The Long Shadows of the Great Inflation: Evidence from Residential Mortgages[*]

Matthew J. Botsch and Ulrike Malmendier[†]

June 26, 2021

Abstract

In most countries, the prevalent long-term mortgage is variable-rate. The US is an outlier, with 80% fixed-rate mortgages. We link the puzzling US market structure to long-lasting effects of the Great Inflation and structurally estimate the welfare implications. First, sentiment towards variable-rate mortgages negatively correlates with past nominal rates. Second, inflation exposure directly affects interest-rate expectations and mortgage choice. Third, we use SCF and RFS data, in combination with interest-rate surveys (PMMS and MIRS), to estimate a structural discrete-choice model and quantify payoff consequences. Our simulations imply that Baby Boomers overpaid by $23bn for fixed-rate mortgages in the late 1980s–1990s.

Keywords: Financial contracts, household finance, experience effects, behavioral finance, inflation expectations, mortgage choice.

JEL Classifications: D14, D83, D84, D91, E31, G41, G51.

[*]We thank workshop participants and discussants at Amherst, Babson, Barcelona, Berkeley, Bowdoin, Cornell, Duke, Haverford, and the New York Fed, as well as the ECB Household Finance conference, the NBER Household Finance Summer Institute, the 2015 World Congress of the Econometric Society, and the ASSA 2016 and 2018 Annual Meetings for helpful comments; and Clint Hamilton, Isaak Heller, Karin Li, Canyao Liu, Junjun Quan, and Jeffrey Zeidel for excellent research assistance.

[†]Botsch: Bowdoin College, mbotsch@bowdoin.edu. Malmendier: UC Berkeley, and NBER, ulrike@berkeley.edu.

1 Introduction

Buying a home is the biggest financial decision in many households, with important consequences for lifetime saving and consumption. Most home buyers take on significant leverage and a financial commitment that commonly stretches over 30 years.

Given the long-term nature of mortgage contracts, it is not surprising that the majority of mortgage markets around the world feature primarily adjustable-rate mortgages (ARMs). Under an ARM contract, the mortgage rate rises and falls with market interest rates, considerably lowering the risk and the cost of financing for the bank, which in turn allows the bank to offer cheaper mortgage products. As the left map in Figure 1 illustrates, ARMs are the standard product across the globe (cf. Lea 2010): in 45 out the 55 countries with available mortgage data, ARMs are a prevalent or even the primary form of mortgage financing.

The map also reveals that the U.S. is an exception. In the U.S., fixed-rate mortgages (FRMs) command a market share of 80%, and households pay a premium of, on average, 170 basis points over equivalent-risk and -term ARMs. The dominance of fixed-rate financing in the U.S. is puzzling since it is costly for consumers and hard to reconcile with standard consumption models. ARMs were introduced nationwide between 1979 and 1982 to provide borrowers with a product that is cheaper, as it entails less interest-risk for the lender, and accessible to households that do not qualify for an FRM. In their seminal work on mortgage choice, Campbell and Cocco (2003, 2015) show that most households are predicted to choose an ARM, particularly if they are younger and more mobile. These predictions are in conflict with the observed dominance of FRMs. Our own calculations below confirm that far more households choose FRMs than the standard economic model predicts, especially Baby Boomers in the wake of the Great Inflation. These generations should have taken out 1 million fewer FRMs in the late 1980s, and half a million fewer in the late 1990s. The costs of their deviations are large. Given expected refinancing behavior and mobility, Baby Boomers overpaid more than \$14 billion on their fixed-rate mortgages in the late 1980s, and almost \$9 billion in the late 1990s.

Why are U.S. consumers so averse to variable-rate mortgages? Market characteristics (such as payment structure, interest deductability, and rental-market regulation) and demographic determinants (such as life-cycle stages, age, fertility, household size, and mobility) play a role but do not suffice to explain the puzzle (cf. Campbell 2013, Guiso and Sodini 2013).

Figure 1. ARM Prevalence and Historical Inflation Around the World

Panel A. Mortgage Markets Panel B. Historical Median Inflation

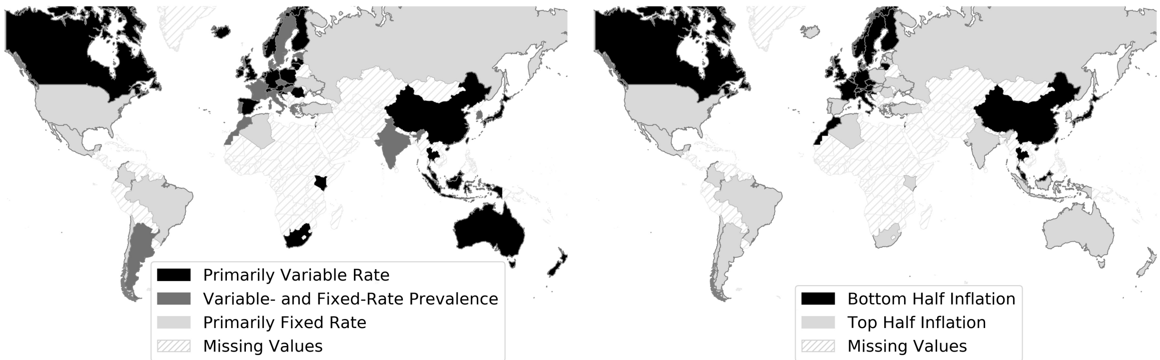

Notes. In Panel A, *Primarily Variable Rate* indicates that at least 75% of all mortgages have variable interest rates for the entire duration of the mortgage or after at most five years; *Variable- and Fixed-Rate Prevalence* indicates that at least 25% and less than 75% of all mortgages have variable interest rates; and *Primarily Fixed Rate* indicates that less than 25% of all mortgages have variable interest. In Panel B, *Bottom Half Inflation* comprises all countries with a median inflation (since 2000) below 2.2%; *Top Half Inflation* includes all countries with a median inflation of at least 2.2%. All data sources for both panels are listed in Appendix A.

In this paper, we build on a notion frequently discussed among practitioners: the idea that the Great Inflation cast "long shadows," which continued to affect the beliefs and fears of householdes decades later. That is, similar to the 2021 discussion of how investors who came of age during a low-inflation period would not be prepared for a high-inflation economy post-pandemic, we hypothesize that mortgage borrowers in the 1980s and 1990s who came of age during a high-inflation period were not prepared for a low-inflation economy.

Empirically, we trace the structure and composition of the U.S. mortgage market back to consumers experiencing dramatic increases in inflation and interest rates in the 1970s and 1980s. These experiences generated a long-lasting aversion to variable-rate borrowing, but also significant cross-sectional differences between cohorts, depending on whether they were more or less exposed to the experience of the Great Inflation. To the best of our knowledge, we are the first to establish a direct link between these experiences and interest-rate expectations and to estimate structurally the payoff consequences of the affected mortgage choices.

While our analysis focuses on the U.S. mortgage market, the international perspective helps to motivate the role of historical inflation rates. In the right map of Figure 1, we graph median inflation rates over the past 20 years for the 55 countries from the mortgage map on the left. The juxtaposition of the two maps reveals the positive correlation between fixed-rate prevalence and high historical inflation rates. This correlation is consistent with a mechanism under which higher historical inflation rates tend to generate a longer-term aversion to ARM usage. [1]

We operationalize the notion of "long shadows of past inflation" for the analysis of the contractual mix and composition of mortgage markets, and provide structural estimates of their magnitude and welfare implications. We show that, after accounting for other determinants of mortgage choice, personal exposure to inflation directly predicts expectations of future interest rates and mortgage choices. Our structural estimates imply that one additional percentage point of experienced inflation increases a borrower's willingness to pay for an FRM by 6–14 basis points. Households who would otherwise have chosen an ARM pay $8,000-$16,000 in year-2000, after-tax dollars for their experience-driven choice of an FRM over their expected tenure in the house.

We start from the historical background of the introduction of ARMs in the U.S. since the late 1970s. We document a significant reversal in the congressional discussion, from outright rejections of all ARM proposals in the 1970s to strong support in 1982. Based on a systematic analysis of all articles published in *The Wall Street Journal* and *The Washington Post* around the time of the nationwide ARM introduction (1971–1984), we construct an index of "ARM sentiment" that measures the direction and prevalence of ARM arguments. We show that ARM sentiment strongly correlates with interest and inflation rates: a rise in rates predicts a decline in public opinion towards ARMs. This negative relationship becomes even more pronounced when accounting for past (lagged) inflation rates.

Motivated by these stylized facts, the core of our analysis focuses on measuring the predictive power of personal exposure to the Great Inflation and other past inflation for interest-rate expectations and mortgage preferences. Our analysis is the first to tease out the role of *interest-rate* beliefs, to construct plausible measures of alternative mortgage products and rates available to the consumer, and to assess the large welfare implication of exposure-induced choices between fixed- and variable-rate contracts.

Our measure of exposure to historical inflation closely resembles the experience-effect measure in Malmendier and Nagel (2016). It differs from prior work on adaptive expectations and extrapolative beliefs (e.g., Koijen, Hemert, and Nieuwerburgh 2009) in accounting for the heterogeneous histories of different individuals' lifetimes. As a result, experience-based beliefs predict cross-sectional differences as well as changes in

[1] See, e.g., *The New York Times* 6/4/2021 ("Inflation is Real Enough to Take Seriously") and *The Wall Street Journal* 6/21/2021 ("What Investors Can Learn From the History of Inflation").

these cross-sectional differences over time. These will be the sources of identification.

We first show that past inflation directly affects interest-rate expectations and mortgage choice using the Survey of Consumer Finances (SCF). Previous work (cf. Malmendier and Nagel 2016) has focused on inflation expectations, but a common criticism of research on consumers' inflation expectations is its supposed lack of influence on actual economic decision-making. We help address these concerns by showing directly that past exposure to inflation affects interest-rate beliefs, which are the intermediating mechanism for the influence of inflation experiences on mortgage choice. Given the tight theoretical and empirical link between inflation and nominal interest rates, the differences between the two rates are small, and we show that consumer expectations move correspondingly. Indeed, we also show that our main results are very similar if we use individuals' personal exposure to past nominal interest rates rather than exposure to past inflation. The correlation between the two corresponding experience measures lies between 70 and 80% (depending on the interest rate used), indicating a common source of variation, the Great Inflation of the 1970s. Our study thus fills another gap in the literature by showing how inflation shocks relate to personal exposure to high interest rates and interest rate experiences, in turn, affect mortgage choice.

The SCF survey waves between 1989 and 2013 elicit respondents' expectations of whether interest rates will rise or fall over the next five years. We construct a measure of personal exposure to past inflation that allows households to overweight their individual lifetime experiences. In the early SCF waves (1989, 1992, etc.), members of the younger cohorts were more likely to expect interest rates to rise on net than members of the older cohorts. These same individuals also have the highest personal exposure to past inflation by our measure. The relative positions of older and younger cohorts reverse in the mid-2000s, as the memory of the Great Inflation is fading and households who have experienced the Great Inflation become older households. At that time new, younger households who put relatively less weight on the Great Inflation enter the sample and have lower expectations.

We then relate the decision to choose an FRM over an ARM to individuals' heterogeneous exposure to historical inflation rates, with beliefs about future interest rates as the intermediating variable. A key challenge is that the SCF elicits interest-rate beliefs as of the time of the survey, not at the time of the (past) mortgage choice. Under the plausible assumption that individuals' interest-rate beliefs are serially correlated, this timing discrepancy introduces non-classical measurement error into the interest-rate belief variable.[2] Instrumenting does not remedy the bias since any instrument that is positively correlated with the time-t interest-rate forecast, such as time-t inflation experiences, has to be negatively correlated with the subsequent change in forecast, so cannot be exogenous. However, we are able to show that, while the univariate probit estimate of mortgage choice on ex-post interest-rate beliefs is attentuated toward zero, the estimate from a bivariate probit model that uses time-t inflation experiences as an instrument for time-t beliefs is inconsistent in the opposite direction: it is amplified rather than attenuated. The univariate and bivariate probit estimates together give us lower and upper bounds on the effect of interest rate beliefs on mortgage choice. With these bounds in place, the SCF analysis confirms that personal exposure to high inflation strongly raises interest-rate expectations, which, in turn, strongly increases the probability of choosing an FRM instead of an ARM.

Our goal is to assess the magnitude of these mortgage-choice effects and move towards an estimation of cohort-specific payoff- and welfare-implications. The SCF does not allow for such an estimation because it lacks contract details and does not provide geographic location, which is crucial for estimating price coefficients in the presence of year fixed effects. To remedy this shortcoming, we turn to a data set that has not been explored in this context, the Census Bureau's Residential Finance Survey (RFS) from 1991 and 2001. Differently from the public-use SCF, the RFS provides geographic location, allowing us to include the menu of fixed and variable rates available to a typical borrower in a give geographic area at a certain point in time while also accounting for origination-year fixed effects.

We estimate the structural parameters of a discrete-choice model over mortgage financing alternatives, and we use these estimates to quantify the effect of past inflation on mortgage choice at the household level, given the ARM and FRM interest rates that a specific household could have qualified for. The estimation faces two challenges: first, we do not observe the contract terms of the alternative that households did not choose. Second, the sample of households that choose a given product is self-selected.

We use a three-step procedure following Lee (1978) and Brueckner and Follain (1988) to overcome these challenges. In Step 1, we estimate a reduced-form model of mortgage choice that only uses *exogenous* explanatory variables. The key explanatory variables in this step are Freddie Mac's Primary Mortgage Market Survey (PMMS) interest rates for standardized FRM and ARM products to a representative, prime borrower in a Census

[2]This applies to any non-longitudinal survey of household expectations, not only the SCF.

region-year. (We also consider the FHFA's Monthly Interest Rate Survey (MIRS) as an alternative explanatory variables, as discussed below.) The resulting estimates are very similar to the SCF results. The replication using such different data, with differences in controls and sample size, provides strong supporting evidence for our hypothesis.

In Step 2, we estimate the fixed- or variable-rate mortgage terms for each household that chose a given product, and use these estimates to predict the out-of-sample interest rate offered to households that did not choose that product, correcting for selection bias. We model in-sample mortgage rates as a function of the respective (FRM or ARM) survey interest rate and household-level attributes associated with risk characteristics and preferences, including marital status, income, urban versus rural location, and mortgage seniority. We implement the semiparametric Newey (2009) series estimator to correct for the selection bias that arises from estimating over the non-random subsample of households that chose a given alternative. The estimator generalizes Heckman (1979) in that it includes polynomial terms of the predicted choice probabilities from the first step, but does not require normally-distributed errors. Identification relies on a pair of cross-equation exclusion restrictions: conditional on the FRM survey rate, the ARM survey rate does not directly influence the FRM rate that an individual household is offered, and vice versa.

In Step 3, we use the predicted, household-characteristic adjusted pairs of mortgage rates for each household to estimate the coefficients of a structural choice model. This model is structural in the sense that the key explanatory variables are pairs of household-varying interest rates, between which each household would choose.

The estimates both of the reduced-form model (Step 1) and the structural mortgage-choice model (Step 3) attest to the lasting legacy of the Great Inflation. Our most conservative estimate is that one in seven households (10–15% of the population) were close enough to indifference between the two alternatives that we can attribute their FRM choice to long-lasting effects of their past exposure to high inflation. This calculation controls for the full information set available to all mortgagors in the origination year via origination-year fixed effects. The fixed effects capture, for example, current inflation as well as the entire history of all past inflation realizations. The choice-model estimates indicate that consumers are willing to pay between 6 and 14 basis points of interest for every additional percentage point of personally experienced inflation.

The identification of household-specific pairs of mortgage rates is a key contribution and ingredient of this analysis. Empirically, borrowers choosing an FRM are likely to differ from those choosing an ARM along both observable and unobservable dimensions and Step 2 provides a borrower-specific markup after controlling for borrower selection. To that end, it is critical, in Steps 1 and 2, to rely on baseline mortgage rates that are not afflicted by borrower heterogeneity. This is the case for the PMMS, which quotes interest rates that would be offered to the same hypothetical borrower, across mortgage products and over time. We also consider another commonly used interest-rate series, the FHFA's Monthly Interest Rate Survey (MIRS). This time series is drawn from *actual* mortgage originations, and so reflects changes in the pool of borrowers across products and over time. We show that, between January 1986 and October 2008, when both series are available, the PMMS tracks the slope of the nominal Treasury yield curve much more closely, revealing the presence of borrower selection in MIRS.

In the last part of the analysis, we assess the dollar cost associated with past inflation experiences and the resulting higher willingness to pay for FRMs. We simulate how much interest an individual would have paid under two standard contracts: a 30-year fully amortizing FRM, and a 30-year 1/1 ARM without caps, i. e., an ARM where the initial rate holds for one year, after which the rate adjusts annually, indexed to the one-year Treasury. We calculate the present value of excess interest paid that is attributable to the individual's inflation-experience coefficient in the structural choice equation. In a typical household, these costs amount to $8,000 (without interest-rate adjustments for household risk characteristics) to $16,000 (with adjustments) in constant year-2000 dollars, accounting for taxes, typical refinancing behavior, and expected tenure given the borrower's age. The estimates imply the potential of significant welfare loss due to the influence of past inflation experiences. The long shadows of the Great Inflation appear to strongly influence mortgage financing choices, and the resulting financial costs to the household are large.

These cost estimates are *ex post* and reflect the actual realization of inflation and interest rates during the "Great Moderation" of the 1980s, 1990s, and early 2000s. We also calculate the *ex ante* cost of exposure-induced choices by simulating different hypothetical inflation environments. In particular, we re-calculate the present value of excess interest that would have been paid under the rising inflation environment beginning in 1971, the falling inflation environment beginning in 1981, and by running a Monte Carlo simulation of different inflation and national average mortgage rate paths over the 30-year lifetime of the mortgage. The Monte Carlo simulations indicate that the ex-ante expected cost of choosing an FRM that is attributable to the individual's inflation-experience coefficient is $6,500 under expected refinancing behavior and tenure. However, the FRM is cheaper for these households in fewer than 25% of replications, and the expected gains are small compared to the expected losses in replications where average inflation is low. Hence, for households choosing an FRM due to their past personal exposure to high inflation, this choice is expensive in expectation.

Our paper contributes to extensive research on residential mortgage choice and consumer welfare. The empirical literature expanded significantly when regulators permitted ARMs in the early 1980s, as indicated by the theoretical and empirical papers cited above. Some of the earlier literature using microdata found that, consistent with the later-formed theoretical predictions, younger households with higher probability of moving and more stable income seem more likely to choose an ARM and emphasized the explanatory power of price variables (see, for example Dhillon, Shilling, and Sirmans (1987), Sa-Aadu and Sirmans (1995), Brueckner and Follain (1988)). Follain (1990) provides an overview of this earlier literature. Among the more recent literature, Paiella and Pozzolo (2007) find that, contrary to theory, most household characteristics cannot explain mortgage choices in Italian microdata, though liquidity constraints and relative prices are driving factors. Bergstresser and Beshears (2010) find a correlation between households' financial literacy and their observed ARM choice. Coulibaly and Li (2009) use the Survey of Consumer Finances (SCF) to show that, besides pricing variables and affordability, mobility expectations, income volatility, and financial risk attitudes influence mortgage choices. Mortgage choice is also a core element of the growing field of household finance, especially since the 2008 financial crisis (Guiso and Sodini 2013; Green and Wachter 2005; Mayer et al. 2009).

Conceptually our analysis builds on the work of Case and Shiller (1988) and Shiller (1999, 2005) as well as an early literature on mortgage financing from the time of the Great Inflation. At that time, researchers first proposed that the resulting change in inflation expectations might distort housing decisions (see, e. g., Kearl (1979), Baesel and Biger (1980), and Alm and Follain (1982)).

A more recent literature on non-standard belief formation has formalized this notion of past realizations affecting beliefs and mortgage borrowing. Koijen, Hemert, and Nieuwerburgh (2009) explain U.S. mortgage choice with an adaptive-expectations "rule of thumb" under which households use only the most recent three years of yield curve data. Extrapolative expectations are also a candidate to explain the house price boom and bust of the mid-2000s (Glaeser and Nathanson 2015). Bailey et al. (2019) and Bailey et al. (2018) consider the role of house-price expectations and its non-standard determinants on mortgage and tenure choice. Armona et al. (2018) show a causal effect of house price beliefs on housing and portfolio choice via a randomized experiment.[3]

One issue with these extrapolative approaches is, as Badarinza, Campbell, and Ramadorai (2018) discuss, that they fail in an international context and are weaker for different time periods. The long shadows cast by past experiences, as we propose in this paper, helps resolve these discrepancies. While consumers do overly rely on the interest-rate realizations of recent years, they also overweight earlier realizations if they have personally experienced them. These early experiences exert a long-lasting, but not permanent, influence, so beliefs are different across different generations and converge slowly over time. In two countries with identical inflation histories but different population age profiles, adaptive-expectation models would predict the same FRM share, whereas our theory predicts different FRM shares.

More broadly, prior research has provided insights into the implications of behavioral factors for mortgage contract design and regulation. For example, Gottlieb and Zhang (2018) study the welfare impact of the option to terminate long-term debt contracts when consumers are present-biased, and other work includes Schlafmann (2016), Ghent (2015), Gathergood and Weber (2017), Atlas et al. (2017), and Bar-Gill (2008).

Beyond the mortgage context, our paper contributes to the broader literature on experience effects. Alesina and Fuchs-Schundeln (2007) relate the personal experience of living in (communist) Eastern Germany to political attitudes post-reunification, and Laudenbach et al. (2018) relate it to households' choice of financial investments, including their persistent aversion to stock-market investment. Oreopoulos et al. (2012) show that the experience of graduating in a recession predicts long-term wage paths. Relatedly, Malmendier and Shen (2015) show that experiences of macroeconomic unemployment conditions predicts lower consumption expenditures, and a higher use of coupons and allocation of expenditures toward lower-end products, for decades afterward.

Much of the recent literature on experience effects has shown that personal experiences of macro-finance outcomes, such as the high inflation of the 1970s, have a lasting impact on individual beliefs and attitudes, often in the context of stock-market participation (cf. Kaustia and Knüpfer (2008), Malmendier and Nagel (2011), Strahilevitz

[3]Other research on extrapolative expectations and house price dynamics includes Glaeser et al. (2008), Mayer and Sinai (2009), Gelain and Lansing (2014), Granziera and Kozicki (2015), Gao et al. (2017), Glaeser and Nathanson (2017), and Guren (2018). On non-standard expectations and house prices more generally, see Piazzesi and Schneider (2009), Case et al. (2012), Favara and Song (2014), Burnside et al. (2016), Suher (2016), Landier et al. (2017), Gao et al. (2018), Kuchler and Zafar (2018), and Nathanson and Zwick (2018).

et al. (2011), Kaustia and Knüpfer (2012), and Knüpfer et al. (2017)). Theoretical treatments include Collin-Dufresne, Johannes, and Lochstoer (2016), Malmendier, Pouzo, and Vanasco (2020), and for the long-lasting effects Schraeder (2015). Most relevant to the analysis of mortgage contracts is the work on inflation experiences. Malmendier and Nagel (2016) first showed that personal inflation experiences predict subjective beliefs about future inflation and, as a result, investment in real estate and mortgage borrowing.[4] Relative to their evidence, our paper provides the first direct evidence relating prior inflation experiences to subjective *interest-rate* beliefs, which is the missing link between the choice of fixed- versus variable-rate instruments and prior exposure to high inflation. Malmendier and Nagel (2016) also relate outstanding mortgage balances in the SCF to lifetime experiences of inflation, though the results on the type of mortgage are weak or insignificant, likely due to data limitations.[5] We overcome these difficulties using the RFS. To the best of our knowledge, we are the first to present quantitative estimates of the direct impact of prior experiences on the choice between FRMs and ARMs and their payoff consequences. Our cost estimates suggest potentially significant welfare consequences. The results aim to be a first stepping stone toward more complete welfare estimations.

2 ARMs: Historical Background and Sentiment Index

To motivate our quantitative choice analysis, we provide more details on the debate surrounding the introduction of ARMs in the U.S. and construct a new "public-sentiment index," which captures attitudes towards adjustable-rate products in the public debate.

The U.S. Mortgage Market. The dominant mortgage in the U.S. is a 30-year, level-payment, self-amortizing, fixed-rate contract with the option to prepay. To foster its popularity, Congress established Fannie Mae (1938) and Freddie Mac (1970) with the mission to purchase long-term FRMs from banks, which otherwise face duration risk from holding these assets.

In the late 1970s and early 1980s, ARMs emerged as an alternative mortgage product. A typical ARM contract also self-amortizes over a long period such as 30 years, but the interest rate resets periodically according to a prespecified margin over an index, typically the one-year Treasury bill or a district cost-of-funds index. As a result,

[4] Past inflation also correlates with homeownership across European countries (Malmendier and Steiny 2016), and influences FOMC members' inflation forecasts and votes (Malmendier et al. 2018).

[5] In unreported results, we replicate the analysis of Malmendier and Nagel (2016) and show that individuals in the RFS with higher lifetime inflation experiences also originate and hold larger balances of fixed-rate liabilities.

monthly payments vary from year to year. More exotic mortgages types became popular during the housing boom of the 2000s, including "hybrid ARMs" whose rates are initially fixed but later may change, and "interest-only" mortgages, under which no principal is paid in early periods to keep initial payments low.

The Census Bureau's RFS data on outstanding residential mortgages reveals the persistent dominance of FRMs, at around 80% market share. Despite their greater liquidity on secondary markets, FRMs are more expensive. According to Freddie Mac's Primary Mortgage Market Survey (PMMS), which controls for risk factors and term length, FRMs are priced at a significant premium of 170 basis points on average between 1984 and 2013 (S.D. = 67 bp).

Historical Debate about Variable-Rate Products. Before ARMs were introduced nationwide in the U.S., the idea of variable-rate mortgages had been contentiously debated for years.[6] The product was first introduced by smaller, state-chartered S&Ls in California in the early 1970s, but larger lenders did not take an interest in them until after interest rates spiked in 1973–4 (Cooper 1975). Regulators and policy-makers including Federal Reserve Chair Arthur Burns supported allowing ARMs nationwide, but Congress rejected all such proposals, mirroring consumers' resistance to the idea (Wall Street Journal 1973; Peek et al. 1990). One prominent senator declared that "Requiring homebuyers to negotiate with sophisticated lenders over the merits of variable-rate mortgages [...] would be 'like putting Baby Snooks in the ring against Muhammad Ali'."[7] In 1979 the tide began to turn. The FHLBB[8] gave federally-chartered S&Ls the authority to originate variable-rate loans, first in California in December 1978, then nationwide in 1979. In 1982 Congress extended permission to all housing lenders (Title VIII of the Garn-St. Germain Depository Institutions Act, "Alternative Mortgage Transactions").

Why did the public and legislators' opinions and attitudes towards ARMs fluctuate so widely? Congress was balancing the interests of two competing groups: lenders whose balance sheets were damaged every time interest rates rose, and borrowers who wanted stable access to housing credit but opposed a transfer of interest-rate risk from

[6]Miller (1986) provides a detailed history, which we draw upon.

[7]Sen. William Proxmire, chairman of the Senate Banking Committee until 1980. Qtd. in *The Wall Street Journal* 8/23/1978 ("Bank Board Mulls Lower Denominations For Variable-Rate Accounts S&Ls Offer'.")

[8]The Federal Home Loan Bank Board (FHLBB) governed the Federal Home Loan Banks (FHLB), the Federal Savings and Loan Insurance Corporation (FSLIC) and nationally-chartered thrifts from 1955 to 1989.

lenders. To investigate why the tide of public sentiment eventually turned and what role interest-rate exposure might have played, we construct a new data set on the public debate about ARMs, construct an index of public sentiment about variable-rate products (*ARM Sentiment Index*), and relate it to inflation and interest rates in the recent past.

First, we identify all articles in *The Wall Street Journal* and *The Washington Post* between 1971 and 1984 that mention "adjustable-rate mortgage" or "variable-rate mortgage" (with and without hyphen). The resulting set of 274 articles captures the press coverage of the political debate from the decade prior to the ARM's nationwide introduction in 1982 and continuing through the early years of its marketing to U.S. consumers.

We extract from these articles all paragraphs that address the concept of an adjustable-rate mortgage, as identified by the phrases "adjustable-rate mortgage" or "variable-rate mortgage" (with and without hyphens), "new mortgage," or "variable interest rate." These paragraphs typically report quotes from political, policy, or business leaders, or discuss incidents related to the ARM concept. We extract 404 paragraphs, 1.47 per article and 28.86 per year.

In the next step, we classify each paragraph's stance towards ARMs, in terms of (1) reported opinions or arguments and (2) the author's own opinion or arguments. For both dimensions we use a scale from -1 to 1, in steps of 0.5. For (1), we categorize a very negative attitude towards ARMs as -1, a neutral report as 0, and strong support for ARMs as 1. A negative value results from the key phrases "fear," "oppose(d)," "boycott," "uproar," "resistance," "difficulty," "problem," "drawbacks," or "worst." A positive value requires the key phrases "need," "should ... adopt," "should be considered," "benefit," or "better off." For (2), we categorize "full disagreement with the reported quote or incident" as -1, neutrality as 0, and "full agreement with the reported quote or incident" as 1. The key phrases "skeptic," "skepticism," "however," as well as colloquial expressions like "pooh-pooh" determine disagreement, while the phrases "indeed," "in addition to the mentioned," "have a point," and "points out correctly" trigger a positive measure. We assign a value of -0.5 ("weakly disagree") when phrasing indicates that others may disagree with the reported argument but does not spell out a counter-argument (e. g., "Advocates of VRMs, for their part, blame the fixed rates for the current mortgage-money famine."). If agreement or disagreement cannot be clearly characterized, we assign a value of 0.

Finally, we combine the two dimensions by adding (1) + (2) if (1) ≥ 0, and sub-

Figure 2. ARM Sentiment Index

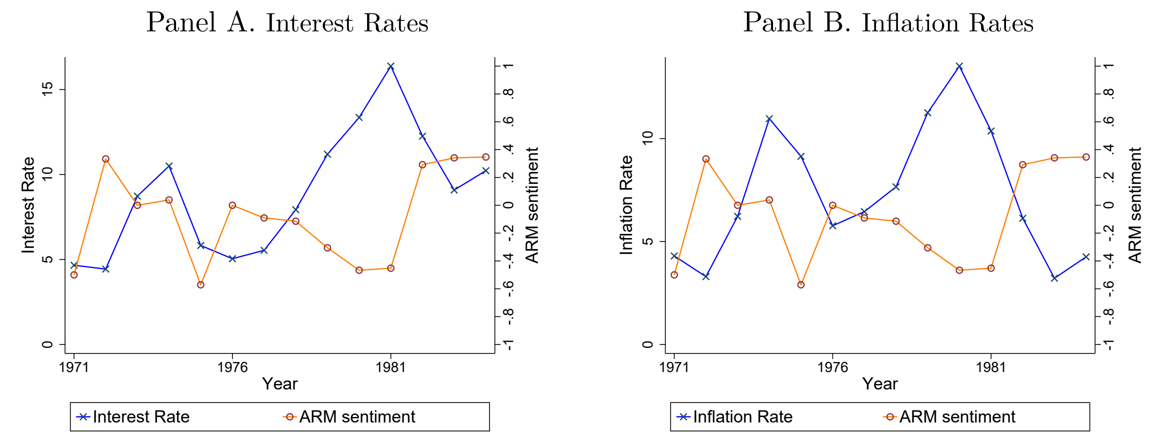

Notes. In Panel A, *Interest Rate* is the yearly average over monthly federal funds rates, obtained from the Federal Reserve Bank of St. Louis. In Panel B, *Inflation Rate* is the annual inflation rate, calculated with the CPI-U data from the BLS. The *ARM Sentiment Index* is calculated as the annual averages of the measure from the data set on the public discussion of ARMs, gathered from all articles in *The Wall Street Journal* and *The Washington Post* from 1971 to 1984 that discuss variable-rate products.

tracting (1) − (2) if (1) < 0, i. e., coding agreement with a negative stance as a negative attitude. The combined measure ranges from −2 to +2, with −2 indicating completely negative sentiment and +2 indicating completely positive sentiment towards ARMs.

In the graphs of Figure 2, we plot the annual averages of this ARM Sentiment Index against inflation and interest rates. In Panel A, we use the yearly averages of monthly federal funds rates from the Federal Reserve Bank of St. Louis, and in Panel B, annual inflation rates, calculated with the BLS CPI-U data.

Both panels reveal a strong negative relationship between the movement of those rates and the ARM Sentiment Index: each rise in interest or inflation coincides with a decline in public opinion about ARMs. In other words, whenever inflation and nominal interest rates are climbing up, the public debate tends to emphasize the downsides of ARM contracts, such as the rate hikes mortgage borrowers may incur. Whenever interest rates are low, these negative aspects dominate the public debate much less. For example, a typical quote from times of higher interest rates (and low sentiment) is: "However, area real estate people and builders were considerably less enthusiastic, with some expressing concern that consumers will be wary of the new mortgages, since rates can rise as fast and as high as general interest rates."[9] Times of lower interest

[9] *The Washington Post* 4/24/1981 ("Panel Allows Home Mortgages With Variable Interest Rates; Mortgage Rule Pleases Lenders Here").We code this paragraph as weakly negative towards ARMs (−0.5) and weak agreement by the newspaper writer (+0.5) for a total sentiment score of −1.

rates and higher sentiment, instead, feature quotes like "A unit of BankAmerica Corp said in a policy statement that the traditional long-term, fixed-interest-rate mortgage loan is too inflexible and should be scrapped as the keystone to home financing. It called for reforms [...] to put home ownership back within reach of more Americans."[10] In other words, exposure to rising inflation and interest rates in 1973–4 and 1978–80 sparked lender enthusiasm for implementing ARMs nationwide but ignited public pushbacks both times; exposure to falling inflation and interest rates made it possible sell consumers on the benefits of a loan that automatically refinances. Not until 1982 did these forces align; the Garn-St. Germain Act passed in the House by a vote of 272–91.

This negative relationship also emerges when looking at lagged inflation rates. In Figure A.1 in the Appendix, we anticipate the construction of measures of personal exposure to past inflation (from the next section) and relate ARM sentiment to inflation over the past years. The resulting graphs are similar or even more pronounced in terms of the observed negative correlation.

Overall, the historical background of the introduction of ARMs in the U.S. and the surrounding debate point to a role for the macroeconomic environment influencing attitudes towards ARMs. The negative correlations between realized rates and attitudes towards ARMs raise the possibility that exposures to historical inflation and nominal interest rates play a role in explaining the popularity and adoption of these products. In our empirical analyses, we investigate this possibility: we relate the limited role of ARMs in the U.S. to the choices of those cohorts that were exposed to the Great Inflation of the 1970s many decades later. If they were more wary of future interest rate hikes, and consequently more skeptical of ARMs, this would be consistent with the notion of "Long Shadows."

3 Measures and Data

The key hypothesis in this paper is that historical rates play a significant role in explaining the puzzling asset composition of the U.S. mortgage market and the sizable costs consumers incur as a result. We investigate the explanatory power of consumers' exposure to historical inflation and interest rates, and in particular the long-lasting consequences of the Great Inflation. In this section, we introduce the key measures and

[10] *The Wall Street Journal* 4/6/1976 ("Bank America Unit Advocates Reforms In Home Mortgages - It Calls for Negotiable-Term, Rollover Loans, Graduated Payments, Other Changes"). We code this paragraph as strongly positive towards ARMs (+1) and weak agreement by the newspaper writer (+0.5) for a total sentiment score of +1.5.

sources of data.

