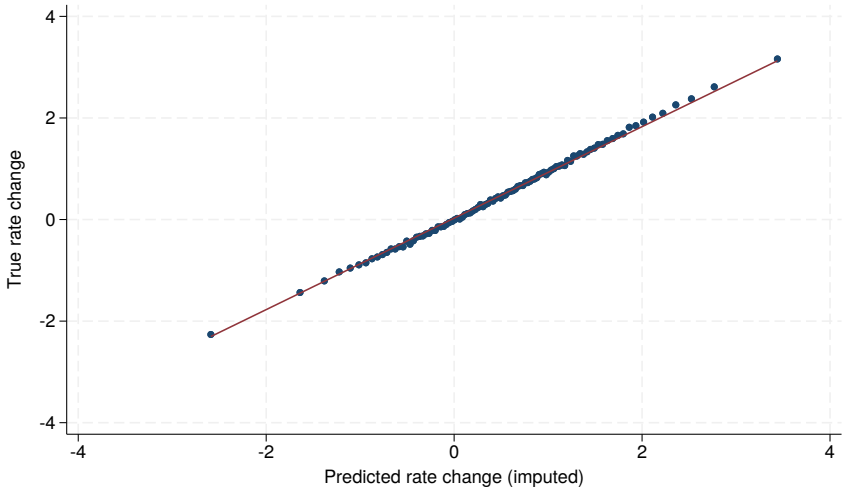


Figure E.1 plots the relationship between the rate gap we predict and the true change in rate among movers who both sell their home and subsequently purchase a new home using a mortgage. The graph is a binscatter with 100 bins, estimated on the subset of movers whose subsequent mortgage we could identify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

Figure E.2: Simulating Missing Moves Due to Rate Changes

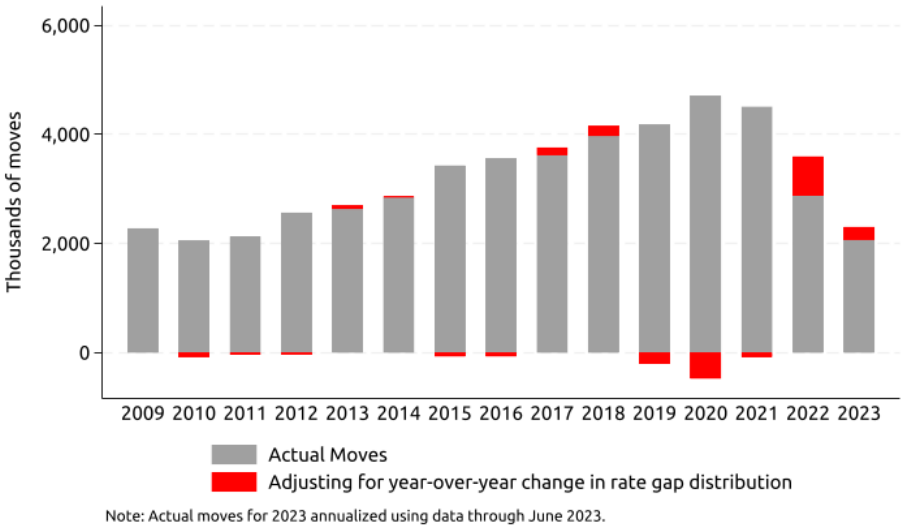


Figure E.2 plots actual moves in each calendar year (gray bars). Red bars represent the additional moves (or fewer moves, if negative) expected if the rate gap distribution had not changed from the prior year, computed as described in 3.3. Taking 2022 as an example, the gray bar shows that about 2.9 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.

