



Household mobility and mortgage rate lock[☆]

Jack Liebersohn^a, Jesse Rothstein^{b,c,d} ^{*}

^a Department of Economics, University of California, Irvine, Irvine, USA

^b Goldman School of Public Policy and Department of Economics, University of California, Berkeley, 2607 Hearst Avenue, Berkeley, CA, 94720-7320, USA

^c NBER, 1050 Massachusetts Avenue, Cambridge, MA, 02138, USA

^d California Policy Lab, 2521 Channing Way, Berkeley, CA, 94720-5555, USA

ARTICLE INFO

Dataset link: [Data for "Household Mobility and Mortgage Rate Lock" \(Reference data\)](#)

JEL classification:

G21
G51
J61
R21
R23

Keywords:

Mortgages
Mobility
Interest rates
Housing lock

ABSTRACT

Rising interest rates can create "mortgage rate lock" for homeowners with fixed rate mortgages, who can hold onto their low rates as long as they stay in their homes but would have to take on new mortgages with higher rates if they moved. We show mobility rates fell in 2022 and 2023 for homeowners with mortgages, as market rates rose. We observe both absolute declines and declines relative to homeowners without mortgages, who are unaffected by mortgage rate lock. Mobility declines are not explained by changes in home values. Overall, our estimates imply that rising interest rates reduced mobility in 2022 and 2023 for households with mortgages by 16% and caused \$20bn of deadweight loss.

1. Introduction

The traditional 30-year, fixed-rate, non-assumable mortgage that is used for most home purchases in the United States is an unusual instrument. Because there are no pre-payment penalties, borrowers can re-finance to take advantage of declines in interest rates. Thus, the rate is in practice adjustable, but only downward. When rates rise, borrowers are protected – the rate does not adjust up – so long as the borrower remains in the house. However, should the original borrower wish to move to a new house, he or she must obtain a new mortgage at the market rate.

This feature can create a very strong disincentive to move for those holding mortgages at rates lower than the currently prevailing rate. Consider a homeowner who took out a fixed rate mortgage in 2016 at 3.5%, a typical rate for that year, and who still owed \$200,000 as of

2023. Suppose that circumstances in her life made it desirable for her to move to a different house of equal value, and that her credit score was excellent, so lenders were eager to offer her a new mortgage at the then-prevailing rate of 7%. Making the move would have increased her monthly payment by 38%, cumulating to over \$110,000 over the remaining life of the loan.¹ This cost could be avoided entirely by remaining in the original house. It thus discourages mobility, and can lead a borrower to avoid moves that would otherwise be desirable (for example, for a new job opportunity).

Interest rates were on a long-term downward trajectory from the early 1980s until the COVID pandemic, so this aspect of mortgages was not empirically relevant for most homeowners for many years. However, between December 2021 and November 2022, average mortgage rates rose from 3.10% to 6.95%, a nearly four point increase. They continued to rise thereafter, peaking at 7.79% in October 2023. Rates

[☆] Toni Whited was the editor for this article. We are grateful to Kirill Borusyak, Damon Clark, Daniel Green, Greg Howard, Jorge de la Roca, Antoinette Schoar, Tejaswi Velayudhan, Toni Whited (editor), and Miguel Zerecero for helpful comments, as well as to seminar participants at PPIC, USC and UC Irvine. We would like to thank Hannah Case for excellent research assistance. We are grateful to Experian and the California Policy Lab for providing the data underlying the University of California Consumer Credit Panel (UC-CCP) on which we rely.

^{*} Corresponding author at: Goldman School of Public Policy and Department of Economics, University of California, Berkeley, 2607 Hearst Avenue, Berkeley, CA, 94720-7320, USA.

E-mail addresses: cjlieber@uci.edu (J. Liebersohn), rothstein@berkeley.edu (J. Rothstein).

¹ To hold as much constant as possible, this calculation assumes that either the new or old loan would be paid off in equal monthly payments over the 23 years remaining of the original loan's term. Applying a discount rate of 7%, the present value of the additional payments is \$55,000.

in 2023 were higher than at any point since 2001. We show below that nearly all mortgage holders in 2022 and 2023 had rates much lower than would have been available to new borrowers, as even older mortgages are likely to have been refinanced during the low-rate period between 2009 and 2020.

We study the effect of rising interest rates on mobility. We show that mobility rates of homeowners with mortgages have fallen dramatically since 2021, and that this has been concentrated among mortgages originated when rates were substantially lower. Our estimates come from a hazard model that considers moving as a function of the gap between a mortgage's origination rate and the current market rate (the "rate gap") and other household characteristics, while controlling for housing tenure.

Each percentage point increase in the currently prevailing rate above a borrower's origination rate is associated with a 7.7% decline in the quarterly mobility probability. From 2022q3–2023q2, the last year we have data, we estimate that rate lock caused between-ZIP-code mobility rate to fall by 16 percent, from 7.3% to 6.1%, for households with a mortgage.² We calculate that rate lock discouraged 800,000 moves over that year. Across several specifications, the effects on inter-state mobility are about half as large as the effects on inter-ZIP mobility.

The change in mobility rates for homeowners without mortgages has been much smaller. We present "difference-in-hazard" estimates that identify the rate lock effect as the differential effect of rates on homeowners with mortgages and on a matched sample of owners without mortgages, allowing for common time effects to control for general equilibrium effects in the housing market. We also estimate instrumental variables specifications that abstract from refinancing decisions by instrumenting for a household's current mortgage rate using the prevailing mortgage rate at origination. Our estimates are similar across specifications.

Rate lock has important financial implications for households and for the banking system (Bolhuis et al., 2024). One effect is an effective transfer to banks from households that moved despite rate lock. We calculate that households who moved despite being "locked in" saw an average increase in annual mortgage payments worth nearly \$5000 as a result of their moves — payments which could have been avoided if they had been able to keep their old mortgages. The present value of this "payment gap" is equal to the change in the discounted value of future mortgage payments caused by a higher interest rate. We find this equaled about \$49,000 per mortgage in the last year of the sample, for a total of \$215bn across all moving households. Rate lock also has economic costs in the form of the deadweight loss caused by forgone moves, relative to a counterfactual where mortgages are assumable. We can value these costs by considering the demand for moving as a function of the present value of mortgage payments. These costs amount to \$20bn in the last year, an average of \$296 per household with a mortgage.

Previous studies have documented the effects of mortgage lock during earlier periods.³ Quigley (1987, 2002) studies lock-in during the 1980s and 1990s, building on the household relocation models in Hanushek and Quigley (1978) and Venti and Wise (1984). Ferreira et al. (2010, 2011) find substantial rate lock-in effects during the 2000s, and also show large effects of negative home equity.

The rapid increase in rates in 2022 and 2023 was much larger than had been seen in other rising-rate episodes in recent decades, and was largely unexpected. This provides a great deal of statistical power

and a unique natural experiment with which to measure rate lock. We use high-frequency credit registry data, along with a strategy that compares homeowners with mortgages to those without, to estimate the effect of rate lock during this period on a representative population of homeowners.⁴

We provide evidence that rate lock mattered more in 2022–2023 than in previous periods. This is due to the interaction of two factors. First, gaps between rates on outstanding mortgages and market rates are much larger in this period; prior to 2022, there are very few observations with gaps larger than about 1 percentage point. Second, in specifications that allow for nonlinear effects of the rate gap, we find very small effects until the rate gap exceeds about 1.5 points.⁵ Thus, our estimates indicate very little effect of rate lock prior to 2022 — an implication that is confirmed when we restrict our sample to the pre-2022 period.

Although the current high-interest-rate regime has been in effect for less than three years, it is already having quantitatively important consequences for aggregate mobility rates. We show that the decline in average moving probabilities since 2021 that is attributable to interest rate lock sped up the secular decline in mobility rates by as much as one year. As many previous authors (e.g. Molloy et al., 2016) have noted, increases in moving costs and declines in mobility have the potential to add substantial friction to the free flow of workers to job opportunities in the labor market and slow recovery from recessions.

Interest rate lock also has consequences for lenders. Insofar as homeowners respond to interest rate increases by reducing mobility, this contributes to the asymmetry between interest rate changes and time-to-mortgage-payoff, reducing mortgage payoffs at exactly the times when it is most costly to the lenders for the mortgages to remain outstanding.

Ours is not the only paper to study rate lock in 2022 and 2023. Two other examples are Fonseca and Liu (2024) and Batzer et al. (2024).⁶ Fonseca and Liu (2024) use credit data similar to ours, while Batzer et al. (2024) use data on home sales. Our work has several advantages over these papers. First, our empirical strategy isolates the variation in rate lock deriving from changes in interest rates over time, rather than from differences among mortgageholders at a point in time in their outstanding loans' interest rates. There are two components of an outstanding loan's rate, both potentially endogenous to unobserved determinants of mobility — the date at which the loan was originated, and the borrower's rate relative to the market rate on that date. The relative rate reflects, among other things, a borrower's creditworthiness, which may be correlated with the propensity to move. The borrower's origination date depends on how recently the borrower refinanced, which may depend on the borrower's mobility plans — a

⁴ Our strategy of comparing mortgage-holders to those without mortgages builds on similar approaches taken in earlier papers, for example Aladangady (2017) and Atalay and Edwards (2022) on housing wealth effects and Chaney et al. (2012) on corporate investment.