3.1 Exposure to Past Inflation

To make our research hypothesis testable, we build on the theoretical and empirical evidence in prior literature that has documented the longlasting effects of personal lifetime experiences of inflation (and other macro-finance variables) on beliefs and financial risk-taking. This literature has estimated a pattern of roughly linearly-declining weights with which individuals weight past lifetime realizations of macro-finance variables when forming beliefs about future realizations of the same variables.[11] We thus calculate the lifetime exposure to inflation as of year t for an individual born in year s as:

$$\pi^e_{s,t} \equiv \sum_{k=0}^{t-s} \frac{t-s-k}{\sum_{j=0}^{t-s}(t-s-j)} \cdot \pi_{t-k}\,. \tag{1}$$

This formula places the highest weight on the most recent observation ($k = 0$), zero weight on the year of birth ($k = t - s$), and connects these endpoints linearly. Thus, while recent realizations receive the highest weight, exposure earlier in life still carry significant weight.

We apply this formula to an extended time series of inflation data, based on the CPI-U from the U.S. Bureau of Labor Statistics (BLS) for 1913–2013. We use the spliced Warren and Pearson series available on Robert Shiller's website to extend this series back over 1876–1912. We calculate annual inflation π_t as the log change in the annual average level of the price index between years $t-1$ and t. We then calculate cumulative inflation exposure $\pi^e_{s,t}$ as of year t for individuals belonging to the cohort born in year s using (1).

Figure 3 illustrates the resulting evolution of $\pi^e_{s,t}$ for two representative households, an "older" household from the 1945 birth-year cohort, and a "younger" household from the 1960 birth-year cohort. The left panel (Panel A) shows the underlying path of historical inflation, where the solid line (filled squares) represents annual CPI-U inflation rates from 1960 to 2013, and the time of the Great Inflation is shaded in gray.[12] We also plot a hypothetical "No-Great-Inflation" path, shown in the lower line (hollow squares), to illustrate how such a scenario would have affected individuals' overall inflation experience. We use a location-scale transformation of actual inflation during the time of the Great Inflation to the No-Great-Inflation mean of 2.5% and S.D. of 1.1%.

[11] Malmendier and Nagel (2011) estimate that individuals apply roughly linearly declining weights to personally experienced past stock-market returns, starting from the current year. Malmendier and Nagel (2016) find that individuals form inflation expectations in a very similar fashion.

[12] Our methodology for dating the Great Inflation is inspired by Scrimgeour (2008); see Appendix D.

Figure 3. Actual and Hypothetical Inflation

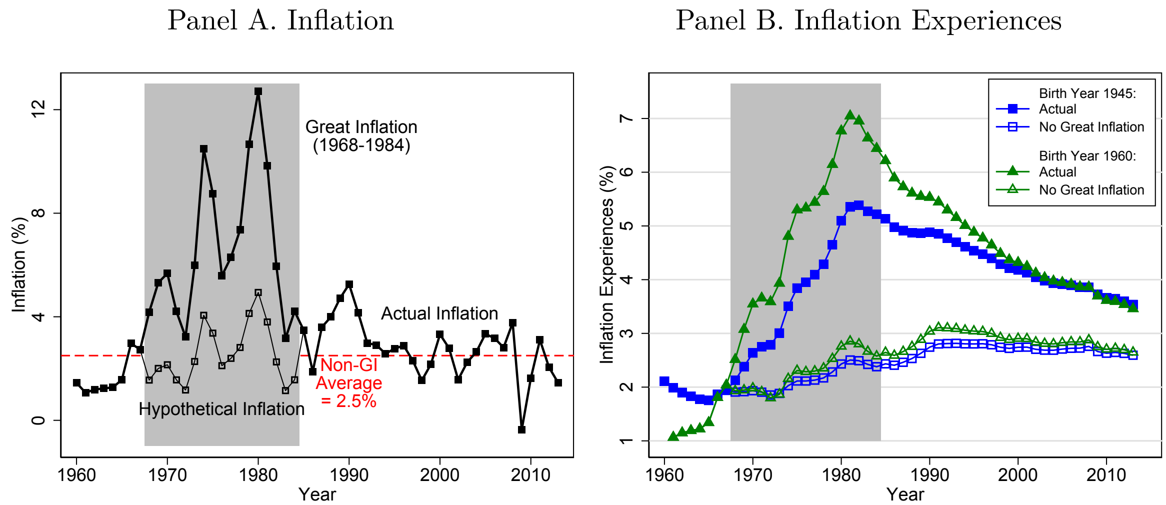

Notes. Data from BLS CPI-U. In Panel A, Actual Inflation (solid boxes) is the annual log change in CPI, and Hypothetical Inflation (hollow boxes) is a location-scale transformation of actual 1968–84 inflation (mean 6.68%, S.D. 2.84%) to the 1960–2013 (excluding 1968–84) mean of 2.55% and S.D. of 1.13%: $\pi_t^H := (\pi_t - 0.0668)/0.0284 \times 0.0113 + 0.0255$. In Panel B, solid symbols show lifetime inflation experiences as in equation (1) using actual inflation. Hollow symbols show the same, but use hypothetical inflation during 1968–84.

Panel B shows the corresponding lifetime weighted-average inflation experiences for both the actual and the hypothetical inflation paths, separately for the "young" and "old" cohorts. The comparison between these lines provide two main insights.

First, young borrowers are particularly affected by inflation shocks. The line for the actual experiences of the young generation (filled triangles) shoots up more steeply than that of the older generation (filled squares) during the time of the Great Inflation, reflecting that members of the younger cohorts have shorter personal histories of inflation and, hence, are more affected by the recent shocks. Notice that, even under the hypothetical "No Great Inflation" scenario, the lifetime average of the younger cohort (hollow triangles) would have increased by 30 basis points more than that of the older cohort (hollow squares) following the second oil crisis in 1979, reflecting the same mechanism. In reality, though, the difference in lifetime inflation exposure between cohorts climbed up significantly more, reaching 170 basis points in 1981. However, by the late 1990s the lifetime exposures of both cohorts are fairly similar in both scenarios.

Second, we also see that, following the actual Great Inflation, the lifetime averages of both cohorts remain higher than under the hypothetical scenario for many years, into the 1990s and 2000s. In other words, inflation shocks have a double effect: an immediate effect on the cross-section and a long-lasting effect on the level. In the empirical analysis we will derive and test the implications of these effects for mortgage choice.

3.2 Data

We rely on two main sources of data: the Survey of Consumer Finances (SCF) and the Residential Finance Survey (RFS).

The Survey of Consumer Finances is a triennial survey run by the Federal Reserve Board to gather detailed, household-level income and balance-sheet data. The SCF allows us to link household inflation experiences, future interest-rate expectations, and mortgage choice in the late 1980s through the late 2000s. Most questions have remained unchanged since 1989, when the SCF underwent a major redesign. Starting with the redesign in 1989, SCF question X302 asked respondents: "Five years from now, do you think interest rates will be higher, lower, or about the same as today?" This question appeared in every survey wave through 2013, after which it was discontinued. We thus confine our analysis to the survey waves between 1989 and 2013.

The typical survey wave interviewed around 4,000 households between 1989 and 2007, and around 6,000 starting in 2010. Since many financial assets are held disproportionately by a small fraction of households, the SCF oversamples the wealthiest households, and the use of sampling weights is important.[13] For example, average household income in the 2013 SCF is \$84k when correctly adjusted for heterogeneous sampling and response probabilities, but \$710k when unweighted. Also, due to the sensitive nature of the questions, data tend to be missing in a non-random fashion. Board statisticians use multiple imputation both to fill in the missing values and to replace non-missing values that might otherwise disclose respondents' identities.[14] The final public-use data contains five simulated "implicates" (imputation-replicates) per household. In our analyses, we calculate all point estimates using SCF sample weights, and we adjust the standard errors for multiple imputation using the standard Rubin (1987) formulas.

Table 1 shows weighted summary statistics from pooling the nine SCF waves between 1989 and 2013. The top panel includes all respondents between the ages of 25 and 74. The mean respondent is 47 years old, has a household income of \$87,560 per year, and reports a net worth of just under \$500k (both in constant 2013 dollars).

[13] For more details on SCF weight construction, see Kennickell and Woodburn 1999.

[14] See Kennickell (1998), and Montalto and Sung (1996) for a user-friendly discussion.

Two-thirds of respondents are home owners, and two-thirds of owners have a mortgage, reducing the sample from 36,266 to 16,824 households. The majority of mortgages have fixed-rate contracts (86%), and only a small fraction of 6% are "jumbo" loans (above the conforming loan limit). More than a quarter of mortgages (28%), however, are "non-conventional," meaning that they carry VA or FHA insurance or guarantees.

The bottom panel of Table 1 restricts the data to observations with a mortgage origination in the survey year. This restriction allows us to approximate a "flow" dataset of recent mortgage decisions. The resulting sample contains 3,257 "new" mortgages originated in one of the survey years between 1989 and 2013 by 3,022 households, just 8% of the original sample. The panel reports the summary statistics for these new mortgages separately for FRMs and ARMs. FRM and ARM holders are similar in age—contrary to the prediction of standard theory that younger households should prefer ARMs—but differ by most other characteristics. Borrowers choosing FRMs tend to have lower income, less net worth, take out smaller loans that are less likely to be "jumbo" and more likely to be non-conventional.

Turning to households' interest-rate expectations, we report the fraction of respondents expecting higher or lower interest rates as well as the net fraction answering "higher" minus the fraction answering "lower." On net, borrowers choosing FRMs are 10 pp more likely to expect higher future interest rates than are borrowers choosing ARMs, 68% versus 58%. The difference is largely driven by FRM holders expecting rising interest rates, and not by ARM holders expecting falling interest rates. This is consistent with the research hypothesis: the aversion to variable-rate borrowing is influenced by worries about future interest rate increases. On the other hand, borrowers choosing FRMs have lower lifetime inflation experiences (4.16% versus 4.34%). ARMs were (paradoxically) more popular in the mid-to-late 1980s, when memories of the Great Inflation carry the largest weight in our inflation experience formula (1). The missing piece here is that, at those times, lenders offered large ARM discounts to encourage takeup of this relatively new product. This points to the importance of performing a within-year calculation.

The SCF uniquely allows us to relate households' past exposure to inflation and interest rates to interest rate expectations. We will also employ the SCF to link past experiences to actual mortgage choice. One shortcoming of the SCF for the analysis of mortgage choice, however, is the relatively small number of respondents with recent mortgage originations, typically about 800 per survey wave. Another limitation is that respondents' geographic locations are not reported in the public data set due to privacy concerns (with the exception of three survey waves in the 1990s). Since our identification strategy requires the inclusion of year fixed effects, the lack of within-survey geographic variation prevents us from estimating some parameters of interest. Finally, the SCF also includes a less extensive list of mortgage contract characteristics than other data.

To remedy these shortcomings, we identify a second source of individual-level data on mortgage financing and demographics, which has not previously been employed in this context, the Residential Finance Survey (RFS). Conducted by the Census Bureau the year after every decennial Census between 1950 and 2000, the RFS is a much larger survey than the SCF: the 1991 homeowner survey interviewed 24,000 households, and the 2001 homeowner survey interviewed nearly 17,000. The unique feature of the RFS is that it consists of two cross-referenced surveys, one of households and one of their mortgage servicers. The household arm of the survey provides demographic and income data, while the lender arm provides the terms of any outstanding loans secured by the property. The sample is drawn from the previous year's Census roster of properties, so it misses newly-constructed housing. The survey oversamples multi-unit properties, particularly rental properties with 5+ units, but is otherwise representative of the stock of outstanding mortgages in the preceding Census year. Property locations are reported at the state level for 12 large states (CA, FL, TX, and NY in both survey years, plus eight additional states in 2001 only) and at the Census region level otherwise. In our final estimation sample we observe the state-level location for 44% of mortgages.

For our primary analysis, we utilize the microdata on mortgages linked to owner-occupied 1–4 unit properties from the 1991 and 2001 waves.[15] Since the sample provides information about outstanding mortgages, rather than flow data of mortgage originations, we do not observe mortgages that were refinanced, repaid in full, or defaulted upon prior to the survey year. To approximate flow data, we restrict the sample to mortgages taken out no more than six years prior to the survey year (1985–1991 and 1995–2001). Mortgagor age at origination is a key input for calculating inflation experiences; we use the age of the self-identified primary owner if the household has multiple members.

Some public-use RFS variables such as income and loan amount are coded to interval means to preserve respondent anonymity, and interest rates are left- and right-censored. We explicitly account for censored dependent variables in our estimation procedure. Also, origination years in the 1991 survey are reported by intervals: 1985–86, 1987–88,

[15] This definition includes second homes and vacation homes as the public-use version of the 1991 RFS does not allow to filter these out.

and 1989–91. To calculate inflation experiences, we assume origination occurred at the beginning of the interval, so as not to include future inflation rates that some borrowers had not yet experienced. When determining conforming versus jumbo status, we use the largest conforming loan limit in each time period, since loans tend to cluster just below this amount. We describe coding decisions for all key variables in Appendix C.

The RFS consistently reports data for three types of mortgage products across both survey waves: FRMs, ARMs, and balloon mortgages. Balloon mortgages are designed to attract borrowers who would not otherwise qualify for a fully-amortizing product. They offer lower monthly payments that are not fully amortizing, so a large lump ("balloon") payment is due at maturity, usually after 7–10 years. Borrowers may be able to refinance upon maturity if their situation has improved, but the mortgages carry greater risk as borrowers have to default if they cannot refinance and cannot afford the balloon payment (MacDonald and Holloway 1996).

Table 2 reports the summary statistics for the RFS data, separately for each type of mortgage. Only 4.8% of mortgages are balloon mortgages, and we will often focus on the comparison of fixed- versus variable-rate contracts. As in the SCF data, FRMs are somewhat smaller than ARMs, and they are also significantly more expensive.

To capture the contemporary economic conditions, we supplement the RFS with data from the Freddie Mac's Primary Mortgage Market Survey (PMMS), a weekly survey of average FRM and ARM rates from a representative nationwide sample of mortgage originators, broken out into five regions. Lenders provide quotes for first-lien, prime, conventional, conforming, home purchase mortgages with an 80% LTV and a 30-year term, for both FRMs and 1/1 ARMs. For ARMs lenders quote both the initial, "teaser" rate and the margin over the one-year Treasury rate after the loan resets. The PMMS provides a useful picture of baseline mortgage rates charged to the same high-quality borrower, across products and over time. The survey rates were reproduced in the Federal Reserve's H.15 "Selected Interest Rate" release until it stopped including non-Fed data in October 2016. Other popular interest rate series such as the FHFA's Monthly Interest Rate Survey (MIRS), which we utilize to consider supply-side implications, are instead drawn from *actual* mortgages and so reflect changes in the pool of borrowers across products and over time. In Appendix E, we show that, between January 1986 and October 2008, when both series are available, the PMMS tracks the slope of the nominal Treasury yield curve much more closely. (The FRM-ARM initial rate spread in the MIRS is about 75 basis points smaller on average, mostly due to higher ARM rates.) To match the PMMS to the RFS, we take annual averages of the weekly PMMS data, then match to borrower locations in the RFS using the Freddie Mac region containing the borrower's state, if reported; else we construct a Census region average by re-weighting the PMMS data from the five Freddie Mac regions to the four Census regions using 1990 Census housing units by state.

As the summary statistics reveal, inflation, the FRM-ARM spread, and the yield spread faced by ARM borrowers tended to be somewhat higher.

The borrower, property, and other loan characteristics, summarized in the lower parts of Table 2, go significantly beyond the level of detail available in the SCF data. When variables are comparable, they are again consistent with the SCF data, especially the comparisons of FRM versus ARM holders. (Some point estimates differ since the RFS time-series coverage stops in 2001 rather than 2013.)

Borrowers choosing ARMs tend to have higher income, are less likely to be first-time homeowners, and are more likely to take out a "jumbo" loan than FRM holders. As in the SCF, there is no significant age difference between FRM and ARM borrowers, contrary to the prediction of standard theory. Prior lifetime exposure to inflation is again a few bps lower for the typical FRM borrower than for the typical ARM borrower (4.74% versus 4.79%). This largely reflects across-year variation in mortgage product shares: FRM takeup is higher in years when current inflation is lower and when FRMs are relatively cheaper, i. e., in the late 1990s. However, this simple comparison pools across all origination years and ignores time-series variation in the relative cost of the two products. As will be seen below, individuals who have experienced higher inflation within an origination year are more likely to choose an FRM. As in the SCF, borrowers choosing FRMs tend to have lower income and take out smaller loans that are less likely to be "jumbo" and more likely to be non-conventional.

We perform the analysis of interest-rate expectations on the SCF data and the analyses of mortgage choice on both the SCF and RFS data sets.

Despite some data limitations, these data provide unique detail into American home buyers' characteristics, loan choices, and expectations in the decades after the Great Inflation. Moreover, while neither the SCF nor the RFS have borrowers' credit scores, we also observe previous access to mortgage credit in the RFS. Combined, these details allow us to make significant progress in constructing hypothetical alternative rates. The datasets are unique in that we can map respondents' beliefs to their choice behavior.

4 Interest Rate Expectations

We begin our analysis by considering beliefs about future interest rates, as they are a key determinant of the ARM-versus-FRM choice. Using the SCF, we test whether consumers' prior exposure to higher or lower inflation over their lives so far influences their beliefs about future nominal rates. We then document the strong influence of interest-rate expectations on mortgage choice, including a first indication of the role of experience-based interest rate expectations.

4.1 Interest Rate Expectations and Inflation Exposure

Prior research has shown that past inflation affects consumer beliefs about future inflation for years (and decades) to come.[16] Thus, by the Fisher equation, $i = r + \mathbb{E}\pi$, it should also affect beliefs about future nominal interest rates, i. In particular, individuals coming of age during periods of high inflation should expect higher nominal interest rates in the future.

To test whether prior exposure to inflation affects households' views on future interest rates, we use the SCF question: "Five years from now, do you think interest rates will be higher, lower, or about the same as today?" While the interest-rate belief variable in the SCF is coarse compared to more modern surveys such as the New York Fed's Survey of Consumer Expectations (launched in 2013), its wording is similar to the question used for many years in Michigan's Survey of Consumers, and it elicits beliefs over a longer forecast horizon (five years rather than one year). For each survey wave between 1989 and 2013, and separately for each cohort, we calculate the net fraction of respondents expecting interest rates to rise as the fraction answering "higher" minus the fraction answering "lower." We then relate this fraction to the personal inflation exposure of the same consumers over their lifetimes so far.

We start from a graphical illustration, shown in Figure 4. For visual purposes, we separate "older" and "younger" cohorts, above and below the sample median of age, respectively. We then calculate the deviation of each group's response from the overall survey-year mean, equivalent to including survey-year fixed effects in a linear regression.

Figure 4 reveals that members of the younger cohorts (in dark red) were more likely to expect interest rates to rise on net than members of the older cohorts (in light blue) during the early SCF years (1989, 1992, etc.). This relationship reverses in the mid-2000s. At that time new, younger households start to have less positive expectations than older households.

[16] Cf. Malmendier and Nagel (2016).

Figure 4. Interest-Rate Expectations and Inflation Experiences

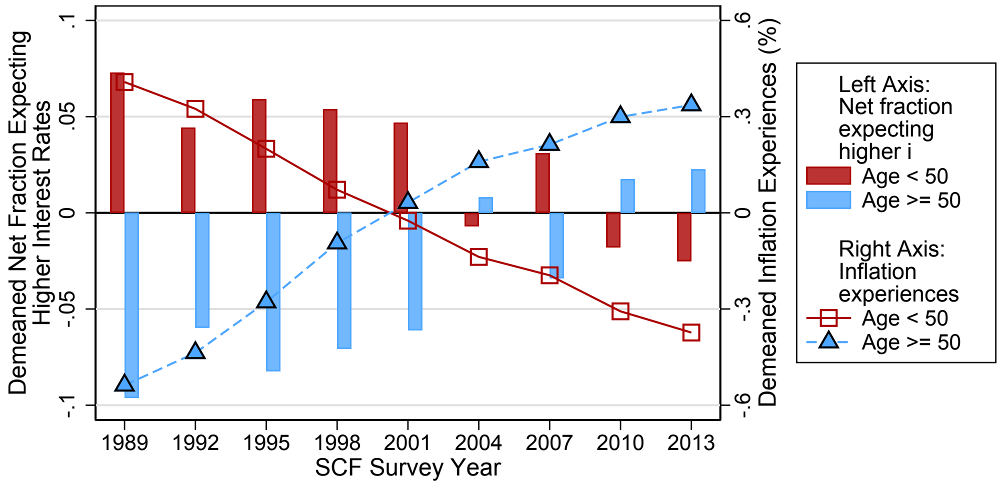

Notes. Values shown are cohort deviations from survey-year mean (average across implicates).

The timing of this reversal in interest-rate expectations coincides almost exactly with the movement of the cross-sectional differences in survey respondents' lifetime exposure to past inflation, calculated using equation (1). As the graph also shows, the lifetime inflation-rate exposure of the younger cohorts (red empty squares) is up to 1 pp higher than that of the older cohorts (blue filled triangles) during the early survey waves, but the relative position of the two groups switches in the mid-2000s. At that point, the prior lifetime of older households starts to hold higher inflation experiences than that of younger households. This reversal happens as new, younger households who did not live through the Great Inflation enter the sample, and the households who did experience the Great Inflation age and become older households. Moreover, the memory of the Great Inflation among older cohorts is fading (i. e., is weighted less).

We confirm the strong visual pattern of the relation between inflation experiences and interest expectations in a formal regression framework, where we account for each cohort-year separately rather than averaging above and below median-age. We estimate

$$\iota_{n,t} = \alpha_{0,t} + \alpha_1 \pi^e_{n,t} + \xi_{n,t}, \tag{2}$$

where $\iota_{n,t}$ is individual n's forecast of future interest rates based on time t information, captured by year fixed effects $\alpha_{0,t}$, as well as the history of n's past inflation experiences, $\pi^e_{n,t}$, and idiosyncratic factors, $\xi_{n,t}$. We note that running the comparison within year is critical. The SCF question elicits respondent beliefs about the *change* in nominal interest rates, $i_{t+5} - i_t$. Looking within year removes the level of the current interest rate i_t. (In our subsequent empirical analysis, we will continue to look "within year" by including year fixed effects in all econometric specifications that also include exposure to past inflation.)

We use three measures of expectations. First, we capture the net fraction expecting higher rates exactly as in the graphical illustration of Figure 4. That is, we code respondents expecting interest rates to rise as $+1$, to fall as -1, and to stay the same as 0. Table 3 shows the corresponding coefficients both under a linear probability model (in column 1) and an ordered probit model (column 2). Second, we employ an indicator for whether respondents expect future interest rates to fall (shown in columns 3 and 4), and third an indicator for respondents expecting interest rates to rise (columns 5 and 6), each time using a linear-probability and a probit model, respectively. The advantage of using LPM is that the coefficients are directly interpretable as marginal effects, though the results are robust to choice of estimation method.

The estimation results shown in Table 3 confirm the visual evidence. Past inflation exposure has a powerful influence on households' beliefs about future interest rate movements for all variants of the independent variable and both estimation methods. Our research hypothesis implies that the coefficient α_1 should be positive in columns 1, 2, 5, and 6, and that it should be negative in 3 and 4. Focusing on the LPM estimates, we see that one additional pp of experienced inflation predicts that a survey respondent is 9.9 pp more likely to expect higher future interest rates on net (col. 1, $p < 0.01$). Columns 3 through 6 show that higher lifetime inflation experiences shift the entire distribution of beliefs to the right. Respondents are 3.9 pp less likely to expect interest rates to fall (col. 3, $p < 0.01$) and 6 pp more like to expect interest rates to rise (col. 5, $p < 0.01$). Thus about 60% of the net effect is driven by more households expecting rising interest rates and about 40% by fewer households expecting falling interest rates.

The robust link between personal exposure to past inflation and interest rate expectations is a first key step towards evaluating the main hypothesis. It lends plausibility to the hypothesized relationship between historical inflationary periods and borrower behavior years later. Individual investors are likely to have interest rates, not inflation on top of their minds, and the above analysis confirms the direct relationship in individuals' beliefs.

4.2 Interest Rate Expectations and Mortgage Choice

Building on the influence of past exposure to inflation on interest-rate beliefs, we turn to an empirical framework that relates interest-rate beliefs to mortgage choice and test this relationship. A key challenge is that the SCF consists of repeated cross-sections every three years, rather than longitudinal data tracking the same households over time, so we observe interest rate beliefs *after* the mortgage was originated instead of *at the same time* as the decision. We derive how this timing discrepancy affects the consistency of OLS and IV estimators, and use these results to estimate lower and upper bounds on the true effect size.

Empirical Framework. Consider the following simple model of mortgage choice. The present value of a fixed-rate mortgage obligation to the borrower is the sum of future payments, which are fixed, divided by a compounded interest rate. Similarly, the present value of an adjustable-rate mortgage obligation is the sum of future payments, which adjust up and down with future interest rates, divided by a compounded interest rate. As just discussed, a borrower who has personally been exposed to higher inflation to date will expect higher inflation in the future, so they will also forecast higher nominal interest rates for any given future path of real interest rates. This means that they will discount any future nominal dollar amount by more. Since future FRM payments are fixed in nominal terms, the perceived present value of the FRM falls as future interest-rate forecasts rise. By contrast, the present value of an ARM is independent of one's beliefs concerning future interest rate movements. A borrower forecasting higher future interest rates will both expect higher future nominal ARM payments and discount them using a larger nominal interest rate. These two effects offset. The bottom line is that borrowers who forecast higher future interest rates will perceive the FRM to be relatively cheaper. Hence, the empirical prediction is that borrowers who have higher interest-rate expectations (e. g., due to their past inflation exposure) are predicted to have a greater inclination to choose an FRM over an ARM.

This reasoning suggests a structural model with three variables of interest and two causal relationships. First, past exposure to high inflation raises an individual's forecast of future nominal interest rates, as captured by a positive coefficient estimate α_1 in estimating equation (2) from above, $\iota_{n,t} = \alpha_{0,t} + \alpha_1 \pi^e_{n,t} + \xi_{n,t}$. And second, individuals who forecast higher future interest rates are more likely to choose an FRM over an ARM, *ceteris paribus*. That is, we estimate a latent utility choice model:

$$U_{n,FRM} - U_{n,ARM} = \delta_0 + \delta_1 \iota_{n,t} + x_n' \delta_x + u_{n,t} \,, \tag{3}$$

where $U_{n,j}$ is the utility that the individual enjoys when choosing mortgage product j; $\iota_{n,t}$ is individual n's forecast of future interest rates based on time t information; and x_n are other socio-demographic factors that influence mortgage contract choice, such as age, income, wealth, currently prevailing FRM and ARM interest rates, etc. (In our dataset each individual n is only observed once, at time t, so we omit the time subscripts on the U's and x.) Our theory indicates that the coefficient δ_1 is also positive.

Equations (2) and (3) formalize our reasoning of how past exposure to high inflation affects mortgage choice, with beliefs about future interest rates as the intermediating variable. Our aim in this section is to jointly estimate the system. This requires a few further assumptions. Real-world inflation is highly persistent: between 1960 and 2013, the autocorrelation parameter on annual log CPI-U inflation was $\phi = 0.8$. To capture this dynamic, we will model inflation as a stationary AR(1) process:

$$\pi_{t+1} = \mu + \phi(\pi_t - \mu) + \epsilon_{t+1}\,,\ 0 \leq \phi < 1\,. \tag{4}$$

Similarly, where there is no new information, an individual's interest-rate forecast should be fairly similar from year to year. We will model this by supposing that the forecast error in (2) follows an AR(1) process:

$$\xi_{n,t+1} = \varphi\xi_{n,t} + \nu_{n,t+1}\,,\ 0 \leq \varphi < 1\,. \tag{5}$$

Given this structure, we need to make some assumptions on the joint error distribution (u, ϵ, ξ, ν) in equations (2) to (5) to guide our estimation. We will assume that the errors are all mean zero; that the inflation innovations ϵ and the forecast innovations ν in (4) and (5) are unpredictable white noise, that the structural model errors in (2) and (3) are contemporaneously orthogonal, and that the regressors in (2) and (3) are pre-determined (a weaker, time-series version of exogeneity). Letting Ω_{t-1} denote all time $t-1$ information:

$$\mathbb{E}[u_{n,t}] = 0\,,\ \mathbb{E}[\nu_{n,t} \mid \Omega_{t-1}] = 0\,, \mathbb{E}[\epsilon_t|\Omega_{t-1}] = 0\,; \tag{6}$$

$$\mathbb{E}[u_{n,t}\xi_{n,t}] = 0\,; \tag{7}$$

$$\text{and } \mathbb{E}[(\pi_t, \pi_{t-1}, \ldots)'\xi_{n,t}] = 0\,,\ \mathbb{E}[(\iota_{n,t}, x_n')'u_{n,t}] = 0\,. \tag{8}$$

Note that Assumption (6) implies that the forecast errors are mean zero, by inverting (5): $\mathbb{E}[\xi] = (1 - \varphi(L))^{-1}\mathbb{E}[\nu] = 0$.

Under these assumptions, equations (2) and (3) can be estimated consistently by single-equation methods. Equation (2), the regression analog to Figure 4, can be estimated by OLS. If we only observe the sign of $\iota_{n,t}$ or the sign of $\iota_{n,t} - i_t$, then (2) is a latent-variable threshold model; if we further make a distributional assumption on the error ξ, then it can be estimated consistently by maximum-likelihood methods such as logit or probit (as we used in Table 3). The same applies to (3) if we know the distribution of u.[17] We refer to these as "OLS-like" estimators, since identification in all these methods relies on the OLS-like orthogonality condition that the regressors are exogenous.

The implementation of this empirical model faces one additional difficulty, arising from the structure of the survey data: the SCF provides information about interest-rates beliefs for mortgage borrowers at the time of the survey, not at the time the mortgage is taken out. That is, instead of observing $\iota_{n,t}$ from the time of the mortgage choice, we only observe an *ex post* interest rate forecast $\iota_{n,t+1}$. This does not impact estimation of (2), in which ι is the dependent variable; but estimation of (3) requires the contemporaneous forecast as an explanatory variable. As is well known, using a mismeasured explanatory variable creates an endogenous-regressors problem: OLS-like estimators are inconsistent.

To derive the asymptotic bias in the mortgage-choice estimation with mismeasured interest rate beliefs, we express $\iota_{n,t+1}$ recursively as

$$\iota_{n,t+1} = \iota_{n,t} + \Delta\iota_{n,t+1}, \tag{9}$$

$$\text{where } \Delta\iota_{n,t+1} = \alpha_1 \Delta\pi^e_{n,t+1} + (\varphi - 1)\xi_{n,t} + \nu_{t+1}\,, \tag{10}$$

with Δ the first-difference operator, following from (2) and (5). Plugging (9) into (3) gives us a feasible second-stage mortgage choice regression:

$$U_{n,FRM} - U_{n,ARM} = \delta_0 + \delta_1 \iota_{n,t+1} + x_n' \delta_x + \underbrace{(u_{n,t} - \delta_1 \Delta\iota_{n,t+1})}_{u^*_{n,t}}. \tag{11}$$

There is now a composite error term $u^*_{n,t}$, consisting of the structural error $u_{n,t}$ minus a "measurement error" term $\delta_1 \Delta\iota_{n,t+1}$.[18]

The comparison of (11) to (3) resembles an errors-in-variables situation. Indeed, if the interest-rate forecast ι were a random walk, then $\Delta\iota$ would be classical measurement error, and OLS-like estimators would attenuated. To see this, observe that the *ex post* interest rate forecast $\iota_{n,t+1}$ is positively correlated with the measurement error term $\Delta\iota_{n,t+1}$ by (9), and this latter term has a negative coefficient in (11), so $\iota_{n,t+1}$ is negatively correlated with the composite error term u^*. In the classical setting, we

[17] E.g., if u has a standard normal distribution, then (3) is a binary-choice probit model. The probability that individual n will choose a fixed rate mortgage is simply the probability that their latent utility difference from the FRM over the ARM is positive: $\Pr(\text{Choose FRM}_n) = \Pr(U_{n,FRM} - U_{n,ARM} \geq 0) = \Pr(-u_{n,t} \leq \delta_0 + \delta_1 \iota_{n,t} + x_n' \delta_x) = \Phi(\delta_0 + \delta_1 \iota_{n,t} + x_n' \delta_x)$.

[18] Similarly, plugging (9) into (2) would give us a feasible first-stage regression equation of *ex post* interest rate expectations on *ex ante* inflation experiences. However, we observe and can use *ex post* inflation experiences as the regressor, so this is not a problem. We show in Appendix F that the OLS estimator otherwise would be attenuated toward zero.

could use past exposure to inflation as of the time of the mortgage choice, $\pi^e_{n,t}$, as an instrument for interest-rate beliefs at the same time, $\iota_{n,t}$, and the resulting IV-like estimator would be consistent.

However, ι is not a random walk, it is the sum of two serially-correlated processes, so the measurement error is non-classical. We derive in Appendix F that, similarly to the classical setting, OLS-like estimators are inconsistent, but differently from that setting, so are IV-like estimators. This is because our measurement error term is negatively correlated with the correct but unobserved regressor: $\mathbb{E}[\Delta\iota_{n,t+1}\iota_{n,t}] < 0$. Intuitively, any instrument that is positively correlated with an individual's time-t interest rate forecast must be negatively correlated with the subsequent change in their forecast, so it cannot be exogenous. But since $\Delta\iota$ has a negative coefficient in (11), an instrument that is negatively correlated with $\Delta\iota$ will be *positively* correlated with the composite error term u^*. Hence, IV-like estimators are amplified rather than attenuated.

We utilize that the OLS- and IV-like estimators are inconsistent in the opposite direction to place bounds on the true effect size by comparing $\hat{\delta}_{1,OLS}$ and $\hat{\delta}_{1,IV}$. The "OLS" and "IV" probability limits are

$$\text{plim } \hat{\delta}_{1,OLS} < \delta_1 < \text{plim } \hat{\delta}_{1,IV} \, . \tag{12}$$

That is, an "OLS" regression of (11) will give a lower bound for δ_1 due to measurement error and attenuation bias; and an "IV" regression using contemporaneous inflation experiences as the instrument will give an upper bound for δ_1.

Finally, plugging (2) into (3) gives us the estimating equation:

$$U_{n,FRM} - U_{n,ARM} = (\delta_0 + \delta_1\alpha_{0,t}) + \delta_1\alpha_1\pi^e_{n,t} + +x'_n\delta_x + (u_{n,t} + \delta_1\xi_{n,t}) \, . \tag{13}$$

Since $\pi^e_{n,t}$ is orthogonal to both structural error terms $u_{n,t}$ and $\xi_{n,t}$, OLS-like estimators of (13) are consistent for the structural coefficient product $\delta_1\alpha_1$. Thus, estimating the effect of inflation experiences on mortgage choice is relatively straightforward, while estimating the intermediating role of interest rate expectations is challenging.

Results. The system of equations (2) and (3) fall under the classic case of using a bivariate probit in a simultaneous equations framework (case 3 of Heckman (1978)) and, assuming the mortgage choice and interest rate forecast errors (u, ξ) are jointly normal, may be estimated by bivariate probit maximum likelihood (Zellner and Lee 1965; Ashford and Sowden 1970). We estimate the system under this parametric assumption in Table 4.

We begin by estimating the feasible probit regression of mortgage choice on *ex* *post* interest rate expectations, equation (11). As just discussed, this equation suffers from errors-in-variables: *ex ante* interest rate expectations are exogenous but *ex post* expectations are not. Using the ex-post regressor attenuates the estimate toward zero. Column 1 estimates equation (11) using the full pooled SCF sample. Here, we expect measurement error to be quite severe for many observations, as the mean gap between the origination year and the survey year is 5 years in the full sample. In column 2 we minimize this problem by restricting the estimation sample to "new" mortgages taken out during the survey year only.