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

Figure E.3: Move rates by tenure at destination

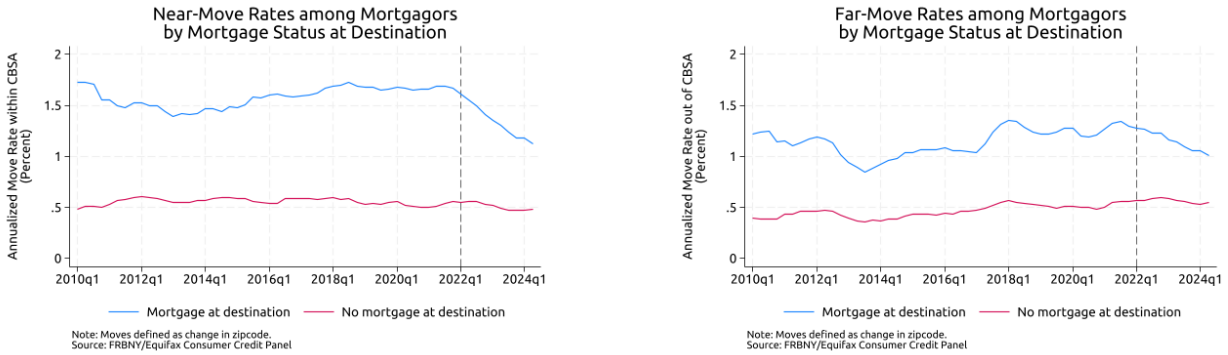


Figure E.3 plots mobility rates in the FRBNY/Equifax Consumer Credit Panel for borrowers with mortgages split by distance of destination and tenure at destination. The left panel shows annualized likelihood of moving within a CBSA split into moves where the borrower had a mortgage at the destination 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 E.4: Effect of Rate Gap and Home Value Change for Within- and Between-CBSA Movers

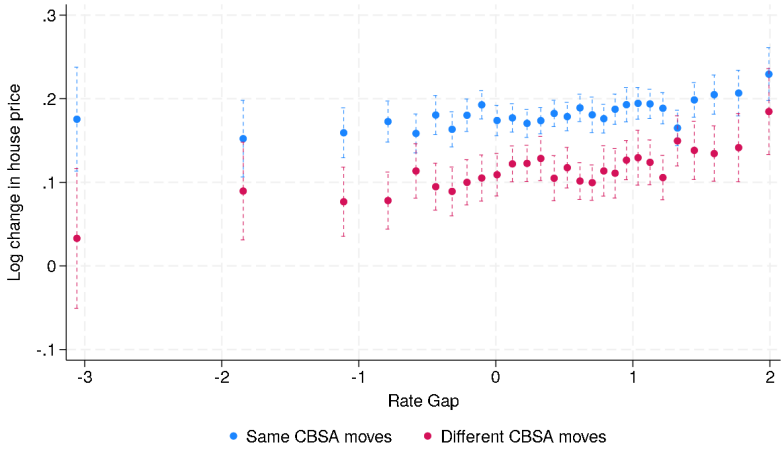


Figure E.4 plots the estimated log change in house value between movers' old and new houses. As in Figure 3, value changes are created using coefficients from Equation 2 evaluated at the mean of control variables. The rate gap is defined as the difference between the rate available to the borrower at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