⁵ Although our primary specifications focus on identifying the effect of positive gaps, where the market rate exceeds the outstanding rate, our nonlinear models allow for effects of negative rate gaps as well. We find some evidence that market rates below the outstanding mortgage rate encourage mobility. This phenomenon is explored in greater depth in Fonseca and Liu (2024). In their model, small negative rate gaps can dissuade moving, but larger negative gaps simply lead to refinancing with no effect on mobility. (Note that their model is developed in terms of the "mortgage delta", defined as -1 times what we call the rate gap.)

⁶ Batzer et al. (2024) postdates the first circulated draft of this paper, in September 2023. The original version of Fonseca and Liu's paper dates to March 2023, but that version uses data only through 2021 and does not incorporate the mortgage holder-non-mortgage holder comparison that we use and that appears in Fonseca and Liu (2024). Both Batzer et al. (2024) and Fonseca and Liu (2024) use as their key variable the difference between an existing mortgage's rate and the current market rate — the opposite of what we define as the rate gap.

² These estimates are consistent with an elasticity of ZIP code mobility with respect to net income between 3 and 4, or of state mobility between 1 and 2. This is in line with other estimates that (Fajgelbaum et al., 2019) review.

³ A large literature studies the effects of negative equity on mobility, a different channel than the one we study. See Andersson and Mayock (2014), Bernstein and Struyven (2022), Foote (2016) and Brown et al. (2019).




Fig. 1. Mortgage interest rates and rate gaps. Notes: Panel A shows the current 30-year FRM interest rate for originating mortgages (from the Federal Reserve's FRED tool, "MORTGAGE30US" series) and quartiles of the interest rates for outstanding 30-year or less, single-family FRMs held by Fannie Mae. Data on outstanding mortgages are compiled from the Fannie Mae Single-Family Loan Performance Data. Panel B shows the average additional monthly payment that mortgage holders would face if their mortgage payments were recalculated using the current 30-year FRM interest rate (assigning zero payment change to any mortgage with a rate above the current rate).

household that is considering moving in the short term may not pay the fixed costs to refinance even when this would lower the rate. Mortgage origination dates may also be correlated with household characteristics that affect migration rates directly, as we discuss in Section 4. By contrast, our preferred estimates include origination date fixed effects and instrument for the current mortgage rate with the market rate when the original mortgage (prior to any subsequent refinancings) was originated. This ensures that rate gap effects are identified solely from the natural experiment generated by the path of interest rates after origination.⁷

Second, we use a hazard framework with a nonparametric baseline hazard, rather than a linear probability model. This is useful for examining an episode in which the key treatment (interest rates above a mortgage's existing rate) is strongly serially correlated. When households are dissuaded from moving in one quarter, this affects the duration composition of outstanding mortgages in the next quarter; because moving rates are sharply higher for people with short tenures in their homes, this directly affects the second quarter's mobility rate. The hazard model with nonparametric controls for duration is better suited for this type of dependence than are the linear probability models used in other work. The hazard framework is also important for computing the cumulative effect of high rates over multiple quarters, without double-counting. We use it to calculate counterfactual aggregate mobility rates by quarter and to calculate total welfare losses coming from all sources of mobility.

Beyond the methodological differences, the credit registry data we use is more granular than what other studies have used. Because we observe moves every quarter, we can construct event studies that trace out the effects of rates on mobility at a high frequency. The event study graphs help with identification because they allow us to provide

evidence for parallel pre-trends and show the timing of the response to interest rate shocks. They also help us identify a valid control group for mortgage holders – cash buyers – not used in other research.⁸ Another use of quarterly data is to identify nonlinear effects of interest rate gaps identified solely from changes in interest rates over time, rather than differences across cohorts; as compared to other work, we find little evidence that negative “rate gaps” affect mobility. We also explore heterogeneous effects by region, race, income and financial factors. The effects are much larger for people who recently refinanced, and about half as large for out-of-state moves as for inter-*zip* moves.

The remainder of this paper is organized as follows. Section 2 discusses institutional details and presents a simple calculation of the contribution of interest rate lock to the cost of moving. Section 3 describes the data we use to obtain high-frequency measures of mobility rates. Section 4 presents our main empirical strategy. Section 5 presents the main results and robustness analysis. Section 6 concludes.

2. Institutional details and motivating framework

Since 2010, the vast majority of U.S. residential mortgages have been fixed rate (known as FRM).⁹ Interest rates are fixed at origination, sometimes with a (pre-established) discount early in the mortgage's life.

Nearly all U.S. residential mortgages are securitized by the home, are not assumable by a new buyer, allow for prepayment without substantial penalties, and must be paid in full if and when the home is sold.¹⁰ These features create what we call “interest rate lock” for FRM borrowers.¹¹ A homeowner who wishes to move must assume

⁷ Fonseca and Liu (2024) instrument for a household's origination rate with the market rate at the time that the mortgage was originated to avoid endogeneity coming from credit scores, but do not address endogeneity coming from origination dates. Moreover, they include calendar year fixed effects, which absorb most of the variation in current market interest rates, so time since origination is their primary source of variation. Batzer et al. (2024)'s preferred specifications use variation in offered rates among mortgages originated at the same time and include calendar time fixed effects. They are thus identified primarily or entirely from the two sources of variation that we see as potentially endogenous.

⁸ Fonseca and Liu (2024) use renters as a comparison for mortgage holders. We found evidence that renters' mobility was not on a parallel trend to mortgage holders prior to the rate hikes, so we focus on non-mortgage homeowners as preferable comparison group.

⁹ The FRM share of mortgage applications rose from about 2/3 in 2004–5 to 95% in 2009, and has been above 90% nearly all of the time since (Goodman et al., 2023).

¹⁰ Mortgages insured by the FHA and VA are assumable, but only under strict conditions, and assumption appears to be rare.

¹¹ A conceptually distinct type of housing lock arises when the market value of the house is insufficient to pay off the remaining balance on the mortgage — when the borrower is “underwater”. This has been studied more — see, e.g., Ferreira et al. (2010).

not only the difference in prices between the old and new house, but also a new interest rate. If market rates are higher at the time of the move than at the time of the original mortgage's origination, his or her payments will go up even if the size of the mortgage is the same. Thus, the rise in interest rates can be seen as imposing a capital loss on the borrower.¹² However, the homeowner can avoid the capital loss by remaining in the old house. He or she thus has an incentive not to move. This distortion can lead to inefficiencies, if homeowners are unable to pursue new job opportunities in different locations or to downsize when life circumstances make that appropriate.¹³

The left panel of Fig. 1 shows the path of mortgage interest rates since 2013. Rates oscillated between about 3.5% and 5% between 2013 and 2020, falling below 3% in the wake of the COVID crisis. In early 2022, however, they began rising sharply, following Federal Reserve monetary policy tightening, and they have been above 6% since September 2022.

We overlay on this graph the 25th, 50th, and 75th percentiles of interest rates on outstanding FRMs, calculated from the Fannie Mae Single-Family Loan Performance sample.¹⁴ These are based on the distribution of rates across loans issued at many different times, and as such move much more slowly than does the current rate series. For example, the decline in rates that began in 2018 does not show up in the outstanding loan rate distribution until 2020. A consequence is that the distribution of rates on outstanding mortgages in 2022 and 2023 largely reflects the pre-2021 low-rate environment. Even the 75th percentile of that distribution was below 4% at the end of 2022, 2.5 percentage points below the rate then being offered on new mortgages.

The cost of taking on a new mortgage is directly related to the gap between the currently offered rate and the rate on the existing mortgage. To fix ideas, consider a homeowner with a mortgage that was taken out in the past at some annual rate R_0 , with remaining principal P and m monthly payments remaining in the term. Suppose that the homeowner is considering moving to a new house of identical value, and converting all of his/her remaining equity into a down payment. This means that he will need to take out a new mortgage with principal P at new interest rate R_1 . For simplicity, assume the remaining term will be the same, m months.

Using standard amortization formulas, the monthly payment for the existing mortgage is $P * f(R_0)$, where $f(R) \equiv \frac{R_0/12}{1 - (1 + R_0/12)^{-m}}$, while the monthly payment for the new mortgage will be $P * f(R_1)$. Note that $f(\cdot)$ is increasing in R , so the new mortgage payment is higher. With discount rate δ , the present value of the cost of trading the former obligation for the latter is

$$P \frac{f(R_1) - f(R_0)}{f(\delta)}.$$

This can be substantial. An increase from $R_0 = 3.5\%$ to $R_1 = 7\%$, with $\delta = 7\%$, raises the present value of future payments by 38%.¹⁵ The right panel of Fig. 1 shows the average monthly payment gap over time — the amount that the average mortgage holder's monthly payment would increase if the mortgage was re-issued at the current market rate, with

¹² Taking out a mortgage is, effectively, issuing a bond. Bond values rise when rates increase and fall when they decrease. Borrowers are short bonds, so take losses and gains, respectively.

¹³ Similar inefficiencies have been noted due to property tax rules that tie the tax bill to the purchase price (Ferreira, 2010) and to rent control regimes that limit rent increases for incumbent tenants (e.g., Munch and Svarer, 2002). A longstanding policy conversation points to declining mobility rates as an indication of reduced dynamism of the U.S. economy (e.g., Molloy et al., 2016).