We observe a positive relationship between expecting higher future interest rates and choosing FRMs in both samples. As expected, attenuation bias is most severe in the full sample, column 1. The coefficient estimate in column 2 is nearly five times larger, 0.24 versus 0.05, and significant at a 1 % level. However, column 2 does not eliminate the measurement error in ι because the survey is still being conducted after the mortgage decision, by as many as 12 months.

In columns 3 and 4 we jointly estimate the system of equations (2) and (3) by bivariate probit, continuing to restrict to new mortgages only. As shown above in (12), the probability limit of the OLS-like estimator in columns 1 and 2 is biased downward, while the probability limit of the IV-like estimator in column 3 is biased upward. The results show that the IV-like estimate in column 3 is significant at a 1% level, five times larger than the OLS-like estimate in column 2, and 25 times larger than the severely attenuated coefficient in column 1. The 95% confidence intervals come close but are non-overlapping: the upper bound in column 2 is 0.42 and the lower bound in column 3 is 0.73.

To assess the economic magnitude of the estimated effect, given these bounds, we can compare the coefficient on interest rate expectations ι to the coefficient on the FRM-ARM spread from the PMMS. Formally, we take the total derivative of utility in (3) and set it equal to zero: $d(U_{n,FRM} - U_{n,ARM}) = \delta_1 \partial\iota_{n,t} + \delta_{Spread} \partial Spread_t = 0$. This generates the following increase in individuals' willingness to pay (WTP) for an FRM:

$$WTP \coloneqq \left. \frac{\partial Spread_t}{\partial \iota_{n,t}} \right|_{d\Delta U_n = 0} = -\frac{\delta_1}{\delta_{Spread}} \,. \tag{14}$$

(See Train 2009, ch. 3.) In column 2 the coefficients are nearly the same in magnitude, indicating that the average individual who expects interest rates to rise, as opposed to stay the same or fall, is willing to pay $-0.245/-0.232 = 1.05$ p.p.s more for an FRM. (For reference, the standard deviation of the PMMS mortgage spread over 1984–2013 is

0.67 p.p.s.) In column 3 the WTP is $-1.351/-0.249 = 5.4$ p.p.s. These estimates of the lower and upper bounds on the true effect size indicate that interest rate expectations have a powerful, and potentially expensive, influence on mortgage choice.

We note that the bivariate probit imposes a strong assumption on the joint distribution of the errors. We test this by running the score test proposed by Murphy (2007), which is asymptotically distributed as $\chi^2_{(9)}$.[19] The sample test statistic is 100.5, as compared to a 5% critical value of 16.9, leading us to reject the null hypothesis of bivariate normality. However, while misspecification of the likelihood function may affect the coefficient estimates, this does not necessarily cause inconsistency in the WTP estimates. Ruud (1983) shows that discrete-choice maximum likelihood (probit or logit) on a misspecified error distribution can still estimate the slope parameters consistently, up to an unknown biasing scale factor, under fairly non-restrictive conditions. The WTP calculation is constructed from the ratios of coefficients, so will eliminate such a biasing scale factor. As a further robustness check, we re-estimated the entire table by LPM; the results showed the same pattern of attenuation and amplification.[20]

Finally, in column 4 of Table 4 we estimate the probit model of mortgage choice on inflation experiences (13). Under our empirical framework assumptions, this equation may be consistently estimated by itself. The results indicate that households with higher lifetime experiences of inflation are significantly more likely to choose an FRM than an ARM ($p = 0.014$).

The estimate provides a first step towards our goal of identifying and quantifying the effect of historical inflation, and in particular the influence of the Great Inflation, on mortgage choice. However, we are unable to calculate a WTP resulting from historical exposure in the SCF for at least two reasons. First, the year fixed effects absorb the time-varying PMMS rates, since the SCF does not provide geographic information. Second, even if we could include these, the PMMS rates are only proxying for the menu of interest rates that an individual household actually faced given its ability to pay and credit characteristics.

In the next section we turn to the RFS and use a three-step estimation procedure to predict these missing, household-level interest rates and estimate a structural model of mortgage choice, which we will then use to assess the economic consequences of high

[19] Chiburis et al. (2012) implement the test in Stata as `scoregof`.

[20] The LPM-2SLS table is available upon request. To summarize: the average marginal effect of expecting higher future interest rates on choosing an FRM is +1.3 p.p.s by OLS on the full sample ($p = 0.096$), +5.9 p.p.s by OLS on the new-mortgage sample ($p = 0.013$), and +32 p.p.s by 2SLS on the new-mortgage sample ($p = 0.002$).

exposure to historical inflation in Sections 6 and 7.

5 Inflation Exposure and Mortgage Choice

We start the analysis of the RFS data with a graphical illustration of the two main implications of our research hypothesis for mortgage choice: (i) Mortgage borrowers who have been exposed to higher historical inflation are more likely to choose FRMs. (ii) Younger cohorts respond more strongly to recent inflation realizations. The first implication reflects the mechanism documented in the previous section: households with personal exposure to high past inflation also expect high future (nominal) interest rates. As a result, they estimate the present value of fixed repayment obligations in real terms to be lower, and future variable rates to be higher. Hence, they are predicted to have a higher willingness to pay for fixed-rate mortgages. The second implication reflects the recency bias embedded in experience-based learning. Younger individuals with shorter lifetime histories so far overweight recent experiences more than those with longer histories, and hence, respond more strongly to recent inflation realizations.

Figure 5. FRM Share and Experienced Inflation by Age Group

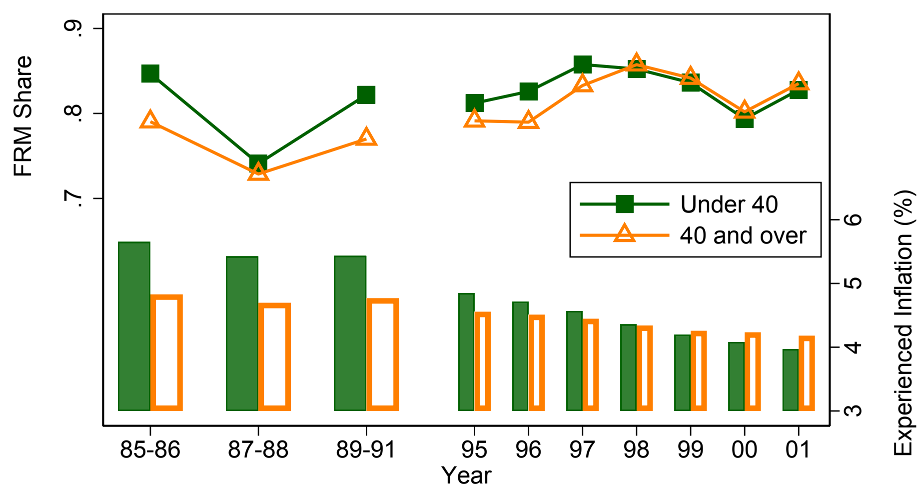

Notes. Data from the 1991 & 2001 RFS and BLS CPI. The 1991 RFS reports origination years in two- or three-year intervals.

Figure 5 illustrates that both predictions hold in the aggregate. Splitting the RFS sample at the median age of 40, we plot the FRM share and inflation experiences of "younger" and "older" borrowers in 1985–1991 and 1995–2001. In the late 1980s, younger cohorts with shorter personal histories (so far) were more affected by the Great Inflation and were more likely to choose fixed rates than older cohorts. In the late 1990s, the inflation experiences of (new) younger and older cohorts converged—memories of the Great Inflation slowly faded for aging, older cohorts, while younger cohorts had no personal memory of it—and so did their mortgage choices. In our main analysis, we test for this pattern formally, in a rich econometric framework.

5.1 Estimation Methodology

Our objective in this section is to estimate the economic importance of personal exposure to past inflationary periods for borrowers' choice of fixed-rate over variable-rate mortgages, as a stepping stone toward calculating the welfare consequences of the Great Inflation in this domain. In the SCF analysis, the PMMS interest rates are absorbed by year fixed effects. In the RFS analysis presented here, we take advantage of within-year geographic variation in the data to overcome these challenges. The section provides an overview of our methodology, and full details are in Appendix G.

We start from generalizing the empirical framework of Section 4 to more than two mortgage types i, namely, $i \in \{FRM,\ ARM,\ Balloon\}$. Suppose that each household n chooses a mortgage once, in year y, and derives indirect utility

$$U_{n,i} = \beta_{0,i,y} + \beta_{R,i} Rate_{n,i} + \beta_{\pi,i}\pi^e_n + \beta_{Inc,i} Income_n + f_i(Age_n) + \upsilon_{n,i} \tag{15}$$

when choosing alternative i. Mortgage preferences depend on a host of demographics and proxies for risk attitudes, including age, mobility, current and expected future income, risk aversion, and beliefs about future short-term interest rates (see, e. g., Stanton and Wallace 1998, Campbell and Cocco 2003, Chambers et al. 2009, and Koijen et al. 2009).[21] Our main observable characteristics are the alternative-specific interest rate $Rate_{n,i}$ offered to borrower n; the borrower's (log) income $Income_n$; and an alternative-specific function of the borrower's age, $f_i(Age_n)$, which we specify as quadratic to capture non-linear life-cycle variation in the attractiveness of a mortgage-contract type. The error term $\upsilon_{n,i}$ accounts for any unobservable factors affecting mortgage choice.

The explanatory variable of interest is borrower n's personal exposure to inflation at the time of the choice situation, π^e_n. As shown in Section 4, inflation experiences only affect households' decision-making indirectly, via their expectations about future inflation and nominal interest-rate movements, so are exogenous.[22] For the interpre-

[21] Koijen et al. (2009) suggest that households use the average of recent short-term Treasury rates to predict future ARM payments after the reset. This decision rule is compatible with learning-from-experiences as it only exploits time-series variation, which is absorbed by time fixed effects in our econometric model.

[22] To see that equation (15) generalizes the empirical framework of Section 4 to $J \geq 2$ alternatives, recall that only differences in utility affect choice behavior. When $J = 2$, a household will choose an

tation of the corresponding coefficient of interest, $\beta_{\pi,i}$, it is essential that we continue to include year fixed effects. The alternative-specific year fixed effects $\beta_{0,i,y}$ capture all aspects of the economic environment at a given time and all information that is common to all households and might enter the rational-expectations forecast of future interest rates ι, including the full history of past inflation (equation (2)). Thus, a borrower's lifetime inflation experiences should not matter unless there is a correspondence between those experiences and borrower beliefs that differs from the baseline rational-expectations forecast. Normalizing $\beta_{\pi,ARM} = 0$, our hypothesis implies $\beta_{\pi,FRM} > 0$, while the standard rational framework predicts $\beta_{\pi,FRM} = 0$.

Alternative i is chosen by household n if

$$D_{n,i} \coloneqq \mathbb{I}\{U_{n,i} > U_{n,j} \quad \forall j \neq i\} \tag{16}$$

equals 1. This could be estimated by standard discrete-choice methods such as probit or logit, except for one major hurdle: interest rates of non-chosen alternatives are not observed. A naïve approach to fill in the missing rates would be to estimate the relation between observed rates and borrower characteristics on the sample of chosen mortgages, and use these estimates to predict mortgage rates offered to all households, as in:

$$Rate_{n,i} = \gamma_{0,i} + \gamma_{R,i} PMMSRate_{y,r,i} + z_n' \gamma_i + \zeta_{n,i}\,. \tag{17}$$

The Freddie Mac survey rate $PMMSRate_{y,r,i}$ represents the baseline price charged to a high-quality borrower in the same year y and Census region r as borrower n, taking out mortgage product i; the other explanatory variables z_n are household risk proxies such as income, first-time homeowner status, marital status, urban/rural property location, and loan size. This model can be estimated separately for each mortgage type i, including the same controls but allowing them to take different values γ_i.

However, since households were not randomly assigned to mortgage types, OLS on (17) will likely be inconsistent due to selection bias. To overcome this, we utilize a three-step procedure suggested by Lee (1978) and Brueckner and Follain (1988).[23] Plugging (17) into (15), we obtain a *reduced-form* choice model that we can estimate:

$$U_{n,i} = \tilde{\beta}_{0,i,t} + \tilde{\beta}_{R,i} PMMSRate_{y,r,i} + \beta_{\pi,i}\pi_n^e + \tilde{\beta}_{Inc,i} Income_n + f_i(Age_n) + \tilde{z}_n' \tilde{\gamma}_i + \tilde{\upsilon}_{n,i}\,. \tag{18}$$

FRM if $U_{n,F} - U_{n,A} > 0$, where $U_{n,F} - U_{n,A} = (\beta_{0,F,y} - \beta_{0,A,y}) + \beta_{R,F} Rate_{n,F} - \beta_{R,A} Rate_{n,A} + (\beta_{\pi,F} - \beta_{\pi,A})\pi_n^e + (\beta_{Inc,F} - \beta_{Inc,A}) Income_n + f_F(Age_n) - f_A(Age_n) + (\upsilon_{n,F} - \upsilon_{n,A})$. Let $\delta_1\alpha_1 \coloneqq \beta_{\pi,F} - \beta_{\pi,A}$, $\delta_{0,y} \coloneqq \beta_{0,F,y} - \beta_{0,A,y}$, $\delta_x \coloneqq \beta_{x,F} - \beta_{x,A}$ for all other sociodemographic variables x_n that only vary by household, and $u_n \coloneqq \upsilon_{n,F} - \upsilon_{n,A}$. Restricting $\beta_{R,F} = -\beta_{R,A}$ gives equation (13).

[23] Lee (1978) confronted a similar problem when estimating the wages of union versus non-union jobs, and Brueckner and Follain (1988) first applied Lee's methodology to a mortgage-choice setting.

We place tildes on coefficients and variables that represent different objects in (18) than in (15). The important takeaway is that we have eliminated the missing data problem by replacing household-level rates $Rate_{n,i}$ with Freddie Mac survey rates $PMMSRate_{y,r,i}$, which do not depend on household characteristics and are always observed for all alternatives. Moreover, since model (17) does not include inflation experiences, the reduced-form model (18) consistently estimates the structural coefficient $\beta_{\pi,i}$.

Our three-step estimation procedure is as follows. First, we estimate the reduced-form choice model (18), where households' decisions depend on region- and time-specific baseline FRM and ARM rates from the PMMS. In the second step, we estimate two mortgage pricing equations (17), where the household's actual FRM (or ARM) interest rate depends on the regional FRM (or ARM) survey rate plus household characteristics that adjust for risk. We use censored least absolute deviations (CLAD, Powell 1984) to account for top-coding in the dataset. To correct for selection bias, we use the predicted choice probabilities from the first step to construct a semiparametric selection correction (SPSC) estimator suggested by Newey (2009), which generalizes the model of Heckman (1979) by using a series approximation for the selection-bias term. Identification of the SPSC model relies on two technical conditions (discussed in Appendix G) and a cross-equation exclusion restriction: conditional on the FRM survey rate, the ARM survey rate does not directly influence the FRM rate a household is offered, and vice versa. That is, we assume that the ARM rate a household is offered is a risk-adjusted markup over the ARM survey rate only. So, the ARM survey rate provides exogenous variation in the probability of choosing an FRM in the first step that may be used to correct for selection bias in the FRM rate equation in the second step, without relying on a distributional assumption. We provide evidence in Appendix E that the PMMS rates are uncorrelated with observed borrower characteristics. This provides reassurance that the exclusion restriction is not violated by the survey rates picking up unobserved borrower characteristics such as credit score and are well suited to the task of predicting selection-corrected, risk-adjusted household mortgage rates.

Since the SPSC control function absorbs the intercept, we follow the suggestion of Heckman (1990) and estimate the intercept of (17) as the median difference between the observed and predicted mortgage rate for households with choice probabilities closest to 1 (i.e., those suffering from the least selection bias). (Schafgans and Zinde-Walsh (2002) show that Heckman's intercept estimator is consistent and asymptotically normal.) This lets us predict mortgage rates for the alternatives a household did not choose, correcting for selection. In the third step, we estimate the structural-choice model over mortgage products (15) using the household-level menu of predicted prices from the second step.

5.2 Choice Model Estimates

We first estimate the reduced-form multinomial choice model in equation (18) using the RFS data. Given that we have already rejected normally-distributed errors for the SCF data, we estimate the model here by multinomial logit (McFadden 1974) and explore alternative error distributions, including probit and semi-nonparametric methods, as a robustness check.

The sample consists of all borrowers aged 25 to 74 at origination for whom all covariates are available. As discussed, we identify $\beta_{\pi,FRM}$ from within-origination year variation in inflation experiences, and from variation in how these differences evolve over time. Multinomial logit coefficients represent the contribution of an attribute or sociodemographic characteristic to the utility of the respective alternative. We normalize $\beta_{\cdot,ARM} \equiv 0$ for all household-level variables, including exposure to past inflation. So, a positive coefficient indicates higher relative utility of, and probability of choosing, an FRM versus the baseline of an ARM. Unlike with the SCF data, we observe borrowers' regional locations, so are able to estimate coefficients on the region- and year-varying PMMS rates while including year fixed effects.

Table 5 presents the estimation results. In column 1, we restrict the coefficients on the FRM rate and the ARM initial rate from the PMMS to be the same (i.e., as in Table 4, households pay attention to the FRM−ARM rate spread). The negative coefficient estimate of $\hat{\tilde{\beta}}_R = -0.483$ indicates that individuals are less likely to choose the FRM when the spread between the FRM and ARM survey rates is higher. Turning to the variable of interest, we estimate a significant, positive coefficient of 0.220 for personal exposure to inflation π^e for the FRM alternative, relative to the baseline ARM alternative. The positive estimate implies that individuals who have lived through periods of high inflation derive greater utility from the FRM alternative, relative to the baseline ARM, than individuals with lower inflation experiences. For completeness, we also show the estimate for balloon mortgages (in the lower half of the table). The coefficient is negative, though less precisely estimated, suggesting that individuals with higher inflation experiences also substitute away from balloon mortgages and into FRMs.

To assess the economic magnitude of the estimated effect, we calculate the additional interest individuals would be willing to pay (WTP) for a fixed-rate mortgage if their lifetime inflation experience were 1 pp higher (see Section 4.2). The estimates in column 1 imply that individuals are willing to pay $-0.220/(-0.483) = 0.456$ pp in the FRM$-$ARM spread due to an additional percentage point of π^e.

In columns 2 to 5, we relax the restriction from column 1 and allow the price coefficients on the FRM and the ARM initial rates to differ. Our coefficient estimate remains very similar to column 1. Using the estimates in column 2, individuals are willing to pay $0.216/3.55 = 0.061$ percentage points more in the FRM rate for every additional percentage point of π^e. This WTP is smaller than in column 1 because we are dividing by a larger FRM rate coefficient, but it is more precisely estimated and statistically distinct from zero at the 5% level using delta-method standard errors.

In column 3, we additionally restrict $\beta_{\pi,Balloon} = \beta_{\pi,ARM}$; in column 4, we control for mortgage characteristics, including junior/senior status, whether it is the refinancing of a previous mortgage, non-conventional status, and discount points paid; and in column 5, we omit the balloon alternative altogether and estimate a binomial choice model between FRMs and ARMs. Under all specifications, personal experiences of higher inflation predict a significant increase in the choice probability of fixed-rate contracts. Since all specifications include origination-year fixed effects, this effect is above and beyond the full-information inflation expectations; rational individuals should place zero additional weight on personal experiences.

Figure 6. Actual and Counterfactual FRM Shares

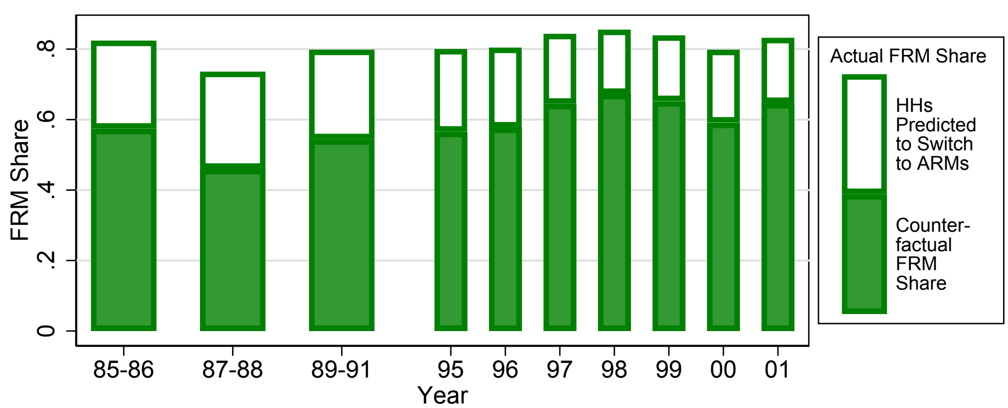

Notes. Data from the 1991 and 2001 RFS. The 1991 RFS reports origination years in two- or three-year intervals. The counterfactual FRM share is based on estimates from Table 2 column 4, with coefficient on experienced inflation set to zero.

To visualize the economic impact of the experience effect on aggregate mortgage choice behavior, Figure 6 shows the fraction of households predicted to switch to an ARM if they were not influenced by personal experiences and ignored π^e. We estimate counterfactual probabilities using the estimates from Table 5, column 4, that include the full battery of mortgage characteristics as controls, except that we replace the estimated coefficient $\hat{\beta}_{\pi,FRM}$ with 0. We then aggregate these probabilities to calculate hypothetical product shares for each origination year. The predicted mortgage shares add the year fixed effect coefficients of the estimation model back in, so adjust for the average level of inflation experiences in each origination year. Since the fixed effects capture all aspects of the economic environment and all information common to all households at the time (including the full history of past inflation), it is sensible to compare the product shares with and without the marginal effect of personal inflation history. The latter captures the prediction of the standard-model mortgage choice determinants, i. e., if $\alpha_1 = 0$ in (2).

As indicated by the shaded parts of the columns relative to their full heights, we predict that the FRM share would have been 25 pp lower in 1985–86, 57% rather than 82%. The effect of personal exposure to inflation diminishes as memory of the Great Inflation recedes, but it does not vanish. By 2001, the counterfactual FRM share is only 18 pp lower than the actual share, 65% rather than 83%. This indicates a sizeable, long-lasting influence of personal experiences on borrower behavior, the double effect previously alluded to in Section 3.1.

In the second step of our three-step procedure, we impute the interest rates of the non-chosen alternatives. Since the balloon alternative occupies such a small market share, we restrict the analysis to FRM and ARM alternatives from here forward.

Table 6 shows censored LAD estimates of the pricing equation (17) for $i \in \{FRM,\ ARM\}$. We use all of the exogenous explanatory variables from Table 5 in the first-step selection model in (18), except for the origination-year fixed effects, which we will include in the final estimation in the third step. Since the first-step choice probabilities are themselves estimates (rather than the true values), we account for the additional uncertainty by bootstrapping the system of equations from steps 1 and 2 and reporting the bootstrapped standard errors.[24]

We show the estimation results both without selection correction (columns 1, 3, 5) and with semiparametric selection correction (SPSC) using Newey (2009)'s series estimator (columns 2, 4, 6). We choose the order K of the approximating power series to the selection-bias term by leave-one-out cross-validation. That is, we run the two-step estimation of equation (see equation (A.21) in Appendix G) for $1 \leq K \leq 4$, on all

[24] Angelis et al. (1993) and Hahn (1995) show that the bootstrap consistently approximates the distribution of LAD-type estimators.

possible leave-one-out subsamples, for both the FRM and ARM rate equations.[25] The mean absolute prediction error is minimized at $K = 4$ for both rates.

Starting with the FRM rate equations in columns 1 and 2, we see that many coefficient estimates are affected by the inclusion of the Newey series correction terms. The biggest difference is in the coefficient on non-conventional status. Nonconventional mortgages carry FHA or VA insurance or guarantees to provide eligible higher-risk households with affordable mortgages, and these borrowers tend to choose FRMs rather than ARMs. Before we correct for sample selection, the coefficient on the non-conventional mortgage dummy is 0.2 basis points (column 1); after correcting for selection, it is −114 basis points (column 2). Intuitively, selection produces positive bias on our estimate of the rate subsidy offered to non-conventional borrowers. Selection correction also has noticeable effects on the coefficients for the PMMS survey rate, joint owners (i.e. marital status), rural county location, loan size/CLL, and jumbo status.

To formally test for the presence of selection bias, we implement a Hausman (1978)-style test suggested by, among others, Donald (1990, ch. 4]) and Martins (2001). The test statistic is a quadratic form of the difference in the coefficients between the two models, excluding the intercept, about the inverse of the covariance matrix of the difference. We bootstrap the distribution of $\hat{\Gamma}_{SC} - \hat{\Gamma}_{noSC}$, since Hausman's simplified variance-covariance matrix is not necessarily applicable. The resulting test statistic is asymptotically chi-squared with degrees of freedom equal to the number of parameters being tested. The sample test statistic reported at the bottom of Table 6 is more extreme than the 5% critical value of 19.7, providing strong evidence in favor of selectivity bias in the FRM pricing equation. We also estimate the selection function from equation (A.20) (in the Online Appendix) as $\hat{g}(\hat{\hat{\eta}}_{n,i,j}) = \widehat{\mathbb{E}}[Rate_{n,i} \mid Z_{n,i}, D_n = 1] - \widehat{\mathbb{E}}[Rate_{n,i} \mid Z_{n,i}]$ and report its mean value within each selected subsample in the bottom row of the table. The result suggests that individuals who selected into the FRM alternative were offered unusually low interest rates given their observable characteristics $Z_{n,i}$.

We repeat this exercise with the ARM initial rate in columns 3 and 4, and with the ARM margin in columns 5 and 6. In the ARM initial-rate pricing equations, the selection bias is weaker. Directionally, inclusion of the selection control function affects the ARM pricing coefficients in a similar manner as the FRM pricing coefficients, but the changes are smaller and the Hausman-style test fails to reject no selection bias

[25] The results of Newey (2009) imply that consistency of the SPSC estimator on a sample of size N requires that the order of the approximating power series be $K = o(N^{1/7})$, which suggests an upper bound of 4 for our sample size.

($p = 0.78$). Somewhat surprisingly, the mean value of $\hat{g}_n$ is positive for those choosing the ARM alternative, although it is smaller in magnitude than for the FRM subsample.

Turning to the ARM-margin estimation, we switch to an ordered-logit estimator (OLOGIT). In unreported specifications, we found that all CLAD estimates other than the junior mortgage dummy are precisely estimated zeros, and the junior mortgage dummy carries the same significant coefficient of +25 bps both without and with the selection correction. That is, CLAD fails to adjust ARM margins for household risk characteristics, possibly because more than half of all ARMs in our sample carry the same margin, 2.75 pp. As an alternative, we discretize the distribution of margins into ten intervals using the 1991 RFS reporting categories[26] and estimate an ordered logit model. This model implicitly accounts for censoring and predicts households' choice probabilities for each interval. We multiply the probabilities by the 2001 RFS medians for each interval to calculate an expected, risk-adjusted margin for each household. Columns 5 and 6 report the marginal effects of each covariate on the expected value of the margin, $\partial \mathbb{E}[Y|X = x]/\partial x$, averaged over all observations, i. e., after calculating $\mathbb{E}[Y|x] = \sum_j \Pr(y \text{ in category } j|x) \times \text{Median}(\text{category } j)$ from the 2001 RFS.

We estimate a slightly inverse relationship between the PMMS initial ARM rate and households' expected margins, suggesting that lenders backload interest when teaser rates are low. The average junior mortgage carries a 30 bp premium over first mortgages (10 bp after correcting for selection effects). Finally, the ordered-logit estimates reveal a big effect of non-conventional status on ARM margins. Most other covariates have small and insignificant marginal effects, and we again fail to reject the null of no selection.

With these estimation results in hand we turn to the structural choice model. Table 7 presents the estimates of (15), where the dependent variable indicates that the household chose an FRM. We use predicted interest rates from the pricing equations (Table 6) for both the chosen and the non-chosen alternative. We adjust standard errors for the first- and second-step estimation by bootstrapping the entire three-step procedure.

A comparison of columns 1, 3, and 5 with columns 2, 4, and 6 reveals the importance of selection correction in the second-stage estimation of (17). Without selection correction, the price coefficients are insignificant and often have the wrong sign. With selection correction, the signs indicate the expected downward-sloping demand.

Columns 1 and 2 include only the FRM and the initial ARM teaser rate predictions

[26]The ten categories are [0, 100), 100, (100, 200), 200, ..., (400, 500), [500, ∞).

from step 2. Columns 3 and 4 add the risk-adjusted ARM margins to the estimation. With the selection correction, the estimated coefficients on the FRM and ARM initial rates are very similar to column 2, while the coefficient on the ARM margin becomes small and insignificant. This suggests that households pay more attention to the up-front costs, and relatively little attention to possible future ARM resets, when deciding between alternatives. To check the robustness of this (auxiliary) finding, we consider an alternate specification. Since the selection correction procedure had the strongest effect on the coefficient estimate of the non-conventional status dummy in the pricing equations in Table 6, we explore whether non-conventional status has an additional, direct effect on mortgage choice. We test this by including non-conventional status as an additional explanatory variable in columns 5 and 6. This generates "correct," negative demand elasticities both with and without the selection correction, indicating that future ARM resets play an important role in households' mortgage contract decisions.

Turning to the variable of interest, we find that borrowers with personal histories of higher inflation are more likely to choose an FRM, independent of how we predict mortgage prices and of the set of controls. We estimate 12-28 additional bp WTP per additional pp of inflation experiences in the structural model, compared to 5-8 bp in the reduced-form model.

Robustness Checks. We employ a battery of alternative estimation approaches and robustness checks to probe our estimation results. These include using alternative data, restricting the data to consumers who are least likely to face supply-side constraints, applying specification tests, and using alternative estimation procedures.

First, we have already estimated the relationship between past exposure to inflation and mortgage choice with the SCF in Section 4. The SCF is conducted at a higher frequency, so we are able to approximate a flow dataset over time while restricting the analysis to only new mortgages (originated in the survey years 1989, 1992, ..., 2013). This should alleviate concerns about selection issues arising from mortgages that are prepaid quickly because the homeowner sells and moves. The replication across two such different datasets provides strong supporting evidence for our hypothesis.

Second, we turn to supply-side constraints. Our baseline analysis assumes that all borrowers have a choice between FRM and ARM contracts. However, some borrowers might have to go for an adjustable-rate contract in order to qualify for a loan due to constraints on the ratio of debt service over income. Conversely, others might not be offered an ARM due to income risk. Appendix H shows that our results are even stronger for borrowers with low loan-to-income ratios, who most likely had "free choice" between FRM and ARM, suggesting that supply-side constraints do not drive our results.

Next, we consider the robustness of our results to different estimation methods. We saw in Section 4 that a parametric probit model is probably not appropriate, at least for the SCF data. To test the validity of the logit choice model for the RFS data, we use the specification test of Horowitz and Härdle (1994). This test compares a parametric regression model to a semiparametric alternative that maintains the same single-index restriction, $E[y|x] = G(x'\beta)$, but allows the link function $G(\cdot)$ to take an unknown form. We describe the implementation of this test in Appendix I.1. The result leads us to reject the logit model, as well. However, visual inspection of the nonparametric estimate of the CDF suggests that deviations from logit are small. To be sure that our results do not depend on a possibly misspecified error distribution, we re-estimate the reduced-form choice model using Gallant and Nychka (1987)'s semi-nonparametric (SNP) estimator, extended to the binary-choice setting by Gabler et al. (1993). The SNP coefficient estimates are very similar to their parametric counterparts after scale normalization; see Appendix I.2. In particular, we estimate a WTP of 5.0 bp for the FRM for every additional percentage point of lifetime inflation experiences by SNP, versus 5.2 bp by logit.

In Appendix I.3, we move in the opposite direction and estimate the three-step model using fully parametric, maximum likelihood methods. We specify the error terms in steps 1 and 2 as multivariate normal. Given the results just discussed, this should be viewed as a simplifying approximation, and the ensuing estimates as *quasi*-maximum likelihood (White 1982). The normality assumption justifies using a Heckit two-step model. To account for the censored dependent variables, we estimate the second-step rate equations by Tobit rather than CLAD, again relying on the normal error distribution assumption. Correcting for selection by fully-parametric methods moves the rate equation coefficients in the same directions as our preferred semiparametric estimator. We find weak statistical evidence of selection bias in both rate equations, again in the wrong direction in the ARM equation. Perhaps reflecting this, the choice of whether or not to use selection-corrected interest rates in the third step is less important for the parametric estimator (both sets of estimates have the correct signs) but increases the precision in step 3. We estimate a 30 bp increase in WTP per pp of inflation experiences, on the high end of our previous estimates.

Finally, we test whether, as an alternative measure of lifetime experiences, we can relate interest rate experiences to mortgage choice behavior. Since Fisher (1930), many macroeconomic models assume that long-run variation in nominal rates is driven by variation in expected inflation ($i = r + \mathbb{E}\pi$).[27] Given this, whether individuals learn from inflation experiences or from nominal interest rate experiences over the course of their lifetimes is not theoretically distinct. Nor do we expect to have much power empirically to distinguish between the mechanisms, since the main source of variation for both is the Great Inflation period (cf. Figure 2).[28] So, rather than running a horse race between the two, we investigate whether this alternative specification generates similar results. In Appendix J, we re-estimate our reduced-form mortgage choice model, replacing π_n^e with i_n^e. As before, we weight historical interest rates using weights that linearly decline to zero in the year that the decision-maker was born. We employ short-term (90-day) T-bill rates as well as long-term (10-year) Treasury rates. Lifetime inflation experiences are highly correlated with both sets of interest rate experiences, $\rho = 0.81$ and 0.69, respectively. As expected, the results are very similar. This finding builds on our evidence in Section 4 that individuals coming of age during the Great Inflation expected higher nominal interest rates than members of earlier or later cohorts, and that this personal history significantly affected their valuation of fixed- versus variable-rate debt contracts for years to come.

Supply-side Implications. Our baseline analysis is of consumer choice and so focuses on demand for fixed- versus variable-rate mortgages. In particular, our identification strategy relies on within-year differences in borrowers' personal exposure to higher inflation. However, in time periods when the memory of the Great Inflation is strongest among home buying cohorts, we would expect to observe increases in equilibrium quantities and prices. Indeed, as Figure 6 illustrates, we calculate that personal exposure to inflation particularly raises the FRM share in the late 1980s, consistent with this theory. We can go one step further and explore the price implications using the two time series of mortgage rates discussed earlier, PMMS and MIRS. In Appendix E, we show that the national average FRM-ARM initial rate spread is indeed higher in years when the average borrower has lived through periods of higher inflation, increasing around half a percentage point for every additional percentage point of inflation experiences (cf. Table A.7). However, such shifts in demand do not necessarily create

[27] The literature testing for a Fisher effect is voluminous; see, e.g., Mishkin (1992), Evans and Lewis (1995), Crowder and Hoffman (1996), King and Watson (1997), and Müller and Watson (2018).

[28] Clarida et al. (2000) find a breakpoint in monetary policy in 1979: the pre-Volcker Fed was "accommodative," allowing nominal rates to rise, but less than one-for-one with expected inflation; whereas post-1979 the Fed became "proactive" and raised nominal rates more than one-for-one.

economic profits for lenders, since they must compete to raise costly funds to finance an increase in supply. Secondary-market investors who are a major source of financing may not view FRMs and ARMs as perfect substitutes, and their pricing will respond to the differing real cash-flow streams and differing probabilities of prepayment and default of these mortgage types.

