FEDERAL RESERVE BANK *of* ATLANTA **WORKING PAPER SERIES**

Mortgage Lock-in: A Review of the Literature

Kristopher Gerardi, Franklin Qian, and David Zhang

Working Paper 2026-11
August 2026

Abstract: This paper reviews the rapidly growing literature on mortgage lock-in and its effects on household mobility, labor reallocation, home sales, owner-occupied prices, and spillovers into rental and commercial markets. We then discuss two future research directions: (1) separately identifying lock-in effects from the direct effects of the existing interest rate on the mortgage, which have different implications for household behavior as market interest rates rise, and (2) assessing the general-equilibrium consequences of reduced mobility on labor markets and productivity. We close by discussing policy responses, including portability, broader assumability, and a Danish-style delivery option.

JEL classification: G18, G21, E52

Key words: mortgage lock-in, fixed-rate mortgage (FRM), household mobility, house prices, rents

https://doi.org/10.29338/wp2026-11

The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Kris Gerardi is with the Federal Reserve Bank of Atlanta. Franklin Qian is with the University of North Carolina. David Zhang is with Rice University. Please address questions regarding content to Kristopher Gerardi, Federal Reserve Bank of Atlanta, 1000 Peachtree Street NE, Atlanta, GA 30309, kristopher.gerardi@atl.frb.org.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's website at www.atlantafed.org. Click "Publications" and then "Working Papers." To receive e-mail notifications about new papers, use atlantafed.org/forms/subscribe.

I. Introduction

Mortgage lock-in refers to situations in which homeowners are paying lower interest rates on their long-term, fixed-rate mortgages (FRMs) compared to the prevailing market rates associated with similar new loans. Since mortgages in the U.S. are not portable or assumable, in such scenarios, homeowners have a financial incentive to postpone moving in response to other types of economic shocks.[1] The economic effects of mortgage lock-in have generated significant recent interest among both researchers and policy makers due to the large increase in mortgage rates that occurred beginning in early 2022, which followed several years of extremely low rates during the COVID-19 pandemic, and left many homeowners with below-market mortgage rates.

Figure 1 illustrates these dynamics. The solid black line shows the path of the average rate of newly originated 30-year FRMs over the past decade according to the Freddie Mac Primary Mortgage Market Survey (PMMS), while the red shaded bars show the share of outstanding 30-year FRMs with an interest rate that is at least 2 percentage points lower than the PMMS rate. Long-term mortgage rates fell from approximately 5 percent in late 2018 to below 3 percent in early 2021 before rising to almost 7 percent in late 2022. Many households purchased homes and/or refinanced existing mortgages to take advantage of the low rates during the pandemic period, which resulted in many locked-in homeowners beginning in early 2022. By early 2023, more than 80 percent of outstanding long-term FRMs had interest rates that were at least 2 percentage points below prevailing market rates. Since mortgage rates have stayed at relatively high levels for the past 4 years, the fraction of locked-in loans has also remained relatively high, at over 45 percent, through the end of 2025 (the most up-to-date data when this article was written).

[1]There are a few types of mortgages in the U.S. that are assumable, including loans insured by the Federal Housing Administration (FHA), Veteran Affairs (VA), and the U.S. Department of Agriculture (USDA). However, these loans account for a minority of U.S. mortgages and are only assumable subject to lender approval.