¹⁴ <https://capitalmarkets.fanniemae.com/credit-risk-transfer/single-family-credit-risk-transfer/fannie-mae-single-family-loan-performance-data>.

¹⁵ We have neglected the possibility that the term could be extended with a new mortgage. This would lower the monthly payment, but (so long as the discount rate is below R_1) not the present value of the stream of payments.

the same term and principal.¹⁶ This is generally low or even zero, when most mortgages are at or above the market rate, but rises above zero when rates increase in 2013, 2018, and especially 2022. By 2023, the average mortgage holder's payment would increase by over \$400 per month if the mortgage was re-issued at the then-prevailing rate. This is the cost that the homeowner would need to pay if he/she wanted to move to a home of equivalent value elsewhere. It creates a large disincentive to move.

Rising rates can also cause lock-in through a second channel. Higher rates can reduce the value of homes, directly by increasing the payment that a prospective buyer would need to pay to finance a mortgage at any given value or indirectly via negative effects on overall economic activity that reduce demand in the housing market. This could push homeowners "underwater", owing more on their mortgage than they could obtain by selling, and thus reduce their ability to finance a move. In the present episode, the rise in rates has not been associated with a large decline in average values. Moreover, our analysis builds in two features that enable us to distinguish interest rate lock from value effects: We compare the change in mobility for mortgage holders to that for non-mortgage-holding homeowners, and we control directly and flexibly for the change in home values in the local area.

3. Data

Our main data source is the University of California Consumer Credit Panel (UC-CCP), developed and maintained by the California Policy Lab at the University of California. UC-CCP is a nationally representative credit registry containing longitudinal information on a 2% random sample of U.S. individuals with a credit history. Quarterly data on households and credit accounts are constructed from records compiled by one of the major credit reporting agencies. Address information includes the ZIP code of residence. Residential locations are updated shortly after moves, as financial institutions stay in close touch with their clients. For this reason, the UC-CCP is ideal for measuring mobility among households with a credit history, a group which includes the mortgage borrowers who are our focus.¹⁷

We identify all unique mortgage originations between the first quarter of 2013 and the end of 2020, identifying a mortgage by the combination of borrower, origination date, and principal amount. We measure whether and when thereafter the borrower relocates to another ZIP code, or in some specifications to another state, within the first ten years (40 quarters) after origination, using data through the second quarter of 2023.¹⁸ We also measure whether the mortgage is closed, which could happen without a move when a mortgage is paid off or refinanced.

We construct a sample of all mortgage originations, treating each as the beginning of a new spell. We think it is important to consider both originations that reflect new purchases and those that reflect refinancings, as the latter make up a large share of the market. However, including both means that a single household-ZIP combination can be represented by several overlapping spells - one begins when the home is purchased, and another begins when it is refinanced. We include both as distinct spells, assigning each a weight of 0.5. Similarly, households that refinance twice receive three spells, each with a 0.33 weight.

The next question is when a spell ends. One option would be to identify a spell with the duration of the original mortgage, considering it to have ended when the mortgage is paid off. However, this would

¹⁶ At the individual mortgage level, this is $\max\{0, P(f(R_1) - f(R_0))\}$. We calculate this in our credit data sample, described below, and then average over the P s, R_0 s, and m s of all outstanding mortgages in each month.

¹⁷ See Holmes (2021) and Holmes and White (2022).

¹⁸ To reflect the possibility that people may not move into a newly purchased house immediately after the mortgage is originated, we identify the house location based on the purchaser's location two quarters after origination, and consider only moves after that point.

mean that households that refinance exit the sample. This would make the sample quite unrepresentative at higher tenures, particularly following the very low rate period in 2019–2020. To ensure that this does not affect the results, households that refinance their mortgages are retained in the sample, not counting as having failed so long as they remain in the original ZIP code. Thus, spells end when the household leaves the origination ZIP, after 40 quarters, or at the end of our data in 2023, whichever comes first.¹⁹

We proxy for a mortgage's interest rate with the market rate at the time that it was originated. This abstracts from idiosyncratic variation across borrowers in the rates they receive, which may reflect unobserved differences in creditworthiness.²⁰ Quarterly market rates come from Freddie Mac's Primary Mortgage Market Survey and pertain to 30-year FRMs. Our primary measure updates to the new market rate whenever a mortgage was refinanced, but we construct a second measure that preserves the original mortgage's origination rate as well.²¹

We construct a second panel of homeowners who do not have a mortgage, who we refer to as "cash buyers". Here, "spells" begin when the household moves into a ZIP code, and end when the household leaves it. This sample includes only households that are classified by the credit panel as homeowners in the quarter they move, excluding households where any household member has an active mortgage in that quarter.²²

Our models control for local home price changes. To measure this, we use ZIP level house price indexes from Zillow. We calculate the change in house prices in a ZIP code since a mortgage was originated.

3.1. Summary statistics

Table 1 presents summary statistics for the main analysis sample. In Panel A, we present statistics at the mortgage level. We have over 1.8 million mortgages in our sample. New purchases are 36% of the originations in our sample, with the remainder being refinancings. The average origination interest rate is 3.8%. 59% of mortgages are closed within five years of origination, about half due to moves out of the zip code and half to refinancings.

Panel B shows statistics at the mortgage-by-quarter level, with 26 million quarterly observations. Our observations on each mortgage begin when it is originated and continue to the earliest of (a) the quarter that the household moves out of the ZIP code, (b) ten years after origination, or (c) the fourth quarter of 2023, when our data end. Appendix Figure B.2 shows the survival curve in our sample. Here, a household is counted as surviving from mortgage origination until the household leaves the ZIP code, even if the mortgage is closed (refinanced or paid off) first; in about 31% of our quarterly observations on surviving spells the original mortgage has been closed. The gap between the mortgage rate in effect for the mortgage and the current rate averages 0.3 percentage points, but for the 53% of mortgages with a positive gap it averages 1.1%. Appendix Table B.1 shows parallel statistics for our sample of homeowners without mortgages.

¹⁹ Another approach to this would be to estimate a competing risks model, where mortgages "fail" either when the household moves or when they are refinanced. We defer this to future work.

²⁰ It also allows us not to rely on mortgage rates computed from credit records, which are measured with substantial error (Shahidinejad, 2023).

²¹ For spells beginning with a refinance, the "origination" rate is that for the date of the focal refinance origination, and the updated rate may differ if the homeowner later refinanced again.

²² We do not condition on not taking out a mortgage later in defining our non-mortgage owner sample. In the rare event where an individual buys a home in cash and moves in, then later takes out a mortgage on the home, he or she will appear in both our non-mortgage owner sample (based on the date he or she arrived in the ZIP) and our mortgage holder sample (based on the date the mortgage was originated). We present robustness checks that define spell starts similarly for mortgage holders.

4. Empirical strategy

Our analysis compares the mobility rates of households facing different interest rate gaps between the fixed rates on their previously issued mortgages and the current market rate. Because our data identify the quarter in which a household moves but not the exact date, we adopt a discrete-time hazard model for data observed at regular intervals. Let Y_i represent the duration (in quarters) from mortgage origination to a household's move out of the zip code, with $Y_i = \infty$ if the household never moves. Let $O^*(i)$ represent the date on which mortgage i was originated, and let $O(i, t) \geq O^*(i)$ be the date on which it was most recently refinanced as of the t th quarter after origination (with $O(i, t) = O^*(i)$ for mortgages that have not been refinanced).

A conventional specification of the survival function is:

$$S(t | X_i) \equiv \Pr(Y_i \geq t | X_i) = S_0(t)^{\exp(X_i\beta)}, \quad (1)$$

where $S_0(t)$ is the baseline survival function and X_i are time-invariant characteristics of unit i .²³ This implies a discrete-time hazard function of the form

$$\lambda(t | X_i) \equiv \Pr(Y_i = t | X_i, Y_i > t - 1) = 1 - [1 - \lambda_0(t)]^{\exp(X_i\beta)}, \quad (2)$$

where $\lambda_0(t)$ is the baseline hazard. This can be rearranged into a simple partially linear form, known as a "complementary log-log" model:

$$\ln(-\ln(1 - \lambda(t | X_i))) = \alpha(t) + X_i\beta, \quad (3)$$

where $\alpha(t) = \ln(-\ln(1 - \lambda_0(t)))$.

In our setting, the variable of interest, the gap between the mortgage interest rate and the current market rate, is time-varying. We assume that the per-period hazard satisfies²⁴:

$$\ln(-\ln(1 - \lambda(t | X_{it}))) = \alpha_t + X_{it}\beta + u_{it}. \quad (4)$$

We use several different specifications for X_{it} . First, we simply include a full set of calendar time indicators. The resulting specification semi-parametrically measures changes in mobility hazards by quarter, controlling for changes in tenure distributions. We plot the resulting estimates to provide graphical evidence for the timing of changes in mobility rates.

Next, we move to a more direct measure of the mortgage rate gap. Let r_t represent the market interest rate at time t . The prevailing rate when mortgage i was last refinanced is then $r_{O(i,t)}$. We measure the rate gap as $g_{it} = \max(0, r_t - r_{O(i,t)})$, and include this in X_{it} .²⁵ Depending on the specification, control variables included in X_{it} are linear calendar time (t) controls, measures of negative rate gaps ($\min(0, r_t - r_{O(i,t)})$), and/or changes in the Zillow home price index in the ZIP code from origination to present for the focal mortgage, $p_{z(i)t} - p_{z(i)O^*(i)}$. We also test for potentially non-linear effects of rate gaps, by including a square or cube of g_{it} or by including indicators for rate gaps in certain ranges (e.g., deciles of the distribution, or values above 1).