In the next section we discuss the welfare implication for switching households—i.e., borrowers who are close enough to indifference between the two alternatives that we can attribute their FRM choice to the long-lasting effects of their past exposure to high inflation. While we will show that choosing an FRM was costly for these households, both *ex post* and *ex ante*, this does not imply that the FRM is a worse deal for all households, nor that lenders are earning positive economic profits in aggregate.

6 Financial Costs and Welfare Implications

Our evidence on mortgage choices is consistent with personal experiences affecting an individual's willingness to pay for the fixed-rate alternative, and the effect on mortgage product shares is economically large. A separate question is how costly these effects are for consumers. Whether exposure to high inflation induce a welfare loss *ex post* depends on realized interest rates; whether they induce a welfare loss *ex ante* depends on the full distribution of possible interest rates that could have occurred.

In this section, we provide estimates of the financial costs of exposure to periods of high inflation on residential mortgage choice, over varying horizons and under varying assumptions about repayment, mobility, and historical as well as simulated interest rates.

6.1 Measurement: Welfare-Relevant Treatment Effect

To assess financial costs, we need to (1) identify whose choice is affected, and (2) calculate whether their exposure-induced choice was costly or beneficial.

As for the first step, some households would have chosen the same mortgage product regardless of whether they overweighted or ignored inflation experiences. The relevant subset are the "switchers:" households who chose an FRM only because personal exposure to inflation figured into their choice function and who would not have chosen the FRM under a full-information Bayesian forecast of future nominal interest rates.

To identify the subset of the population who are affected by their inflation experiences, we define each household's *switching probability* as

$$h_n = \Pr(D_n = 1|\beta_\pi = \beta_\pi) - \Pr(D_n = 1|\beta_\pi = 0)\,, \tag{19}$$

where D_n is an indicator for choosing the FRM. We obtain an estimator of h_n by comparing choice probabilities with the coefficient on inflation experiences in the choice model set to its "true," estimated value in Table 5 or 7 versus zero, leaving all other estimated coefficients the same. For example, if a household's true probability of choosing an FRM is 90% and the counterfactual probability (ignoring experienced inflation) is 70%, then for every 100 observationally-equivalent households, we expect 70 of them to choose an FRM no matter what, 10 to choose an ARM no matter what, and 20 to switch from the FRM to the ARM.

As for the second step, there are periods when locking in a low nominal fixed-rate was advantageous *ex post*. The historical PMMS data show that the FRM-ARM initial rate spread is always positive, so individuals with a sufficiently short time horizon will usually benefit from the ARM's low teaser rate, but over longer time horizons the resets could make the ARM more expensive. For example, an individual taking out an FRM in 1993 would lock in a nominal rate of 7.31% for the life of the loan. An individual taking out a 1/1 ARM with no reset caps and a 2.75 margin over the one-year Treasury rate would pay only 4.58% in 1993, but this would reset to 8.06% in 1994, 8.70% in 1995, etc. Resets would keep the subsequent ARM rate above the 1993 FRM rate every year until 2001.

To establish the counterfactual (hypothetical) mortgage payments, we use our pricing estimates in Table 6 and simulate the monthly payments each household would make under an FRM and an ARM. For ease of comparison, all mortgages carry a 30-year term, are self-amortizing, paid on time (no late penalties or prepayments), and originated on January 1.

We consider three interest-rate scenarios. Each makes progressively greater adjustments for risk characteristics, at the cost of increasing sensitivity to modeling assumptions. In Scenario 1, we assign everyone the Freddie Mac PMMS mortgage rate, varying only by region. This sidesteps the issue of estimating individual-level pricing equations, but may over- or understate the financial costs by not correcting for household risk characteristics. In Scenario 2, we use the selection-corrected CLAD estimation to predict risk-adjusted FRM rates and ARM teasers (Table 6, columns 2 and 4), while ARM margins are adjusted for seniority only. In Scenario 3, we use ordered logit to predict ARM margins based upon household-level characteristics (Table 6, column 6).

Under all three scenarios, individuals choosing an ARM receive the teaser rate for one year, after which annual resets are based on the appropriate margin over the average value of a 1-year constant maturity Treasury for that year: plus 2.75 percentage points

(Scenario 1), plus 2.75 if first-lien and 3.00 if second- or third-lien (Scenario 2), or plus a risk-adjusted margin from the selection-corrected ordered logit estimation results (Scenario 3). We do not roll the ARM over into the new PMMS initial ARM rate after year 1, since these discounted rates are available only to new mortgage borrowers and doing so would bias in favor of the ARM being less expensive. Our Scenario 1 reset margin of 2.75 p.p.s from the PMMS is very close to the average ARM margin in the RFS, confirming its validity (see Table 2).

For each scenario, we simulate the full path of future interest payments that a household would make under both mortgage types. Letting $Y_{n,1}$ be interest payments under the FRM and $Y_{n,0}$ under the ARM alternative, $\Delta Y_n \equiv Y_{n,1} - Y_{n,0}$ is the *ex post* financial cost of choosing the FRM (if positive) or benefit (if negative) for household n.

Our summary measure, the *Welfare-Relevant Treatment Effect (WRTE)*, is the weighted sum of the simulated ΔY_n across all households, using their estimated switching probabilities as weights:

$$\widehat{WRTE} \coloneqq \sum_{n=1}^{N} \Delta\hat{y}_n \left(\frac{\hat{h}_n}{\sum_n \hat{h}_n} \right). \tag{20}$$

We show in Appendix G.2 that the WRTE is equivalent to the expected difference between FRM and ARM payments for households that chose an FRM because of their inflation experiences.[29] We can now calculate the cost of experience-induced FRM choices.

6.2 Costs over Different Holding Periods

We begin by calculating the WRTE as of the RFS survey years (1991 and 2001). Since we know that a mortgage exists as of the RFS survey year—the household has not defaulted or moved—we can provide a lower bound on the true WRTE with very few modeling assumptions. In this spirit, we run this simulation under Scenario 1, with PMMS rates and switching probabilities from the reduced-form choice model. On average, borrowers in the 1991 RFS had already paid \$4,700 in cumulative extra interest as of year-end 1991, and borrowers in the 2001 RFS had already paid \$1,700 cumulative extra as of year-end 2001, due to experienced inflation. Moreover, for all but one origination year (1998, when FRM rates were unusually low), overweighting personal exposure to inflation and taking out an FRM proved to be *ex post* costly.

Turning to longer holding periods, we need to make a few additional assumptions

[29] We choose the name "WRTE" in reference to Heckman and Vytlacil (2007), who formulate a "policy-relevant treatment effect" (PRTE) using the same weighted average.

regarding the refinancing behavior. Most mortgages in the U.S. allow refinancing without paying a penalty. To accurately gauge the *ex post* financial cost of holding a fixed- versus adjustable-rate mortgage over longer time periods, we consider households' likely refinancing behavior.

Refinancing Scenarios. We consider three sets of assumptions about refinancing. First, we assume that households hold the original fixed-rate mortgage until maturity, as if the contracts prohibited prepayment. This is a worst-case scenario for an FRM in a dis-inflationary environment, and provides an upper bound to our cost estimates.

Second, we assume that households refinance whenever the difference between the old and the new interest exceeds a threshold that accounts for the fixed cost of refinancing and the option value of waiting. Such optimal refinancing is a best-case scenario for fixed-rate mortgagors. Agarwal, Driscoll, and Laibson (2013, hereafter ADL) provide a closed-form solution for this threshold. We simulate the new interest rate a household would be offered using the estimates in Table 6 and updated PMMS rates for each year, then plug the differential into ADL's square-root rule approximation to the optimal threshold.

Third, we calculate costs based on "expected refinancing," which provides an intermediate case between the two extremes of no refinancing and optimal refinancing. An extensive literature documents that mortgagors do not exercise this real option optimally.[30] They sometimes refinance too early, before the rate differential has crossed the optimal threshold, or too late, waiting months or years after the differential has crossed the threshold. To calculate a household's expected mortgage payments, we use estimates from Andersen et al. (2015) that describe the probability a household will refinance every period as a function of the interest rate differential. Iterating these refinancing probabilities forward starting in year 2 of the mortgage gives us a set of probabilities describing, t years after origination, the probability that the household holds a mortgage last (re-)financed s years after origination, $0 \leq s \leq t$. We use these probabilities to calculate the household's expected FRM payments across the entire distribution of possible time-t interest rates. See Appendix G.3 for further details.

Simulation Results. Turning to the full simulation results, we show the cost estimates for all switching households in Table 8. All calculations are presented for holding periods up to 15 years (i. e., up to year 2016 for mortgages originated in 2001). Positive numbers indicate financial costs from choosing the FRM.

[30] Cf. Green and Shoven (1986), Stanton (1995), Green and LaCour-Little (1999), Bennett et al. (2000), Agarwal et al. (2015), Andersen et al. (2015), Bajo and Barbi (2015), and Keys et al. (2016).

In the top panel, we display the simulation results under Scenario 1 (using unadjusted PMMS rates) for the three refinancing assumptions. The first row shows how costly it would be to continue holding the mortgage beyond the survey year if switching households never refinanced. We see that the WRTE doubles over five years, from \$2,400 to \$5,500 per household. After 15 years, the WRTE exceeds \$17,000 per household in after-tax, present value terms. Allowing households to refinance ameliorates this cost, to approximately \$10,000 per household under "Expected Refi," and \$8,000 under "Optimal Refi."

The middle and bottom panels report Scenarios 2 and 3, in which we adjust the FRM rate, the initial ARM rate, and (in Scenario 3) the ARM margin for risk characteristics. Scenarios 2 and 3 provide similar and significantly larger estimates at every holding period; e.g., after 15 years, from \$18,000 if households refinance optimally, to \$27,000 if they never refinance.

To generate bottom-line numbers for all three scenarios, we calculate each household's expected tenure as a function of age. We obtain five-year non-mover rates from the Current Population Survey Annual Social and Economic Supplement (CPS ASEC) for 2000–05 and 2005–10 in the general U.S. population that is at least 20 years old. We convert these staying probabilities into one-year moving probabilities and fit them to a fourth-order polynomial function of age. This generates moving probabilities that slope downward in age. For example, we estimate that a 25-year old household has a 17.4% probability of moving in the next year. This declines to 13.1% by age 30 and 5.1% by age 50. See Appendix K for more details.

We assume that moving events are exogenous and unanticipated by the household, arriving according to the empirical distribution we have just estimated. Upon moving, the household sells the house and the stream of mortgage payments stops. Using these probabilities, we re-calculate the present discounted value of the difference between FRM and ARM interest payments, weighting each difference by the probability that the household has not yet moved. These results are reported in the final column of Table 8, labeled "$\mathbb{E}$[tenure|age]". The order of magnitude resembles our estimates for a 10-year holding period even though we now put positive probability on the entire holding period (through the end of our data). We estimate a bottom-line cost based on expected refinancing of \$8,000 under Scenario 1 and \$15,000 under Scenario 3. To put these numbers in perspectives, our *ex ante* WTP estimates imply an expected 30-

year cost of \$3300-\$7600 in PDV terms.[31] This underscores that, for most switching households, taking out an FRM was likely a very costly mistake *ex post*.

Robustness: Discount Points. Our baseline methodology to estimate expected tenure in the house is completely nonparametric and relies only on the borrower's age. Alternatively, a literature dating back to Dunn and Spatt (1988) suggests that borrowers reveal private information about their expected tenure in the house by purchasing discount points. Discount points allow borrowers to pay the lender upfront and purchase a lower future interest rate. Each discount point costs 1% of the amount borrowed, and reduces the mortgage interest rate by approximately a 25 basis point. Common investment advice is to purchase enough points such that, over the expected tenure in the house, the lower monthly payments just offset the upfront cost.[32] However, households might pay fewer points if they are risk averse or face liquidity constraints at the time of mortgage origination. Moreover, Agarwal et al. (2017) show that in practice borrowers do not pay points optimally, calling into question the rational interpretation of borrowers' empirically observed menu choices. In our data, only 16.5 percent of households pay discount points, with a median of 2 points paid.

Nevertheless, we check the robustness of our results to utilizing discount points for the estimation of geographic mobility. As detailed in Appendix K, we estimate each household's expected tenure in the house as the number of years until the household breaks even in present-value terms. We then fit these break-even horizons to two plausible parametric distributions of moving times: a negative exponential distribution, which assumes a constant hazard of moving, and a Weibull distribution, which allows the hazard of moving to decrease over time.

As anticipated, the resulting estimates of implied tenure are very low. Since most borrowers do not pay any discount points, the average of households' median tenure is 3.6 years under the negative exponential distribution, and 4.7 years using the Weibull distribution, versus 12.5 years based on household age. Hence, households do not appear make the purchase decision of a risk-neutral rational agent without liquidity constraints.

If we ignore these discrepancies and nevertheless assume risk-neutral optimal pur-

[31] Using an estimate of 6–14 bp per percentage point of lifetime inflation experiences for an average-sized mortgage, this amounts to \$60–\$140 per year, or \$700–\$1,600 over 30 years (discounting at 8%) per pp of lifetime inflation experiences, times 4.75 pp.

[32] Cf. https://www.investopedia.com/articles/pf/06/payingforpoints.asp or https://www.bankrate.com/finance/mortgages/mortgage-points.aspx. In theory, a risk-neutral household should purchase points until the expected tenure exactly equals the break-even time it will take to recover the upfront payment.

chase decisions without liquidity constraints, the WRTE estimates are still significant, albeit 40% to 45% lower: $9,106 under Scenario 3 interest rates, expected refinancing behavior, and negative exponential distribution, and $8,275 under the same scenario with a Weibull distribution.

If, instead, we acknowledge that households choose less than the optimal number of points for one of the reasons discussed above, then our estimates of occupancy time are too short—expected tenure will exceed the break-even horizon. We model some adjustments in Appendix K, which raises the average median time of occupancy to 6.4 years, and reduces the gap between the dollar costs estimated under the two methodologies. Now the cost estimate rises to $11,176 ($11,629), only 25% (20%) lower than our baseline estimates.

In principle, we could use other additional methodologies to back out moving probabilities, but the evidence in this section suggests that our results are robust to a wide array of assumptions.

6.3 Different Inflation Environments

An important limitation to our *ex post* estimates is that they rely on the actual realization of inflation and interest rates after each origination. This ignores the range of other possible inflation environments that might have occurred. To estimate the *ex ante* value of choosing an FRM versus an ARM, we re-simulate interest payments for switching households under other inflation environments.

Historical Environments of Rising versus Falling Inflation. The expected path of future inflation affects the slope of the nominal yield curve, and thus the FRM-ARM spread today. We first use prior historical inflation and term structure data to engage in a thought experiment: what would be the WRTE for the households in our sample had they originated their mortgages in a different historical inflation environment?

We choose two points in time that represent a rising versus a falling inflation environment: 1971, just as the Great Inflation took off; and 1981, the year that inflation began to subside (and FRM rates peaked). We assume that the households are completely identical in every respect, including their lifetime inflation experiences, except that they are facing a hypothetical FRM/ARM interest rate schedule of 1971 or 1981 (and subsequent years), if both contracts had been available.[33] We use Scenario

[33]The ARM was not available nationwide until 1982. Since the PMMS initial ARM rate series begins in 1984, we impute the survey rate for 1971 and 1981 by assuming that it would have taken its

3 estimates to simulate each household's interest payments over the lifetime of both mortgage alternatives, estimating the probability that the homeowner sells the house and moves based on head of household age. Our goal is to isolate the effect of different inflation realizations after the mortgage is originated, so we continue to use the same switching probabilities as weights in calculating the WRTE. That is, for the purposes of this thought experiment, the only component of equation (20) that we change is the simulated interest payments $\Delta\hat{y}_n$.

In a rising inflation environment such as the one that followed 1971, the WRTE is *negative*, indicating that households who choose an FRM instead of an ARM due to their inflation experiences end up paying *less*. The average switching household is better off by \$8,423 under optimal refinancing behavior, compared to \$7,406 under expected refinancing behavior and \$8,833 if they never refinance. This economic environment represents a best-case scenario for choosing an FRM. Due to rising inflation over the 1970s, it is never optimal for any of the households in our sample to refinance during the first twenty years of the mortgage's life.

By contrast, in a falling inflation environment such as the one that followed 1981, choosing an FRM can be extremely costly—even if a household refinances close to optimally. We estimate that the average switching household would pay \$18,346 more over its expected lifetime in the house, given optimal refinancing behavior, compared to \$20,304 if they refinance as expected and \$44,463 if they never refinance.

This exercise illustrates that, historically, there are plausible scenarios when the choice of an FRM paid off, even though the embedded inflation insurance was rarely in the money during the Great Moderation of the 1990s and 2000s. Hypothetical best-case payoffs are on the order of 50–60% of our empirical cost estimates, whereas the hypothetical worst-case loss is about one-third larger (130%) than our estimates.

Simulated Inflation Environments. To take the *ex ante* analysis a step further, we turn to a wider range of possibilities and run a Monte Carlo simulation of different possible inflation environments. Our setup follows Campbell and Cocco (2003). The simulation environment simplifies certain aspects of the real world, but it is rich enough to capture the key dynamics of how expected and realized inflation impact mortgage cost. Each replication has two independent sources of variation: a 30-year sequence of inflation rates and of short-term real interest rates. All other variables are derived by exact, linear relationships.

average value over the 1-year constant-maturity Treasury rate of 1.5 percentage points.

First, as in Section 4, we assume that inflation follows an AR(1) process, $\pi_t = \mu + \phi(\pi_{t-1} - \mu) + \epsilon_{\pi,t}$, with serially-independent innovations $\epsilon_{\pi,t} \sim \mathcal{N}(0, (1-\phi^2)\sigma_\pi^2)$. One-year log real interest rates are serially uncorrelated: $r_t = \rho + \epsilon_{r,t}$, where $\epsilon_{r,t} \sim$ indep. $\mathcal{N}(0, \sigma_r^2)$ that are mutually-independent to the inflation innovations: $\epsilon_{r,\cdot} \perp \epsilon_{\pi,\cdot}$. Short-term nominal (log) interest rates equal the real interest rate plus actual inflation: $y_t^1 = r_t + \pi_t$. Long-term nominal rates follow the expectations hypothesis with a term premium: $y_t^T = \frac{1}{T}\sum_{s=1}^{T} \mathbb{E}_t y_{t+s-1}^1 + \theta^T$, where $\mathbb{E}_t y_{t+s}^1 = \rho + \phi^s(\pi_t - \mu) + \mu$. We acknowledge that the real-world time-series dynamics of inflation may be more complicated than an AR(1) process, that short-term real interest rates in the U.S. may exhibit some serial correlation, and the rich literature modeling the term premium would not exist if it were simply a constant. However, this simplified environment is sufficient to generate realistic autocorrelations in short-term and long-term nominal interest rates due to the mean reversion of inflation, which will carry over to the path of mortgage rates.

Second, we assume that ARM rates equal the one-year nominal bond rate plus a term premium: the initial ARM rate (in year 1) is $y_1^A = y_t^1 + \theta^{A,1}$; and the ARM reset rate (years 2–30) is $y_t^A = y_t^1 + \theta^A$. The FRM rate (all years) equals the ten-year nominal bond rate plus a term premium: $y_t^F = y_t^{10} + \theta^F$. Hence, as we document in Appendix E, the FRM-ARM spread closely tracks the nominal bond yield spread. Because the expectations hypothesis holds, today's FRM rate depends on expected short-term nominal bond rates over the next ten years, so except for differences due to bond premia, households cannot time the market and profit from choosing one mortgage instrument over another.

We calculate the simulation parameters from average U.S. values over the longest subset of 1960–2013 available, to capture the average economic environment before, during, and after the Great Inflation. Table 9 gives the values and sources for all the simulation parameters. Of particular importance are the premia. We set the term premium as the average constant maturity U.S. Treasury ten-year minus one-year spread, 1%. Using PMMS data, we calculate an average FRM markup of $\theta^F = 1.7\%$ over the ten-year nominal bond rate, an average initial ARM markup of $\theta^{A,1} = 1.5\%$ over the one-year nominal bond rate, and a subsequent ARM markup of $\theta^A = 2.75\%$. By contrast, Campbell and Cocco (2003) use the same term premium, a ten-basis-point higher FRM premium of 1.8%, and a constant ARM markup of 1.7%, so in expectation their ARM is 110 basis points cheaper over every ten year period.[34] Our assumptions are less

[34] Campbell and Cocco (2003) observe that their ARM premium "may be biased downward" due to teaser rates (p. 1466).

favorable to the ARM: we specify that the ARM is five basis points more expensive in expectation over every ten-year period starting after year 1, $\theta^A - (\theta^F - \theta^{10}) = 0.05\%$. This is offset by the initial teaser ARM rate a household pays in year 1 and the cost of refinancing an FRM should rates fall, realistically capturing that younger and more mobile households should prefer the ARM.

In each replication, we draw 30 years of inflation and nominal mortgage rates, which we treat as the national baseline rates. We plug these simulated rates in place of the PMMS rates into our selection-corrected, Step 2 FRM and ARM rate equation estimates to obtain a set of counterfactual mortgage rates for each household in each replication (Scenario 3). We then simulate the full set of mortgage payments each household would face in that replication, under various assumptions concerning refinancing behavior and expected tenure in the house. We calculate the WRTE by (20) under the same three, "no"/"expected"/ "optimal" refinancing assumptions. We can again make different assumptions concerning mobility, including age and discount points paid. For space we focus on our preferred age-based mobility metric. As in the historical exercise, we use the same household switching probabilities $\hat{h}_n$ that we estimated using the RFS data, in the actual origination year, changing only $\Delta\hat{y}_n$, the present value of the excess payments a household would face by choosing an FRM over an ARM in that replication. We repeat the process 100 times to obtain a set of 100 different inflation environments, mortgage rate paths, and WRTE realizations.

Table 10 summarizes the simulation results. The results indicate that choosing an FRM due to personal exposure to high past inflation is costly in expectation. As in Table 8, the expected refinancing case is intermediate between the unrealistic extremes of no refinancing and optimal refinancing: the expected WRTE is about \$8000 under no refinancing, \$6500 under expected refinancing behavior, and \$5100 under optimal refinancing behavior. The WRTE is positive, indicating that the FRM is more costly, in over 75% of replications regardless of refinancing behavior. This is despite the simulated FRM-ARM rate spread being nearly symmetric around zero: the baseline FRM rate path is only 5 basis points higher than the baseline ARM rate path on average (over all 30 years, including the initial teaser rate in year 1 and the rate resets in years 2–30). Even in realizations that are extremely favorable for the FRM, with high average inflation rates, it appears that the expected savings from choosing an FRM are minimal: in only 10% of cases do they exceed \$4700 given expected refinancing.[35]

[35]Note that the S.D. of average inflation involves a long-run variance, because the within-replication inflation data are dependent, and so exceeds $\sigma_\pi/\sqrt{30}$. Average inflation is never negative, but we

Figure 7. Average Inflation and E[WRTE] in Monte Carlo Simulation

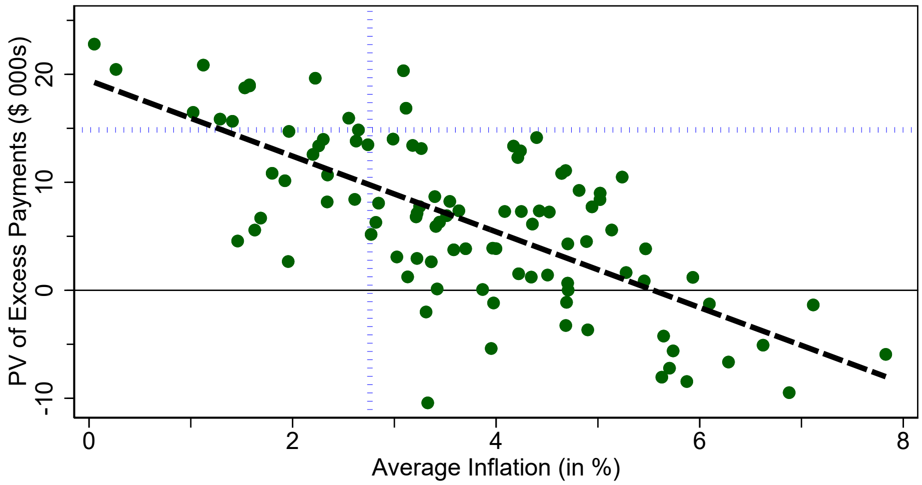

Notes. The horizontal axis is average 30-year inflation and the vertical axis is the WRTE in a replication. Each point is a replication, and thick dashed line is OLS regression line across all 100 replications. Horizontal and vertical crossing lines indicate average values of inflation over 1986–2013 and WRTE from the RFS. Calculations based on Scenario 3 estimates, expected refinancing behavior, and age-based mobility.

To explore how the inflation environment affects the realized cost of choosing an FRM further, Figure 7 plots the simulated WRTEs against average inflation in each of the 100 replications, again using our preferred age-based mobility metric and expected refinancing behavior. The horizontal and vertical crossing lines indicate the average value of each variable from the RFS data for reference: average inflation of 2.8% over 1986–2013 and the Scenario 3 WRTE of just under $15,000 given expected refinancing and age-based mobility, from Table 8. As expected, we observe a strong inverse correlation between realized inflation and the *ex post* cost of a fixed-rate mortgage. Every additional percentage point of average inflation over the 30-year simulation reduces the ex-post cost of the FRM by $3,573 (s.e. 311), controlling for initial interest rate conditions. In replications with average inflation exceeding about 5.5%, the expected WRTE becomes negative, indicating that the FRM is *ex post* cheaper. Overall, the simulation results indicate that the embedded inflation-insurance of an FRM is costly on average, it rarely pays out, and when it does the payout is small.

Finally, the simulation suggests that FRMs were not unusually expensive given actual subsequent economic conditions in the 1990s and 2000s. We assume that inflation reverts to a long-run mean of 3.8%, based on an historical average that includes the

observe deflation in 58% of replications and just over 10% of simulated years, 311 out of 3,000.

Great Inflation. However, U.S. inflation averaged only 2.8% over 1986–2013. In replications with similarly low values of average inflation, we predict that the expected WRTE is $9,718 (s.e. 632). The bottom line is that while the low-inflation experience of the 1990s was disadvantageous to FRM holders, choosing an FRM is predicted to be expensive even in "average" time periods, particularly for those who are making their decisions due to their past personal exposure to high inflation.

7 Discussion: The Long-Lasting Effects of the Great Inflation

The cost estimates in this paper leave us with a striking conclusion about the long-run consequences of the Great Inflation, both in terms of the composition of asset markets and in terms of welfare implications. Suppose, as shown in Figure 3, that the Great Inflation had not occurred. Our structural choice model can be used to determine what share of FRM choices are attributable to this experience: if there had not been a Great Inflation, the FRM share would have been 5.5 percentage points lower across all the households in our sample. Our model estimates also specify that this effect was concentrated among younger households taking out mortgages in the late 1980s—essentially, the Baby Boom generation, many of whom were entering the housing market and buying their first homes at this time. According to our structural model estimates, these individuals would have take out 1 million fewer FRMs if not for the Great Inflation, lowering their FRM share by 8.1 percentage points (Table 11). A decade later, differences between the inflation experiences of Boomers and earlier generations recede, but these older generations continue to overweight the 1970s vis-a-vis younger Gen Xers. We estimate that the memory of the Great Inflation raises the FRM share among Baby Boomers' mortgage originations in the late 1990s by 3.6 percentage points, or half a million additional FRMs. In other words, the long shadow of the Great Inflation has significantly altered the composition of one of the largest asset markets in the U.S., and we can pinpoint the cohorts that are particularly affected.

These decisions are costly. Based on the aggregate of our interest rate estimates, using expected refinancing behavior and mobility, Baby Boomers likely ended up overpaying over $14bn on their FRMs in the late 1980s, and almost $9bn in the late 1990s (under risk-adjusted, Scenario 3 interest-rate predictions). Even under Scenario 1, i. e., assigning each borrower the average PMMS mortgage rate rather than risk-adjusting, the dollar figures are still substantial, about half as large. These calculations underscore the point that young borrowers' beliefs are particularly affected by macroeconomic shocks, since they have the shortest personal histories of lifetime experiences. Such changes in beliefs can produce long-lasting effects that only temper many years later.

Our results are, however, not restricted to the Great Inflation period. While a large share of the identifying variation in this paper stems from the 1970s, the above cited papers on inflation experiences among U.S. consumers in the Michigan Survey of Consumers (MSC) and among European consumers in the European Household Finance and Consumption Survey (HFCS) document similar magnitudes of experience-based learning. This paper is the first to pinpoint the effects on contract choice, quantify those effects, and provide cost estimates. Higher lifetime inflation experiences are the determining factor in choosing an FRM for between 10 and 20 percent of outstanding mortgages, and households exhibit an *ex ante* willingness to pay of between 6 and 14 basis points on the FRM mortgage contract. *Ex post* (as of the RFS survey year), the average switching household would have been better off by $8,000 to 16,000 after accounting for expected refinancing behavior and years of occupancy in the home.

Looking ahead, we can ask whether the experience of the mortgage crisis from 2007–2010 will have similar long-lasting effects and welfare implications for members of the Gen-X and Millennial generations who were first-time homeowners then.

Which policies could help ameliorate the costs for consumers? The answer to this question depends on the extent to which these decisions represent a mistake (biased beliefs) versus increased demand for insurance due to non-standard (instable) preferences. Our evidence on the influence of experiences on interest-rate beliefs point to a mistake. Moreover, the cost of this mistake is amplified by other well-known financial household mistakes, such as the failure to refinance optimally, which we accounted for in our analysis. Policy proposals to alleviate the cost of that mistake (including borrower counseling surrounding the refinancing decision and the marketing of FRMs that refinance automatically—cf. Keys et al. 2016) would also help here.

Another important dimension are general-equilibrium effects. Increased demand for FRMs by households with biased beliefs about future inflation and interest rates raises the FRM-ARM spread, so the higher mortgage rates paid by behavioral FRM borrowers help finance lower rates paid by non-behavioral ARM borrowers (Gabaix and Laibson 2006). Any policy that encourages greater ARM takeup would raise borrowing costs for these non-behavioral households, unless the reduced need for bank risk management resulted in large cost savings that could be passed through to all borrowers. Nevertheless, to the extent that such cross-subsidization is undesirable, our study suggests that the ARM's low reputation among borrowers is deserving of rehabilitation.

References

Agarwal, S., I. Ben-David, and V. Yao (2017). Systematic mistakes in the mortgage market and lack of financial sophistication. *Journal of Financial Economics 123*(1), 42–58.

Agarwal, S., J. C. Driscoll, and D. Laibson (2013). Optimal Mortgage Refinancing: A Closed Form Solution. *Journal of Money, Credit and Banking 45*(4), 591–622.

Agarwal, S., R. J. Rosen, and V. Yao (2015). Why Do Borrowers Make Mortgage Refinancing Mistakes? *Management Science 62*(12), 3494–3509.

Alesina, A. and N. Fuchs-Schundeln (2007). Good-bye Lenin (or Not?): The Effect of Communism on People's Preferences. *American Economic Review 97*(4), 1507–1528.

Alm, J. and J. R. Follain (1982). Alternative Mortgage Instruments, the Tilt Problem, and Consumer Welfare. *Journal of Financial and Quantitative Analysis 19*(1), 113–126.

Andersen, S., J. Y. Campbell, K. M. Nielsen, and T. Ramadorai (2015). Inattention and inertia in household finance: Evidence from the Danish mortgage market. Working Paper, National Bureau of Economic Research.

Angelis, D. D., P. Hall, and G. Young (1993). Analytical and bootstrap approximations to estimator distributions in L^1 regression. *Journal of the American Statistical Association 88*(424), 1310–1316.

Armona, L., A. Fuster, and B. Zafar (2018). Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment. *The Review of Economic Studies*, Forthcoming.

Ashford, J. and R. Sowden (1970). Multi-variate probit analysis. *Biometrics*, 535–546.

Atlas, S. A., E. J. Johnson, and J. W. Payne (2017). Time preferences and mortgage choice. *Journal of Marketing Research 54*(3), 415–429.

Badarinza, C., J. Y. Campbell, and T. Ramadorai (2018). What Calls to ARMs? International Evidence on Interest Rates and the Choice of Adjustable Rate Mortgages. *Management Science 64*(5), 2275–2288.

Baesel, J. and N. Biger (1980). The Allocation of Risk: Some Implications of Fixed Versus Index-Linked Mortgages. *Journal of Financial and Quantitative Analysis 15*(2), 457–468.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2019). The Economic Effects of Social Networks: Evidence from the Housing Market. *Journal of Political Economy*, Forthcoming.

Bailey, M., E. Davila, T. Kuchler, and J. Stroebel (2018). House Price Beliefs and Mortgage Leverage Choice. Working paper.

Bajo, E. and M. Barbi (2015). Out of Sight, Out of Mind: Financial Illiteracy and Sluggish Mortgage Refinancing. Working paper.

Bar-Gill, O. (2008). The law, economics and psychology of subprime mortgage contracts. *Cornell Law Review 94*, 1073.

Bennett, P., R. Peach, and S. Peristiani (2000). Implied Mortgage Refinancing Thresholds. *Real Estate Economics 28*(3), 405–434.

Bergstresser, D. and J. Beshears (2010). Who selected adjustable-rate mortgages? Evidence from the 1989-2007 Surveys of Consumer Finances. *Evidence from the 2007*, 10–083.

Brueckner, J. and J. Follain (1988). The Rise and Fall of the ARM: An Econometric Analysis of Mortgage Choice. *The Review of Economics and Statistics 70*(1), 93–102.

Brueckner, J. K. (1994). Borrower mobility, adverse selection, and mortgage points. *Journal of Financial Intermediation 3*(4), 416–441.

Burnside, C., M. Eichenbaum, and S. Rebelo (2016). Understanding booms and busts in housing markets. *Journal of Political Economy 124*(4), 1088–1147.

Campbell, J. Y. (2013). Mortgage Market Design. *Review of Finance 17*(1), 1–33.

Campbell, J. Y. and J. F. Cocco (2003, November). Household Risk Management and Optimal Mortgage Choice. *The Quarterly Journal of Economics 118*(4), 1449–1494.

Campbell, J. Y. and J. F. Cocco (2015). A model of mortgage default. *The Journal of Finance 70*(4), 1495–1554.

Case, K. E. and R. J. Shiller (1988). The behavior of home buyers in boom and post-boom markets. *New England Economic Review* (Nov), 29–46.

Case, K. E., R. J. Shiller, and A. K. Thompson (2012). What Have They Been Thinking?: Homebuyer Behavior in Hot and Cold Markets. *Brookings Papers on Economic Activity 2012*(2), 265–315.

Chambers, M. S., C. Garriga, and D. Schlagenhauf (2009, July). The Loan Structure and Housing Tenure Decisions in an Equilibrium Model of Mortgage Choice. *Review of Economic Dynamics 12*(3), 444–468.

Chiburis, R. C., J. Das, and M. Lokshin (2012). A practical comparison of the bivariate probit and linear IV estimators. *Economics Letters 117*(3), 762–766.

Clarida, R., J. Gali, and M. Gertler (2000). Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly journal of economics 115*(1), 147–180.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016). Asset pricing when 'this time is different'. *The Review of Financial Studies 30*(2), 505–535.

Cooper, R. (1975, January). Sliding Scale: Four Big California S&L Associations Quietly Adopt Variable-Rate Mortgages. *Wall Street Journal*, 24.

Coulibaly, B. and G. Li (2009). Choice of mortgage contracts: Evidence from the survey of consumer finances. *Real Estate Economics 37*(4), 659–673.