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

Figure E.5: Geographic Variation in Exposure to Lock-In

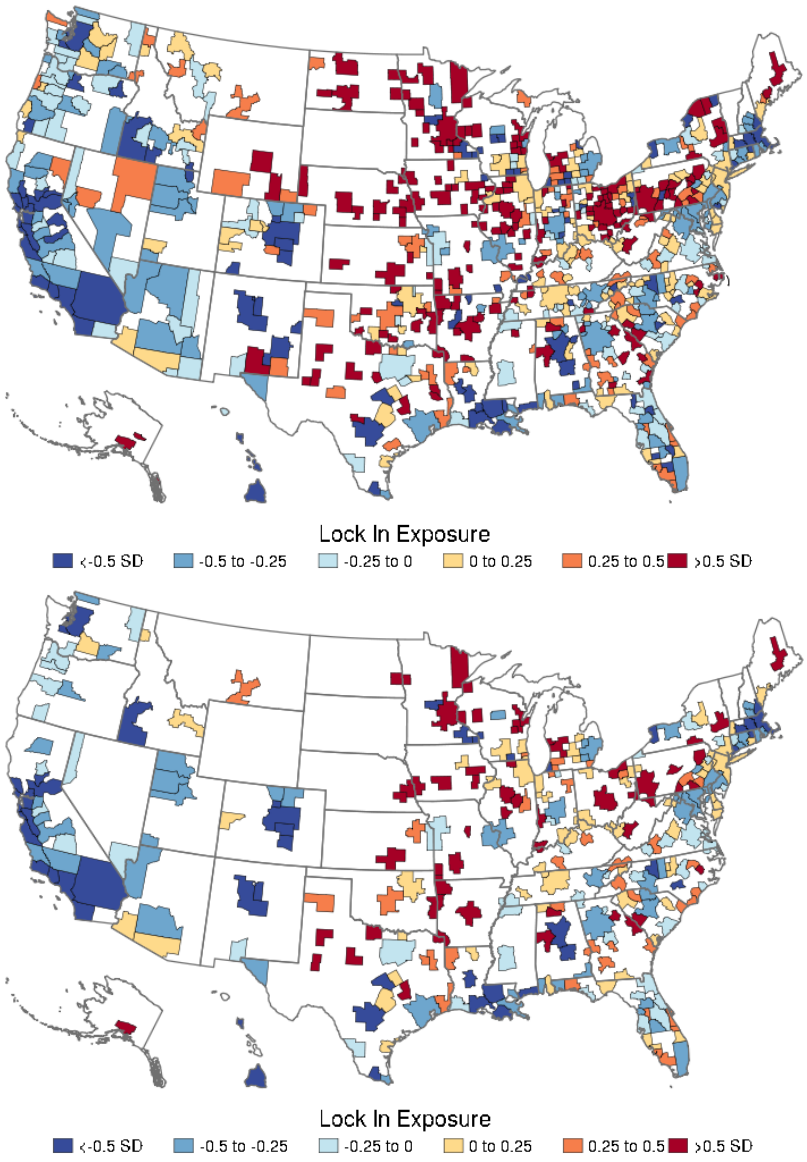


Figure E.5 plots estimated CBSA-level exposure to the Lock-In shock, as described in Section 4. Exposure is measured in standard deviations from the mean CBSA. The top map includes all metros in our sample; the bottom includes just the 300 most populated CBSAs. Exposure thresholds are consistent across the two maps.

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

Figure E.6: Lock-in Driven House Price Growth is Strongest in Tighter, Larger Markets

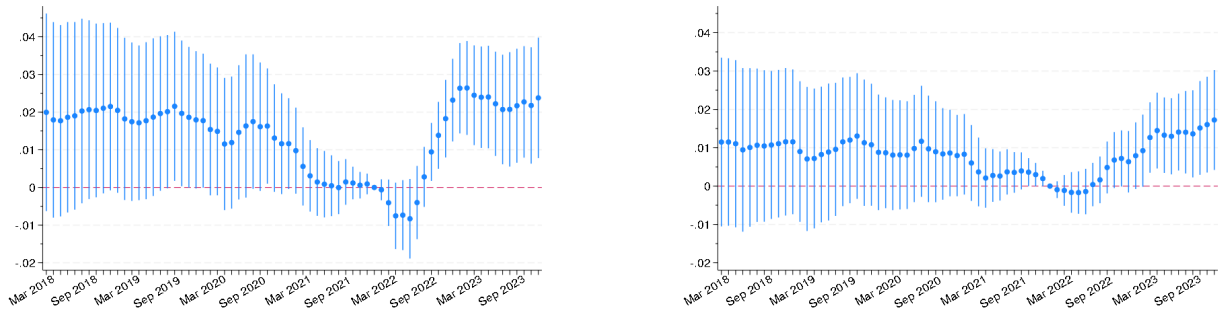


Figure E.6 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses log house prices at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. This figure includes only the top 300 most-populated CBSAs. The 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 E.7: Lock-in Driven House Price Growth is Weakest in Looser, Smaller Markets