²³ A survival function of this form can be derived by assuming that continuous-time data is described by the standard proportional hazards model, but is observed only at discrete time intervals.

²⁴ With time-varying X_{it} , this hazard does not generate a simple closed form for the survival function, which in general will depend on the history of X from the mortgage's origination to t , but the hazard remains well-defined. When we interpret the β coefficients in terms of their implications for survival, we compute the running product of the implied $1 - \lambda_{it}$ s. Note also that modeling survival would require accounting for right-censoring of spells that are ongoing at the end of the sample, but this is not needed when fitting the hazard directly.

²⁵ Fonseca and Liu (2024) and Batzer et al. (2024) examine what they call a "rate delta" that is defined as the difference between the outstanding mortgage's rate and the current market rate — that is, the opposite sign as our rate gap.

Table 1
Mortgage summary statistics.

Panel A: Mortgage-level sample					
	Mean	Median	Standard deviation	Min	Max
Origination year	2017.393	2017	2.934	2013	2022
Refinance	0.641	1	0.480	0	1
Principal (thousands of dollars)	260.689	213.75	249.124	−0.02	100 000
Origination rate	3.796	3.78	0.691	2.65	7.08
Quarters observed (max = 40)	21.831	22	11.667	0	40
Credit score	746.284	759	63.025	348	850
Age	46.136	45	14.147	18	128
Closed within 5 years	0.590	1	0.492	0	1
Moved within 5 years	0.314	0	0.464	0	1
Closed by end of panel	0.588	1	0.492	0	1
Moved by end of panel	0.478	0	0.500	0	1
Panel B: Mortgage-by-quarter panel					
	Mean	Median	Standard deviation	Min	Max
Year	2019.065	2019	2.698	2013	2023
Rate gap (g)	0.338	0	1.199	−13.86	4.01
Rate gap (g), conditional on positive	0.567	0	13	0	4.01
Rate gap (g), conditional on negative	−0.228	0	0.414	−13.86	0
Positive rate gap (g)	0.480	0	0.500	0	1
Rate gap vs. origination (g*)	0.313	0.04	1.196	−2.18	4.01
Rate gap vs. origination (g*), conditional on positive	1.110	0.66	1.085	0	4.01
Rate gap vs. origination (g*), conditional on negative	−0.582	−0.49	0.449	−2.18	−0
Positive rate gap vs. origination (g*)	0.530	1	0.499	0	1
Log ZHVI change	0.208	0.16	0.197	−2.33	3.83
Mortgage still open	0.690	1	0.462	0	1
Mortgage closed this quarter	0.026	0	0.160	0	1
Household moved this quarter	0.034	0	0.180	0	1

Notes: N = 1,870,171 mortgages in panel A; N = 26,509,578 mortgage-quarter observations in panel B. Statistics pertain to mortgage sample; mortgage-quarter sample is restricted to quarters before a household moves.

Of course, individual-level hazard rates are not observed directly. To implement our analysis, we divide our sample into cells defined by the interaction of origination quarter, duration, mortgage type (purchase mortgage or refinance), and ZIP code home price appreciation bins.²⁶ The resulting cells each include thousands of mortgages. Let c index cells, where g_c is the rate gap for cell c and X_c is the vector of controls.²⁷ We compute the empirical hazard for cell c , λ_c , as the share of households in the cell who move out of the ZIP code at duration d (which we write as $d(c)$ to reflect that duration is fixed for each cell). Because cells are large, we can measure this hazard accurately. Per Eq. (3), this gives rise to a simple regression:

$$\ln(-\ln(1 - \lambda_c)) = \alpha_{d(c)} + g_c \gamma + X_c \beta + u_c. \quad (5)$$

We estimate this via weighted least squares, weighting by the number of mortgages in cell c and allowing separate baseline hazards for purchase mortgage and refinanced mortgages. Because λ_c is generally small, $\ln(1 - \lambda_c) \approx -\lambda_c$, so the left hand side of (5) is approximately equal to $\ln(\lambda_c)$. This means that γ can be interpreted as the percentage change in the hazard per one-unit increase in X_c .

As noted, we keep mortgages in our panel even if they are refinanced. Keeping refinanced mortgages is important for identification, since refinancing is likely to depend on the current rate gap. However, if households that anticipate moving soon are less likely to refinance to obtain lower rates, this could create an endogenous relationship between a household's propensity to move and the measured rate gap. To address this, we use an instrumental variables strategy, instrumenting for g_{it} with an alternative rate gap that uses the rate at the time the focal

mortgage was originated rather than the most recent refinance, $g_{it}^* \equiv \max(0, r_t - r_{0^*(i)})$. Palmer (2022) discusses the estimation of instrumental variables hazard models via a control function approach. Our cell-based approach allows a simpler estimator: We simply estimate (5) with two stage least squares, again weighting by cell size.

We interpret the coefficient on g as the causal effect of the rate gap on moving hazards, and use our estimates to calculate counterfactual survival curves given different levels of rate lock. The large and unexpected nature of the 2022 interest rate shock helps to identify the effects of rate gaps. Estimates from earlier years might be confounded by slow-moving macroeconomic variables, like demographic change or secular trends in migration. In 2022, event study graphs show large changes in mobility, right at the time that rates went up, exactly for the groups we expect.

Our preferred specification controls for origination-cohort fixed effects, ensuring that the estimates are identified by changes in rates after each mortgage was originated. An alternative strategy, pursued by Fonseca and Liu (2024), would use variation in mortgage origination dates to identify the effects of mortgage rate lock — homeowners who originated at lower-rate times face more of a penalty for moving than do other homeowners who originated when rates were higher, so the contrast in the two groups' mobility rates could be used to measure the effect of rate lock. This approach could result in biased estimates, however. Households who know they are likely to move may be less likely to refinance, creating a correlation between time-since-origination and current mobility. Origination quarter is also likely correlated with household characteristics since some types of households are systematically more likely to refinance when rates fall (Andersen et al., 2020; Gerardi et al., 2023). For example, younger households are more likely to refinance when rates are low and are also more mobile, biasing downward estimates that do not control for origination date.²⁸

²⁶ The home price appreciation bins group ZIP codes with similar price trends from origination to the end of our sample. For each origination date, we divide ZIPs into ten deciles.

²⁷ g_{it} may vary within cells, as households with the same origination date may have refinanced at different times. Per (4), we use the average of g_{it} within the cell for g_c . Below, we discuss an instrumental variables strategy that uses a gap constructed from the origination rate $r_{0^*(i,t)}$; this is constant within cells, as are our other X variables.

²⁸ In Appendix Table B.2, we estimate specifications identified by cohort variation in mortgage origination rates, as in Fonseca and Liu (2024), rather

The main challenge to our interpretation is that omitted variables that are correlated with interest rate movements may affect household mobility. An obvious candidate is the COVID-19 pandemic. Interest rates rose exactly when the U.S. economy was recovering from the pandemic, which might have affected mobility directly. Our strategy here is to compare mortgage holders to other households that also were experiencing any pandemic effects but were not directly affected by interest rates. We show that mobility of households with mortgages is much more sensitive to interest rate gaps than is the mobility of mortgage-free homeowners.

In comparing mortgage holders to non-mortgage homeowners (“cash buyers”), a concern is that they may be different in other relevant characteristics (e.g., age) that influence sensitivity to market conditions. Thus, as a first step we reweight the cash buyer sample to match our main sample on observable characteristics, including the calendar quarter, origination cohort, ZIP code group, and homeowner age, credit score, and outstanding debt. Details of this reweighting are discussed in the appendix. We construct cell observations for cash buyers from the reweighted data. We combine these with the mortgage holders to estimate “difference-in-hazard” specifications. These interact our duration controls and all of our X_c variables with indicators for the type of owner, and identify the mortgage lock effect from the extra effect of the rate gap on mortgage holders relative to its effect on cash buyers.²⁹ We add to the specifications time fixed effects that capture any time-varying market conditions (e.g., liquidity), on the assumption that these conditions have similar effects on mortgage holders and on observably similar cash buyers.

5. Main results

5.1. Empirical hazard rates by cohort

The left panel of Fig. 2 shows empirical estimates of the hazard for mobility out of the zip code, grouping mortgages into four two-year origination “cohorts”. For all cohorts, the mobility hazard is high in the first two years after origination, then declines to a low but stable level thereafter. However, we see that each cohort mobility hazard turns down sharply near the end of the available data (though for the 2021–22 originations this downturn is simultaneous with the post-origination decline). For each cohort, the timing of the downturn corresponds to observations from calendar year 2022 or 2023, which in event time represent quarters 4–11 for 2021:Q1 originations but quarters 20–27 for 2017:Q1 originations.