Crowder, W. J. and D. L. Hoffman (1996). The long-run relationship between nominal interest rates and inflation: the Fisher equation revisited. *Journal of money, credit and banking 28*(1), 102–118.

De Luca, G. (2008). SNP and SML estimation of univariate and bivariate binary-choice models. *Stata Journal 8*(2), 190–220.

den Haan, W. J. and A. T. Levin (1997). A practitioner's guide to robust covariance matrix estimation. In G. Maddala and C. Rao (Eds.), *Handbook of statistics*, Volume 15, Chapter 12, pp. 299–342. Elsevier.

Dhillon, U. S., J. D. Shilling, and C. F. Sirmans (1987). Choosing Between Fixed and Adjustable Rate Mortgages. *Journal of Money, Credit and Banking 19*(2), 260–267.

Donald, S. G. (1990). *Estimation of heteroskedastic limited dependent variable models.* Ph. D. thesis, University of British Columbia.

Dunn, K. B. and C. S. Spatt (1988). Private information and incentives: Implications for mortgage contract terms and pricing. *The Journal of Real Estate Finance and Economics 1*(1), 47–60.

Dynarski, M. (1985). Housing demand and disequilibrium. *Journal of Urban Economics 17*(1), 42–57.

Evans, M. D. and K. K. Lewis (1995). Do expected shifts in inflation affect estimates of the long-run Fisher relation? *The Journal of Finance 50*(1), 225–253.

Fan, Y. and Z. Liu (1997). A Simple Test for a Parametric Single Index Model. *Journal of Quantitative Economics 34*(3), 186–195.

Favara, G. and Z. Song (2014). House price dynamics with dispersed information. *Journal of Economic Theory 149*, 350–382.

Federal Register (2019). Designation of Replacement for Federal Housing Finance Agency's ARM Index. *84*(31), 32738–32742. July 1, 2019, No. 2019-N-5.

Fisher, I. (1930). *The Theory of Interest as Determined by Impatience to Spend Income and Opportunity to Invest It.* New York: Macmillan. Reprinted 1974, Augustus M. Kelly Publishers, Clifton, NJ.

Follain, J. (1990). Mortgage Choice. *Real Estate Economics 18*(2), 125–144.

Gabaix, X. and D. Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics 121*(2), 505–540.

Gabler, S., F. Laisney, and M. Lechner (1993). Seminonparametric Estimation of Binary-Choice Models With an Application to Labor-Force Participation. *Journal of Business & Economic Statistics 11*(1), 61–80.

Gallant, A. R. and D. W. Nychka (1987). Semi-Nonparametric Maximum Likelihood Estimation. *Econometrica 55*(2), 363–390.

Gao, Z., M. Sockin, and W. Xiong (2017). Economic consequences of housing speculation. Working paper.

Gao, Z., M. Sockin, and W. Xiong (2018). Learning about the Neighborhood: The Role of Supply Elasticity for Housing Cycles. Working paper.

Gathergood, J. and J. Weber (2017). Financial literacy, present bias and alternative mortgage products. *Journal of Banking & Finance 78*, 58–83.

Gelain, P. and K. J. Lansing (2014). House prices, expectations, and time-varying fundamentals. *Journal of Empirical Finance 29*, 3–25.

Ghent, A. (2015). Home ownership, household leverage and hyperbolic discounting.

Real Estate Economics 43(3), 750–781.

Glaeser, E. L., J. Gyourko, and A. Saiz (2008). Housing supply and housing bubbles. *Journal of Urban Economics 64*, 198–217.

Glaeser, E. L. and C. G. Nathanson (2015). Housing bubbles. In *Handbook of regional and urban economics*, Volume 5, pp. 701–751. Elsevier.

Glaeser, E. L. and C. G. Nathanson (2017). An extrapolative model of house price dynamics. *Journal of Financial Economics 126*(1), 147–170.

Gottlieb, D. and X. Zhang (2018). Long-Term Contracting with Time-Inconsistent Agents. Working paper.

Granziera, E. and S. Kozicki (2015). House price dynamics: Fundamentals and expectations. *Journal of Economic Dynamics and control 60*, 152–165.

Green, J. and J. B. Shoven (1986). The Effects of Interest Rates on Mortgage Prepayments. *Journal of Money, Credit and Banking 18*(1), 41.

Green, R. K. and M. LaCour-Little (1999). Some Truths about Ostriches: Who Doesn't Prepay Their Mortgages and Why They Don't. *Journal of Housing Economics 8*, 233–248.

Green, R. K. and S. M. Wachter (2005). The American mortgage in historical and international context. *Journal of Economic Perspectives 19*(4), 93–114.

Guiso, L. and P. Sodini (2013). Household finance: An emerging field. In G. M. Constantinides, M. Harris, and R. M. Stulz (Eds.), *Handbook of the Economics of Finance*, Volume 2, pp. 1397–1532. Elsevier.

Guren, A. M. (2018). House Price Momentum and Strategic Complementarity. *Journal of Political Economy 126*(3), 1172–1218.

Hahn, J. (1995). Bootstrapping quantile regression estimators. *Econometric Theory 11*(1), 105–121.

Hall, R. E. (1988). Intertemporal substitution in consumption. *Journal of political economy 96*(2), 339–357.

Härdle, W. and O. Linton (1994). Chapter 38 Applied nonparametric methods. Volume 4 of *Handbook of Econometrics*, pp. 2295–2339. Elsevier.

Hausman, J. A. (1978). Specification tests in econometrics. *Econometrica: Journal of the econometric society*, 1251–1271.

Heckman, J. (1979). Sample Selection Bias as a Specification Error. *Econometrica 47*(1), 153–161.

Heckman, J. (1990). Varieties of Selection Bias. *American Economic Review 80*(2), 313–318.

Heckman, J. J. (1978). Dummy Endogenous Variables in a Simultaneous Equation System. *Econometrica 46*(4), 931–959.

Heckman, J. J. and E. J. Vytlacil (2007). Chapter 71 Econometric Evaluation of Social Programs, Part II: Using the Marginal Treatment Effect to Organize Alternative Econometric Estimators to Evaluate Social Programs, and to Forecast their Effects in New Environments. *Handbook of Econometrics 6*(07), 4875–5143.

Horowitz, J. L. and W. Härdle (1994). Testing a parametric model against a semiparametric alternative. *Econometric Theory 10*(5), 821–848.

Jansson, M. (2004). The error in rejection probability of simple autocorrelation robust tests. *Econometrica 72*(3), 937–946.

Kaustia, M. and S. Knüpfer (2008). Do investors overweight personal experience? Evidence from IPO subscriptions. *Journal of Finance 63*(6), 2679–2702.

Kaustia, M. and S. Knüpfer (2012). Peer performance and stock market entry. *Journal of Financial Economics 104*(2), 321–338.

Kearl, J. R. (1979). Inflation, Mortgages, and Housing. *Journal of Political Economy 87*(5), 1115–1138.

Kennickell, A. B. (1998). Multiple imputation in the Survey of Consumer Finances. *Proceedings of the Section on Survey Research Methods*. 1998 Annual Meetings of the American Statistical Association, Dallas, Texas.

Kennickell, A. B. and R. L. Woodburn (1999). Consistent Weight Design for the 1989, 1992 and 1995 SCFs, and the Distribution of Wealth. *Review of Income and Wealth 45*(2), 193–215.

Keys, B. J., D. G. Pope, and J. C. Pope (2016). Failure to refinance. *Journal of Financial Economics 122*(3), 482–499.

Kiefer, N. M. and T. J. Vogelsang (2005). A new asymptotic theory for heteroskedasticity-autocorrelation robust tests. *Econometric Theory*, 1130–1164.

King, R. G. and M. W. Watson (1997). Testing Long-Run Neutrality. *Federal Reserve Bank of Richmond Economic Quarterly 83*(3), 69.

Knüpfer, S., E. Rantapuska, and M. Sarvimäki (2017). Formative experiences and portfolio choice: Evidence from the Finnish Great Depression. *Journal of Finance 72*(1), 133–166.

Koijen, R. S., O. V. Hemert, and S. V. Nieuwerburgh (2009, August). Mortgage Timing. *Journal of Financial Economics 93*(2), 292–324.

Kuchler, T. and B. Zafar (2018). Personal experiences and expectations about aggregate outcomes. *Journal of Finance, Forthcoming*.

Landier, A., Y. Ma, and D. Thesmar (2017). New Experimental Evidence on Expectations Formation. CEPR Discussion Paper No. DP12527.

Laudenbach, C., U. Malmendier, and A. Niessen-Ruenzi (2018). The Long-lasting Effects of Experiencing Communism on Financial Risk-Taking. Working Paper.

Lazarus, E., D. J. Lewis, J. H. Stock, and M. W. Watson (2018). HAR inference: Recommendations for practice. *Journal of Business & Economic Statistics 36*(4), 541–559.

Lea, M. (2010). International comparison of mortgage product offerings. Research Institute for Housing America Research.

Lee, L.-F. (1978). Unionism and Wage Rates: A Simultaneous Equations Model with Qualitative and Limited Dependent Variables. *International Economic Review 19*(2), 415–433.

Luce, R. D. and P. Suppes (1965). Preferences, Utility, and Subjective Probability. In *Handbook of Mathematical Psychology*. New York: Wiley.

MacDonald, G. and T. Holloway (1996). Early Evidence on Balloon Performance. *The Journal of Real Estate Finance and Economics 12*(3), 279–293.

Malmendier, U. and S. Nagel (2011, April). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking? *The Quarterly Journal of Economics 126*(1), 373–416.

Malmendier, U. and S. Nagel (2016). Learning from inflation experiences. *The Quarterly Journal of Economics 131*(1), 53–87.

Malmendier, U., S. Nagel, and Z. Yan (2018). The Making of Hawks and Doves. *Working Paper*.

Malmendier, U., D. Pouzo, and V. Vanasco (2020). Investor Experiences and Financial Market Dynamics. *Journal of Financial Economics*.

Malmendier, U. and L. S. Shen (2015). Experience Effects in Consumption. Working paper, UC-Berkeley.

Malmendier, U. and A. Steiny (2016). Rent or buy? the role of lifetime experiences of macroeconomic shocks within and across countries. Working Paper, Working Paper.

Martins, M. F. O. (2001). Parametric and semiparametric estimation of sample selection models: an empirical application to the female labour force in Portugal. *Journal of Applied Econometrics 16*(1), 23–39.

Mayer, C., K. Pence, and S. M. Sherlund (2009, March). The Rise in Mortgage Defaults. *Journal of Economic Perspectives 23*(1), 27–50.

Mayer, C. and T. Sinai (2009). US house price dynamics and behavioral finance. *Policy Making Insights from Behavioral Economics. Boston, Mass: Federal Reserve Bank of Boston*.

McFadden, D. (1974). Conditional Logit Analysis of Qualitative Choice Behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. New York: Academic Press.

Miller, D. W. (1986). Adjustable Rate Mortgages: A Proposed Statutory Reform. *Santa Clara L. Rev. 26*, 253.

Mishkin, F. S. (1992). Is the Fisher effect for real?: A reexamination of the relationship between inflation and interest rates. *Journal of Monetary economics 30*(2), 195–215.

Montalto, C. P. and J. Sung (1996). Multiple imputation in the 1992 Survey of Consumer Finances. *Journal of Financial Counseling and Planning 7*.

Müller, U. K. (2014). HAC corrections for strongly autocorrelated time series. *Journal of Business & Economic Statistics 32*(3), 311–322.

Müller, U. K. and M. W. Watson (2018). Long-Run Covariability. *Econometrica 86*(3), 775–804.

Murphy, A. (2007). Score tests of normality in bivariate probit models. *Economics Letters 95*(3), 374–379.

Nathanson, C. G. and E. Zwick (2018, 12). Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market. *The Journal of Fi-*

nance 73(6), 2587–2633.
Newey, W. K. (2009, January). Two-Step Series Estimation of Sample Selection Models. *Econometrics Journal 12*, S217–S229.
Newey, W. K. and K. D. West (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica 55*(3), 703–708.
Oreopoulos, P., T. von Wachter, and A. Heisz (2012). The Short- and Long-Term Career Effects of Graduating in a Recession. *American Economic Journal: Applied Economics 4*(1), 1–29.
Paiella, M. and A. F. Pozzolo (2007). Choosing between fixed-and adjustable-rate mortgages. In *Household Credit Usage*, pp. 219–236. Springer.
Peek, J. et al. (1990). A call to ARMs: adjustable rate mortgages in the 1980s. *New England Economic Review*, 47–61.
Piazzesi, M. and M. Schneider (2009, May). Momentum Traders in the Housing Market: Survey Evidence and a Search Model. *American Economic Review 99*(2), 406–11.
Powell, J. L. (1984). Least Absolute Deviations Estimation for the Censored Regression Model. *Journal of Econometrics 25*(3), 303–325.
Quigley, J. M. (1987). Interest Rate Variations, Mortgage Prepayments and Household Mobility. *Review of Economics and Statistics 69*(4), 636–643.
Rubin, D. B. (1987). *Multiple Imputation for Nonresponse in Surveys.* New York: John Wiley & Sons.
Ruud, P. A. (1983). Sufficient Conditions for the Consistency of Maximum Likelihood Estimation Despite Misspecification of Distribution in Multinomial Discrete Choice Models. *Econometrica 51*(1), 225–228.
Sa-Aadu, J. and C. F. Sirmans (1995). Differentiated contracts, heterogeneous borrowers, and the mortgage choice decision. *Journal of Money, Credit and Banking 27*(2), 498–510.
Schafgans, M. M. and V. Zinde-Walsh (2002). On intercept estimation in the sample selection model. *Econometric Theory 18*(1), 40–50.
Schlafmann, K. (2016). Housing, mortgages, and self control. Working paper.
Schraeder, S. (2015). Information Processing and Non-Bayesian Learning in Financial Markets. *Review of Finance*, 1–31.
Scrimgeour, D. (2008). The Great Inflation Was Not Asymmetric: International Evidence. *Journal of Money, Credit and Banking*, 799–815.
Shiller, R. J. (1999). Human behavior and the efficiency of the financial system. *Handbook of macroeconomics 1*, 1305–1340.
Shiller, R. J. (2005). *Irrational Exuberance, Second Edition.* Princeton University Press.
Stanton, R. (1995). Rational Prepayment and the Value of Mortgage-Backed Securities. *Review of Financial Studies 8*(3), 677–708.
Stanton, R. and N. Wallace (1998). Mortgage Choice: What's the Point? *Real Estate*

Economics 26(2), 173–205.

Strahilevitz, M., T. Odean, and B. Barber (2011). Once Burned, Twice Shy: How Naive Learning, Counterfactuals, and Regret Affect the Repurchase of Stocks Previously Sold. *Journal of Marketing Research 48*, 102–120.

Suher, M. (2016). Inferring expectations from homeownership decisions: House price forecasts during a boom and bust. Working paper.

Train, K. (2009). *Discrete Choice Methods with Simulation* (Second edition). Cambridge: Cambridge UP.

Wall Street Journal (1973, January). Variable-Rate Mortgages. pp. 10.

White, H. (1982). Maximum Likelihood Estimation of Misspecified Models. *Econometrica 50*(1), 1–25.

Yogo, M. (2004). Estimating the elasticity of intertemporal substitution when instruments are weak. *Review of Economics and Statistics 86*(3), 797–810.

Zellner, A. and T. H. Lee (1965). Joint estimation of relationships involving discrete random variables. *Econometrica: Journal of the Econometric Society*, 382–394.

Table 1: SCF Summary Statistics

All SCF HHs			
	Mean	**SD**	**N**
Respondent age	47.2	13.4	36,266
HH income (2013 $)	87,560	259,644	36,266
HH net worth (2013 $)	477,242	2,888,737	36,266
Homeowner	0.68	0.47	36,266
Mortgage | Howeowner	0.68	0.47	26,219
FRM | Mortgage	0.86	0.35	16,824
Loan/CLL | Mortgage	0.46	0.41	15,708
Jumbo | Mortgage	0.06	0.24	15,446
Non-conventional | Mortgage	0.28	0.45	16,824
Has second mortgage | Mortgage	0.09	0.28	16,824
Has second home	0.19	0.39	36,266
Fraction expecting higher *i*	0.70	0.46	36,266
Fraction expecting lower *i*	0.06	0.25	36,266
Net fraction expecting higher *i*	0.63	0.60	36,266
Inflation experiences (%)	4.09	0.67	36,266

New Mortgages			
	FRM	**ARM**	**FRM - ARM**
N =	*2,538*	*725*	
Respondent age	44.0	44.8	-0.8
HH income (2013 $)	136,756	205,648	-68,892*
HH net worth (2013 $)	712,055	1,516,472	-804,417*
Loan / CLL	0.47	0.68	-0.21*
Jumbo loan	0.06	0.17	-0.11*
Non-conventional	0.30	0.16	0.14*
Junior mortgage	0.09	0.14	-0.05*
Second home	0.15	0.25	-0.09*
Fraction expecting higher *i*	0.74	0.65	0.08*
Fraction expecting lower *i*	0.06	0.07	-0.02
Net fraction expecting higher *i*	0.68	0.58	0.10*
Inflation experiences (%)	4.09	4.18	-0.09*

Notes. The table reports summary statistics for respondents to the 1989-2013 waves of the SCF. The top panel is all respondents aged 25-74; each observation is a household. The bottom panel is new mortgages that were originated in the survey year only (1989, 1992, ..., 2013); each observation is a mortgage. Age, income, net worth, and inflation experiences are as of the survey year. Calculations use SCF "revised consistent" sampling weights (X42001), rescaled so that each survey wave receives equal weight. We adjust for multiple imputation following Rubin (1987). All statistics are based on available cases. * $p < 0.05$.

Table 2: RFS Summary Statistics

	FRM	ARM	Balloon	FRM - ARM
N =	12,416	2,245	735	
Contract Characteristics				
Current rate (bps)	972.7	924.5	870.8	48.2*
Initial rate (bps)	"	876.2	"	96.4*
Margin (bps)	n.a.	282.7	n.a.	n.a.
Years since origination	2.6	2.8	2.1	-0.2*
Original Term (years)	23.2	26.1	8.9	-2.9*
Loan Amount (2000 $k)	102.0	140.3	89.9	-38.3*
Prepayment penalty	0.061	0.091	0.058	0.0*
Economic Conditions (all in %)				
Inflation	3.24	3.35	3.45	-0.12*
FRM - ARM spread	1.75	1.86	1.69	-0.11*
Default spread	2.09	2.09	2.06	0.00
Yield spread	0.90	0.99	0.84	-0.09*
Borrower Characteristics				
Primary owner age	41.4	41.8	42.8	-0.4
Inflation experiences (%)	4.74	4.79	4.68	-0.05*
Non-white	0.136	0.099	0.121	0.037*
Hispanic	0.508	0.580	0.516	-0.071*
Veteran	0.226	0.216	0.245	0.010
Joint owners	0.703	0.694	0.660	0.009
First-time owner	0.413	0.348	0.347	0.065*
Has investment income	0.282	0.302	0.256	-0.021
Has business income	0.094	0.106	0.135	-0.012
Total income (2000 $)	75,177	84,165	71,479	-8,989*
Property Characteristics				
Central city of MSA	0.257	0.258	0.214	0.000
Rural county	0.143	0.162	0.310	-0.018*
Second home	0.012	0.017	0.017	-0.005
Mobile home	0.032	0.020	0.049	0.012*
Condo	0.071	0.118	0.057	-0.047*
Other Loan Characteristics				
Junior mortgage	0.129	0.086	0.233	0.043*
Non-conventional	0.211	0.061	0.049	0.150*
Refi	0.256	0.244	0.294	0.012
Loan / income	1.73	2.04	1.54	-0.31*
Loan / value × 100	81.7	90.0	80.2	-8.3*
Loan / CLL	0.426	0.554	0.386	-0.128*
Jumbo loan	0.043	0.127	0.056	-0.084*
Points paid (bps)	39.6	42.1	14.9	-2.5

Notes. The table reports summary statistics for respondents to the 1991 and 2001 RFS of homeowner properties, with origination at most 6 years before the survey year (1985-1991, 1995-2001) and primary-owner age between 25 and 74 years at origination. All statistics are as of the origination year, based on available cases. Investment income, second home status, and buydown indicator only available for 2001. "FRM - ARM spread" is from Freddie Mac PMMS, by origination year and Census region. "Default spread" is Moody's seasoned corporate BAA rate minus 10-year CM Treasury. "Yield spread" is the 10-year CM Treasury minus the 1-year CM Treasury rates. All other variable definitions are in Appendix C. * $p < 0.05$.

Table 3: Interest Rate Expectations and Inflation Experiences

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable is:	Net Expectation (-1/0/1)		Expect Lower i ? (0/1)		Expect Higher i ? (0/1)	
Estimation method:	LPM	Ordered Probit	LPM	Probit	LPM	Probit
Inflation experiences (%)	0.0990***	0.200***	-0.0385***	-0.283***	0.0605***	0.175***
	(0.01)	(0.03)	(0.01)	(0.04)	(0.01)	(0.03)
Survey Year FE?	YES	YES	YES	YES	YES	YES
Number of Households	36,264	36,264	36,264	36,264	36,264	36,264
R^2	0.023		0.008		0.024	
Pseudo R^2		0.016		0.016		0.020
F-stat on π^e	55.54	55.07	43.95	49.21	37.57	38.57

Table 4: Interest Rate Expectations and Mortgage Choice

	(1)	(2)	(3)	(4)	(5)
Dependent variable is:	FRM	FRM	FRM	Expect higher i	FRM
Estimation method:	Probit	Probit	Bivariate Probit		Probit
			(2nd Stage)	(1st Stage)	(RF)
Expect higher i (0/1)	0.0535*	0.245***	1.351***		
	(0.03)	(0.09)	(0.32)		
Inflation experiences (%)				0.290**	0.388**
				(0.13)	(0.16)
FRM - ARM spread (%)	-0.158***	-0.232***	-0.249***		*(Absorbed*
	(0.02)	(0.06)	(0.06)		*by FE)*
Mortgage controls	YES	YES	YES		YES
Sociodemographic controls	YES	YES	YES		YES
Origination year FE				YES	YES
Sample	All Mtgs.	New Mtgs.	New Mtgs.		New Mtgs.
Number of Mortgages	21,330	3,123	3,123		3,123
Pseudo R^2	0.033	0.051	0.097		0.086
ρ			-0.68***		

Notes. Tables 3 and 4 report linear probability model and probit coefficient estimates of equations (2) and (3), relating past inflation experiences to future nominal interest rate forecasts and future interest rate forecasts to mortgage choice. "Net expectations" codes households expecting higher interest rates as +1, lower interest rates as -1, and about the same as 0. "FRM" is an indicator equal to 1 if the household chose an FRM and 0 if it chose an ARM. All other variable definitions are in Appendix C. Mortgage controls are Refi dummy, Junior Mortgage dummy, Non-conventional dummy, Loan / CLL, and Jumbo dummy. Sociodemographic controls are log(Income), log(Net Worth), Age, Age^2, and Married dummy. Table 3: the sample is all SCF households with a respondent between the ages of 25 and 74 in survey waves 1989-2013. Table 4: the "All Mtgs." sample consists of all SCF households with a mortgage, conditional on having a respondent between the ages of 25 and 74 in survey waves 1989-2013; the "New Mtgs." sample further restricts the sample to households with mortgages originated in the survey year only (1989, 1992, ..., 2013); each observation is a mortgage. Regressions use SCF "revised consistent" sampling weights (X42001), rescaled so that each survey wave receives equal weight. We adjust for multiple imputation using the Rubin (1987) methodology. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 5: Reduced-Form Logit Model of Mortgage Choice

	(1)	(2)	(3)	(4)	(5)
Freddie Mac PMMS index	-0.483**				
rate (%)	(0.237)				
FRM Alternative-Specific Characteristics					
Freddie Mac PMMS FRM		-3.55***	-3.56***	-3.33***	-3.59***
index rate (%)		(0.549)	(0.549)	(0.575)	(0.816)
Inflation experiences (%)	0.220**	0.216**	0.292***	0.254***	0.187*
	(0.095)	(0.095)	(0.083)	(0.086)	(0.098)
Log(Income)	-0.0069	-0.0062	-0.0063	0.0276**	0.0278**
	(0.014)	(0.014)	(0.014)	(0.012)	(0.012)
Age	-0.019	-0.019	-0.017	0.019	0.017
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Age^2 / 100	0.020	0.020	0.020	-0.018	-0.017
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
ARM Alternative-Specific Characteristics					
Freddie Mac PMMS ARM		-0.861***	-0.865***	-0.768***	-0.844***
initial rate index (%)		(0.243)	(0.243)	(0.250)	(0.314)
Balloon Mortgage Alternative-Specific Characteristics					
Inflation experiences (%)	-0.308*	-0.303*			
	(0.168)	(0.168)			
Log(Income)	-0.0342*	-0.0346*	-0.0349*	0.0054	
	(0.020)	(0.020)	(0.020)	(0.020)	
Age	-0.0204	-0.0213	-0.0184	-0.0298	
	(0.027)	(0.027)	(0.027)	(0.029)	
Age^2	0.02420	0.02520	0.02820	0.03250	
	(0.02990)	(0.02990)	(0.02960)	(0.03080)	
Alternative-specific constants	YES	YES	YES	YES	YES
Origination year FE	YES	YES	YES	YES	YES
Mortgage controls				YES	YES
Socidemographic controls				YES	YES
Number of Choice Situations	15,051	15,051	15,051	15,051	14,337
Number of Alternatives	3	3	3	3	2
Pseudo R2	0.018	0.020	0.019	0.071	0.069
$-\beta_{\pi, FRM} / \beta_{Rate, FRM}$	0.456	0.061**	0.082***	0.076***	0.052*
(S.E. by delta method)	(0.295)	(0.028)	(0.027)	(0.029)	(0.030)

Notes. The table reports coefficient estimates for a reduced-form, multinomial logit model of mortgage choice among FRM, Balloon, and ARM alternatives in the 1991 and 2001 RFS. Cols. 1-4 include all three alternatives, while Col. 5 reports binomial logit coefficients, excluding the balloon alternative. The sample is mortgages originated ≤ 6 years prior to the survey year, with primary owner age between 25 and 74 years. The omitted category for sociodemographic variables is ARM. Separate coefficients for all mortgage / sociodemographic controls are estimated for each alternative. Mortgage controls are Refi dummy, Junior Mortgage dummy, Non-conventional dummy, Loan / CLL, Jumbo dummy, and Points Paid. Sociodemographic controls are First-time Owner dummy, Joint Owners dummy, and Rural county dummy. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 6: Selection-Corrected Mortgage Rate Equations

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable is:	FRM Rate		ARM Initial Rate		ARM Margin	
Estimation Method	CLAD	SPSC CLAD	CLAD	SPSC CLAD	OLOGIT	SPSC OLOGIT
Freddie Mac PMMS index	84.21***	96.61***	77.28***	86.71***	-11.83***	-6.513**
rate (%)	(0.79)	(2.98)	(3.35)	(6.45)	(2.26)	(2.66)
Log(Income)	-0.411	-2.056*	1.559	-0.00414	-1.516	-1.608
	(0.84)	(1.14)	(2.25)	(2.60)	(1.20)	(1.14)
First-time owner	7.209***	6.734	16.74**	13.16	1.849	0.505
	(2.41)	(4.52)	(8.16)	(9.57)	(5.18)	(5.17)
Joint owners	-4.273*	-17.59***	8.587	-1.483	0.413	-3.729
	(2.47)	(5.22)	(8.34)	(10.85)	(5.10)	(5.14)
Rural county	12.43***	33.49***	55.44***	73.96***	-10.1	-4.308
	(3.55)	(7.73)	(10.79)	(12.84)	(7.78)	(8.90)
Refi	-25.71***	-35.34***	13.13	-0.751	3.542	-1.14
	(2.94)	(5.05)	(8.80)	(12.18)	(5.35)	(6.10)
Junior mortgage	171.5***	141.9***	194.5***	175.8***	30.5	10.74
	(9.52)	(13.54)	(15.46)	(28.72)	(18.86)	(22.03)
Non-conventional	0.201	-114.0***	-45.61**	-47.4	-60.11***	-160.4***
	(2.62)	(28.80)	(19.74)	(56.07)	(10.81)	(36.29)
Points paid (pctg points)	-1.194*	-0.396	-7.850**	-8.548*	0.522	1.26
	(0.70)	(1.43)	(3.37)	(4.50)	(1.72)	(1.74)
Loan / CLL	-54.43***	1.202	-97.21***	-62.47**	-19.52**	-10.94
	(6.27)	(14.92)	(15.46)	(25.99)	(9.21)	(13.45)
Jumbo loan	35.85***	67.76***	60.70***	71.47***	-2.891	-13.02
	(7.81)	(17.94)	(17.99)	(19.11)	(9.73)	(10.02)
Constant[a]	156.2***	187.2***	256.5***	156.1**	-	-
	(11.71)	(22.98)	(33.68)	(73.56)	-	-
Margin reference rate dummies					YES	YES
Observations	12,155	12,155	1,410	1,410	1,490	1,490
Pseudo R2	0.219	0.221	0.270	0.276	0.026	0.031
χ^2 test of H0: no selection bias[b]		21.49		7.201		14.510
[p-value]		[0.029]		[0.783]		[0.339]
Average Selection Bias[c]		-116.9		50.5		-

Notes. The table reports two-step censored least absolute deviation (CLAD) estimates and CLAD semiparametric selection-corrected (SPSC) estimates of the mortgage rate pricing equations. The sample is mortgages originated $\leq$ 6 years ago as of 1991 and 2001 Residential Finance Surveys, with primary owner age between 25 and 74 years. Dependent variables are FRM, ARM initial, and ARM margin rates expressed in bps. Standard errors (in parentheses) are analytic, robust standard errors in columns 1 and 3, bootstrapped standard errors, adjusted for first-step estimation, from 200 repetitions in columns 2, 4, and 6, bootstrapped standard errors from 200 repetitions in column 5.

a. SPSC absorbs the intercept into the control function. As suggested by Heckman (1990), we estimate the intercept as the median of $Rate_n - Z_n\hat{\Gamma}_i$ in the subsample of observations n with choice probabilities for alternative i above the 90$^{\text{th}}$ percentile. Cols 5-6 are marginal effects, so no intercept is reported.

b. Test statistic for no selection bias is a quadratic form for the difference in slope parameters: $(\hat{\Gamma}_{SC} - \hat{\Gamma}_{noSC})'\hat{V}^{-1}(\hat{\Gamma}_{SC} - \hat{\Gamma}_{noSC}) \sim \chi^2(L)$, where $L = \text{length}(\Gamma)$ (11, 11, and 13, respectively). We calculate V by bootstrapping the difference 200 times. In column 6, the test statistic is calculated on the underlying ordered logit slope coefficients.

c. Average Selection Bias is average value of the selection poynomial in the subsample choosing alternative i. *** p<0.01, ** p<0.05, * p<0.1

Table 7: Structural Logit Model of Mortgage Choice

	(1)	(2)	(3)	(4)	(5)	(6)
Step 2 Selection Correction?	No	Yes	No	Yes	No	Yes
FRM Rate Offered	0.764	-1.474**	-0.127	-1.272***	-0.575	-0.692*
	(0.74)	(0.58)	(0.60)	(0.45)	(0.45)	(0.41)
Initial ARM Rate Offered	-0.368	1.280**	0.838	1.196***	0.184	0.593
	(0.62)	(0.54)	(0.55)	(0.38)	(0.35)	(0.39)
ARM Margin Offered			-2.364***	-0.302	3.738***	2.600**
			(0.55)	(0.47)	(1.03)	(1.22)
Inflation experiences (%)	0.237**	0.181*	0.222**	0.180*	0.181*	0.192**
	(0.09)	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
Log(Income)	0.00221	-0.00875	-0.0572	-0.0171	0.0798*	0.0916
	(0.02)	(0.03)	(0.04)	(0.04)	(0.05)	(0.06)
Age	-0.015	0.004	-0.007	0.004	0.007	0.015
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Age^2 / 100	0.018	-0.005	0.010	-0.004	-0.006	-0.014
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Joint owners	0.144	-0.074	0.035	-0.062	0.101	0.183
	(0.12)	(0.13)	(0.15)	(0.12)	(0.16)	(0.20)
Rural county	-0.053	-0.776**	-0.860**	-0.761***	0.106	-0.375
	(0.32)	(0.35)	(0.36)	(0.28)	(0.33)	(0.40)
Non-conventional					3.744***	4.736**
					(0.59)	(2.16)
Alternative-specific constants	YES	YES	YES	YES	YES	YES
Origination year FE	YES	YES	YES	YES	YES	YES
Number of Choice Situations	14,337	14,337	14,337	14,337	14,337	14,337
Pseudo R2	0.023	0.059	0.041	0.060	0.064	0.066
$-\beta_{\pi, FRM} / \beta_{Rate, FRM}$	-0.31**	0.123*	1.75	0.142*	0.315*	0.277*
(S.E. by delta method)	(0.129)	(0.067)	(1.836)	(0.078)	(0.186)	(0.149)

Notes. The table reports binomial logit coefficient estimates for the structural model of mortgage choice between FRM and ARM alternatives in the 1991 and 2001 RFS. Estimates are produced by a three-step procedure, in which interest rates for both alternatives are predicted (step 2) after correcting for sample selection (step 1) using the estimates from Tables 6 and 5, respectively. The sample is mortgages originated ≤ 6 years prior to the survey year, with primary owner age between 25 and 74 years. The dependent variable equals 1 if an FRM is chosen, and 0 for ARMs. Bootstrapped standard errors in parentheses, adjusting for first- and second-step estimation, from 200 repetitions. *** p<0.01, ** p<0.05, * p<0.1

Table 8: Additional Interest Paid Due to Inflation Experiences

	Scenario 1: Primary Mortgage Market Survey rates				
Time Horizon:	Survey Year	5 years	10 years	15 years	E[tenure | age]
After-tax PDV: (all in $)					
No Refi	2,386	5,542	11,148	17,085	13,052
Expected Refi	-	5,422	7,681	9,924	7,827
Optimal Refi	-	4,805	6,213	7,993	6,493
% switching households	13.8	13.8	13.8	13.8	13.8

	Scenario 2: Risk-adjusted rates, seniority-adjusted ARM margins				
Time Horizon:	Survey Year	5 years	10 years	15 years	E[tenure | age]
After-tax PDV: (all in $)					
No Refi	5,674	10,124	19,126	27,345	20,819
Expected Refi	-	10,056	15,886	20,505	15,769
Optimal Refi	-	9,455	14,460	18,639	14,475
% switching households	13.5	13.5	13.5	13.5	13.5

	Scenario 3: Risk-adjusted rates and ARM margins				
Time Horizon:	Survey Year	5 years	10 years	15 years	E[tenure | age]
After-tax PDV: (all in $)					
No Refi	5,355	9,635	18,193	26,176	19,964
Expected Refi	-	9,556	14,915	19,261	14,854
Optimal Refi	-	8,947	13,474	17,374	13,543
% switching households	14.3	14.3	14.3	14.3	14.3

Notes. The table reports the "welfare-relevant treatment effect" (WRTE) on switching households, measured as the differential after-tax interest + refinancing costs paid by a household choosing an FRM instead of an ARM due to overweighting their inflation experiences. All dollar figures are in constant year-2000 units. Positive values indicate that the FRM is more expensive than the ARM. To calculate the WRTE on switching households, each household is weighted by their decline in probability of choosing an FRM contract when the experienced inflation coefficient is turned off in the choice model (scenario 1 = Table 5 col. 5, scenario 2 = Table 7 col. 2, scenario 3 = Table 7 col. 6). PDV calculations assume a nominal discount rate of 8% / year ($r = .04$, $\pi = .04$). Predicted interest rates in scenario 1 are from the PMMS, and in scenarios 2 and 3 from Table 6, cols. 2, 4, and 6. In the "No Refi" row, the household holds the initial FRM until maturity. In the "Expected Refi" row, the household is assumed to refinance probabilistically, according to a probit function of the differential between the current FRM rate i_0 and the refinanced rate i, estimated in Andersen et al. (2015) Table 8, column 1. (The timing of principal repayment is the same as in Optimal Refi row.) In the "Optimal Refi" row, the household refinances deterministically whenever $i_0 - i > OT$, where OT is the square-root rule approximation to the optimal threshold for refinancing, derived by Agarwal et al. (2013). The mortgage interest deduction is calculated assuming a 25% marginal tax rate. Refinancing costs $2,000 and is not tax-deductible. "E[tenure | age]" indicates that probability of moving every year estimated as a 4th-order polynomial in head of household's age, using 5-year migration / geographic mobility data from CPS ASEC 2005 and 2010.