Figure 1: Share of Rate-Locked FRMs

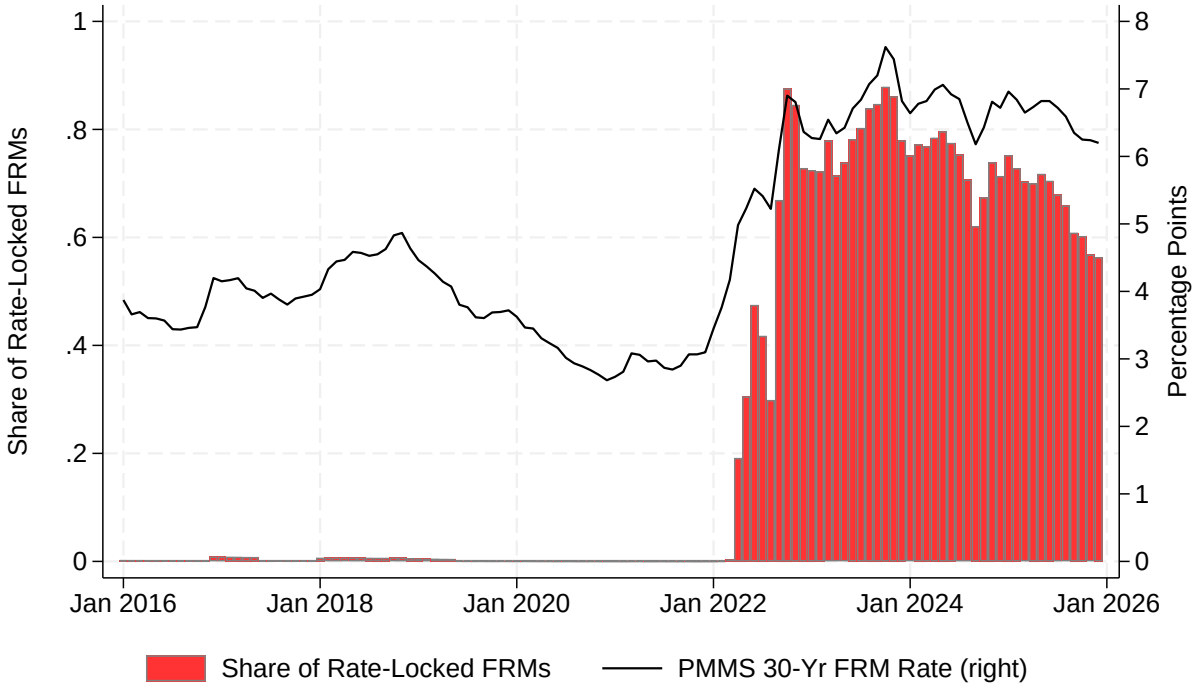

Notes: This figure shows the fraction of Outstanding 30-Year FRMs with an interest rate that is at least 200 basis points below the prevailing average rate for newly originated FRMs according to the Freddie Mac Primary Mortgage Market Survey (PMMS). Underlying data on mortgage rates come from Intercontinental Exchange's McDash servicer database.

A burgeoning academic literature has focused on the economic effects of mortgage lock-in. In this chapter, we review this literature, beginning with the relatively sparse pre-pandemic literature that focused exclusively on the mobility effects of lock-in. We then turn to a larger post-pandemic literature, which, in addition to revisiting the effects of lock-in on household mobility, has studied how lock-in affects labor market reallocation patterns and how it impacts housing market liquidity and prices. After briefly reviewing this literature, we pivot to a discussion of the direction in which we believe future research should turn to, focusing specifically on two challenges: cleaner identification of causal effects and general equilibrium considerations for labor markets and firm production. Finally, we close the chapter with a brief discussion of potential policy solutions. The main driver of lock-in is the primacy of the long-term fixed-rate mortgage in the U.S. mortgage market. The simplest way to eliminate lock-in would be to shift the market toward adjustable-rate loans, but given the extreme popularity of the FRM, we focus on more pragmatic solutions that have successful track records in other countries.

II. Rate-lock Effects on Mobility

a. The Early Literature

Until the early-to-mid 2020s, the academic literature on mortgage lock-in was sparse. The limited number of studies on the topic focused exclusively on household mobility effects. To our knowledge, Quigley (1987) was the first paper to analyze the effects of lock-in on mobility rates. The paper estimated a Cox proportional hazard model of household mobility using data from the Panel Study of Income Dynamics PSID from 1979-1981, a period characterized by rapidly rising interest rates.[2] To measure lock-in it constructed the mortgage "premium" for each household, defined as the difference between the present value of remaining mortgage payments at the contract interest rate and the present value of remaining payments at the

[2]Average mortgage rates rose from approximately 9 percent to 15 percent between 1977 and 1982.

prevailing market interest rate. The estimated models yielded highly statistically significant, negative effects of increased lock-in on household mobility rates. In a follow-up study, Quigley (2002) again used PSID data and a similar empirical methodology to estimate the effect of lock-in on household mobility in the early 1990s, a period characterized by declining interest rates that made mortgage lock-in less severe. The results were largely consistent with those of Quigley (1987). Finally, in a related analysis, Potepan (1989) showed that mortgage lock-in partially explained the significant increase in home improvement expenditures during the late 1970s and early 1980s. The paper also used data from the PSID and documented that lock-in increases the likelihood that households choose to stay and make home improvements to their current homes rather than move and upgrade their properties.[3]

b. Post-Pandemic Literature

As lock-in has become a pervasive issue in the U.S. mortgage market in the post-pandemic period (Figure 1), a second wave of studies has emerged that re-examines its economic effects. Whereas the early literature discussed above primarily used survey data, the newer literature has used administrative data from credit bureaus and mortgage servicers to create much larger samples.

Recent studies on mortgage lock-in have considered numerous economic outcomes, but in this subsection, we focus narrowly on mobility effects. These mobility estimates are collated in Table 1. Fonseca and Liu (2024) and Liebersohn and Rothstein (2025) build samples from large nationally representative consumer credit panels, the GCCP (Gies Consumer and Small Business Credit Panel) and the UC-CCP (University of California Consumer Credit Panel), respectively, that track borrowers across moves at an annual or quarterly frequency. Aladangady et al. (2025), Gerardi et al. (2024), and Katz and Minton (2026) instead use the CRISM (Credit Risk Insight Servicing McDash) database, which links Equifax credit

[3] Schwab (1985) developed a theoretical model of a rate-locked homeowner's choice between moving to a different home and renovating their current house.

records to mortgage servicing records providing detailed borrower-level loan characteristics. Aladangady et al. (2025) and Katz and Minton (2026) supplement CRISM with deeds data, while Gerardi et al. (2024) additionally merge it with the New York Fed Consumer Credit Panel.

The point estimates in the last column of Table 1 mostly line up in terms of a per-percentage point interest rate effect on mobility (Fonseca and Liu, 2024; Liebersohn and Rothstein, 2025; Aladangady et al., 2025), clustering between roughly 8% and 11%. The two papers that aggregate their per-period estimates over the 2022–2023 rate-hike episode also reach almost identical conclusions about the effect of lock-in: Liebersohn and Rothstein (2025) attribute a 16% reduction in mortgage-holder mobility to lock-in over 2022Q3–2023Q2, while Gerardi et al. (2024), working from a net present value measure, ΔNPV, that is similar to the one developed in Quigley (1987) rather than Δr, the simple difference between the borrower's rate at origination and the prevailing market rate, arrive at a 17% aggregate reduction over the post-2021 period. One clear outlier is Katz and Minton (2026), which shows that following the November 2016 rate hike, a 1 pp lower outstanding mortgage rate leads to a $\sim$42% lower 4-year sale probability. Fonseca and Liu (2024)'s estimated effect of Δr over the 2022–2024 period is also significantly larger than the estimates of Liebersohn and Rothstein (2025) when accounting for the over 3 percentage point increase in mortgage rates over 2022. Some of these differences are likely due to different approaches to identifying the causal effects of lock-in, which we discuss below in Section 4.