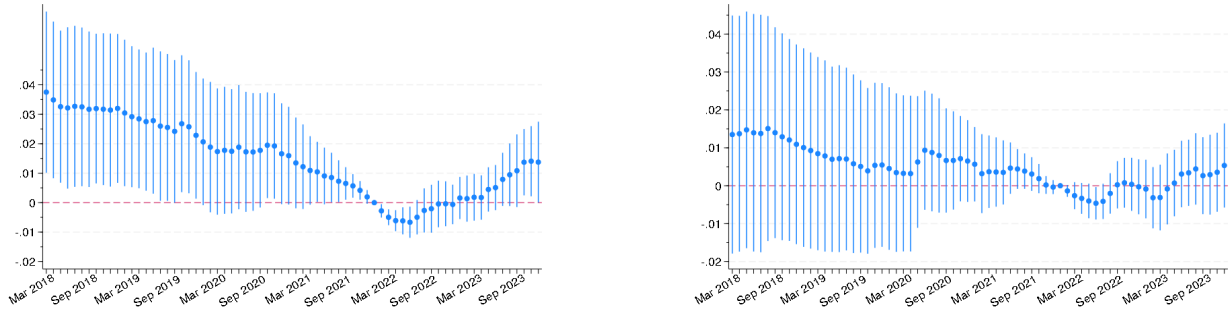


Figure E.7 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses log house prices at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. This figure includes only CBSAs outside the 300 most-populated. 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 E.8: Effect of Exposure to Lock-In on CBSA House Price Growth
Overall vs. Top300 CBSAs vs. All Others

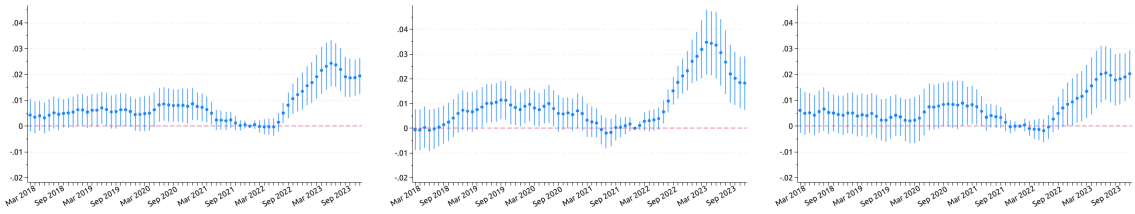


Figure E.8 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses ***log year-over-year house price changes*** at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. Effect sizes are for a one standard 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 E.9: Effect of Exposure to Lock-In on (log) CBSA House Prices
Overall vs. Top300 CBSAs vs. All Others

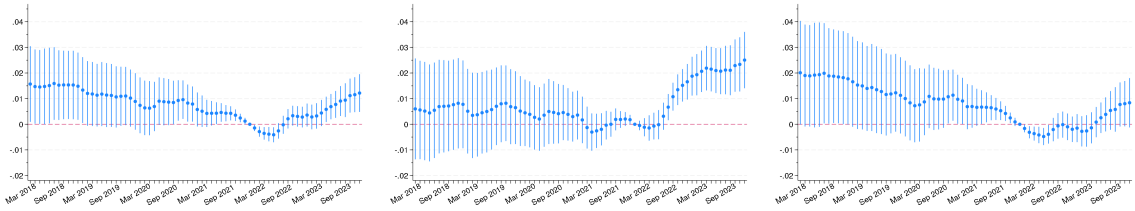


Figure E.9 plots coefficients from a two-way fixed effect difference-in-difference style specification, which regresses ***log house prices*** at the CBSA level (according to CoreLogic's CBSA House Price Index) on CBSA-level exposure to lock-in. Effect sizes are for a one standard deviation increase 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