The right panel of Fig. 2 shows hazards for a different outcome, closing the mortgage. This can precede moves when mortgages are refinanced or simply prepaid. The profile here is different. There is a prominent peak in each series that corresponds to calendar times around 2020. We interpret this as reflecting large-scale refinancing in the low interest rate environment of 2019–2021. The hazard of mortgage closing then falls in 2022 and 2023.

than time variation in market rates on a potential new mortgage. The coefficients are less than half the size of those in our preferred specifications and are much noisier — likely because including current quarter effects removes much of the identifying variation.

²⁹ Specifically, we include a rate gap main effect and rate gap-mortgage holder interaction, and report the latter coefficient. In IV specifications, our instruments are g^* (using zero for cash buyer observations) and a g^* -mortgage holder interaction.

5.2. Estimates of moving hazard

To aggregate the different series in Fig. 2 into a single quantitative estimate of the time profile of mobility, we fit a semi-parametric complementary log–log hazard model, (5).

We begin with a very flexible model that includes only a full set of calendar time (t) indicators, along with separate duration indicators (baseline hazards) for purchase loan and refinance observations. The upper left panel of Fig. 3 plots the estimated calendar time coefficients. (The estimated baseline hazard function $\lambda_0(d)$ is plotted in Appendix Figure B.3.³⁰) This shows a sharp decline of about 20 percentage points in the mobility hazard at the beginning of 2022. The timing of this decline lines up neatly with the rise in interest rates and the increase in predicted interest rate lock in Fig. 1. The mobility hazard declines further in late 2022 and 2023. There is also a temporary upward discontinuity in mobility in mid-2020, reflecting COVID effects.

The upper right panel shows coefficients from a similar model estimated on refinanced loans. The series is a bit noisier here, but shows a decline of similar magnitude in early 2022. The lower left panel shows a model for purchase loans (i.e., mortgages taken out on the purchase of the home — the complement of the refinance sample). This series is noisier still, and the 2020 jump is more notable. Nevertheless, we still see a sharp drop in mobility in 2022.

The final panel of Fig. 3 plots estimates for the sample of homeowners without mortgages (“cash buyers”). Here and in subsequent analyses, this sample is reweighted to match the mortgage holder sample on calendar time, cohort, zip code group, and homeowner age, credit score, and outstanding debt, as discussed in Section 4. There is no sharp change in mobility rates at the start of 2022 for this sample, though there is a small decline in mid-2022. A plausible interpretation of this delayed decline in mobility is that lower housing turnover affected even households that who did not need a mortgage, underscoring the importance of using homeowners as a comparison group rather than including renters, who would not be affected by changes in owner-occupied housing turnover.³¹

The pattern in Fig. 3 clearly points to interest rate changes in 2022–23 as drivers of the decline in mobility of mortgage holders in this period. To explore this, we move to a more parametric model that replaces calendar time effects with the rate gap measure defined above, plus controls.

The first three columns in the first row of Table 2 present coefficient estimates from (5). In column 1, we include just season fixed effects as controls (along with the baseline hazard, allowed to vary freely for purchase loans and refinances). Column 2 adds a linear time control, while column 3 adds origination cohort (quarter) fixed effects as well as controls for negative interest rate gaps (set to zero when the gap is positive) and a flexible polynomial in the change in log home values in the ZIP code from mortgage origination to present (specified as separate cubics in positive and negative changes). We see substantial negative effects on mobility rates. The -0.05 coefficient in column 1 implies that a one percentage point increase in the rate gap (e.g., a rise in interest rates from 4% to 5%) reduces the probability that a rate-locked

³⁰ The baseline hazard for refinance mortgages falls off surprisingly quickly after the first year after origination. This plausibly reflects misclassification of locations. Recall that we assign a mortgage’s location based on the mortgage holder’s address two quarters after origination, and classify a mortgage as a refinance if the holder does not move around the origination. If an individual buys a home with a mortgage, but does not move in until 3 or more quarters after the origination, we classify the mortgage as a refinance located in the pre-move location, and count the homeowner as moving out when in fact they move in. Our results are robust to estimating our mobility models excluding all quarters within as little as one quarter or as much as one year of origination.

³¹ When we repeat the analysis for renters, we do not find a discontinuity in mobility when rates rise, suggesting that renters’ mobility depends on different factors than homeowners.

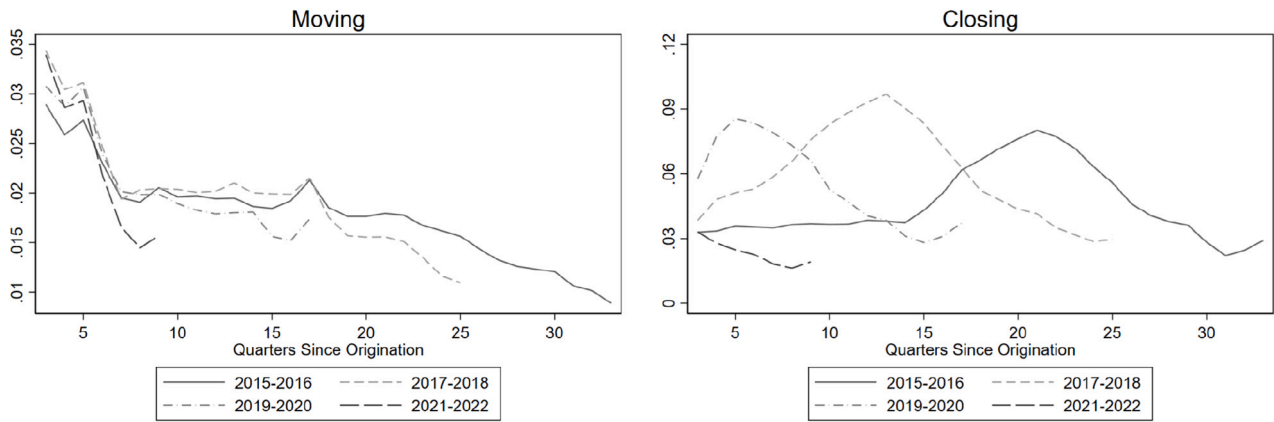


Fig. 2. Empirical hazards of moving and closing mortgage, by time since mortgage initiation. Notes: Panel A shows ZIP code moving hazards by quarter since origination. The share of moving households at time t is calculated for each cohort as the share of households moving between $t-1$ and t divided by the share that have not moved at $t-1$. Panel B shows mortgage closing hazards, calculated similarly. Mortgages may be closed when the household moves or by prepayment or refinancing.
Source: University of California (UC-CCP).

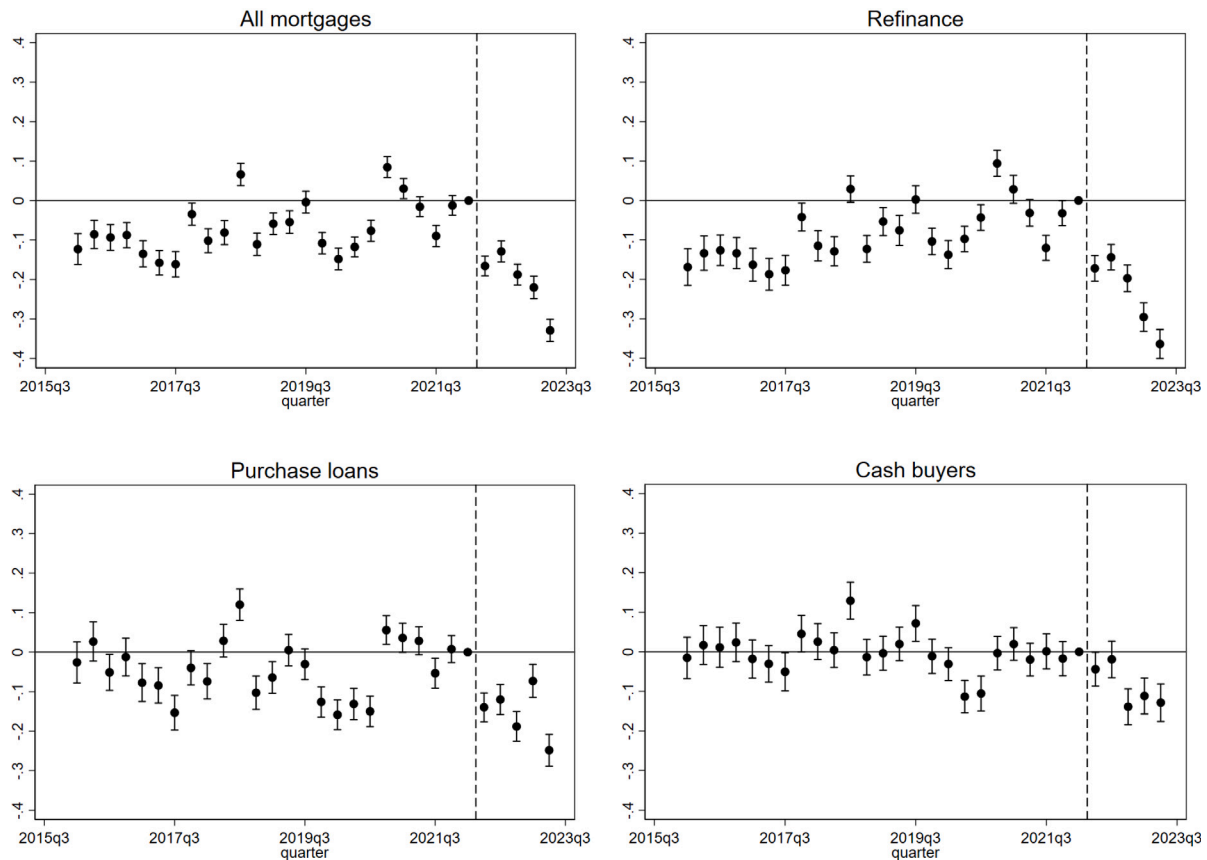


Fig. 3. Calendar time effects on mobility from complementary log-log hazard model, varying samples. Notes: Calendar time effects are from estimates of complementary log-log hazard models where failure is mobility out of the ZIP code. All models control for nonparametric baseline hazard in the elapsed time since the mortgage was originated ($t - O(i)$). Calendar time fixed effects are seasonally adjusted by subtracting the seasonal mean over the sample period.
Source: University of California (UC-CCP).

homeowner (with a mortgage rate below 4%) moves in a quarter by about 5%. This grows to nearly 8% with the additional controls.