Table 1: Recent empirical literature on mortgage rate lock-in and household mobility

Paper	Sample period	Data	Primary measure of lock-in	Effect on mortgage-holder mobility
Fonseca and Liu (2024)	2010–2024	GCCP (1% credit-bureau panel)	Δr = origination rate − market rate	1 pp lower $\Delta r \Rightarrow$ 9% lower annual mobility on average; 1 pp lower $\Delta r \Rightarrow$ 16% lower mobility in 2022–2024.
Liebersohn and Rothstein (2025)	2013–2023	UC-CCP (Experian, 2% credit panel)	Rate gap = market rate − origination rate (positive ⇒ locked in)	1 pp higher rate gap ⇒ 7.7% lower quarterly mobility; 16% reduction in mortgage-holder mobility over 2022Q3–2023Q2.
Aladangady et al. (2025)	2009–2023	CRISM + Cotality deeds	Borrower-specific imputed rate gap (imputed origination rate − imputed market rate), with control for maximum forgone rate gap to date	1 pp lower rate gap ⇒ ≈11% lower annualized mobility.[a]
Gerardi et al. (2024)	2005–2024	CRISM–NY Fed CCP + Cotality deeds/MLS	ΔNPV: NPV of remaining mortgage-payment savings at origination rate vs. market rate	A \$10,000 decline in ΔNPV lowers overall quarterly mobility by ≈7% (IV); the post-2021 average shock of −\$23,556 implies a ≈17% reduction in mortgage-holder mobility over that period.
Katz and Minton (2026)	2016–2023	CRISM + CoreLogic deeds	Prepayment cost: mortgage balance − PV of remaining payments at the outstanding rate, scaled by property value	Around the November 2016 rate hike, 1 pp lower outstanding mortgage rate ⇒ ≈42% lower 4-year sale probability.

Notes: GCCP = Gies Consumer and Small Business Credit Panel. CRISM = Equifax Credit Risk Insight Servicing McDash. UC-CCP = University of California Consumer Credit Panel. CCP = NY Fed/Equifax Consumer Credit Panel. All effects in the last column are signed as mobility reductions in response to greater lock-in. [a] Computed from Table E.4 of Aladangady et al. (2025): between rate gap midpoints of −1.081 and +1.057, the annualized move probability rises from 6.788 to 8.304, a slope of ≈0.71 pp per 1 pp rate gap, or ≈11% relative to the paper's unconditional mean annualized move hazard of 6.4%.

Beyond the overall mobility estimates, the post-pandemic literature report interesting heterogeneity in lock-in responses. Fonseca and Liu (2024) document an asymmetry in the response at extreme positive values of the gap between the origination rate and the market rate: the effect on mobility is much weaker for high positive levels of Δr, consistent with the fact that borrowers whose origination rate exceeds the market rate can capture the

savings by refinancing rather than moving. Liebersohn and Rothstein (2025) find that the rate-gap response is strongly nonlinear, with essentially no effect until the rate gap exceeds about 1.5 percentage points and a substantial effect above that threshold; this nonlinearity is central to their argument that lock-in had little bite before 2022, when few borrowers had gaps that large. Aladangady et al. (2025) find that the rate-gap response is concentrated in within-CBSA and short-distance (under 30-mile) moves, while cross-CBSA and long-distance moves are largely insensitive to rate gaps. Gerardi et al. (2024) show that the lock-in effect is concentrated entirely among moves that involve a home sale, with no detectable effect on moves where the borrower retains their mortgage, and that younger borrowers are roughly twice as sensitive to lock-in as older borrowers in raw magnitudes, reflecting their higher baseline mobility. Katz and Minton (2026) document that the sale-probability response to lock-in is sharper for borrowers with larger outstanding balances and longer remaining terms, for whom the present-value cost of giving up a below-market rate is largest.

c. Implications for Labor Reallocation and Housing Mismatch

The post-pandemic literature has further explored the effect of reduced mobility on labor reallocation and housing mismatch, with a different emphasis on each. Focusing on labor reallocation, Fonseca and Liu (2024) finds an effect of lock-in on migration to nearby areas with wage shocks as well as on self-employment. Feng et al. (2026) further studies the effect of lock-in on job changes and the patenting behavior of inventors. In contrast, Aladangady et al. (2025) argues that because the cross-CBSA effects of mortgage lock-in are relatively low, the labor reallocation effects of lock-in are limited, and the more welfare-relevant effect of lock-in is on housing mismatch.[4] Liebersohn and Rothstein (2025) use a revealed preference approach to estimate the welfare losses from mortgage lock-in due to a combination of labor reallocation and housing mismatch effects, and find them to be substantial at $20 billion between 2022Q3 and 2023Q2.

[4] Gerardi et al. (2024) shows that for younger borrowers, the cross-CBSA effects of lock-in are still significant, suggesting that it may have an effect on labor reallocation for households early in the job ladder.