Table 9: Simulation Parameters

Parameter	Description	Value	Source
μ	Mean log inflation	0.038	CPI-U, 1960-2013
σ_π	Standard deviation of log inflation	0.027	CPI-U, 1960-2013
ϕ	Log inflation autoregression parameter	0.811	CPI-U, 1960-2013
ρ	Mean log real interest rate	0.02	Campbell & Cocco (2003)
σ_r	Standard deviation of log real interest rate	0.022	Campbell & Cocco (2003)
θ^{10}	Ten-year nominal term premium	0.01	Average of ten-year minus one-year constant maturity U.S. Treasury yields, 1960-2013
$\theta^{A,1}$	ARM initial premium over one-year nominal bond (year 1 only)	0.015	Average spread between PMMS initial rate and CM U.S. Treasury, 1984-2013
θ^{A}	ARM reset margin over one-year nominal bond (years 2-30)	0.0275	Average PMMS margin, 1987-2013
θ^{F}	FRM premium over ten-year nominal bond	0.017	Average spread between PMMS rate and CM U.S. Treasury, 1971-2013

Table 10: Monte Carlo Simulation Results

	Mean	**SD**	**10th Pctl.**	**90th Pctl.**	**Pct. > 0**
WRTE (after-tax $)					
No Refi	7,926	9,474	-4,214	19,963	78
Expected Refi	6,481	7,659	-4,661	16,213	82
Optimal Refi	5,118	7,586	-6,038	15,014	76
Economic Conditions (Years 1-30)					
Average inflation (%)	3.70	1.54	1.60	5.67	100
Average FRM-ARM spread (%)	0.05	0.95	-1.37	1.14	50

Notes. The table reports summary statistics from a Monte Carlo simulation of different possible inflation and mortgage rate paths, $T = 30$ years each, using the parameters in Table 9, based on 100 replications. WRTE is calculated using Scenario 3 household-level interest rates given simulated baseline rates and mobility given borrower age. We use the same experience-induced switching probabilities as in Table 8, from the actual RFS origination years. All other calculation details are the same as in Table 8.

Table 11: Aggregate Cost of the Great Inflation

(1)	(2)	(3)	(4)	(5)
Survey Year - Cohort	% Switching HHs	E[WRTE] per switching HH ($)	# of switching HHs (1000s)	Total Cost ($m)
1991 - G.I. & Silent Gens.	6.6	15,869	322.3	5,115
1991 - Baby Boomers	8.1	14,433	1,018.1	14,694
2001 - G.I. & Silent Gens.	3.3	15,314	129.7	1,987
2001 - Baby Boomers	3.6	17,769	502.7	8,933
2001 - Gen Xers	2.9	12,495	248.8	3,108

Notes. The table reports the aggregate additional interest paid (in 2000 $) by members of each generation who chose an FRM instead of an ARM because of their inflation experiences during 1968-84, among mortgages originated $\leq$ 6 years prior to survey year. The G.I. and Silent Generations are individuals born prior to 1946; Baby Boomers are born between 1946 and 1964; and Gen Xers are born after 1964. Column (2) shows the predicted change in the FRM product share if the Great Inflation had not occurred, as shown in Figure 3. Column (3) shows the Scenario 3 WRTE under expected refinancing and mobility given age. Column (4) assumes that every sample household represents 2,599 population HHs in 1991 and 3,655 population HHs in 2001. Column (5) = Column (3) $\times$ Column (4).

Online Appendix

A International Mortgage and Inflation Data

Country	Mortgage Type	Source (Mortgage Data)	Median Inflation	Source (Inflation Data)
Algeria	Fixed	Ehlers and Villar (2015)	4.09	IMF (2020)
Argentina	Mixed	Ehlers and Villar (2015)	8.48	World Bank (2020)
Australia	Variable	Lea (2010)	2.47	OECD (2020)
Austria	Variable	Albertazzi et al. (2019)	2.00	OECD (2020)
Belgium	Fixed	Albertazzi et al. (2019)	2.01	OECD (2020)
Brazil	Fixed	Ehlers and Villar (2015)	6.27	OECD (2020)
Canada	Variable	Lea (2010)	1.98	OECD (2020)
Chile	Fixed	Ehlers and Villar (2015)	3.03	OECD (2020)
China	Variable	Warnock and Warnock (2007)	2.00	OECD (2020)
Colombia	Fixed	Ehlers and Villar (2015)	4.65	OECD (2020)
Croatia	Variable	Reuters (2017)	2.14	World Bank (2020)
Cyprus	Mixed	Ehrmann and Ziegelmeyer (2014)	2.30	World Bank (2020)
Czech Republic	Mixed	Ehlers and Villar (2015)	2.03	OECD (2020)
Denmark	Mixed	Lea (2010)	1.76	OECD (2020)
Estonia	Mixed	Swedish Bankers' Association (2018)	3.43	OECD (2020)
Finland	Variable	Scanlon and Whitehead (2004)	1.13	OECD (2020)
France	Mixed	Lea (2010)	1.65	OECD (2020)
Germany	Variable	Lea (2010)	1.51	OECD (2020)
Greece	Mixed	Albertazzi et al. (2019)	2.90	OECD (2020)
Hungary	Mixed	Kubas (2018)	4.07	OECD (2020)
Iceland	Variable	Bjarnason (2014)	3.99	OECD (2020)
India	Mixed	Campbell et al. (2012)	5.83	OECD (2020)
Indonesia	Variable	Ehlers and Villar (2015)	6.21	OECD (2020)
Ireland	Variable	Lea (2010)	1.95	OECD (2020)
Israel	Variable	Ehlers and Villar (2015)	1.10	OECD (2020)
Italy	Mixed	Albertazzi et al. (2019)	1.91	OECD (2020)
Japan	Variable	Lea (2010)	-0.03	OECD (2020)
Kenya	Variable	The World Bank (2011)	9.23	World Bank (2020)
Latvia	Variable	Albertazzi et al. (2019)	2.73	OECD (2020)
Lithuania	Variable	Scanlon and Whitehead (2004)	1.85	OECD (2020)
Luxembourg	Mixed	Albertazzi et al. (2019)	2.15	OECD (2020)
Malaysia	Variable	Endut and Hua (2009)	1.92	World Bank (2020)
Malta	Variable	Central Bank of Malta (2018)	1.86	World Bank (2020)
Mexico	Fixed	Ehlers and Villar (2015)	4.13	OECD (2020)
Morocco	Mixed	Dübel et al. (2016)	1.39	World Bank (2020)
Netherlands	Variable	Lea (2010)	1.70	OECD (2020)
New Zealand	Variable	Fitch Ratings (2020)	2.20	OECD (2020)
Norway	Variable	Almklov and Tørum (2007)	2.17	OECD (2020)
Poland	Variable	Ehlers and Villar (2015)	2.22	OECD (2020)
Portugal	Mixed	Swedish Bankers' Association (2018)	2.32	OECD (2020)
Romania	Variable	Hegedüs and Struyk (2005)	5.69	World Bank (2020)
Russia	Fixed	Hegedüs and Struyk (2005)	9.34	OECD (2020)
Singapore	Variable	Ehlers and Villar (2015)	0.98	World Bank (2020)
Slovakia	Variable	Kubas (2018)	2.73	OECD (2020)
Slovenia	Variable	Albertazzi et al. (2019)	2.13	OECD (2020)

Country	Mortgage Type	Source (Mortgage Data)	Median Inflation	Source (Inflation Data)
South Africa	Variable	Everything Overseas (2016)	5.51	OECD (2020)
South Korea	Mixed	Ehlers and Villar (2015)	2.40	OECD (2020)
Spain	Variable	Albertazzi et al. (2019)	2.62	OECD (2020)
Sweden	Mixed	Scanlon and Whitehead (2004)	1.26	OECD (2020)
Switzerland	Variable	Lea (2010)	0.64	OECD (2020)
Thailand	Variable	Ehlers and Villar (2015)	1.85	World Bank (2020)
Turkey	Fixed	Ehlers and Villar (2015)	8.87	OECD (2020)
Ukraine	Fixed	Cerutti et al. (2015)	11.46	World Bank (2020)
United Kingdom	Variable	Lea (2010)	2.05	OECD (2020)
United States	Fixed	Lea (2010)	2.20	OECD (2020)

Notes. Mortgage Type is one of three classifications: *Variable* indicates that at least 75% of all mortgages in that country have variable interest rates throughout or after an initial period of at most five years; *Mixed* indicates that at least 25% but less than 75% of all mortgages have variable interest rates; and *Fixed* indicates that less than 25% of all mortgages have variable interest rates after at most five years. Median Inflation is the median inflation in a given country from 2000 to present.

B ARM Sentiment & Experienced Inflation Rates

Figure A.1. ARM Sentiment Index and Experienced Past Inflation Rates

Panel A. Inflation over Past Two Years

Panel B. Inflation over Past Three Years

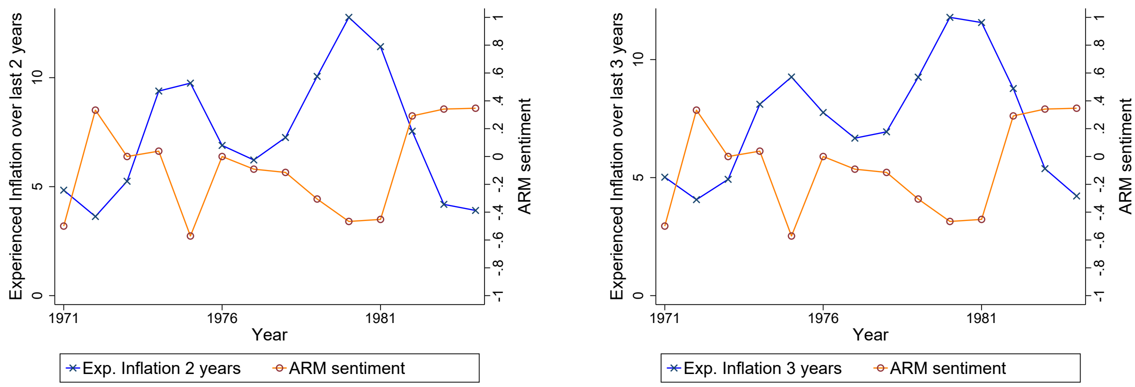

Notes. In Panel A, *"Experienced Inflation (past 2 years)"* is the weighted inflation rate over the past two years with the highest weight on the most recent observation, zero weight on the observation of t-2 and a linear connection of these endpoints. In Panel B, *"Experienced Inflation (past 3 years)"* is the weighted inflation rate over the past three years with the highest weight on the most recent observation, zero weight on the observation of t-3 and a linear connection of these endpoints. Both experienced inflation rate graphs are calculated according to the methodology of Equation (1). The data used for calculating annual inflation rates is obtained from the CPI-U of the BLS. The *ARM Sentiment Index* is calculated as the annual averages of the combined measure from the data set on the public discussion of ARMs. The information is gathered from all articles in *The Wall Street Journal* and *The Washington Post* that discuss variable-rate products from 1971 to 1984.

C Variable Definitions

SCF Variables (Federal Reserve Board)

Variable	Units	Description	SCF Source
Expect Higher i	$\{0, 1\}$	=1 if expects higher interest rates five years from now	X302=1
Expect Lower i	$\{0, 1\}$	=1 if expects lower interest rates five years from now	X302=2
Net expectation	$\{-1, 0, 1\}$	= Expect Higher i − Expect Lower i	X302
Married	$\{0, 1\}$	=1 if married	X8023=1
Age	years	Age of survey respondent	X8022
Total Income	const. year 2013 $	1989,1992: total HH income in survey year -1. 1995-2013: "normal" income (i.e., permanent income) in survey year -1. Bottom-coded to $1 in log specifications.	1989,1992: Summary Extract Data. 1995-2013: =X7362 if X7650 in (1,2), =X5729 o/w.
Net Worth	const. year 2013 $	HH net worth in survey year. Bottom-coded to $1 in log specifications.	Summary Extract Data
Home owner status	$\{0, 1\}$	0 if rents, 1 if owns (including ranch, farm, mobile home, house, condo), missing otherwise	same as Summary Extract SAS code (FRB website)
Has mortgage	$\{0, 1\}$	=1 if has a mortgage on primary or secondary residence (excluding land contracts).	Any of X723=1, X830=1, X1711=1, X1811=1
Refi status	$\{0, 1\}$	(First mortgage on primary residence only) 1989,1992: =1 if origin. year > purchase year. 1995-2013: =1 if reports taking out this loan to refinance a previous loan.	1989,1992: origin year is X802; purchase year is X606, X626, X630, X634, or X720. 1995-2013: X7137 in (1, 3).
Junior mortgage	$\{0, 1\}$	=1 if second mortgage (primary residence only)	X830=1
Non-conventional	$\{0, 1\}$	=1 if the first or main mortgage is federally guaranteed (includes FHA, VA, and "other programs").	X724=1
Loan amount	$	Original loan amount	X804, X904, X1714, X1814

RFS Variables (Census Bureau)

Variable	Units	Description
FRM Rate, ARM Initial Rate, ARM Margin	% or bps	Contractual interest rates charged to mortgage borrowers, top-and bottom-censored. 1991 RFS rates are also interval-censored; we code these to interval midpoints.
Total Income	const. year 2000 $	Real total household income in origination year. We impute total household income in Census year (1990 or 2000) back to origination year using peak-to-peak log growth rate in U.S. nominal median household income over 1980-2001 from CPS Historical Table H-6 (4.14% / year), then inflate to constant year 2000 dollars. For 1991 RFS, income is imputed back to interval midpoints (1985.5 for 1985-86, 1987.5 for 1987-88, and 1990 for 1989-91). Real income is bottom-coded to $1 in log specifications.

Variable	Units	Description
Age	years	Primary owner's age in origination year = age in survey year - (survey year - origination year). For 1991 RFS, age is coded to average within each origination year interval.
Joint owners	$\{0,1\}$	=1 if number of property owners exceeds one.
Rural county	$\{0,1\}$	=1 if property is located outside of an MSA.
Junior mortgage	$\{0,1\}$	=1 for second or third mortgage on a property.
Non-conventional	$\{0,1\}$	=1 if mortgage is FHA-, VA-, or FmHA/RHS-insured or guaranteed.
LTI ratio	fraction	Face amount of loan at origination / total household income in origination year. Ratio is symmetrically 1% Winsorized in pooled RFS sample of all FRM / ARM / balloon mortgages.
LTV ratio	fraction	Face amount of loan at origination / property value at origination (2001 RFS) or purchase price (1991 RFS). Ratio is symmetrically 1% Winsorized in pooled RFS sample of all FRM / ARM / balloon mortgages.
Loan / CLL	fraction	Face amount of loan at origination / Conforming Loan Limit for properties with same number of units. The CLL is updated every October. For 1991 RFS, we use the maximum CLL within each origination year interval (generally the last year). Ratio is symmetrically 1% Winsorized in pooled RFS sample of all FRM / ARM / balloon mortgages.
Jumbo loan	$\{0,1\}$	=1 if Loan / CLL > 1.
Points paid	% or bps	Discount points paid as interest at inception of first mortgage, excluding loan origination and non-interest fees.

Other Variables

Variable	Units	Description	Source
Inflation Experiences	%	Weighted average inflation (log change in annual average CPI-U) over respondent or primary owner's lifetime, using linearly decreasing weights starting from current year: for year $k \in [s,t]$, weight $w_k \propto k - s$, where s is the birth year and t is origination year. For the 1991 RFS, we use inflation experiences as of the first year in each origination year interval (1985, 1987, and 1989).	BLS CPI-U & Robert Shiller's website / authors' calculations
PMMS Index Rates	% or bps	Average rate on an FRM, or average first-year "teaser" rate on a 1/1 ARM, offered to a first-lien, prime, conventional, conforming mortgage borrower with an LTV of 80% and a 30-year term. Annual average of weekly data, re-weighted from five Freddie Mac regions to four Census regions using 1990 Census housing unit counts by state. We use the corresponding Freddie Mac regional rate if borrower's home state is reported, and the Census region rate otherwise.	Freddie Mac PMMS
CLL	$	Conforming Loan Limit	Fannie Mae

D Dating the Great Inflation

We determine the dates for the Great Inflation in a data-driven manner, proposed by Scrimgeour (2008). We first extract the trend component of BLS CPI-U log annual

Figure A.2

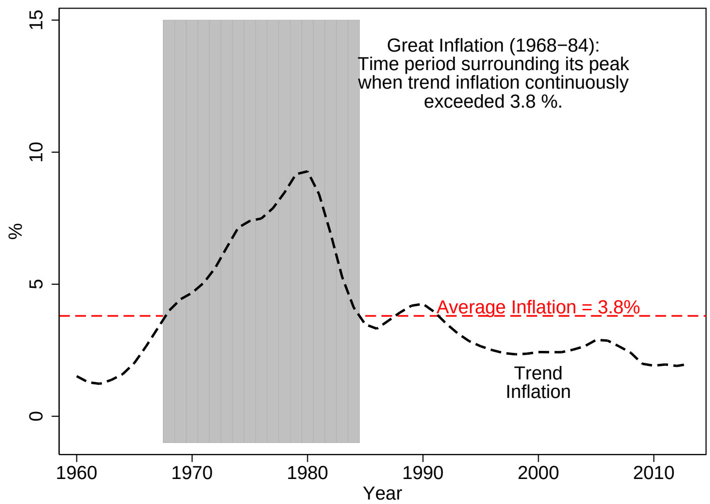

inflation using a triangular moving-average filter:

$$\pi_t^{trend} = \sum_{j=-h}^{h} \frac{h-|j|}{h^2} \pi_{t+h}, \tag{A.1}$$

with half-width $h = 4$ years. We then identify those years surrounding the mid-1970s when trend inflation *continuously* exceeded a pre-determined threshold, its 1960–2013 mean of 3.8%. This methodology determines that the U.S. Great Inflation began in 1968 and lasted through 1984. Scrimgeour (2008) calculates dates of 1969–1983 using the GDP deflator and a 4% threshold. Other authors suggest a starting dates as early as 1965; see the references cited in Scrimgeour.

E Comparison of Mortgage Rates in PMMS and MIRS

This overview of the PMMS and MIRS draws on summaries in Koijen et al. (2009) Appendix A, a 2019 Federal Register filing by the FHFA (84 Fed. Reg. 32,738), and Freddie Mac's website.[36]

Freddie Mac's PMMS is based on a representative nationwide survey of lenders (including thrifts, credit unions, commercial banks, and mortgage companies) collected Monday through Wednesday and released every Thursday. Lenders provide quotes for first-lien, conventional, non-jumbo, home purchase mortgages with an 80% LTV and a 30-year term that a prime borrower would receive that week, so the quotations hold borrower and loan characteristics constant both across products and over time. As of 2009 the PMMS included around 125 lenders per week; as of 2019 the sample size

[36]See http://www.freddiemac.com/pmms/about-pmms.html.

Figure A.3. FRM and ARM Rates in PMMS and MIRS

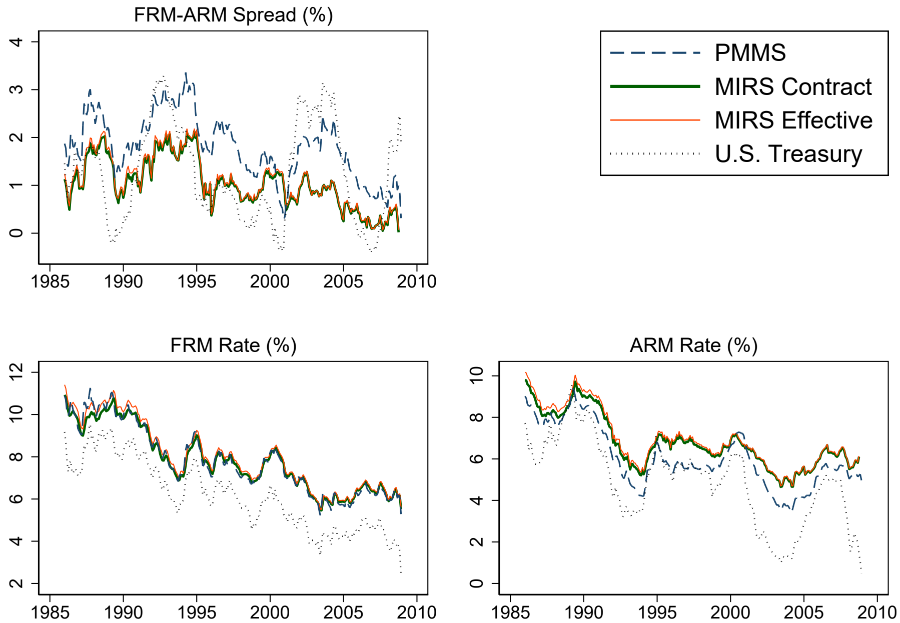

Notes. The figure shows monthly data from the FHFA's Monthly Interest Rate Survey (MIRS) and Freddie Mac's Primary Mortgage Market Survey (PMMS) (monthly average of weekly data), January 1986–October 2008. U.S. Treasury rates are 10-year minus 1-year spread (top left), 10-year rate (bottom left) and 1-year rate (bottom right).

was around 80 lenders per week. Interest rates are a weighted average based on lender size. The PMMS has added and subtracted products over time as the mortgage market has evolved, including the 30-year FRM since its inception (April 1971) and adding a 15-year FRM (August 1991–present); a 1/1 ARM (January 1984–December 2015); and a 5/1 "hybrid" ARM (January 2005–present). Data for five regions of the U.S. were also broken out through December 2015.

The MIRS was launched by the Federal Home Loan Bank Board (FHLBB) in the 1960s, then taken over by the Federal Housing Finance Board (FHFB) in 1989 and by the Federal Housing Finance Agency (FHFA) in 2008. Lenders provide information "on the terms and conditions on all conventional, single-family, fully amortized, purchase-money mortgage loans closed during the last five working days of the preceding month" (Federal Register 2019). Similar to PMMS, the MIRS excludes refinancings, FHA- or VA- insured or guaranteed loans, and multifamily properties; unlike PMMS, MIRS includes non-conforming (jumbo) loans. Koijen et al. (2009) report that the June 2006 MIRS had data from 74 lenders. By 2018 the sample size had shrunk to 20 per month. In May 2019 a single respondent accounting for more than half of the loans informed

FHFA that it was dropping out, leading to the survey's discontinuation. Breakouts of interest rates and other loan terms by property type, by lender type, and by region are available at various frequencies. Separate interest rate data for FRMs and ARMs are available between January 1986 and October 2008, when the FHFA stopped reporting ARM data due to insufficient observations. Averages were weighted by lender size and type through 2011, and unweighted starting in 2012.

This summary highlights at least three important differences. First, MIRS tends to track PMMS with a lag, since MIRS reflects originations of mortgages that had their rates quoted and locked in several months earlier. The FHFA's 2019 analysis found that an 11-week lag provides the best fit when constructing a transition index from MIRS to PMMS. Second, MIRS is a survey of originations and so reflects changing borrower and loan characteristics across products and over time (including term, LTV, and credit score) whereas PMMS attempts to hold these characteristics fixed. Third, MIRS includes "hybrid" ARMs with initial fixation periods longer than a year in its ARM summary data—e.g., rates on the 1/1 ARM and the 5/1 ARM are averaged together—while PMMS breaks these rates out separately. Hybrid ARMs will tend to carry higher initial rates than 1/1 ARMs since they provide insurance against rising interest rates for a longer initial time period.[37]

Table A.5: PMMS and MIRS Summary Statistics

Variable	Source	Mean	SD	10th Pctl.	90th Pctl.
FRM Rate	PMMS	7.85	1.58	5.92	10.27
FRM Rate, Contract	MIRS	7.84	1.47	6.03	10.05
FRM Rate, Effective	MIRS	8.03	1.58	6.09	10.43
ARM Rate	PMMS	6.04	1.46	4.18	8.39
ARM Rate, Contract	MIRS	6.79	1.26	5.33	8.84
ARM Rate, Effective	MIRS	6.94	1.34	5.40	9.07
FRM - ARM Spread	PMMS	1.81	0.66	0.88	2.71
FRM - ARM Spread, Contract	MIRS	1.05	0.52	0.37	1.84
FRM - ARM Spread, Effective	MIRS	1.09	0.53	0.40	1.89

Notes. The table reports summary statistics for the FRM and ARM initial rates and spreads reported in PMMS and MIRS, monthly averages over January 1986–October 2008. All variable are in percentage points.

[37]Between 2005 and 2015, when the PMMS reports both, the 5/1 ARM initial rate is 38 basis points higher on average.

Table A.5 reports that the average "contract" FRM rate in MIRS is very close to the PMMS, and that the average "effective" rate in MIRS that includes origination points and fees is only a little higher, about 18 basis points. By contrast, the average ARM rates in MIRS are 75 to 90 basis points higher than the average ARM rate in PMMS. Figure A.3 plots the rates and spreads over time; it is visually apparent that FRM-ARM spreads are consistently lower in MIRS, and that this difference is largely due to consistently higher initial ARM rates in MIRS. One possible explanation is the rise in popularity of hybrid ARMs in the early aughts, with initial fixation periods as long as 10 years (cf. Koijen et al. 2009, Figure 5); these carry higher rates and are included in MIRS but not PMMS. Were this the cause, we would expect the MIRS−PMMS difference to increase in the early aughts. Inspection of the data underlying Figure A.3 reveals a local maximum difference in ARM rates in 2004–5, consistent with this hypothesis, but also in 1992–4 and 1996–7, well before the explosion in popularity of hybrid ARMs. Moreover, Figure 5 in Koijen et al. (2009) also indicates a spike in hybrid ARM popularity in 2000, but in this year the MIRS ARM contract rate actually fell *below* the PMMS rate (bottom-right panel of Figure A.3). This suggests that hybrid ARMs are not the entire story. Despite these level differences, the series track each other quite closely. The correlation coefficients among the three spread series exceed 80% (Table A.6).

Table A.6: Correlations Among FRM-ARM Spreads in PMMS and MIRS and 10 Year−1 Year Treasury Yield Spread

	Treasury	PMMS	MIRS Contract	MIRS Eff.
10Y−1Y Treasury Yield Spread	1.000			
PMMS FRM-ARM	0.643	1.000		
MIRS FRM-ARM, Contract	0.411	0.819	1.000	
MIRS FRM-ARM, Effective	0.397	0.812	0.998	1.000

Notes. The table reports Pearson correlation coefficients among the U.S. Treasury 10-year minus 1-year yield spread and the FRM-ARM spreads from PMMS and MIRS, monthly averages over January 1986–October 2008.

Table A.6 also reveals that the PMMS tracks the U.S. Treasury yield curve much more closely than MIRS. The correlation between the 10-year minus 1-year Treasury yield spread and the PMMS FRM-ARM spread is 0.643,indicating that the yield spread explain $0.643^2 \approx 41\%$ of the variation in PMMS. By comparison, the yield spread only explains $0.411^2 \approx 17\%$ and $0.397^2 \approx 16\%$ of the MIRS spreads. (We round these figures to 40% and 15% in Section 1.) This suggests that non-interest factors drive the remaining variation in MIRS to a greater extent than PMMS, possibly due to changes in the borrower and loan pool changes over time.

To explore this further, we convert the data to annual averages and regress each mortgage spread on the Treasury yield spread and average borrower and loan characteristics in that origination year, from the most recent SCF wave starting with 1989. We cannot run monthly regressions because SCF suppresses the origination month in the public-use dataset. We exclude 2008 because we do not have a full year of MIRS data (the FHFA stopped reporting ARM data in November). Moreover, disruptions to the mortgage market brought by the financial crisis and the first round of quantitative easing, in which the Fed directly bought $1.7 trillion of mortgage-backed securities, may make 2008 unrepresentative. This leaves us with 22 complete years of data.

Table A.7: Determinants of the National Annual Average FRM-ARM Spread, 1986–2007

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable is	PMMS Spread (%)		MIRS Contract Spread (%)		MIRS Effective Spread (%)	
10Y - 1Y Treasury Yield Spread (%)	0.384***	0.560***	0.176***	0.254**	0.172***	0.253**
	(0.04)	(0.11)	(0.05)	(0.08)	(0.05)	(0.08)
Inflation Experiences (%)	0.586***	0.111	0.592***	0.410**	0.634***	0.453**
	(0.10)	(0.20)	(0.09)	(0.18)	(0.09)	(0.18)
Mortgage controls		YES		YES		YES
Sociodemographic controls		YES		YES		YES
Observations	22	22	22	22	22	22
R^2	0.771	0.938	0.695	0.935	0.715	0.936
[p-value: Mtg. controls all 0]		[0.791]		[0.064]		[0.081]
[p-value: Sociodem. controls all 0]		[0.205]		[0.058]		[0.067]

Notes. The table reports OLS coefficient estimates of the determinants of the FRM-ARM initial rate spread reported in PMMS and MIRS, annual averages over 1986–2007. Inflation experiences and controls are annual averages by origination year in the most recent SCF wave (1989, ..., 2007) using "revised consistent" sampling weights (X42001). We adjust for multiple imputation following Rubin (1987). Mortgage controls ($K_1 = 5$) are Refi dummy, Junior Mortgage dummy, Non-conventional dummy, Loan / CLL, and Jumbo dummy. Sociodemographic controls ($K_2 = 5$) are log(Income), log(Net Worth), Age, Age^2, and Married dummy. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

We report regression results in Table A.7. For each mortgage spread, we regress the spread on the Treasury yield spread and average borrower inflation experiences in that origination year, plus the same ten controls for average mortgage and borrower characteristics from SCF as in Table 4 (leaving $22 - 12 - 1 = 9$ residual degrees of freedom).

The table shows that the yield spread is a significant predictor of the mortgage rate spread in all specifications, however, its marginal effect on the PMMS spread is about twice as large as its effect on the MIRS spreads. E.g., without additional controls a one percent steepening of the yield curve increases the PMMS FRM-ARM spread by 0.38 p.p.s (column 1) versus 0.18 and 0.17 p.p.s for the MIRS FRM-ARM spreads (columns 3 and 5). With additional controls, the yield spread has an even larger effect on all three spreads, but again the effect size is about twice as large on PMMS: 0.56 p.p.s in column 2 versus 0.25 p.p.s in columns 4 and 6.

The table also shows that the mortgage and borrower controls are jointly significant predictors of the MIRS spreads at a 10% level, but they do not have a significant impact on the PMMS spread. We find this a reassuring confirmation that the design of the PMMS questionnaire is succeeding at holding these characteristics constant both across products and over time.[38]

Finally, the table shows that average borrower inflation experiences have a significant and powerful effect on average national mortgage loan pricing: when borrowers have lived through on average one percentage point higher inflation, the gap between FRMs and ARMs rises by between 0.4 and 0.6 p.p.s depending on the specification (columns 1 and 3–6). This is consistent with higher inflation experiences raising borrower demand for FRMs; as the demand curve for FRMs shifts right, both the price and the quantity of FRMs increase. The one exception is the PMMS regression with additional mortgage and borrower controls (column 2), in which the effect of inflation experiences is indistinct from zero. This is consistent with the idea that PMMS asks for quotes to the same representative borrower over time.

For robustness, we also tried running the PMMS specification in column 2 over the full 1984–2015 period for which we have data, even though we cannot make a direct comparison to MIRS. The results confirm that the Treasury yield spread remain a significant predictor of the mortgage rate spread ($\hat{\beta} = 0.406$, S.E. $= 0.12$, $p < 0.01$) and inflation experiences remain insignificant in the presence of controls ($\hat{\beta} = 0.369$, S.E. $= 0.27$, $p = 0.19$). The additional mortgage and borrower controls from the SCF remain jointly insignificant.

Table A.7 reports Huber-Eicker-White heteroskedasticity robust standard errors.

[38] In terms of individual coefficients: only log(net worth) has a marginally significant, negative impact on the PMMS spread. In columns 4 and 6: the spread is lower in years with more married couples ($p < 0.05$) and more junior loans ($p < 0.10$); and it is higher in years with more nonconventional loans ($p < 0.01$).

We also re-ran all tests using Newey-West (1987) standard errors that are robust to serially-correlated errors. Following the guidance of Lazarus et al. (2018), we set the maximum lag truncation parameter to $1.3T^{1/2} \approx 6$ and used the "fixed-b" asymptotic critical values of Kiefer and Vogelsang (2005).[39] These are larger than conventional t critical values to adjust for the tendency of Newey-West tests to overreject the null, particularly when the true data generating process is not highly correlated.[40] The intuition is that using a longer maximum lag length reduces the bias of the covariance matrix estimator but increases the size distortion, so non-standard, larger critical values are necessary to control the size of the test. E.g., the two-sided 5% t_{19} critical value is 2.093 and t_9 critical value is 2.262, versus 2.943 using fixed-b asymptotics with $b = (6+1)/22 = 31.8\%$ ($6+1$ is the maximum number of time-series lags plus the current observation). The Newey-West results are very similar to our baseline results in Table A.7. In particular, the coefficient on inflation experiences remains significant in all specifications except column 2; the additional controls in column 2 remain jointly insignificant (with slightly larger p-values), while those in columns 4 and 6 remain jointly marginally significant (with slightly smaller p-values).[41]

This analysis highlights a second cost of overweighting personal exposure to higher inflation. Our main analysis focuses on consumer choice and holds relative prices fixed, which makes sense from the individual household level. But at the aggregate level, if a large fraction of borrowers are demanding more FRMs, this has supply implications: FRM prices will increase relative to ARMs. I.e., demanding extra insurance against future interest rate increases is a particularly costly mistake when many other individuals are making the same mistake.[42]

[39] Tim Vogelsang provides Stata programs `neweyfixedb` and `testfixedb` to implement the tests on his website: `https://sites.google.com/view/tim-vogelsang-msu/code`.

[40] For Monte Carlo evidence, see den Haan and Levin (1997) Table 3 and Kiefer and Vogelsang (2005) Figure 1. Jansson (2004) provides a theoretical explanation for the error in rejection probability of Newey-West tests. Müller (2014) surveys the literature.

[41] It is not surprising that this adjustment does not make a difference. Although this is a time series setting, we do not expect the population errors to exhibit serial correlation. Even if all the spread series are serially correlated, as we expect they are, the *difference* between the FRM-ARM spread and the Treasury spread should be unpredictable, white noise.

[42] This also suggests a second plausible counterfactual in Section 6: the fraction of switching households in our WRTE calculation might also affect the counterfactual mortgage payments that the switching households face. This would be particularly pressing if we used the MIRS rates as our baseline mortgage rates in Step 2, but Table A.7 indicates it is a less pressing concern with PMMS rates.