The second row of Table 2 repeats these specifications, this time examining moves out of the state rather than out of the zip code. These coefficients are a bit smaller and more sensitive to the inclusion of a linear time trend, but with that control we find that a one percentage point increase in the rate gap reduces out-of-state mobility by about 4.3%.

Columns 4–6 of the table repeat the model using our instrumental variables specification, instrumenting for the current rate gap with the one that would apply if the focal mortgage had not been subsequently refinanced.³² This leads to very slightly larger effects, but generally does not change the results meaningfully.

³² First stage and reduced form specifications are reported in Appendix Table B.3.

Table 2

Estimates of the effects of interest rate gaps on mobility, varying samples.

Source: University of California Consumer Credit Panel (UC-CCP).

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Main sample						
Out-of-ZIP moves	−0.050 (0.003)	−0.076 (0.004)	−0.077 (0.005)	−0.052 (0.003)	−0.077 (0.004)	−0.079 (0.005)
Out-of-state moves	−0.011 (0.005)	−0.041 (0.005)	−0.043 (0.007)	−0.014 (0.005)	−0.044 (0.005)	−0.050 (0.007)
Controls						
Season FE	×	×	×	×	×	×
Linear time		×	×		×	×
Origination cohort FE			×			×
ZHVI & negative rate gap			×			×
Alternative sample (out-of-ZIP moves)						
Mortgage purchase	−0.026 (0.004)	−0.018 (0.004)	−0.023 (0.006)	−0.025 (0.004)	−0.018 (0.004)	−0.028 (0.006)
Refinance	−0.062 (0.004)	−0.105 (0.005)	−0.100 (0.008)	−0.065 (0.004)	−0.107 (0.005)	−0.101 (0.008)
Cash buyers	−0.034 (0.003)	−0.022 (0.003)	−0.022 (0.005)			
Stacked specification, main sample vs. cash buyers (out-of-ZIP moves)						
With time FEs	−0.014 (0.006)	−0.054 (0.007)	−0.056 (0.010)	−0.014 (0.006)	−0.050 (0.007)	−0.054 (0.012)
With time-cohort-location FEs	−0.018 (0.006)	−0.055 (0.007)	−0.066 (0.012)	−0.022 (0.007)	−0.055 (0.008)	−0.061 (0.015)

Notes: In this table, columns 1–3 present OLS estimates from equation 5. Columns 4–6 present IV estimates where the current rate gap is instrumented using the rate that would prevail if the focal mortgage had remained unrefinanced. Interest rate gap is the difference between the current market rate at time t and the market rate at the time the mortgage was last refinanced, g_{it} . “Mortgage Purchases” are mortgages taken out to finance a new purchase (identified from households who move into the ZIP code around the time of origination). “Cash buyers” are homeowners who do not have a mortgage at the time they move into a zip code, and are reweighted to match mortgage holders (purchases and refinances) on observables. Final panel includes both mortgage holders and cash buyers in the sample and interacts all variables (including the baseline hazard) with indicators for the subsample. Coefficients reported here are for the rate gap-mortgage holder interaction, controlling for the rate gap main effect. In IV specifications, the main effect and interaction are instrumented with the no-refinance instrument (set to zero for cash buyers) and its interaction with type. Standard errors in all panels are clustered at the origination quarter by zip home price index group level.

The second panel of the table shows estimates for several alternative samples. We first divide the main sample into purchase mortgages and refinances. Effects are notably larger for refinances, perhaps reflecting greater financial sophistication of households that have refinanced or nonlinear effects combined with larger rate gaps for households that refinanced to very low rates around 2019. (We explore this possibility further below.) In our preferred specification (column 5), a one percentage point increase in the rate gap reduces mobility of purchase mortgages by 2% and that of households holding refinanced mortgages by 11%.

The next row shows results for our sample of (reweighted) cash buyers. Here, we present only OLS specifications, as the IV strategy does not apply when refinancing is not possible. The coefficients are around −0.022, much smaller than for refinances or for our main sample (but comparable to those for mortgage purchases).

The final panel of the table shows difference-in-hazard specifications that contrast the mobility hazards of mortgage holders and cash buyers (after reweighting the latter to have similar observables as the mortgage holders, as discussed above). The first row includes calendar time (year-quarter) fixed effects, while the second further adds time-by-origination cohort-by ZIP code group fixed effects. Results are a bit smaller than what we find in our main sample, but still highly significant in specifications including time controls. This suggests that the dynamics we identify in our main results are not driven by secular changes in mobility or other aspects of the housing market (e.g., changes in home values) that would affect mortgage holders and owners without mortgages similarly.

Appendix Figure B.4 shows difference-in-hazard event study graphs, similar to Fig. 3 but showing the difference in fixed effects between mortgage purchasers and cash buyers. The first panel shows estimates without controls, equivalent to the difference between the first and last panels of Fig. 3. In the second panel we add controls as in Column 3 of

Table 2, excluding controls that are collinear with the time fixed effects we estimate to create the figure. We are reassured by the fact that the estimated decrease in mobility post-rate hikes is similar across panels, starting around −0.1 and falling to around −0.2 after three quarters. The rise in mobility after the start of the COVID-19 pandemic is visible in both figures, following what we see in the first panel of Fig. 3. Adding controls in the second panel of Figure B.4 reduces the noise in the pre-COVID period, making the effects of higher rates more prominent.

We have also explored whether the effects of rate gaps vary with homeowner or neighborhood characteristics. Table 3 presents estimates of heterogeneity along a number of dimensions. For each indicated characteristic, we divide our main sample in half, and estimate our main OLS and IV specifications in each. For example, the first row shows estimates for homeowners in zip codes with minority shares below 21.5% (“Low”) and above it (“High”). We find slightly more responsiveness in low-minority-share ZIP codes. In general, we see relatively little evidence of heterogeneous effects, though it does appear that older households are less affected (in the IV specification), as are those with low loan-to-value (proxied by having a principal that is above 84% of the average home value in the ZIP code).

5.2.1. Robustness tests

Our base models allow for the rate gap to have different effects when positive (and thus the household is potentially locked in) and when negative (which would make refinancing a potentially attractive option but should not affect mobility). But they otherwise force the effect to be linear. Fig. 4 loosens this. We divide the rate gap distribution (including both positive and negative values) into 20 ventiles, and then include indicators for 19 of them in our regression. The omitted category is the ventile corresponding to zero rate gap. For ventiles corresponding to rate gaps above about 1 percent, we see a clear negative effect of higher rate gaps on mobility. The trend is harder

Table 3

Heterogeneous effects of interest rate gaps on mobility by borrower characteristics.

Source: University of California Consumer Credit Panel (UC-CCP).

	OLS		IV	
	Low	High	Low	High
	(1)	(2)	(3)	(4)
Minority share of zip	-0.081	-0.069	-0.082	-0.071
High: share ≥ 0.215	(0.008)	(0.007)	(0.007)	(0.006)
Home values	-0.067	-0.080	-0.069	-0.083
High: value $\geq \$209,900$	(0.007)	(0.007)	(0.007)	(0.007)
Homeowner share	-0.067	-0.081	-0.070	-0.082
High: share ≥ 0.672	(0.007)	(0.007)	(0.007)	(0.007)
Median incomes	-0.064	-0.080	-0.066	-0.082
High: value $\geq \$61,774$	(0.006)	(0.007)	(0.006)	(0.007)
Age	-0.054	-0.099	-0.055	-0.100
High: age ≥ 45	(0.007)	(0.007)	(0.007)	(0.007)
Urban share	-0.080	-0.074	-0.082	-0.076
High: share ≥ 0.85	(0.007)	(0.007)	(0.006)	(0.007)
Credit Score at Origination	-0.070	-0.075	-0.072	-0.077
High: score ≥ 753	(0.006)	(0.009)	(0.006)	(0.009)
Mortgage term (months)	-0.089	-0.065	-0.092	-0.066
High: term (months) ≥ 360	(0.010)	(0.007)	(0.010)	(0.006)
Principal	-0.082	-0.070	-0.084	-0.072
High: principal $\geq \$213,750$	(0.006)	(0.008)	(0.006)	(0.008)
Principal to average zip home value	-0.090	-0.059	-0.092	-0.060
High: ratio ≥ 0.842	(0.007)	(0.007)	(0.007)	(0.007)
Season FE	X	X	X	X
Linear Time Controls	X	X	X	X

Notes: In this table, columns 1 and 2 present OLS estimates from Eq. (5) while columns 3 and 4 present IV estimates. Samples are divided according to the indicated variable; the dividing point is the median for all variables except urban share (where we use 0.85, and about 2/3 of observations are in the “high” category) and mortgage term (where we use 360 months, and about 3/4 of observations are in the “high” category). Standard errors are clustered at the origination quarter by zip home price index group level.

to see for lower rate gaps — there may still be a negative slope, but it is definitely smaller.