III. Lock-in Effects on Housing Markets

A second, more recent strand of the literature has shifted attention from household mobility to the consequences of mortgage lock-in for housing markets themselves: the volume and composition of transactions, the level and dispersion of prices, the time it takes to sell, and the spillovers to rental and commercial markets. The clearest finding, on which the literature converges, is that lock-in substantially reduces sales and listings. The harder and more contested question is whether lock-in raises prices, since the same friction operates on both the demand and supply sides of the market. In this section, we briefly summarize the main findings of these studies, and then turn to outcomes that go beyond prices to spillovers into adjacent markets.

a. Lock-In Reduces Listings and Sales

The most robust empirical fact in this literature is that mortgage rate lock-in substantially reduces home sales and listings, in addition to household mobility. The papers that measure home sales reach this conclusion using different data, designs, and units, but their estimates are qualitatively similar. Fonseca and Liu (2024) shows a significant effect of county-level rate gaps on county-level listings. Using individual-level data on home sales, Gerardi et al. (2024) decompose moves into those involving a sale and those not involving a sale, and show that the lock-in effect is concentrated almost entirely in the former. Aladangady et al. (2025) shows a significant effect of lock-in on home sales and market tightness. Katz and Minton (2026) exploit a discontinuity in mortgage rates created by the spike in Treasury yields after the November 2016 U.S. presidential election, finding that a one-percentage-point increase in the contract rate raises the four-year sale probability by more than one-third.

b. Does Lock-In Raise Owner Prices?

While there is a broad consensus that lock-in lowers property listings and sales, a more contested question is whether it raises owner-occupied house prices. The nature of this debate is the fact that the fixed-rate mortgage friction operates on both the demand and supply sides of the market: the owner who refrains from selling is, in many cases, also the household that would have bought a different home. If sellers and buyers were perfectly symmetric, the supply and demand effects of lock-in would cancel, leaving prices unaffected. There are two potential channels: first, lock-in could affect own-to-rent transitions, which would build in natural asymmetry since some of the sellers affected by lock-in could not re-buy. Second, lock-in could affect the size of housing demanded.

The first account, developed by Katz and Minton (2026) and Amromin and Eberly (2024), emphasizes incomplete substitution between owner-occupied and rental housing: some locked-in households would have transitioned from owning to renting absent lock-in, and these missing exits remove owner supply without removing owner demand. Katz and Minton (2026)'s November 2016 event study shows that a one-percentage-point higher fixed rate raises the four-year overall sale probability by 3.1 percentage points (a 33-percent semi-elasticity) and the four-year own-to-rent transition probability by 1.3 percentage points (a 36-percent semi-elasticity). Since own-to-rent transitions are a subset of sales, roughly 42 percent (1.3/3.1) of lock-in-suppressed sales are sellers who would have become renters, contracting owner supply without a matching contraction in owner demand. To estimate equilibrium price effects, the authors instrument local mortgage value with a twin-birth IV that exploits finite-sample variation in family-size shocks to predict accumulated mobility, and find that cross-sectional variation in mortgage value explains approximately 78 percent of 2021–2023 local price growth. Their structural dynamic model, which incorporates both the user-cost effect of higher rates and the supply-restricting effect of lock-in, predicts a substantially smaller decline in real house prices from the post-2022 tightening than user-cost models

without lock-in. Amromin and Eberly (2024) reach a parallel conclusion through their own structural-calibration approach, with own-to-own moves canceling out in their market-clearing equation so that the price effect operates entirely through the own-to-rent transition margin. A model with full turnover (i.e., no listing friction, so any owner can sell each period) predicts that 2022 mortgage rates should have lowered prices by roughly 37 percent from their peak. But introducing a Calvo-style moving friction (in which only a random fraction of owners list each period, calibrated to match observed 2022 listings) shrinks the predicted decline to about 15 percent, and updating the model to reflect the high incomes of actual 2023 buyers shrinks it further to roughly 5 percent. This account is also consistent with the out-of-sample model tests of Gerardi et al. (2024).

The second account, due to Mabille et al. (2024), focuses on the segmentation in housing size rather than in tenure. The paper presents reduced-form evidence showing that the same one-percentage-point reduction in Δr that lowers overall moves by 9.6 percent lowers downsizing moves by 15.4 percent and upsizing moves by only 3.2 percent, with essentially no effect on exits from homeownership, and size measured in physical units. The missing movers are disproportionately downsizers, so the homes withheld from the market exceed in value the homes those households would otherwise have demanded. A spatial housing-ladder model with six segments: rental, starter, and trade-up housing in low- and high-cost areas, translates this asymmetry into equilibrium prices that are roughly 2 percent higher for low-cost starter homes and 5 percent higher for high-cost trade-up homes under lock-in than in a counterfactual without it, with rents likewise rising in high-cost areas but not in low-cost ones.

The two accounts differ in terms of the data used to motivate them and have different implications for modeling. We hope that future research will unify these accounts.

c. Beyond Prices: Search, Matching, and Life-Cycle Migration

Most of the literature above treats prices as the central outcome and aggregate supply contraction as the central mechanism. Gerardi et al. (2024) push the analysis past supply-and-demand accounting into housing-market microstructure by combining micro-BLP demand estimation with a multi-market search-and-matching equilibrium across over 150 CBSAs. Their framework brings industrial-organization tools to bear on lock-in: rich seller heterogeneity, estimated from microdata, is embedded in a continuous-time matching model with endogenous listing, Nash bargaining, and re-buy decisions. Local markets, defined as Core-Based Statistical Areas crossed with census-tract income terciles, are then linked through endogenous seller-buyer flows. Counterfactually setting the lock-in coefficient to zero, the analogue of moving to a Danish-style system in which borrowers can repurchase their loan at market value, lowers mean prices by 9.7 percent and raises time-on-market by 8.7 percent. Lock-in raises prices while compressing time-on-market because it shrinks the seller pool more than the rebuyer-driven buyer pool, leaving a tighter market in which elevated reservation values translate into higher transacted prices.