F Empirical Framework Derivations

In section 4.2, we argue that we can place bounds on the true effect size of the individual inflation forecast on mortgage choice in equation (11) by comparing $\hat{\delta}_{1,OLS}$ and $\hat{\delta}_{1,IV}$. We show this more carefully here.

Recall first that the just-identified (for ease of exposition) IV estimator of $y_i = x_i'\beta + \upsilon_i$ with positively-correlated instruments z_i takes the form

$$\hat{\beta}_{IV} = \left(\frac{1}{N}\sum_i z_i x_i'\right)^{-1}\left(\frac{1}{N}\sum_i z_i y_i\right) \xrightarrow{p} \beta + \left(\mathbb{E}[z_i x_i']\right)^{-1}\mathbb{E}[z_i \upsilon_i]\,. \tag{A.2}$$

In OLS, the instruments and regressors are the same, $z_i \equiv x_i$. If $\mathbb{E}[z_i\upsilon_i] = 0$, then the estimator is consistent for β. Otherwise, if the moment expression $\mathbb{E}[z_i\upsilon_i]$ is positive (negative), the probability limit is greater than (less than) the structural parameter β.

Preliminary Result: Non-classical measurement error ι. The source of endogeneity in our empirical framework is measurement error in the interest rate forecast ι. This arises due to a timing discrepancy: interest rate expectations are observed after the mortgage is taken out instead of contemporaneously. Because expectations are serially correlated, the measurement error term is "non-classical" and is negatively correlated with the level of the initial forecast: $\mathbb{E}[\iota_{n,t}\Delta\iota_{n,t+1}] < 0$. We show this first, as a preliminary result; then we use this to analyze the probability limits of the OLS and IV estimators given by (A.2).

Letting L be the lag operator, observe that personal inflation experiences (1) are obtained by applying an absolutely-summable linear filter $w_{s,t}(L) = \left(\sum_{j=0}^{t-s}(t-s-j)\right)^{-1} \cdot \sum_{k=0}^{t-s}(t-s-k)L^k$ to a stationary AR(1) process: $\pi^e_{n,t} = w_{s,t}(L)\pi_t$. In fact, $\pi^e_{n,t}$ is a weighted sample mean, and it follows that $\mathrm{Var}(\pi^e_{n,t+1}) < \mathrm{Var}(\pi^e_{n,t})$, $\lim_{n\to\infty}\mathrm{Var}(\pi^e_{n,t}) = 0$, and $\pi^e_{n,t} \xrightarrow{p} \mathbb{E}[\pi]$.

Further, inflation experiences may be written recursively as

$$\pi^e_{n,t+1} = \phi_{t+1}\pi^e_{n,t} + (1-\phi_{t+1})\pi_{t+1}\,, \tag{A.3}$$

$$\text{so } \Delta\pi^e_{n,t+1} = (\phi_{t+1}-1)(\pi^e_{n,t} - \pi_{t+1})\,, \tag{A.4}$$

where $\phi_{t+1} = t/(t+2)$ for an individual born at $s = 0$.

Equation (10) states that the *ex post* interest rate forecast differs from the *ex ante* forecast by three factors: the evolution of personal inflation experiences $\Delta\pi^e$, the mean-reverting component of the forecast error term ξ, and a new forecast innovation ν that is white noise. Given this, the change in the interest rate forecast $\Delta\iota$ in (9) is not pure white noise:

$$
\begin{aligned}
\mathbb{E}[\iota_{n,t}\Delta\iota_{n,t+1}] &= \mathbb{E}[(\alpha_{0,t}+\alpha_1\pi^e_{n,t}+\xi_{n,t})(\alpha_1\Delta\pi^e_{n,t+1}+(\varphi-1)\xi_{n,t}+\nu_{t+1})] \\
&= \alpha_{0,t}\left(\alpha_1\underbrace{\mathbb{E}[\Delta\pi^e_{n,t+1}]}_{=0\text{ by (A.4)}}+\underbrace{\mathbb{E}[(\varphi-1)\xi_{n,t}+\nu_{t+1}]}_{=0\text{ by (5) and (6)}}\right) \\
&\quad+\alpha_1\left(\alpha_1\mathbb{E}[\pi^e_{n,t}\Delta\pi^e_{n,t+1}]+(\varphi-1)\underbrace{\mathbb{E}[\pi^e_{n,t}\xi_{n,t}]}_{=0\text{ by (8)}}+\underbrace{\mathbb{E}[\pi^e_{n,t}\nu_{t+1}]}_{=0\text{ by (6)}}\right) \\
&\quad+\alpha_1\underbrace{\mathbb{E}[\xi_{n,t}\Delta\pi^e_{n,t+1}]}_{\substack{=(1-\phi_{t+1})\mathbb{E}[\xi_{n,t}\pi_{t+1}]\text{ by (A.4) and (8)}\\ =0\text{ by (4), (6), and (8)}}}+(\varphi-1)\mathbb{E}[\xi^2_{n,t}]+\underbrace{\mathbb{E}[\xi_{n,t}\nu_{t+1}]}_{=0\text{ by (6)}} \\
&= \alpha_1^2\underbrace{\mathbb{E}[\pi^e_{n,t}\Delta\pi^e_{n,t+1}]}_{<0}+(\varphi-1)\mathbb{E}[\xi^2_{n,t}]<0\,. \qquad \text{(A.5)}
\end{aligned}
$$

By the Cauchy-Schwarz inequality,

$$
Cov(\pi^e_{n,t+1},\pi^e_{n,t}) \le SD(\pi^e_{n,t+1})SD(\pi^e_{n,t}) < SD(\pi^e_{n,t})^2 = Var(\pi^e_{n,t})\,. \qquad \text{(A.6)}
$$

Add $(\mathbb{E}[\pi^e])^2$ to both sides to get that $\mathbb{E}[\pi^e_{n,t}\pi^e_{n,t+1}] < \mathbb{E}[(\pi^e_{n,t})^2]$, or $\mathbb{E}[\pi^e_{n,t}\Delta\pi^e_{n,t+1}] < 0$. So, (A.5) is negative, as we have claimed.

For further intuition, use (A.4) to rewrite the final line of (A.5) as

$$
\mathbb{E}[\iota_{n,t}\Delta\iota_{n,t+1}] = (\phi_{t+1}-1)\alpha_1^2\mathbb{E}[(\pi^e_{n,t})^2-\pi^e_{n,t}\pi_{t+1}]+(\varphi-1)\mathbb{E}[\xi^2_{n,t}]\,.
$$

If it were the case that $\phi_{t+1}=1$ and $\varphi=1$, then the interest rate forecast would be a random walk with white noise innovations ν and $\mathbb{E}[\Delta\iota_{t+1}\iota_t]=0$. However, because $0\le\varphi<1$ and $0\le\phi_{t+1}<1$, the forecast ι is serially correlated, so the change in the forecast $\Delta\iota$ is negatively correlated with the level of the lagged forecast.

Result 1: Probability limit of OLS. Suppose that we were to ignore the presence of measurement error and naïveley run an OLS-like regression of (11). Consistency of this estimator relies upon the orthogonality condition $\mathbb{E}[\iota_{n,t+1}u^*_{n,t}]=0$. Expanding the moment expression gives us

$$
\begin{aligned}
\mathbb{E}[\iota_{n,t+1}u^*_{n,t}] &= \mathbb{E}[\iota_{n,t+1}u_{n,t}]-\delta_1\mathbb{E}[\iota_{n,t+1}\Delta\iota_{n,t+1}] \\
&= \underbrace{\alpha_1}_{>0}\underbrace{\mathbb{E}[\pi^e_{n,t+1}u_{n,t}]}_{\text{ambig.}}-\underbrace{\delta_1}_{>0}\underbrace{\left(\mathbb{E}[\iota_{n,t}\Delta\iota_{n,t+1}]+\mathbb{E}[\Delta\iota^2_{n,t+1}]\right)}_{\ge 0}\neq 0 \qquad \text{(A.7)}
\end{aligned}
$$

in general. (We show how to derive the second line below.) For comparison, in a classical errors-in-variables setting, the first and second expectation terms in the final line would drop out, leaving $\mathbb{E}[\iota_{n,t+1}u^*_{n,t}] = -\delta_1\mathbb{E}[\Delta\iota^2_{n,t+1}] < 0$, since $\delta_1 > 0$. So the mis-measured regressor ι_{t+1} would be negatively correlated with the composite error term, leading to attenuation bias in the coefficient estimates.

However, (A.7) contains two additional terms. First, since $\mathbb{E}[\iota_{n,t}\Delta\iota_{n,t+1}] < 0$, adding this term attenuates the attenuation bias. Using the Cauchy-Schwarz inequality and the fact that ι is serially-correlated, the entire expression in parentheses remains positive, so the magnitude of bias is smaller but the direction remains negative. Second, $\mathbb{E}[\pi^e_{n,t+1}u_{n,t}]$ will be nonzero if $\mathbb{E}[\pi_{t+1}u_{n,t}] \neq 0$, i.e., if future inflation is predictable using any unobserved or omitted factors affecting current mortgage choice. Empirically, the nationwide FRM share in year t is negatively correlated with inflation in year $t+1$.[43] This would suggest $\mathbb{E}[\pi^e_{n,t+1}u_{n,t}] < 0$, and since $\alpha_1 > 0$, makes the entire expression (A.7) more negative and the attenuation bias more severe.

Derivation of the OLS moment expression. To derive the second line of (A.7), we use (2) and (5) to expand $\iota_{n,t+1} = \alpha_{0,t+1} + \alpha_1\pi^e_{n,t+1} + (\varphi\xi_{n,t} + \nu_{n,t+1})$ and simplify the first right-hand term:

$$\begin{aligned}\mathbb{E}[\iota_{n,t+1}u_{n,t}] &= \alpha_{0,t+1}\underbrace{\mathbb{E}[u_{n,t}]}_{=0 \text{ by } (6)} + \alpha_1\mathbb{E}[\pi^e_{n,t+1}u_{n,t}] + \varphi\underbrace{\mathbb{E}[\xi_{n,t}u_{n,t}]}_{=0 \text{ by } (7)} + \underbrace{\mathbb{E}[\nu_{n,t+1}u_{n,t}]}_{=0 \text{ by } (6)} \\ &= \alpha_1\mathbb{E}[\pi^e_{n,t+1}u_{n,t}]\,, \end{aligned} \tag{A.8}$$

then rewrite the second right-hand term as $\mathbb{E}[\iota_{n,t+1}\Delta\iota_{n,t+1}] = \mathbb{E}[(\iota_{n,t} + \Delta\iota_{n,t+1})\Delta\iota_{n,t+1}]$.

Result 2: Probability limit of IV. Now, suppose we were to address the endogeneity between $\iota_{n,t+1}$ and $u^*_{n,t}$ via instrumental variables. A common empirical technique in rational expectations models where the researcher only observes an *ex post* outcome is to use lagged values of variables as instruments (e.g., Hall (1988) on consumption, Yogo (2004) on real interest rates). The structure of our two-equation model (2) and (3) suggests such an instrument. Contemporaneous, time-t inflation experiences $\pi^e_{n,t}$ are correlated with the contemporaneous, time-t interest rate forecast $\iota_{n,t}$, and they have no direct effect on mortgage choice, except through their impact on an individual's interest rate forecast. However, we need the instrument to be orthogonal not only to the structural error term u in (3), but to the composite error term u^* in

[43] Using the most recent SCF survey to calculate the FRM share in every origination year between 1987–2013, $\rho = -0.72$.

the feasible regression equation (11). Expanding the exogeneity moment expression,

$$\begin{aligned}\mathbb{E}[\pi^e_{n,t}u^*_{n,t}] &= \mathbb{E}[\pi^e_{n,t}u_{n,t}] - \delta_1\mathbb{E}[\pi^e_{n,t}\Delta\iota_{n,t+1}] \\ &= \quad 0 \quad - \underbrace{\delta_1\alpha_1}_{>0}\underbrace{\left(\mathbb{E}[\pi^e_{n,t}\pi^e_{n,t+1}] - \mathbb{E}[(\pi^e_{n,t})^2]\right)}_{<0} > 0\,. \end{aligned} \tag{A.9}$$

$\delta_1 > 0$ and $\alpha_1 > 0$ by economic theory, and the sign of the term in parentheses is found by applying the Cauchy-Schwarz inequality, as we show below.

Derivation of the IV moment expression. To derive the second line of (A.9), consider in turn the two terms on the right-hand side of the first line. Theory tells us that contemporaneous inflation experiences have no direct effect on mortgage choice, except through their impact on an individual's interest rate forecast:

$$\mathbb{E}[\pi^e_{n,t}u_{n,t}] = \mathbb{E}[\alpha_1^{-1}(\iota_{n,t} - \alpha_{0,t} - \xi_{n,t})u_{n,t}] = 0 \tag{A.10}$$

by (6), (7), and (8). It remains to evaluate

$$-\delta_1\mathbb{E}[\pi^e_{n,t}\Delta\iota_{n,t+1}] = -\delta_1\mathbb{E}[\pi^e_{n,t}(\alpha_1\Delta\pi^e_{n,t+1} + (\varphi-1)\xi_{n,t} + \nu_{t+1})] \tag{A.11}$$

$$= -\delta_1\left(\alpha_1\mathbb{E}[\pi^e_{n,t}\Delta\pi^e_{n,t+1}] + (\varphi-1)\underbrace{\mathbb{E}[\pi^e_{n,t}\xi_{n,t}]}_{=0\text{ by }(7)} + \underbrace{\mathbb{E}[\pi^e_{n,t}\nu_{t+1}]}_{=0\text{ by }(6)}\right) \tag{A.12}$$

$$= \underbrace{-\delta_1\alpha_1}_{<0}\underbrace{\left(\mathbb{E}[\pi^e_{n,t}\pi^e_{n,t+1}] - \mathbb{E}[(\pi^e_{n,t})^2]\right)}_{<0} > 0\,. \tag{A.13}$$

$\delta_1 > 0$ and $\alpha_1 > 0$ by economic theory, and the sign of the term in parentheses is found by the Cauchy-Schwarz inequality (A.6). This gives the second line of (A.9).

Result 3: Feasible first-stage regression. Plugging (9) into (2) gives

$$\iota_{n,t+1} = \alpha_{0,t} + \alpha_1\pi^e_{n,t} + \underbrace{(\xi_{n,t} + \Delta\iota_{n,t+1})}_{\xi^*_{n,t}} \tag{A.14}$$

The feasible first-stage regression (A.14) adds measurement error to the dependent variable rather than an independent variable. If the interest rate forecast were a random walk, then $\Delta\iota$ would be random noise and OLS would be consistent.

The OLS orthogonality expression for the first-stage feasible regression (A.14) is

$$\begin{aligned}\mathbb{E}[\pi^e_{n,t}\xi^*_{n,t}] &= 0 + \mathbb{E}[\pi^e_{n,t}\Delta\iota_{n,t+1}] \\ &= \left(\mathbb{E}[\pi^e_{n,t}\pi^e_{n,t+1}] - \mathbb{E}[(\pi^e_{n,t})^2]\right) < 0\,, \end{aligned} \tag{A.15}$$

by the same argument as equations (A.11) to (A.13). So, the probability limit of the OLS estimator of α_1 in (A.14) is attenuated, as we saw in the "All Mortgages" versus "New Mortgages" columns of Table 4. Moreover, this is easily resolved by using the "correct" regressor $\pi^e_{n,t+1}$ so there is no timing discrepancy.

G Methodology in Detail

G.1 Estimation Methodology

Our key prediction is that relatively high lifetime experiences of inflation are a significant factor in explaining the tilt in mortgage financing toward fixed-rate contracts. As the main estimation approach we utilize a discrete choice model over mortgage products using a three-step procedure suggested by Lee (1978) and Brueckner and Follain (1988):

1. Estimate a reduced-form model of mortgage choice using only *exogenous* explanatory variables (equation (18)).
2. Predict FRM and ARM mortgage rates at the household level, correcting for selection bias (equation (17)).
3. Estimate a structural model of mortgage choice using individual-level predicted mortgage rates (equation (15)).

We begin by assuming that a household n derives utility $U_{n,i} = x'_{n,i}\beta_i + \upsilon_{n,i}$ when choosing alternative i from a menu of J alternatives, $i \in \{FRM,\ ARM,\ Balloon\}$, depending on observed components $x'_{n,i}\beta_i$ and unobserved components $\upsilon_{n,i}$. Each household lives in Census region r and chooses a mortgage only once, in year y (unless they take a junior mortgage), so we omit time subscripts for notational simplicity. Observed components may include attributes of the alternative, such as its cost, as well as household characteristics that sway the decision toward one alternative. The latter includes our variable of interest, namely past lifetime experiences such as living through the Great Inflation. Alternative i is chosen by household n if

$$\begin{aligned} D_{n,i} &\coloneqq \mathbb{I}\{U_{n,i} > U_{n,j} \quad \forall j \neq i\} \\ &= \mathbb{I}\{\upsilon_{n,j} - \upsilon_{n,i} < x'_{n,i}\beta_i - x'_{n,j}\beta_j \quad \forall j \neq i\} \end{aligned} \tag{A.16}$$

equals 1.[44] Marley (cited by Luce and Suppes (1965)) and McFadden (1974) show necessary and sufficient conditions on the distribution of the unobserved utility components υ_{ni} for the implied choice probabilities $\Pr(D_{n,i} = 1) = F(x'_{n,i}\beta_i - x'_{n,1}\beta_1, \ldots, x'_{n,i}\beta_i - x'_{n,J}\beta_J)$ to be described by a logit formula. This likelihood function is globally concave in β, so that the utility parameters can be estimated by maximum likelihood (up to

[44]Since utility is continuous, ties are of probability zero and are broken at random.

scale).[45] The definition of $D_{n,i}$ implies that, if explanatory variables do not vary across alternatives within household ($x_{n,i} = x_n \ \ \forall i$), as is the case for sociodemographic characteristics, then β_i can only be estimated for $J - 1$ of the J alternatives. We normalize $\beta_{\cdot,ARM} \equiv 0$ for all sociodemographic characteristics, including experienced inflation.

Theoretically, the mortgage product preferred by a household depends on a host of demographics and proxies for risk attitudes, including age, mobility, current and expected future income, risk aversion, and beliefs about future short-term interest rates (see, among others, Stanton and Wallace (1998), Campbell and Cocco (2003), Chambers et al. (2009), and Koijen et al. (2009)). Our main observable characteristics are the alternative-specific interest rate offered to the borrower, $Rate_{n,i}$; the borrower's (log) income, $Income_n$; and an alternative-specific function of the borrower's age, $f_i(Age_n)$. Our baseline age specification is quadratic, to capture possibly non-linear life-cycle variation in the attractiveness of a given mortgage contract type. The explanatory variable of interest is borrower n's lifetime experience of inflation at the time of the choice situation, π^e_n. We obtain the following estimating equation ((15) in the paper):

$$U_{n,i} = \beta_{0,i,y} + \beta_{R,i} Rate_{n,i} + \beta_{\pi,i}\pi^e_n + \beta_{Inc,i} Income_n + f_i(Age_n) + \upsilon_{n,i}\,, \tag{A.17}$$

with the error term capturing any unobservables. Since each borrower is only observed once, we omit the time subscripts on all borrower characteristics, even though some characteristics such as income are time-varying. Note that our model includes alternative-specific year fixed effects $\beta_{0,i,y}$. These control for the desirability of a given alternative in a given year. They capture all aspects of the economic environment at a given time and all information that is common to all households and might enter the rational-expectations forecast, including the full history of past inflation. They are essential for the interpretation of our coefficient of interest, $\beta_{\pi,i}$. In the presence of year fixed effects, a borrower's lifetime inflation experiences should not matter unless there is a correspondence between those experiences and borrower beliefs that differs from the baseline rational-expectations forecast. Specifically, the experience-effect hypothesis implies $\beta_{\pi,FRM} > 0$, while the standard rational framework predicts $\beta_{\pi,FRM} = 0$.

The main estimation difficulty is that the interest rates of the non-chosen alternatives are not observed. If households were randomly assigned to mortgage types, we could simply estimate the correlation between borrower characteristics and interest rates using the subsample of borrowers who chose each alternative. Specifically, we

[45] That is, the ratios of utility slope coefficients are identified, but the levels are not. We follow the usual practice of normalizing the variance of the υ's to $\pi^2/6$ before estimating the coefficients.

would use the subset of households n choosing alternative i to estimate the following equation ((17) in the paper) for all J all alternatives:

$$\begin{aligned} Rate_{n,i} &= \gamma_{0,i} + Z'_{n,i}\Gamma_{n,i} + \zeta_{n,i} \\ &= \gamma_{0,i} + \gamma_{R,i} PMMSRate_{y,r,i} + z'_n\gamma_i + \zeta_{n,i} \,. \end{aligned} \tag{A.18}$$

The equation decomposes the the explanatory variables $Z_{n,i}$ into $(PMMSRate_{y,r,i}, z'_n)'$, where the Freddie Mac survey rate $PMMSRate_{y,r,i}$ represents the baseline price charged to a high-quality borrower in the same year y and Census region r as borrower n, taking out mortgage product i; and the other explanatory variables z_n control for household-varying risk proxies such as income, first-time homeowner status, marital status, urban/rural property location, and loan size. The specification includes the same controls in each rate equation but allow them to have different slope coefficients γ_i. The error term $\zeta_{n,i}$ captures all remaining, unobserved factors that affect the interest rate for alternative i being offered to household n.

The goal of estimating equation (A.18) is to predict interest rates for households who did not choose product i. However, since households were not randomly assigned to mortgage types, OLS will likely be inconsistent due to selection bias. Specifically, households might have been offered an unusually low rate for the alternative they chose, so we expect the mean pricing error to be negative rather than zero: $\mathbb{E}\left[\zeta_{n,i}|Z_{n,i}, D_{n,i} = 1\right] = f(Z_{n,i}) < 0$. Our estimation must account for a correlation between the explanatory variables $Z_{n,i}$ and factors affecting sample selection. Otherwise our out-of-sample predictions will also be biased and inconsistent.

An additional wrinkle is that mortgage rates are top-coded in the public-use RFS files (at 14.1% in the 1991 survey and at 20% in 2001), and censoring of the dependent variable leads to inconsistent OLS estimators. Moreover, parametric methods such as Tobit do not perform well in the presence of non-normal errors. Powell (1984) first observed that estimators based on a conditional *median* restriction $\mathbb{E}\left[\text{sgn}(\zeta_{n,i})|Z_{n,i}\right] = 0$, rather than the usual conditional mean restriction $\mathbb{E}\left[\zeta_{n,i}|Z_{n,i}\right] = 0$, are robust to top- and bottom-censoring of the dependent variable, without further assumptions on the distribution of the errors. We thus use a censored least absolute deviations (CLAD) estimator as our benchmark estimator of equation (A.18).

Although our coefficient estimates from (A.18) do not provide us directly with predicted rates, we can plug them into (A.17) and obtain a *reduced-form* choice model that we can estimate ((18) in the paper):

$$\begin{aligned} U_{n,i} &= \tilde{x}'_{n,i}\tilde{\beta} + \tilde{\upsilon}_{n,i} \\ &= \tilde{\beta}_{0,i,t} + \tilde{\beta}_{R,i} PMMSRate_{y,r,i} + \beta_{\pi,i}\pi^e_n + \tilde{\beta}_{Inc,i} Income_n + f_i(Age_n) + \tilde{z}'_n\tilde{\gamma}_i + \tilde{\upsilon}_{n,i} \,. \end{aligned} \tag{A.19}$$

We place tildes on coefficients and variables that represent different objects than in equation (A.17). For example, the coefficient on the PMMS rate in equation (A.19) is the structural coefficient from equation (A.17), scaled by the partial correlation between household interest rates and PMMS rates from equation (A.18): $\tilde{\beta}_{R,i} := \beta_{R,i}\gamma_{R,i}$. We write $\tilde{z}_n$ to represent the subset of variables in z_n from equation (A.18) that do not appear directly in (A.17) (e.g., excluding household income). The pricing errors from (A.18), ζ_{ni}, are absorbed into the unobserved component of latent utility: $\tilde{\upsilon}_{n,i} := \upsilon_{n,i} + \beta_{R,i}\zeta_{n,i}$.

The important takeaway is that we have eliminated the missing data problem by replacing household-level interest rates $Rate_{n,i}$ with the Freddie Mac survey rates $PMMSRate_{y,r,i}$, which do not depend on an individual household's characteristics and are always observed for both alternatives. Moreover, since lifetime inflation experiences do not appear in equation (A.18), we can consistently estimate the structural coefficient $\beta_{\pi,i}$ in the reduced-form choice model.

We now have all of the pieces in hand to run our three-step estimator and obtain structural mortgage choice estimates. We work backward, estimating (A.19) first, (A.18) second, and (A.17) third. Model (A.19) can be consistently estimated by standard maximum likelihood methods, since it only depends on exogenous characteristics that are observed for all households. We then use the predicted choice probabilities to correct for any selection bias in the FRM and ARM rate equations (A.18) semiparametrically. Specifically, let $\tilde{\eta}_{n,i,j} := \tilde{x}'_{n,i}\tilde{\beta}_i - \tilde{x}'_{n,j}\tilde{\beta}_j$ denote the difference in the observed components of utility for the i^{th} and j^{th} alternatives. We can decompose the rate equation error in equation (A.18) as

$$\begin{aligned} \zeta_{n,i} &= \mathbb{E}\left[\zeta_{n,i} \mid Z_{n,i}, D_{n,i} = 1\right] + w_{n,i} = \mathbb{E}\left[\zeta_{n,i} \mid Z_{n,i}, \tilde{\upsilon}_{n,j} - \tilde{\upsilon}_{n,i} < \tilde{\eta}_{n,i,j} \; \forall j \neq i \,\right] + w_{n,i} \\ &= g(\tilde{\eta}_{n,i,1}, \ldots, \tilde{\eta}_{n,i,J}) + w_{n,i} \,, \end{aligned} \tag{A.20}$$

where $w_{n,i}$ is a mean-zero error that is independent of $(Z'_{n,i}, D_{n,i})'$. This decomposition states that, conditional on selection, the mean of the pricing error depends on $Z_{n,i}$ only through the $J-1$ choice indices $\tilde{\eta}_{n,i,1}, \ldots, \tilde{\eta}_{n,i,J}$.

Newey (2009) analyzes the case $J = 2$ and suggests a semiparametric selection correction (SPSC) estimator that uses a series approximation for the selection bias term: $g(\tilde{\eta}_{n,i,j}) \approx \sum_{k=0}^{K} \tau_k \cdot p(\tilde{\eta}_{n,i,j})^k$, where $p(\cdot)$ is some function, and τ_k is the coefficient on the k^{th} polynomial term. Consistency of the two-step series estimator requires that the order K of the approximating power series grows with sample size N according to $K = o(N^{1/7})$. Plugging the approximation terms into equation (17), we obtain

$$Rate_{n,i} \approx \gamma_{R,i} PMMSRate_{y,r,i} + z_n' \gamma_i + \sum_{k=0}^{K} \tau_k \cdot p(\tilde{\eta}_{n,i,j})^k + w_{n,i}\,. \tag{A.21}$$

In the special case where $K = 1$ and $p(\cdot)$ is the inverse of Mill's ratio, equation (A.21) is the familiar Heckman (1979) two-step selection model. Newey (2009) establishes the consistency and root-N asymptotic normality of this semiparametric, two-step series estimator $\hat{\Gamma}_{n,i}$ when $K \to \infty$, without requiring joint normality of the pricing and selection equation errors.

Note that specification (A.21) drops the intercept $\gamma_{0,i}$ from (A.18) since the series approximation includes a possibly non-zero constant (for $k = 0$). Thus, unlike in Heckman's two-step model, the model intercept $\gamma_{0,i}$ is not separately identified from the selection control function $g(\cdot)$.

Identification of the slope parameters requires a "single-index restriction" on the first-step selection process: $\Pr(D_{n,i} = 1 \mid \tilde{x}'_{n,i}, \tilde{x}'_{n,j}) = \Pr(D_{n,i} = 1 \mid \tilde{\eta}_{n,i,j})$, which a binomial logit or probit model satisfies; additive separability of the selection function in the second step; and an exclusion restriction. To satisfy the final condition, we assume that the PMMS survey rate for the non-chosen alternative does not directly influence the rate for the chosen alternative, except via the probability of being selected. So the ARM survey rate is absent from the FRM pricing equation, and the FRM survey rate from the ARM pricing equation. We also exclude borrower age, age^2, and experienced inflation from the second-stage pricing equations.

In the third step, we impute pairs of interest rates for each household using our selection-corrected estimates of the pricing equation coefficients, and use these predicted explanatory variables to estimate the structural-choice model in (A.17). As mentioned, the pricing equation intercept $\gamma_{0,i}$ is not identified in the two-step series estimator (A.21). However, Heckman (1990) suggests estimating it by calculating the mean or median difference between the dependent variable and the predicted value conditional on all other explanatory variables, $Rate_{n,i} - Z'_{n,i}\hat{\Gamma}_{n,i}$, using only those observations whose selection probabilities for alternative i are close to 1. Intuitively, these individuals are likely to have chosen the i due to observed factors. They suffer from little selection bias, and their mean or median pricing error should be close to zero. Schafgans and Zinde-Walsh (2002) show that Heckman's intercept estimator is consistent and asymptotically normal. We estimate the intercept as the median difference within the top 10% of observations from each selected subsample, sorted by their predicted choice probabilities.

G.2 Derivation of the WRTE

In each scenario, we can describe the cost of choosing an FRM over an ARM for switching households using the language of potential treatments and potential outcomes. We focus on the binary choice problem and number the FRM alternative as 1 (and the ARM alternative as 0). In every choice situation n, the household faces two potential outcomes: mortgage payments $Y_{n,1}$ under the FRM and mortgage payments $Y_{n,0}$ under the ARM. The observed set of mortgage payments in our data is $Y_n = D_n Y_{n,1} + (1 - D_n) Y_{n,0}$, where $D_n \in \{0, 1\}$ is the mortgage choice of household n ("treatment status"). As defined in equation (16), the value of D_n depends on the difference in latent utility in equation (15) between the alternatives: the FRM is chosen if the difference in observed components of latent utility exceed the difference in unobserved components, $-(\upsilon_{n,1} - \upsilon_{n,0}) < x'_{n,1}\beta_1 - x'_{n,0}\beta_0$. Observed latent utility may include alternative characteristics, such as prices, as well as household characteristics, and experienced inflation. The coefficients in Table 7 are estimates of their effects.

Let $D_n(b_\pi)$ be the potential choice individual n would make given experienced-inflation coefficient b_π ("potential treatment"). We can rewrite the choice observed in our data as

$$D_n = \int A_n(\beta_\pi) D_n(b_\pi) db_\pi \,, \tag{A.22}$$

where $A_n(\cdot) = \mathbb{I}\{b_\pi = \cdot\}$ and β_π is the true experienced-inflation coefficient, representing the additional weight placed on π^e beyond the full-information Bayesian optimum. The household's actual choice, under the true utility model, is $D_n(\beta_\pi) \in \{0, 1\}$; and the welfare-relevant counterfactual is the choice the household would have made in the same choice situation if placing no additional weight on experienced inflation: $D_n(0) \in \{0, 1\}$. If $D_n(\beta_\pi) = D_n(0)$, then "assignment" (experience-based learning) was irrelevant and experienced inflation did not influence the mortgage choice. If $D_n(\beta_\pi) \neq D_n(0)$, then the household would switch out of an FRM into an ARM under the counterfactual model.[46]

[46]Households only switch in one direction because we model $\Pr(D_n = 1 | b_\pi \pi^e)$ as a logit function, so that expected household choice is monotonic in $b_\pi \pi^e$, and $\pi^e > 0$.

Using this notation, the expected financial cost (or benefit) for switching households is

$$\mathbb{E}[Y_{n,1} - Y_{n,0} | D_n(\beta_\pi) = 1, D_n(0) = 0] \,, \tag{A.23}$$

i. e., the expected difference between FRM and ARM payments for households that chose an FRM because of their inflation experiences. Positive numbers represent overpayment, and negative numbers underpayment. The conditioning set restricts us to the subset of mortgagors for whom experienced inflation was the determining factor in their mortgage choice.

If we observed the actual realizations of these differences $\Delta Y_n = Y_{n,1} - Y_{n,0}$ across switching households, we could calculate the average and obtain a measure of the expected ex-post financial cost. While we can replace these unknown realizations with estimates, we still cannot directly estimate equation (A.23), because we do not observe households' counterfactual choices $D_n(0)$. However, Bayes' rule lets us rewrite (A.23) as

$$\begin{aligned} \mathbb{E}\left[\Delta Y_n | D_n(\beta_\pi) = 1, D_n(0) = 0\right] &= \int \Delta y \cdot f(\Delta y | D_n(\beta_\pi) = 1, D_n(0) = 0) d\Delta y \\ &= \frac{\int \Delta y \cdot h(D_n(\beta_\pi) = 1, D_n(0) = 0 | \Delta y) f(\Delta y) d\Delta y}{g(D_n(\beta_\pi) = 1, D_n(0) = 0)} \,. \end{aligned} \tag{A.24}$$

The first line of equation (A.24) gives the definition of a conditional expectation, using $f(\Delta y|\cdot)$ to notate the density of payment differences Δy conditional on the household being a switcher. This conditional density is unknown and cannot be estimated directly. The second line replaces the unknown density function with a probability mass function, $h(\cdot|\Delta y)$, giving the *probability* that a household facing payment difference Δy would switch to an ARM were it not for the presence of personal inflation experiences in its choice function. Multiplication by the unconditional density $f(\Delta y)$ indicates that we need to integrate over all payment differences Δy according to how often they occur in the population; and division by the unconditional mass function g merely ensures that the densities integrate to 1.

Thus, we have replaced households' unknown counterfactual choices with switching probabilities that we can estimate. Intuitively, the second line of equation (A.24) is the weighted average difference in FRM versus ARM mortgage payments, using households' switching probabilities as weights. We can estimate the probability h that a household facing payment difference Δy is a switcher, by comparing two predicted choice probabilities: the "true" probability that uses all of the coefficient estimates, and a "counterfactual" probability that sets $\beta_\pi = 0$ but uses all of the other coefficients as estimated:

$$h(D_n(\beta_\pi) = 1, D_n(0) = 0|\Delta y) = \Pr(D_n = 1|b_\pi = \beta_\pi, \Delta y) - \Pr(D_n = 1|b_\pi = 0, \Delta y)\,. \tag{A.25}$$

For example, if a household's true probability of choosing an FRM is 90% and the counterfactual probability (ignoring experienced inflation) is 70%, then for every 100 observationally-equivalent households, we expect 70 of them to choose an FRM no matter what, 10 to choose an ARM no matter what, and 20 to switch from the FRM to the ARM. These choice probabilities can be obtained by calculating predicted values from the estimates in Table 5 or 7. We can replace β_π, the unknown population coefficient on lifetime inflation experiences, with the logit estimate $\hat{\beta}_\pi$ from either the reduced-form or the three-step estimation, since both are consistent. Finally, we replace the actual FRM−ARM payment difference Δy_n with predicted differences $\Delta \hat{y}_n$ obtained from the selection-corrected pricing equations estimated in Table 6.