Appendix Figure B.6 presents a similar graph using only 10 deciles, allowing more precision but less flexibility. Here, we can see a clear downward trend as the rate gap moves from -1.5 to -0.5 , little effect of the rate gap between -0.5 and 1.0 , and then the same negative effect above 1.0 that we see in Fig. 4. Fonseca and Liu (2024) estimate a qualitatively similar pattern, with flattening out in the middle of the distribution.³³

Appendix Table B.4 presents more parametric models that allow for nonlinear effects of the rate gap via polynomial terms and indicators for gaps in specified ranges (e.g., for gaps above 2 percentage points). They indicate that the downward trend in mobility as the rate gap increases from -1.5 to $+1$ that is visible in Fig. 4 is statistically significant, but that even when this is included there is still a substantial, significant additional effect of rate gaps above $+1$.

An important question is whether the relationship between rate gaps and mobility changed in the recent period. The simple nonlinearity seen in Fig. 4 would mean different effects in 2022 and 2023 even

³³ Fonseca and Liu (2024) emphasize an apparent kink in the relationship for rate gaps below about -2 (in our notation). As Fig. 4 indicates (see also Appendix Figure B.5), we have very few gaps that negative in our sample. This reflects differences in measurement. We compute the rate gap using the market rate at the time of origination, while Fonseca and Liu (2024) main estimates use the imputed actual rate on the mortgage. Using the market rate reduces the variation, but guards against bias from unobserved characteristics that influence the origination rate. By using a linear regression, we also avoid the bias in the binscatter command described by Cattaneo et al. (2024).

with fixed coefficients, as the distribution of rate gaps in this period is so different from what has been seen before, but there could be even larger changes if the response function has changed as well. Appendix Figure B.6 shows estimates both for the full sample and the pre-2022 subsample. As it indicates (see also Appendix Figure B.5), there are essentially no observations before 2022 with rate gaps corresponding to the ventiles with large effects in Fig. 4. However, in the range where the two distributions overlap, there is little indication of changes in the response function in 2022.

Appendix Table B.5 presents a number of additional robustness tests. The first row repeats our primary specification. The second row shows results where we cluster by origination quarter (rather than origination quarter-by-price growth group, as in the main results). Rows 3, 4, and 5 present alternative specifications of the hazard model (i.e., of the link function in (3) — first a Poisson model, then using the log hazard or the hazard itself as the dependent variable in (5).³⁴ These yield broadly similar results.³⁵

Rows 6 and 7 vary the controls — 6 removes entry quarter (cohort) effects, while 7 adds year-by-quarter fixed effects. The former has minimal effect, indicating that our estimates are driven by variation in the contemporaneous rate, not by variation in the origination rate.

³⁴ We do not present an IV version of the Poisson model.

³⁵ The rate gap coefficient on the untransformed hazard is -0.0012 . This specification identifies an effect in percentage points, where our main specifications estimate percentage effects. The sample average hazard is 0.034 (Table 1), so a -0.12 percentage point effect corresponds to a percentage effect of -0.035 at the sample mean — smaller than in our other estimates, but not wildly different.




Fig. 4. Nonlinear effects of rate gap. Notes: Figure shows coefficients on indicators for 19 ventiles of the rate gap, including both positive and negative values. Specification otherwise matches Table 2, column 3. The excluded category is zero rate gap, accounting for 19% of the sample. Spikes show 95% confidence intervals. Source: University of California (UC-CCP).

In the latter specification, all time series variation in r_t is absorbed by the fixed effects, and identification comes only from differences in rates among households that originated (or refinanced) at different times. Not surprisingly, this has a large effect – standard errors more than triple, and point estimates shrink – we cannot reject either zero effects or substantial ones. There is simply little variation in origination rates with which to identify the effect. As variation in rate lock coming from different origination times could be correlated to other differences in cohort characteristics, such as differences in economic experiences or demographic changes, we find it reassuring that our estimates are not identified by cross-cohort variation.

The next panel of the table explores alternative samples. We first limit to mortgages originated before the COVID pandemic, then consider only moves in 2015 and thereafter, and then exclude moves in 2022 and 2023. The first two have only minor impacts on the results. The final change eliminates the effect, while tripling the standard errors, indicating that our identification comes primarily from the recent increase in rates (consistent with our above discussion).

Finally, we present two sets of estimates that vary the way that we calculate an observation's elapsed duration. These address the concern that we calculate duration somewhat differently for cash buyers and mortgage holders in our main results. In the first row of the last panel, we re-calculate duration for refinance observations, measuring it as time since the household moved into the ZIP rather than as time since the focal refinance. This makes the refinance measurement more similar to the way we measure duration for cash buyers. It has little effect on the results. In the last row, we return to our cash buyer sample, but expand it to include households who moved into the ZIP code between 2004 and 2013 (where previously we included only post-2013 entrants in this sample, corresponding to post-2013 mortgage originations in our main sample). Again, this has little effect. This provides assurance that measurement differences are not driving the contrasting responsiveness of mortgage holders and cash buyers.

5.3. Discussion

In this section, we explore the aggregate implications of mortgage rate lock. We are particularly interested in two questions. First, how would aggregate mobility have been different if no households were

locked in — for example, because there had been no rate increase, or because mortgages were assumable? Second, how large are the costs of mortgage lock, and how are the costs and benefits distributed? We answer these questions using the estimates from the hazard model shown in Table 2 and considering the implications for the population of mortgage borrowers.

To calculate the effects of rate lock on aggregate mobility, we use Eq. (5) to estimate how much higher cohort-specific hazard rates would have been if mortgage lock were zero. Denoting as λ_c^0 cell c 's counterfactual hazard rate with no rate lock,

$$\lambda_c^0 = 1 - (1 - \lambda_c) \exp(-\hat{\gamma} g_c) \quad (6)$$

where $\hat{\gamma}$ is the estimated coefficient on the rate gap. We calculate λ_c^0 using the effect on ZIP code mobility from Column (3) of Table 2 for $\hat{\gamma}$. Then we aggregate the cell-specific counterfactual mobility rates to calculate how overall mobility would have been different if rate lock were zero for all cohorts.³⁶ Appendix Figure B.1 shows the actual and counterfactual mobility hazards by quarter.³⁷ The effects of rate lock on mobility hazards vary over time. In 2021q2, when the average rate gap was essentially zero, both the actual and counterfactual quarterly ZIP code mobility hazard in our sample were around 1.5 percent. In 2022q2, the counterfactual quarterly mobility hazard was 1.6 percent, which was 0.16 percentage points above the actual quarterly mobility hazard of 1.44 percent. By 2023q2, the counterfactual and actual mobility hazards were 1.72 and 1.45 percent per quarter respectively. Over the entire last year of our sample, which extends from 2022q3 to 2023q2, rate lock reduced cumulative mobility from about 7.3% to 6.1%, a decrease of about 1.2 percentage point or 16 percent. This amounts to 800,000 moves across ZIP codes that were prevented.

³⁶ This aggregation benefits from the use of a hazard model rather than a linear probability model for the effect of rate lock — we are able to account for the fact that a household that is dissuaded from moving in one quarter remains in the risk pool the following quarter.

³⁷ One implication of the results in Appendix Figure B.1 is that mobility would have risen in 2022 above its 2020–21 level, but for rate lock. Batzer et al. (2024) also find increases in counterfactual mobility that were masked by rate lock, though their rate lock effects are much larger than ours.

Aggregating over the longer period from the time interest rate hikes began in mid-2021, around 1,000,000 fewer people with a mortgage moved than would have if there had been no rate lock.

The magnitude of these estimates implies that aggregate mobility was reduced by interest rate lock. Because approximately one-quarter of adults hold mortgages, a decline in quarterly mobility over the last year of the sample from 1.85% to 1.60% (a 0.25 percentage point decline) reduces *overall* mobility by 0.06 percentage points. On average, mobility has declined by a bit more than half over the last half century, or by about 0.017 percentage points per quarter in the quarterly rate. The decline that we attribute to mortgage lock since 2022q3 accelerated this secular decline by about one year. These calculations leave out potentially important equilibrium effects coming from lower housing transaction volumes and higher search costs.

Another way to understand the magnitude of our results is to compare them to estimates of the elasticity of migration with respect to local net-of-tax rates. Head and Mayer (2021) tabulate estimates of this parameter in their Table B.1. They report a median estimate of 1.63; estimates for state-to-state mobility in the U.S. are 1.21 (Suárez Serrato and Zidar, 2016), 1.81 (Moretti and Wilson, 2017), 1.73 (Fajgelbaum et al., 2019), and 2.69 (Bryan and Morten, 2019). As shown in Fig. 1, by October 2022, the rate gap in dollar terms reached about 3 percentage points and over \$400 per month. For households moving between July 2022 and June 2023, the average rate gap was about 2.7 percentage points, which corresponded to a difference in annual payments of \$4928. Assuming \$5000 is 5% of household income, our estimates imply an elasticity of ZIP code mobility of 15% / 5%, or about 3. If we instead use our estimated effects on between-state mobility, which are a bit more than half as large, we obtain an elasticity around 1.5, very much in line with other studies.