Several findings from this framework are new to the literature. First, the average price impact is larger than other recent quantitative estimates, in part because the model fits a sizable own-to-rent channel out of sample using the matched CRISM–CCP data, which in turn is close to the Census Bureau's Survey of Income and Program Participation (SIPP) data. Second, the cross-market heterogeneity across over 150 CBSAs interacted with their income terciles is far richer than what aggregate or single-market models can deliver, which has important welfare implications: many markets experience price or time-on-market effects opposite in sign to the average, reflecting general-equilibrium reallocation of seller-rebuyers across destinations. Average price impacts rise monotonically across tract-income terciles (3.8, 7.8, and 12.3 percent), and the same monotonicity holds for time-on-market. In superstar cities (as definied in Gyourko et al. (2013)), the magnitudes are roughly twice as large, at 22.6

percent for prices and 24 percent for time-on-market, reflecting both inelastic supply and a buyer pool concentrated in exactly the demographic group whose moves lock-in most distorts. Borrowers under 35 are roughly 83 percent more likely to move within a CBSA and 23 percent more likely to move across CBSAs than borrowers aged 55-65, and conditional on moving are significantly more likely to upgrade tract income. The youngest borrowers therefore gain most from removing lock-in, both because they forgo more moves and because the forgone moves are higher-value upgrades to higher-opportunity neighborhoods. The model also matches untargeted moments on the own-to-own versus own-to-rent decomposition of moves with a sale, lending external validity to the counterfactual exercises.

d. Spillovers Beyond the Owner-Occupied Sales Market

De la Roca et al. (2024) use detailed rental-listings data for Los Angeles County, matched to nearby owner-occupied mortgage histories at high spatial resolution. For each rental listing, the paper constructs a local payment gap measuring the percentage increase in the monthly mortgage payment a buyer would face today relative to the typical legacy contract within a half-mile radius. A one-percent increase in this local payment gap raises asking rents by about 0.35 percent and reduces rental time-on-market—a combination that rules out a landlord-cost channel, which would raise asking rents and lengthen vacancies simultaneously. Using the payment gap as an instrument for nearby starter-home prices, the authors find that approximately one-third of the lock-in-induced increase in local sales prices spills over to local rents. The effect is largest in low-socioeconomic-status tracts, where rental demand is least elastic, and in larger multifamily buildings; it is muted in single-family rentals (with rent sensitivity to the payment gap roughly 60 percent smaller than in large multifamily buildings), in part because some locked-in owners may convert their unsold homes to rentals: a one-standard-deviation higher local payment gap raises the likelihood that a rental listing is single-family by roughly 8 percent.

Flynn et al. (2025) extend the analysis to commercial real estate using transaction-level

data on properties financed with fixed-rate CMBS loans. The setting is interesting because commercial-property sellers do not typically become buyers: the buyer-equals-seller offset that complicates identification of price effects in residential housing largely disappears. A 100-basis-point increase in the financing advantage embedded in the existing loan reduces the monthly probability of sale by 5.3 basis points (approximately 17 percent of the unconditional probability over a 10-year loan life) and raises transaction prices by roughly 4 percent. A market-wide spillover analysis indicates that unlocked properties in markets with high fixed-rate exposure trade at prices about 9 percent higher than otherwise similar properties in low-exposure markets.

IV. Challenges and Future Research Directions

a. Identification of origination interest rate effects

An important challenge that the post-pandemic literature confronted is the fact that interest rates are not randomly assigned — and thus, exogenous variation in the origination interest rate is needed to identify the effect of lock-in on various outcomes. To this end, Fonseca and Liu (2024) developed an instrument based on the gap between the market rate at origination and the current market rate, which eliminates borrower-specific variation in mortgage rates that might be correlated with variation in mobility. Liebersohn and Rothstein (2025) instead used the gap between the current market rate and the market rate at origination directly, rather than as an instrument for borrower-level variation in origination rates. Aladangady et al. (2025) further controlled for the refinancing tendency of the borrower by including the maximum refinancing rate incentive the borrower has experienced as a covariate. Generally speaking, the literature has used the *timing of origination* as a source of identification. Katz and Minton (2026) pointed out that using variation in the timing of mortgage originations might also introduce endogeneity bias and thus explored an entirely new approach, using unanticipated, rapid increases in U.S. Treasury yields in November 2016 and twin births as

sources of exogenous variation in origination rates. The Katz and Minton (2026) approach to identification yields significantly larger magnitudes of the lock-in effects, as shown in Table 1.[5]

Importantly, as Fonseca and Liu (2024) and Katz and Minton (2026) both emphasize, balancing in terms of observable covariates is important to validate these instrumental variables approaches. There is, of course, some tension on what observables to balance on — for example, given that the interest rate at origination directly affects the debt-to-income (DTI) ratios of borrowers holding FRMs, it is theoretically impossible to simultaneously hold fixed borrower income, borrower DTI, and borrower loan amount while allowing the interest rate at origination to vary. Therefore, an interesting direction for future research would be to develop more theoretically grounded choices of what set of covariates to balance on and what instruments to use to obtain plausibly exogenous variation in the origination interest rate, which, given its large apparent impact on the resulting estimates, will likely have implications for a significant number of lock-in related outcomes.

b. Separate identification of lock-in effects vs. interest rate effects

Assuming that the effect of the mortgage rate at origination can be identified, the separation of lock-in effects from the direct effects of origination rates remains a foundational challenge. Consider a model in which an outcome y (e.g., mobility, labor supply, consumption) is related to the origination rate r_0, the prevailing market rate r_t, and the rate gap:

$$y \;=\; a\, r_0 \;+\; b\, r_t \;+\; c\,(r_0 - r_t) \;+\; \varepsilon, \tag{1}$$

where c isolates the pure lock-in effect, over and above any level effect of either rate. The identification challenge is that even an ideal experiment that independently randomizes r_0 and r_t cannot, in general, recover c. Exogenous variation in r_0, for example, from origination-

[5]While Katz and Minton (2026) reports their estimate in 4-year sale terms, as this estimate result from the cumulative hazard of the hazard rate in each of the individual years, its relative magnitude should still be comparable to the annual level estimates.

timing instruments, identifies $a + c$, while exogenous variation in r_t identifies $b - c$.

This problem can be characterized as a "trilemma," since there are three unknown parameters but only two variables that can be potentially randomized. The problem is compounded by the fact that, based on the literature, the level of r_0 plausibly has a direct effect on the outcomes of interest: for example, Zator (2025) and Das et al. (2026) show that higher mortgage payments themselves induce households to supply more labor, so $a \neq 0$ may be a confound for any lock-in estimate built on origination-rate variation. Similarly, r_t is likely correlated with both aggregate and idiosyncratic factors that impact mobility and labor outcomes: for example, high interest rates are known to reduce home purchases (Bosshardt et al., 2024) and can slow firm hiring, which suggests that $b \neq 0$ is also likely.

To further illustrate this issue, Table 2 studies the effect of mortgage interest rate on divorce. We measure r_0 as the Primary Mortgage Market Survey (PPMS) rate at the month of origination. We measure $\Delta r = r_0 - r_t$, where r_t is the average interest rate in year t. We then regress divorce using a linear probability hazard model, which is equal to 1 in the year of an individual's marriage before their marriage status in the Rice University Experian credit panel switches from "Married" to "Single". Single households are excluded from this analysis. Butler et al. (2025) uses this data to study the effects of debt relief on divorce. We include origination year fixed effects, current year fixed effects, and county fixed effects, and so the variation in Δr is generated by within-year variation in the mortgage market interest rate at origination, r_0. In principle, the Δr captures the lock-in effect, whereas r_0 captures the direct effect of the origination rate, which may apply due to the effect of financial stress induced by a higher mortgage rate at origination.

Table 2 shows a large effect of interest rates on divorce, with a 1 percentage point higher market mortgage interest rate at the month of origination leading to a 0.36 percentage point higher probability of divorce, which is 10% of the sample mean of 3.61 percentage points. However, comparing Column (1) and Column (2) of Table 2 shows a difficulty with

interpretation: the estimated effect of the rate gap Δr on divorce in Column (1) is identical to the estimated effect of the origination interest rate r_0 on divorce in Column (2). This is because they use the same variation, which ultimately comes from the intra-year variation in the mortgage interest rate in the year of origination. Which specification should we prefer? Is the effect of mortgage interest rates on divorce a lock-in effect or an origination interest rate effect?

Table 2: Divorce: a lock-in effect or an origination interest rate effect?

	(1) divorced	(2) divorced
Δr	0.00365***	
	(0.000273)	
r_0		0.00365***
		(0.000273)
Mean of Dep. Var.	0.03614	0.03614
N	3978126	3978126
Origination Year FE	Yes	Yes
Current Year FE	Yes	Yes
County FE	Yes	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table's data came from the Rice 2010–2022 Experian credit panel. Only married individuals are kept in the sample. "divorced" is a dummy variable equal to 1 if the individual's marriage status switches from "Married" to "Single" in the following year. Standard errors are clustered by county.

One could certainly theorize one way or another. For example, lock-in could reduce divorce by reducing the incentive of couples to sell their homes. Alternatively, a lower interest rate could directly reduce divorce, for example, by alleviating financial pressures in a household.[6] Absent a way to tell the two mechanisms apart, it is difficult to separate lock-in effects from interest rate effects, but the two effects do have different predictions for when interest rates rise: the lock-in channel predicts there'll be fewer divorces, whereas the direct effect of interest rates on household finances channel does not. As the two effects are equivalent in

[6]As mentioned in Section 4.1, since it is mathematically impossible to simultaneously hold fixed borrower income, borrower DTI, and borrower loan amount while allowing the interest rate at origination to vary, a higher interest rate must be reflected in one of these variables.

Table 2, it might be reasonable to present the regression in Column (2) of Table 2 as the main specification, and speculate that lock-in is one of the channels through which the origination interest rate affects a given outcome. This example illustrates the importance of separating lock-in effects and interest rate effects as a direction for future research, which could be done across many outcomes.

We now discuss several possibilities for separating lock-in effects from the direct effects of the origination interest rate. In principle, one could lean on a nonlinearity specific to lock-in, such as a hypothesized kink at $r_0 - r_t = 0$, to separately identify c from a and b. This is, roughly, the logic that allows for the identification of refinancing responses to interest rate gaps, where a sharp kink at a gap of zero is visible in the data (Berger et al., 2021; Gerardi et al., 2023). In the lock-in setting, however, no comparable kink appears in mobility outcomes. Furthermore, the theoretical model of Fonseca and Liu (2024) does not predict such a kink at 0, which precludes this method of identification. The existence of a natural kink that is specific to refinancing implies that the effect of interest rate gaps on refinancing is more straightforward to identify than the effect of interest rate gaps on mobility and other outcomes.