Using the mortgage amortization formula and the characteristics of outstanding mortgages, we calculate how much higher each household's principal balance would have to be for their monthly payments to be the same as if they refinanced their actual mortgage at existing rates. We find that a 2.7 percentage point rate gap is equivalent to a difference in average mortgage balance value – the “price” of moving – of \$49,400. Our approach to calculating the price of moving implicitly uses market rates' embedded expectations about refinancing rates and the path of future Fed Funds rate movements. An alternative approach to calculate the costs of rate lock would be to simulate the path of future interest rates while making assumptions about refinancing behavior. A disadvantage of this latter approach is that it requires assumptions about what households will do. Given that fewer than half of households refinance optimally (Agarwal et al., 2016), we prefer an approach without that assumption.

Mortgage lock creates deadweight loss when households are deterred from moving by rate gaps. One way to calculate the deadweight loss is to consider the demand for moving as a linear function of the present value of the payment gap. By integrating under the demand curve for moving, we can calculate the welfare losses. We assume that moving demand is linear and decreasing in wealth. Our hazard estimates imply that in the last year of the sample, around 1.2 percentage point more people would have moved if there had been no rate lock. Together with the wealth decrease we calculate, the decrease in moving implies a deadweight loss of about \$296 per household in 2022q3–2023q2, the last year of the sample. In aggregate, these losses amount to around \$20bn in that year. Our deadweight loss calculation implicitly includes the costs of all reasons that a household might want to move yet be deterred from doing so. This is preferable to an approach based on the effects of a particular motivation for moving, e.g., for a new job, as that would underestimate the overall costs of rate lock.

Another effect of rate lock is that households who do move must pay higher rates. The \$49,400 difference in present values is a pure transfer from households to the banking system which would be saved if households could retain their mortgage after moving. Since 6.1% of sample

households moved in the last year of the sample, the monthly per-household cost was around \$315 per household, implying a present-value balance transfer to the banking system worth \$3160 per household. In aggregate, the costs amount to \$21bn per month in mortgage payments, or a balance transfer worth \$215bn in total.

6. Concluding remarks

We estimate that interest rate lock has a substantial effect on individuals' propensity to move ZIP codes. Our preferred specification is a hazard specification that models ZIP code moving probability as a function of the gap between the rate a household is paying for its mortgage and the current prevailing mortgage rate. The hazard model implicitly controls for the baseline hazard rate, which is modeled as a function of the time since a household has a mortgage.

Our preferred estimates come from instrumental variables models which instrument for the interest rate on a mortgage using the prevailing rate at the time of mortgage origination. The IV specifications show that each percentage point increase in the gap between the mortgage's rates and prevailing rates reduces mobility between 5% and 8%. When we repeat the specifications for homeowners without mortgages, the estimates are much smaller. Therefore we think that macroeconomic conditions or other omitted variables do not explain our results.

CRedit authorship contribution statement

Jack Liebersohn: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jesse Rothstein:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jfineco.2024.103973>.

Data availability

Data for “Household Mobility and Mortgage Rate Lock” (Reference data) (Mendeley Data)

References

- Agarwal, S., Rosen, R.J., Yao, V., 2016. Why do borrowers make mortgage refinancing mistakes? *Manage. Sci.* 62 (12), 3494–3509.
- Aladangady, A., 2017. Housing wealth and consumption: evidence from geographically linked microdata. *Amer. Econ. Rev.* 107 (11), 3415–3446.
- Andersen, S., Campbell, J.Y., Nielsen, K.M., Ramadorai, T., 2020. Sources of inaction in household finance: Evidence from the danish mortgage market. *Amer. Econ. Rev.* 110 (10), 3184–3230.
- Andersson, F., Mayock, T., 2014. How does home equity affect mobility? *J. Urban Econ.* 84, 23–39.
- Atalay, K., Edwards, R., 2022. House prices, housing wealth and financial well-being. *J. Urban Econ.* 129, 103438.
- Batzer, R., Coste, J., Doerner, W., Seiler, M., 2024. The Lock-In Effect of Rising Mortgage Rates. FHFA Staff Working Paper Series.

- Bernstein, A., Struyven, D., 2022. Housing lock: Dutch evidence on the impact of negative home equity on household mobility. *Am. Econ. J.: Econ. Policy* 14 (3), 1–32.
- Bolhuis, M.A., Cramer, J.N.L., Schulz, K.O., Summers, L.H., 2024. The Cost of Money is Part of the Cost of Living: New Evidence on the Consumer Sentiment Anomaly. Working Paper 32163, National Bureau of Economic Research, <http://dx.doi.org/10.3386/w32163>, URL <http://www.nber.org/papers/w32163>.
- Brown, J.R., Cookson, J.A., Heimer, R.Z., 2019. Growing up without finance. *J. Financ. Econ.* 134 (3), 591–616.
- Bryan, G., Morten, M., 2019. The aggregate productivity effects of internal migration: Evidence from Indonesia. *J. Polit. Econ.* 127 (5), 2229–2268.
- Cattaneo, M.D., Crump, R.K., Farrell, M.H., Feng, Y., 2024. On binscatter. *Amer. Econ. Rev.* 114 (5), 1488–1514.
- Chaney, T., Sraer, D., Thesmar, D., 2012. The collateral channel: How real estate shocks affect corporate investment. *Amer. Econ. Rev.* 102 (6), 2381–2409.
- Fajgelbaum, P.D., Morales, E., Suárez Serrato, J.C., Zidar, O., 2019. State taxes and spatial misallocation. *Rev. Econ. Stud.* 86 (1), 333–376.
- Ferreira, F., 2010. You can take it with you: Proposition 13 tax benefits, residential mobility, and willingness to pay for housing amenities. *J. Public Econ.* 94 (9–10), 661–673.
- Ferreira, F., Gyourko, J., Tracy, J., 2010. Housing busts and household mobility. *J. Urban Econ.* 68 (1), 34–45.
- Ferreira, F., Gyourko, J., Tracy, J., 2011. Housing Busts and Household Mobility: An Update. Working Paper 17405, National Bureau of Economic Research.
- Fonseca, J., Liu, L., 2024. Mortgage Lock-In, Mobility, and Labor Reallocation. Working Paper.
- Foote, A., 2016. The effects of negative house price changes on migration: Evidence across US housing downturns. *Reg. Sci. Urban Econ.* 60, 292–299.
- Gerardi, K., Willen, P.S., Zhang, D.H., 2023. Mortgage prepayment, race, and monetary policy. *J. Financ. Econ.* 147 (3), 498–524.
- Goodman, L., Ratcliffe, J., Neal, M., Choi, J.H., Zhu, L., Walsh, J., Young, C., Pang, D., Zinn, A., Visalli, K., Mehrotra, A., Rincon, A., Neeley-Bertrand, D., Hill, T., Barcus, A., 2023. Housing Finance At A Glance: A Monthly Chartbook, May 2023. Technical Report, Urban Institute, URL <https://www.urban.org/research/publication/housing-finance-glance-monthly-chartbook-may-2023>.
- Hanushek, E.A., Quigley, J.M., 1978. An explicit model of intra-metropolitan mobility. *Land Econom.* 54 (4), 411–429.
- Head, K., Mayer, T., 2021. The United States of Europe: a gravity model evaluation of the four freedoms. *J. Econ. Perspect.* 35 (2), 23–48.
- Holmes, N., 2021. CalExodus: Are People Leaving California? Policy Brief, California Policy Lab.
- Holmes, N., White, E., 2022. Pandemic Patterns: California Is Seeing Fewer Entrances and More Exits. Policy Brief, California Policy Lab.
- Molloy, R., Trezzy, R., Smith, C.L., Wozniak, A., 2016. Understanding declining fluidity in the US labor market. *Brook. Pap. Econ. Act.* 2016 (1), 183–259.
- Moretti, E., Wilson, D.J., 2017. The effect of state taxes on the geographical location of top earners: Evidence from star scientists. *Amer. Econ. Rev.* 107 (7), 1858–1903.
- Munch, J.R., Svarer, M., 2002. Rent control and tenancy duration. *J. Urban Econ.* 52 (3), 542–560.
- Palmer, C.J., 2022. An IV Hazard Model of Loan Default with an Application to Subprime Mortgage Cohorts. Working Paper.
- Quigley, J.M., 1987. Interest rate variations, mortgage prepayments and household mobility. *Rev. Econ. Stat.* 636–643.
- Quigley, J.M., 2002. Homeowner mobility and mortgage interest rates: new evidence from the 1990s. *Real Estate Econ.* 30 (3), 345–365.
- Shahidinejad, A., 2023. Are (Nonprofit) Banks Special? The Economic Effects of Banking with Credit Unions. Working Paper.
- Suárez Serrato, J.C., Zidar, O., 2016. Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms. *Amer. Econ. Rev.* 106 (9), 2582–2624.
- Venti, S.F., Wise, D.A., 1984. Moving and housing expenditure: transaction costs and disequilibrium. *J. Public Econ.* 23 (1–2), 207–243.