In our view, there are two fruitful avenues for future research on this topic. One method of breaking the "lock-in trilemma" in Equation (1) is via functional form: to introduce more non-linearity in the definition of the lock-in effect, for example through a ΔNPV measure, defined based on the borrower's origination mortgage rate, r_i^o, the market rate r_i^t, and average prepayment hazards $\overline{\omega}$. Formally, it can be defined as follows:

$$\Delta NPV = (m_i^o - m_i^t) \times \frac{[1 - (1 - i - \overline{\omega})^{(TM_i - k_i)}]}{i + \overline{\omega}} \tag{2}$$

where

$$m_i^o = P_i^t \times \frac{r_i^o/12(1+r_i^o/12)^{(TM_i-k_i)}}{(1+r_i^o/12)^{(TM_i-k_i)}-1}$$

$$m_i^t = P_i^t \times \frac{r_i^t/12(1+r_i^t/12)^{(TM_i-k_i)}}{(1+r_i^t/12)^{(TM_i-k_i)}-1}$$

and P_i^t is the borrower's remaining principal balance at time t, TM_i is the mortgage term (360 months for all loans since we restrict our sample to 30-year FRMs), k_i is the age/seasoning of the mortgage in terms of months since origination as measured in the first month of the quarterly observation, i is the rate at which borrowers discount future cash flows (assumed to be 4% annualized), and $\overline{\omega}$ is the average monthly prepayment rate calculated across all borrowers in the sample. The coefficient in front of ΔNPV can be estimated controlling for loan amount and remaining term effects — and even loan amount by time, remaining term by time, and the origination rate by time effects if warranted, which then uses the exact functional form of Equation (2) for identification while controlling for the time-varying effects of each of its components. The combination of using a ΔNPV measure plus controls for time-varying effects of its components (including origination rate by time effects, which in effect controls for Δr) is therefore one possible method of separating lock-in effects from origination interest rate effects, and it would be interesting for future research to explore their implications for various outcomes.

A second potentially promising method of separately identifying lock-in effects from origination interest rate effects is to use a "triple difference" approach by comparing to a country where mortgage lock-in effects plausibly do not exist. That is, if we can assume $c = 0$ in another country, we can identify a in that country, and apply a to the U.S. setting. Berger et al. (2026) does this for mobility in Denmark. The paper finds that the coefficient in front of $a + c$ is close to zero in Denmark, where lock-in effects are close to zero by construction due to the ability of households to repurchase their mortgages at their market values when

interest rates rise, which implies that a is close to zero in Denmark. In turn, this suggests that the existing literature attributing the effect of the origination interest rate on mobility as the effect of lock-in, by implicitly assuming $a = 0$, may be plausible. That said, given the significant differences in mobility patterns, labor markets, and geographic topology between the United States and Denmark, we view alternative identification strategies that resolve the lock-in trilemma as an important direction for future research.[7] Liebersohn and Rothstein (2025)'s comparison with cash owners over the 2022-2023 interest rise period, which identifies the effect of mortgage lock-in by relying on a sample of households for whom lock-in effects are theoretically equal to zero, is another example of the "triple difference" approach. Other approaches of this style might be plausibly produced for each particular outcome under study.

c. General Equilibrium

A second, equally unresolved set of questions concerns the general equilibrium implications of mortgage lock-in on labor markets and firm productivity. Most of the empirical literature estimates partial-equilibrium responses at the household level, while the aggregate welfare and policy consequences depend on how these household-level responses interact through labor and housing markets. A first-order issue is the specificity of locked-in workers for firms. If locked-in workers are imperfect substitutes for non-locked-in workers—because of occupation, skill, or location-specific human capital differences—then reduced mobility can translate into meaningful labor misallocation and wage dispersion effects.[8] If, by contrast, workers are close substitutes at the relevant margins, equilibrium adjustment through within-market reallocation or more mobile workers may absorb most of the lock-in shock, resulting in only

[7] For example, one potential way forward is to compare areas of Denmark and the U.S. that are more similar to one another, instead of relying only on aggregate comparisons.

[8] One class of workers for whom this is plausible is inventors, for whom Feng et al. (2026) finds significant effects of lock-in on patenting.

relatively minor misallocation costs. Letting firm j's effective labor input be CES:

$$L_j \;=\; \left(\sum_i \theta_{ij}\, \ell_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

the welfare cost of lock-in is likely increasing in the dispersion of match-specific productivities θ_{ij} and decreasing in the elasticity of substitution σ. Disciplining both objects in the presence of lock-in would be an interesting empirical exercise. On a related note, the rise of remote work potentially attenuates lock-in's real consequences by decoupling job location from residence, which may be interesting to incorporate in future models. A quantitative general-equilibrium framework that jointly incorporates worker substitutability and remote work with mortgage lock-in strikes us as a natural next step for this literature.

V. Policy Solutions

Mortgage lock-in occurs due to the combination of a large stock of low-rate, long-term FRMs and relatively high prevailing market rates. Thus, the lock-in problem will ameliorate when interest rates fall back to (or near) their previously low levels. That may take a long time, however, as persistently high inflation continues to prevent the Federal Reserve from returning short-term interest rates to the low levels that characterized the pandemic period, when many of the outstanding low-rate FRMs were originated.

More permanent solutions to lock-in will likely need to involve contract design changes to the long-term, fixed-rate mortgage that has dominated the product space in the U.S. for the past 30 years. There are at least three types of changes that would be effective in eliminating lock-in. First, allowing for portability in the U.S. mortgage market would break the link between a loan and a specific property. Portable mortgages, which are common in Canada and the United Kingdom, allow a borrower to transfer their existing loan to their next property when they move. Currently, nearly all mortgages in the U.S. have a "due-on-sale"

clause, which forces homeowners to pay off their loans when they sell and effectively prohibits portability. The Trump administration has considered proposals to make loans portable,[9] but such a reform would be complicated because it would significantly change prepayment behavior by extending mortgage durations, and thus lead to a transfer from investor to borrowers, while further increasing the convexity of mortgage prepayment behavior which likely lead to higher interest rates as investors would demand higher compensation for the increased extension risk.

Similarly, increasing the share of assumable mortgages, which allow a homebuyer to take over the seller's existing loan, could lessen the effects of lock-in. Currently, only government-insured loans (FHA, VA, and USDA) are assumable. However, even within this segment of the market, the option to assume is rarely exercised for a couple of reasons that would be difficult to overcome. First, lenders must provide their approval, and are often reluctant to do so because they make higher profits originating new, higher-rate loans. Second, assuming a loan typically requires the buyer to make a significantly higher down payment than they would if they simply obtained a new mortgage. This is due to the realized amortization of the assumed loan and the fact that the buyer is typically paying a higher price than what the seller originally paid for the property.[10]

A potentially promising solution would be to adopt the "delivery" option that is embedded in Danish fixed-rate loans. The option gives mortgage borrowers the ability to prepay their loans at their current market value rather than at par. When interest rates rise, FRMs that were originated at lower rates decrease in value. Thus, borrowers who would like to sell their properties in high-rate environments are able to settle their mortgage debt by paying an amount that is less than the outstanding principal balance and pocketing the difference between the par and market values of the loan. This option theoretically prevents mortgage lock-in because it eliminates the financial incentive to hold onto a below-par mortgage rate.

[9] See an article describing the proposal HERE.

[10] See Berndt et al. (2024) for a detailed analysis of the impact of assumable financing on the U.S. housing market.

Indeed, Berger et al. (2026) document that the delivery option is frequently exercised in Denmark and show that, as a result, household mobility in Denmark is largely insensitive to market rates moving above rates on outstanding Danish loans. However, this option is not without its costs, and Berger et al. (2026) estimates a roughly 20 basis points increase in interest rates for this option.

References

Aladangady, A., J. Krimmel, and T. Scharlemann (2025). Locked in: Rate hikes, housing markets, and mobility.

Amromin, G. and J. Eberly (2024). Macro shocks and housing markets. Technical report, Working Paper.

Berger, D., J. Jeong, J. Marx, M. V. Olesen, and F. Tourre (2026). A danish fix for us mortgage lock-in?

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

Berndt, A., J. Zheng, and M. Zhu (2024). Breaking the lock-in effect: How assumable mortgages shape housing prices and supply dynamics. *Available at SSRN 5064022*.

Bosshardt, J., M. Di Maggio, A. Kakhbod, and A. Kermani (2024). The credit supply channel of monetary policy tightening and its distributional impacts. *Journal of Financial Economics 160*, 103914.

Butler, A. W., I. Spyridopoulos, Y. Tellez, and B. Xu (2025). Till debt do us part? the effects of debt relief on household stability. *The Effects of Debt Relief on Household Stability (April 17, 2025)*.

Das, M. M., J. Hambur, K.-P. Hellwig, and J. A. Spray (2026). Labor supply effects of monetary policy: Evidence from australian mortgage holders. Technical report, International Monetary Fund.

De la Roca, J., M. Giacoletti, and L. Liu (2024). Mortgage rates and rents: Evidence from local mortgage lock-in effects. Technical report, Working Paper.

Feng, I., J. S. Xue, and D. Zhang (2026). Unlocking innovation: Household mobility, inventor careers, and productivity. Technical report, Working Paper.

Flynn, S., R. J. Kurtzman, and A. Tchistyi (2025). Does mortgage rate lock-in dampen commercial real estate busts?

Fonseca, J. and L. Liu (2024). Mortgage lock-in, mobility, and labor reallocation. *Journal of Finance*.

Gerardi, K., F. Qian, and D. Zhang (2024). Mortgage lock-in, lifecycle migration, and the welfare effects of housing market liquidity.

Gerardi, K., P. S. Willen, and D. H. Zhang (2023). Mortgage prepayment, race, and monetary policy. *Journal of Financial Economics 147*(3), 498–524.

Gyourko, J., C. Mayer, and T. Sinai (2013). Superstar cities. *American Economic Journal: Economic Policy 5*(4), 167–199.

Katz, J. and R. Minton (2026). Rising interest rates, mortgage rate lock, and house price fluctuations. *Working paper*.

Liebersohn, J. and J. Rothstein (2025). Household mobility and mortgage rate lock. *Journal of Financial Economics 164*, 103973.

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

Potepan, M. J. (1989). Interest rates, income, and home improvement decisions. *Journal of Urban Economics 25*(3), 282–294.

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

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

Schwab, R. M. (1985). Renovation and mobility: An application of the theory of rationing. *Southern Economic Journal*, 203–215.

Zator, M. (2025). Working more to pay the mortgage: Household debt, interest rates, and family labor supply. *The Journal of Finance 80*(2), 1171–1207.