In reference to Heckman and Vytlacil (2007)'s formulation of the "policy-relevant treatment effect" (PRTE), who use the same weighted average that we have derived above, we denote our estimator of the weighted average of the difference in mortgage payments as the Welfare-Relevant Treatment Effect (WRTE):

$$\begin{aligned}\widehat{WRTE} &:= \widehat{\mathbb{E}}\left[Y_{n,1} - Y_{n,0}|D_n(\beta_\pi) = 1, D_n(0) = 0\right] \\ &= \sum_{n=1}^{N} \Delta\hat{y}_n \cdot \left\{ \frac{\widehat{\Pr}(D_n(\hat{\beta}_\pi) = 1|\Delta\hat{y}_n) - \widehat{\Pr}(D_n(0) = 1|\Delta\hat{y}_n)}{\sum_n \left(\widehat{\Pr}(D_n(\hat{\beta}_\pi) = 1|\Delta\hat{y}_n) - \widehat{\Pr}(D_n(0) = 1|\Delta\hat{y}_n)\right)} \right\}\,,\end{aligned} \tag{A.26}$$

where the weights are proportional to the difference in probability of choosing an FRM under the estimated ("true") and counterfactual experienced-inflation coefficients. Note that the WRTE (and PRTE) differ from standard objects reported in the treatment literature. For example, an Average Treatment Effect (ATE) is estimated as an unweighted average of the difference in expected payments, $\mathbb{E}[Y_n|b_n = \beta_n] - \mathbb{E}[Y_n|b_n = 0] = \sum_{i=0}^{1} \Pr(D_n(\beta_\pi) = i) \cdot Y_{n,i} - \sum_{i=0}^{1} \Pr(D_n(0) = i) \cdot Y_{n,i}$, using the actual versus the counterfactual choice probabilities.[47]

G.3 Modeling Refinancing Behavior

Optimal Refinancing. Agarwal, Driscoll, and Laibson (2013, hereafter ADL) provide a closed-form solution for this threshold. We use their square-root rule approximation

[47]By this logic, our "welfare-relevant treatment effect" is a Local Average Treatment Effect for the subset of the population for whom assignment is deterministic.

to the optimal threshold:

$$OT_{n,t} \approx -\sqrt{\frac{\sigma\kappa}{M_{n,t}(1-\tau)}\sqrt{2(\rho+\lambda_{n,t})}}\,, \tag{A.27}$$

where σ is the annualized standard deviation of movements in the FRM rate, κ is the fixed cost of refinancing, M is the outstanding mortgage balance, τ is the household's marginal tax rate, ρ is the household's intertemporal discount rate, and λ is the Poisson arrival rate of exogenous prepayment events. We follow ADL in parameterizing $\sigma = 0.0109$, $\kappa = \$2000$, and $\rho = 0.05$; and we continue to set the marginal tax rate $\tau = 0.25$. (ADL use the next bracket up, 28%.) The mortgage prepayment process parameterized by $\lambda_{n,t}$ is derived from three exogenous sources of principal repayment:

$$\lambda_{n,t} = \mu + \frac{i_n}{\exp(i_n(T-t))-1} + \pi\,. \tag{A.28}$$

The first term, μ, represents the hazard of moving and selling the house; this could in principle vary across households, but we follow ADL and set $\mu = 0.10$ (corresponding to an expected residency of $1/\mu = 10$ years). The second term represents the annual scheduled repayment of principal for a self-amortizing FRM carrying interest rate i_n with $T-t$ years remaining. The third term represents declines in the real value of future mortgage payments due to inflation. This could also vary over time with actual inflation, but for simplicity we set $\pi = 0.04$ (the mean CPI inflation rate over 1960–2013).

Expected Refinancing. To calculate a household's expected mortgage payments, we borrow estimates from Andersen et al. (2015) that describe the probability of refinancing as a function of the "incentive to refinance" embedded in the difference between the optimal threshold and the actual rate differential. Their baseline estimate of the probability that a household n will refinance in month m in year y is

$$\Pr(\mathit{Refi}_{n,y,m}|i_0) = \Phi\left(-1.921 + \exp(-1.033)\times(OT_{n,y} - (i_{n,y} - i_0))\right)\,, \tag{A.29}$$

where i_0 is the interest rate on the outstanding fixed-rate mortgage and $i_{n,y}$ is the interest rate on a new mortgage issued if the household refinances in year y.[48] We convert from a monthly to an annual horizon by assuming that monthly refinancing events are i. i. d. within a year: $\Pr(\mathit{Refi}_{n,y}|i_0) = 1 - (1 - \Pr(\mathit{Refi}_{n,y,m}|i_0))^{12}$. The refinancing probability may be interpreted as a transition probability between two "states": the state of holding a year-$(\mathit{OrigYr}_n + s)$ mortgage and the state of holding a year-$(\mathit{OrigYr}_n + t)$ mortgage, where s and t denote the number of years between origination and the pre-

[48]From Andersen et al. (2015), Table 9, col. 1, based on a sample of Danish households from 2008 to 2012.

vious refinancing or today, respectively. If i_0 is the rate $s \geq 0$ years after origination, and today is $t > s$ years after origination, then

$$P_n(S_t = t | S_{t-1} = s) := \Pr(\mathit{Refi}_{n, \mathit{OrigYr}_n + t} | i_0 = i_{n, \mathit{OrigYr}_n + s}) \cdot \mathbb{I}\{s < t\}\,. \tag{A.30}$$

$S_t \in \{0, 1, 2, ..., t\}$ denotes the household's current, time-t "state," i. e., the time of the most recent refinancing. To obtain the set of unconditional probabilities that, at time t, household n will hold a mortgage last refinanced at time s, $\{P_n(S_t = s), 0 \leq s \leq t \leq 29\}$, we begin with the initial condition that $P_n(S_0 = 0) = 1$ and solve forward iteratively.[49]

H Robustness Check: Supply-Side Constraints

Throughout the analysis, we take the supply side as fixed (i. e., the spread between FRM and ARM rates does not vary when households make counterfactual choices), and we assume that all borrowers have a choice between the FRM and ARM. However, lenders might impose constraints on some borrowers. Borrowers with high loan-to-income (LTI) ratios may face debt servicing constraints and need to get an ARM in order to qualify for a mortgage loan at all, or, conversely, they may not be offered an ARM due to income risk. Borrowers with low LTI ratios are more likely to have "free choice" between the two contract types.

To address supply-side confounds, we test whether our results persist in the subsample of unconstrained borrowers with low LTIs. In Table A.8, we re-estimate the reduced-form binomial choice model separately on above- and below-median LTI subsamples. We that estimated experience effect is even stronger in the unconstrained, low-LTI subsample (column 2). Among high-LTI borrowers, instead, who might not have a choice between the alternatives, inflation experiences play a weaker and insignificant role (column 1). As an additional test, we estimate the choice model on the full sample while flexibly controlling for the possibility of borrower constraints by including a fifth-order polynomial in LTI, in column 3. This does not substantially affect the coefficient on lifetime inflation experiences (cf. Table 5, column 5). We conclude that supply-side constraints in the mortgage lending process are not driving our results.

[49]The calculations also need to keep track of the household's outstanding mortgage balance at the beginning of each year. This state variable depends on the entire path of prior interest rates. There are $2^{29} \approx 500$ million such paths for every mortgage. To simplify matters, we assume that the timing of principal repayment in the "Expected Refinancing" case is the same as in the "Optimal Refinancing" case.

Table A.8: Supply-Side Constraints

	High LTI Subsample	Low LTI Subsample	Full Sample
	(1)	(2)	(3)
Freddie Mac PMMS FRM	-3.939***	-3.291***	-3.916***
index rate (%)	(1.18)	(1.18)	(0.84)
Freddie Mac PMMS ARM	0.969**	0.896*	1.005***
initial rate index (%)	(0.45)	(0.46)	(0.32)
Experienced inflation in %	0.118	0.319**	0.188*
	(0.13)	(0.16)	(0.10)
Log(Income)	0.002	0.062	-0.031
	(0.04)	(0.04)	(0.06)
Age	0.016	0.009	0.012
	(0.02)	(0.02)	(0.02)
$Age^2/100$	-0.015	-0.010	-0.012
	(0.02)	(0.03)	(0.02)
Number of Choice Situations	6,965	6,966	13,931
Pseudo R2	0.092	0.047	0.073
$-\beta_{\pi,\text{ FRM}} / \beta_{\text{Rate, FRM}}$	0.03	0.097*	0.048*
(S.E. by delta method)	(0.035)	(0.057)	(0.028)
Origination year FE	YES	YES	YES
Mortgage controls	YES	YES	YES
Sociodemographic controls	YES	YES	YES
5^{th}-order polynomial in LTI			YES

Notes. This table reports binomial logit coefficient estimates of choice between FRM, and ARM in the 1991 and 2001 RFS for mortgages originated ≤ 6 years ago, for subsamples split by borrower loan-to-income (LTI) ratios above or below the sample median. The dependent variable is an indicator equal to 1 if the household took out an FRM. Mortgage controls are Refi dummy, Junior Mortgage dummy, Nonconventional dummy, Loan / CLL, Jumbo dummy, and Points Paid. Sociodemographic controls are First-time Owner dummy, Joint Owners dummy, and Rural county dummy. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

I Robustness Check: Alternate Estimation Methods

I.1 Specification Test for the Parametric Choice Model

In our three-step estimation procedure, we estimate the first and third steps parametrically, by logit: $\Pr(D_n = 1) = F(x_n'\beta)$, where D_n is an indicator variable that equals 1 if the individual chose an FRM (and 0 otherwise), x_n is the set of explanatory variables in equation (15), and $F(v) = e^v/(1 + e^v)$ is the logit function. Horowitz and Härdle (1994) (HH) describe a specification test of a parametric conditional moment model versus semiparametric alternatives,

$$H_0 : E[D_n|x_n] = F(x_n'\beta) \text{ versus } H_1 : E[D_n|x_n] = G(x_n'\beta), \tag{A.31}$$

where F is the known (logit) CDF and G is an unknown CDF. Both the null and the alternative hypotheses maintain the *single-index restriction* that households' choice probabilities depend on the explanatory variables only via the one-dimensional index $v(x_n, \beta) = x_n'\beta$. This restriction is common in semiparametric models in order to avoid the "curse of dimensionality."

The HH test statistic is

$$HH := h^{1/2} \sum_{n=1}^{N} \omega_n \cdot (D_n - F(x_n'\hat{\beta})) \cdot (\hat{F}(x_n'\hat{\beta}) - F(x_n'\hat{\beta})) \,. \tag{A.32}$$

Intuitively, this statistic compares the average distance between the parametric link function F and a nonparametric estimate $\hat{F}_n = \hat{E}[D_n | x_n'\hat{\beta}]$, weighted by the parametric-model residuals. $\hat{F}_n$ must be independent of D_n for every n and asymptotically unbiased; h is the bandwidth used to estimate $\hat{F}$; and ω_n are a set of non-negative weights chosen to maximize power against the alternative hypothesis: $E[HH|H_1] = E[\omega_n \cdot (G_n - F_n)^2] =: \mu > 0$. (Note that the alternative is one sided.)

Under H_0, $\hat{F}_n - F_n$ is an asymptotically mean-zero, root-Nh consistent estimator, so by the appropriate Central Limit Theorem, $HH \xrightarrow{d} \mathcal{N}(0, V)$, with

$$V = 2 \int K(u)^2 du \cdot \int \omega(z)^2 \sigma^4(z) dz \,. \tag{A.33}$$

K is the kernel used to estimate $\hat{F}$ nonparametrically. A consistent estimator for $\text{Var}(HH)$ under H_0 is

$$\hat{V} = 2 \int K(u)^2 du \cdot \frac{1}{N} \sum_{n=1}^{N} \omega_n^2 \frac{[(F(x_n'\hat{\beta}))(1 - F(x_n'\hat{\beta}))]^2}{\hat{f}(x_n\hat{\beta})} \,. \tag{A.34}$$

The first term, $\int K^2$, is non-random and depends only on the choice of kernel function. The second term replaces an unknown population moment $E[\omega^2\sigma^4/f]$ with its sample analogue. The expression for $\sigma_n^2 = \text{Var}(D_n | x_n'\hat{\beta})$ relies on the observation that D_n is Bernoulli and uses the parametric model to estimate its conditional variance. The density of $x'\hat{\beta}$ is estimated using the same kernel and bandwidth as for $\hat{F}$.

We require that $\hat{F}_n$ be independent of D_n and asymptotically unbiased for $E[D_n | x_n'\beta = v_n]$. The former is achieved by using a leave-one-out kernel regression estimator, and the latter is achieved by using a bias-reducing kernel. Higher-order ($r > 2$) kernels reduce the asymptotic bias of $\hat{F}$ to order h^r, at the cost of possibly poor finite-sample performance because they take both positive and negative values. See, e.g., Härdle and Linton (1994) for further details on bias reduction and bandwidth selection for kernel estimators.

To implement this test, we must choose weights, a kernel function, and a bandwidth.

For the weights, we follow the suggestion of Fan and Liu (1997) and set $\omega_n = \hat{f}(x_n'\hat{\beta})$. (The other standard choice is a window function that equals 1 between the α and $1-\alpha$ quantiles of $x'\hat{\beta}$, and 0 everywhere else; e.g., $\alpha = 0.01$ or $\alpha = 0.05$.) We use a fourth-order kernel:

$$K^{(4)}(u) \coloneqq \frac{15}{8}\left(1 - \frac{7}{3}u^2\right) \times K^{(2)}(u) \times \mathbb{I}\{|u| \leq 1\}, \tag{A.35}$$

$$\text{where } K^{(2)}(u) \coloneqq \frac{3}{4}\left(1 - u^2\right) \times \mathbb{I}\{|u| \leq 1\}. \tag{A.36}$$

$K^{(2)}$ is the standard second-order Epanechnikov kernel. For our variance calculation, we note that this kernel has $\int [K^{(4)}(u)]^2 du = 5/4$. We choose the bandwidth for $\hat{F}$ by least-squares cross-validation: $h_{N,CV} \coloneqq \arg\min_{h \in H} N^{-1} \sum_n (D_n - \hat{F}^{(2)}(v_n; h))^2$, where $\hat{F}^{(2)}(v_n; h)$ is the leave-one-out estimator using $K^{(2)}$. We then plug $h_{N,CV}$ into $K^{(4)}$.

Figure A.4. HH Specification Test

Panel A. **Cross-validation function** Panel B. **P and NP Estimates of F**

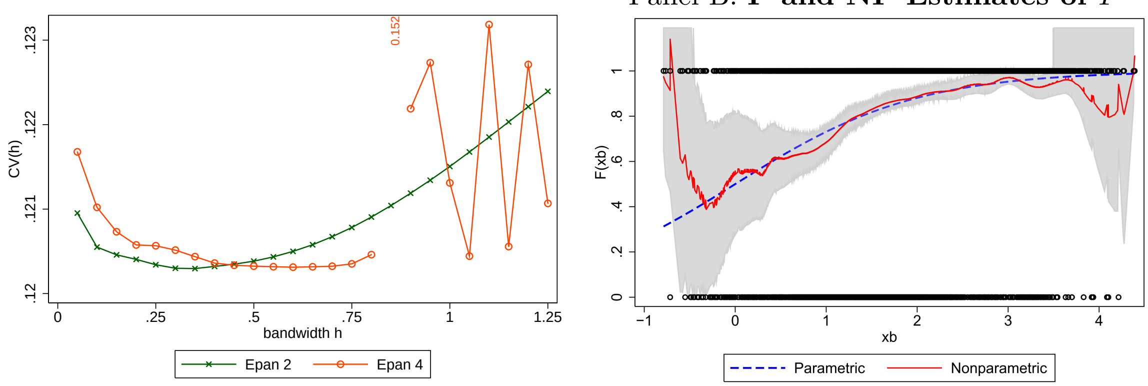

Notes. Panel A shows the cross-validation function for the reduced-form binomial mortgage choice model (18), with the Epanechnikov kernel ($r = 2$) and its fourth-order analogue ($r = 4$). Panel B shows the parametric and nonparametric estimates of the link function $\Pr(D_n = 1 \mid x_n'\hat{\beta})$ conditional on the reduced-form logit coefficients estimated in Table 5, column 5. The nonparametric estimator $\hat{F}$ is calculated using bandwidth $h = 0.35$ and $K^{(4)}$. Shaded area is a uniform 2-SE confidence interval for the nonparametric estimator, constructed using the Bonferoni correction for multiple testing.

Figure A.4(A) illustrates the bandwidth selection procedure. The index $x'\hat{\beta}$ is calculated using the reduced-form, binomial logit choice model coefficients reported in Table 5, column 5. We calculate the CV function on a grid over $h \in [0.05, 1.25]$ in increments of 0.05. For $K^{(2)}$, the criterion is minimized at $h_{N,CV} = 0.35$. The analogous grid search using $K^{(4)}$ in the CV function has a minimum at 0.60. Using the second-order crossvalidated bandwidth in conjunction with a fourth-order kernel guarantees that we will undersmooth asymptotically, as required to eliminate bias in $\hat{F}$.

Figure A.4(B) shows the two competing estimates of $F(x'\beta)$. The Horowitz and Härdle test statistic for the logit specification is $HH = 0.576$, with $\hat{V} = 0.017$. The associated Z-statistic is 4.40, well above the one-sided 1% critical value of 2.33, meaning that we reject the logit model. Results are similar for other values of the bandwidth ($h \in \{0.25, 0.45, 0.65\}$).

I.2 Semi-Nonparametric ML Estimation of the Choice Model

Given our rejection of the logistic distribution, we consider whether our results are affected by allowing the errors to come from a more general family of distributions. Maximum likelihood on a misspecified error distribution can still estimate the slope parameters of a discrete choice model consistently up to scale; see Ruud (1983) for sufficient conditions.

The "semi-nonparametric" (SNP) estimator of Gallant and Nychka (1987) is a pseudo-maximum likelihood estimator for models with the form $y_n = v(x_n'\beta) + e_n$. The single-index restriction $E[y_n|x_n] = v(x_n'\beta)$ is maintained, and the unknown density $g(e_n)$ is approximated by multiplying the standard normal density φ by Hermite polynomials:

$$g^*(e) = P(e)\varphi(e) = \left(\sum_{r=0}^{R} H_r(e)\right)\varphi(e)\,. \tag{A.37}$$

The approximate density g^* is substituted for the unknown density g into the log-likelihood. Gabler et al. (1993) extend the SNP estimator to binary-choice models, and De Luca (2008) implement it in Stata. Estimation proceeds by maximum likelihood with respect to the model coefficients, β, plus $R-2$ additional coefficients in front of the polynomial terms. The first two Hermite coefficients are fixed by location and scale normalizations, so the SNP estimator nests the probit estimator when $R = 2$.

Table A.9 presents estimates of the parametric probit model versus the semi-nonparametric model for $R = 3$ and 4. Model selection criteria such as Schwartz's BIC prefer $R = 3$, or just one additional parameter beyond the probit model. For comparability across models with different scale normalizations, we rescale the coefficient on the PMMS FRM rate to -1, so coefficients may be interpreted as WTPs in terms of the FRM rate. The probit coefficients in column 1 are almost identical to logit coefficients presented in Table 5, column 5, after rescaling. Our estimates are mostly unaffected by the switch from parametric to semi-nonparametric estimation in columns 2 and 3. In particular, we estimate a WTP of 5.0 basis points for every additional percentage point of lifetime inflation experiences in the SNP model with $R = 3$, versus 5.4 basis points in the probit model, and 5.2 in our baseline logit model.

Table A.9: Semi-Nonparametric Estimation of the Reduced-Form Choice Model

	(1)	(2)	(3)
Estimation Method:	*Probit*	*SNP (R=3)*	*SNP (R=4)*
Freddie Mac PMMS FRM	-1	-1	-1
index rate (%)	-	-	-
Freddie Mac PMMS ARM	0.239***	0.218***	0.219***
initial rate index (%)	(0.055)	(0.058)	(0.057)
Experienced inflation (%)	0.054*	0.050*	0.049*
	(0.029)	(0.030)	(0.029)
Log(Income)	0.007**	0.011***	0.011***
	(0.004)	(0.004)	(0.004)
Number of Choice Situations	14,337	14,337	14,337
Number of Alternatives	2	2	2
Pseudo-log likelihood	-5,701.2	-5,663.6	-5,663.5
Schwartz's BIC	11,641.7	11,585.7	11,595.0
Alternative-specific constants	YES	YES	YES
Origination year FE	YES	YES	YES
Mortgage controls	YES	YES	YES
Socidemographic controls	YES	YES	YES

Notes. The table reports estimates of the reduced-form model for households' choice betwee the FRM and ARM alternatives in the 1991 and 2001 RFS, for mortgages originated ≤ 6 years ago. All three columns rescale the coefficients so $b_{\mathrm{FRMRate}} = -1$ for comparability. The dependent variable is an indicator equal to 1 if FRM (and 0 if ARM). In columns 2 and 3, SNP indicates the semi-nonparametric pseudo-ML estimator of Gallant and Nychka (1987), where R is the order of the Hermite polynomial approximation to the unknown error density. Mortgage controls are Refi dummy, Junior Mortgage dummy, Nonconventional dummy, Loan / CLL, Jumbo dummy, and Points Paid. Sociodemographic controls are Age, Age^2, First-time Owner dummy, Joint Owners dummy, and Rural county dummy. Robust standard errors by the delta method in parentheses. *** p<0.01, ** p<0.05, * p<0.1

I.3 Three-Step Estimation by Fully Parametric Methods

We report estimates of the structural choice model using fully-parametric predicted values in Tables A.10 and A.11. We continue to estimate a powerful correlation between individuals experiencing higher levels of lifetime inflation and their propensity to choose an FRM, with a WTP of 30–45 basis points for every additional percentage point of lifetime inflation experiences.

Table A.10: Fully-Parametric Choice Model, Step 2 (Selection-Corrected Mortgage Rate Equations)

	(1)	(2)	(3)	(4)
Dependent variable:	FRM Rate		ARM Initial Rate	
Estimation Method:	Tobit	Heckit-Tobit	Tobit	Heckit-Tobit
Freddie Mac PMMS	54.8***	73.3***	77.5***	84.9***
index rate (%)	(1.65)	(10.90)	(3.31)	(4.90)
Log(Income)	-6.19***	-9.26***	0.827	0.201
	(2.24)	(3.23)	(2.09)	(2.12)
First-time owner?	12.6*	8.86	16.9**	13.9
	(6.50)	(8.78)	(8.28)	(8.80)
Joint owners?	-19.7***	-41.2***	13.3	5.42
	(7.30)	(15.60)	(8.66)	(9.91)
Rural?	25.6***	62.5**	50.9***	62.2***
	(9.86)	(25.00)	(10.60)	(12.10)
Refi?	-35.7***	-54.7***	17*	10
	(7.23)	(14.50)	(9.64)	(10.80)
Junior mortgage?	191***	153***	180***	165***
	(12.00)	(25.50)	(18.70)	(20.40)
Nonconventional?	-53.8***	-175**	-72.4***	-125***
	(6.55)	(71.30)	(16.60)	(32.90)
Points paid (pctg points)	-11.1***	-10.9***	-4.79	-4.38
	(1.46)	(2.52)	(3.54)	(3.70)
Loan / CLL	-66.9***	28.1	-101***	-71.7***
	(16.20)	(58.80)	(16.20)	(22.30)
Jumbo loan?	103***	181***	43.5***	54.6***
	(23.30)	(55.40)	(16.70)	(18.00)
Constant	587***	617***	281***	96.1
	(27.90)	(40.20)	(33.20)	(100.00)
Inverse of Mill's ratio		-601*		91.3*
		(349)		(47.8)
Observations	12,155	12,155	1,410	1,410
Pseudo-R2	0.008	0.008	0.041	0.041

Notes. The table reports fully-parametric estimates of the mortgage rate pricing equations, assuming joint normality of the first- and second-step errors. The sample is mortgages originated $\leq$ 6 years ago as of the 1991 and 2001 Residential Finance Surveys, with primary owner age between 25 and 74 years. The dependent variable is the interest rate in bps. In columns 2 and 4, the first step is a binomial probit model of mortgage choice on the same explanatory variables as in Table 5, column 5. Standard errors, in parentheses, adjusted for first-step estimation by mult-eqn. GMM formulas. *** p<0.01, ** p<0.05, * p<0.1

Table A.11: Fully-Parametric Choice Model, Step 3 (Structural Logit Model of Mortgage Choice)

	(1)	(2)	(3)	(4)
Step 2 Selection Correction?	*No*	*Yes*	*No*	*Yes*
FRM Rate Offered	-0.568*	-0.636***	-0.589***	-0.434***
	(0.30)	(0.22)	(0.14)	(0.17)
Initial ARM Rate Offered	0.606**	0.482**	0.93***	0.666***
	(0.31)	(0.20)	(0.15)	(0.24)
Experienced inflation in %	0.211**	0.184*	0.196**	0.192*
	(0.10)	(0.10)	(0.10)	(0.10)
Log(Income)	-0.0418	-0.038	-0.0273	-0.0194
	(0.03)	(0.02)	(0.03)	(0.02)
Age	-0.023	0.00466	0.0039	0.00975
	(0.02)	(0.02)	(0.02)	(0.02)
$Age^2/100$	0.0231	-0.00583	-0.00242	-0.00891
	(0.02)	(0.02)	(0.02)	(0.02)
Joint owners?	-0.105	-0.128	-0.091	-0.0502
	(0.12)	(0.08)	(0.11)	(0.08)
Outside MSA?	-0.399**	-0.225*	-0.568***	-0.422**
	(0.16)	(0.13)	(0.15)	(0.17)
Nonconventional Dummy			1.82***	1.5***
			(0.19)	(0.43)
Origination year FE	YES	YES	YES	YES
Number of Choice Situations	14,337	14,337	14,337	14,337
Pseudo-R2	0.022	0.050	0.063	0.065
$-\beta_{\pi,\text{ FRM}} / \beta_{\text{Rate, FRM}}$	0.371	0.289**	0.332*	0.441**
(S.E. by delta method)	(0.268)	(0.122)	(0.185)	(0.189)

Notes. The table reports binomial logit coefficient estimates for the fully parametric, structural model of mortgage choice between FRM and ARM alternatives in the 1991 and 2001 RFS. The dependent variable is an indicator equal to 1 if FRM, and 0 if ARM. Estimates are produced by a three-step procedure, in which interest rates for both alternatives are predicted (step 2) after correcting for sample selection (step 1) using first-step probit and second-step Heckit-Tobit. The sample is mortgages originated *le* 6 years prior to the survey year, with primary owner age between 25 and 74 years. Standard errors in parentheses, adjusted for first- and second-step estimation by mult.-eqn. GMM formulas. *** p<0.01, ** p<0.05, * p<0.1

J Learning from Nominal Interest-Rate Experiences

We re-estimate our reduced-form mortgage choice model, replacing π_n^e with i_n^e using short-term and long-term nominal interest rates from the CRSP U.S. Treasuries and Inflation Indexes database and the Historical Statistics of the United States (HSUS). Since these series only begin in 1926 and 1918, respectively, and 1915 is the earliest birth year in our dataset, we drop these few early years and re-normalize the weights to construct i_n^e. The resulting measures are highly correlated with inflation experiences ($\rho = 0.81$ and 0.69 for short-term and long-term rates, respectively).

Table A.12: Learning from Nominal Interest Rates

	(1)	(2)
Freddie Mac PMMS FRM index	-3.590***	-3.588***
rate (%)	(0.816)	(0.816)
Freddie Mac PMMS ARM index	0.845***	0.845***
rate (%)	(0.313)	(0.313)
Short-term interest rate experiences (%)	0.144*	
	(0.0770)	
Long-term interest rate experiences (%)		0.162*
		(0.0931)
Origination Year FE	YES	YES
Mortgage controls	YES	YES
Sociodemographic controls	YES	YES
Number of Choice Situations	14,337	14,337
Pseudo R2	0.069	0.069

Notes. This table reports binomial logit coefficients from a reduced-form choice model with the same sample and control variables as in Table 5, column 5. Interest rate experiences are constructed using linearly-declining weights from the current year to the year of birth, as in (1). Short-term nominal rates (column 1) are average annual returns on the 90-day Treasury bill from the CRSP US Treasuries and Inflation Indexes database (1926-2001). Long-term nominal rates (column 2) are U.S. government long-term bond yields (HSUS series Cj1192) between 1919-1961, and 10-year constant-maturity Treasury yields (Fed Release H.15) beginning in 1962. Mortgage controls and sociodemographic controls are the same as in Table 5, col. 5. Robust standard errors in parentheses. * p<0.10, ** p<0.05 and *** p<0.01.

Estimation results are in Table A.12. The direction and magnitude are similar to our baseline but statistically weaker. Model-selection criteria point to inflation beliefs as the preferable independent variable, followed by the short-term interest-rate-experiences model (e.g., Schwartz's BIC = 11,625.07 for the baseline model, followed by 11,625.18 for the short-term interest rate model and 11,625.68 for the long-term interest rate model).

K Estimating Geographic Mobility

Moving Probabilities based on Age. To estimate household moving probabilities, we obtain CPS-ASEC five-year geographic mobility estimates for the time periods 2000–05 and 2005–10 from the Census Bureau. We choose these time periods in order to capture both an expansion and a recession, so that we may smooth over business-cycle frequency variation in mobility rates. The Census's survey question asks respondents whether they lived in the same house or apartment five years ago, and classifies movers by type of move (within county, state, division, region, or from abroad). Since even a local move necessitates terminating the mortgage, we use the total mobility rate. The data does not break out renters versus homeowners, so our mobility-rates estimates are based on the entire population.

Figure A.5. Age and Mobility

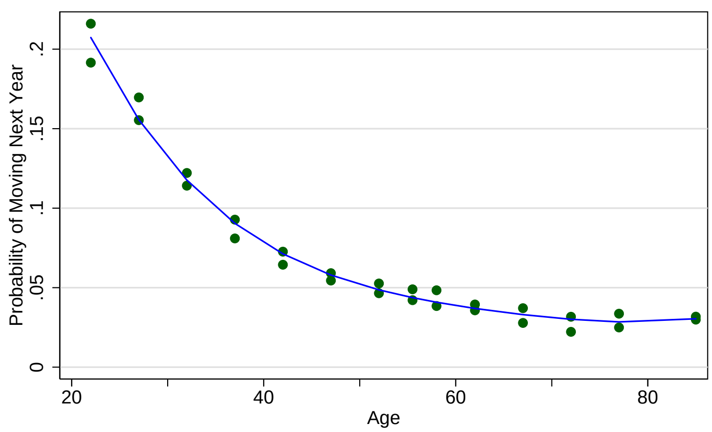

Notes. The data source is the CPS ASEC from 2005 and 2010. Fitted values are calculated using fourth-order polynomial function of age.

We convert five-year moving frequencies into one-year (ex-ante) probabilities as follows. First, since respondents are grouped into five-year age ranges, we code individuals' ages at the interval medians. So for example, individuals in the 35–39 year interval are coded as 37 years old today, and as 32 years old five years ago. We further top-code the highest interval (85+ years) at 85, and we drop respondents who were minors five years ago (i. e., aged less than 22 years at the time of the survey). We then convert the five-year moving probabilities to one-year moving probabilities by using an "independent-increments" (Poisson) assumption:

$$\text{MoveProb}_a^{1y} \equiv 1 - (1 - \text{MoveProb}_a^{5y})^{1/5} = 1 - \left(\frac{N_{a+5}(\text{Nonmovers})}{N_{a+5}(\text{Total})}\right)^{1/5},$$

where y is year(s), a an age bracket, and $N_a(\cdot)$ the number of individuals in age bracket a in the CPS data. We plot these one-year moving probabilities in Figure A.5. Mobility declines with age, leveling off in the mid-to-late 40s, and increasing again slightly in the late 80s.

We model the relationship between mobility and age by regressing one-year moving rates against a fourth-order polynomial in householder age:

$$\widehat{\text{MoveProb}}^{1y}(age) = \underset{(0.077)}{0.696} - \underset{(0.007)}{0.0355} \times age + \underset{(0.0002)}{0.000752} \times age^2 - \underset{(2.80 \cdot 10^{-6})}{7.40 \cdot 10^{-6}} \times age^3 + \underset{(1.30 \cdot 10^{-8})}{2.80 \cdot 10^{-8}} \times age^4 \,. \tag{A.38}$$

(Standard errors are in parentheses.) We finally use these coefficients to estimate the probability that a householder of age a today will still be in the house after T years:

$$\text{StayProb}(a, T) = \prod_{s=0}^{T-1} \left(1 - \widehat{\text{MoveProb}}_{1y}(a+s)\right) \,. \tag{A.39}$$

Moving Probabilities Based on Discount Points Paid. Discount points represent a trade-off between an upfront cost and a future benefit. Each discount point costs 1% of the amount borrowed, and buys approximately a 25 basis point reduction in the mortgage interest rate. The exact point-interest rate schedule may vary by bank and over time, but inspection of our data suggests that a quadratic function is a good Ndescription of the average schedule: $r(p) = r_0 - 0.0027p + 0.0002p^2$. This is the same order of magnitude that Brueckner (1994) finds for the early 1990s.

To estimate moving probabilities, we calculate the break-even horizon τ^* for each household, given the number of points paid and predicted future interest rate savings, discounting at an annual rate of 8%. If households are risk-neutral and face no liquidity constraints, then they will expect to reside in the house for exactly τ^* years. Assuming a constant hazard rate of moving (homogeneous Poisson), then years until moving $\tau \sim N.E.(\lambda)$ with intensity parameter $\lambda = 1/E[\tau] = 1/\tau^*$. Alternately, to model a hazard rate that decreases with time due to community attachment (Dynarski 1985, Quigley 1987), we let moving times follow a Weibull(λ, α) distribution with shape parameter $\alpha = 0.7$.[50] Finally, we allow for the possibility that individuals choose fewer than the optimal number of points due to risk aversion or liquidity constraints by fitting the intensity parameter to the median, rather than the mean: $F_\lambda^{-1}(\tau^*) = 0.5$.

Table A.13 reports the estimation results. We see significantly lower estimates of median tenure relative to our previous age-based calculations (bottom row of each panel).

[50]The negative exponential distribution equals the Weibull distribution with $\alpha = 1$.

This discrepancy reflects that most households do not pay any discount points (see Section 6.2). The discrepancy is exacerbated under the Weibull distribution (columns 4–5) and ameliorated when we fit each household's break-even horizon to the median rather than the mean (columns 3 and 5). Ignoring these concerns, we estimate the WRTE to be somewhat lower using the points-paid methodologies: under expected refinancing behavior and Scenario 3 interest rates, between $8 and 12 thousand, as compared to $15 thousand using age-based estimates of mobilitiy.

Table A.13: Moving Probabilities Based on Discount Points Paid

	(1)	(2)	(3)	(4)	(5)
P(Moving) based on:	*Age*	*Discount Points Paid*			
Distribution:		Neg. Exp. (λ)		Weibull(λ, 0.7)	
Break-Even Year (τ*):		τ*=E[τ]	F(τ*)=0.5	τ*=E[τ]	F(τ*)=0.5
	Scenario 1: Primary Mortgage Market Survey rates				
After-tax PDV [in $]:					
No Refi	13,052	6,603	8,805	6,222	9,815
Expected Refi	7,827	5,095	6,136	4,636	6,325
Optimal Refi	6,493	4,368	5,172	3,999	5,316
Av. Median Tenure (years)	12.5	4.9	6.6	3.6	6.6
	Scenario 2: Risk-adjusted rates, seniority-adjusted ARM margins				
After-tax PDV (all in $):					
No Refi	20,819	10,953	14,265	10,160	15,574
Expected Refi	15,769	9,630	11,848	8,724	12,315
Optimal Refi	14,475	8,945	10,937	8,123	11,357
Av. Median Tenure (years)	12.5	4.7	6.4	3.6	6.4
	Scenario 3: Risk-adjusted rates and ARM margins				
After-tax PDV (in $):					
No Refi	19,964	10,436	13,607	9,720	14,912
Expected Refi	14,854	9,106	11,176	8,275	11,629
Optimal Refi	13,543	8,416	10,256	7,668	10,661
Av. Median Tenure (years)	12.5	4.7	6.4	3.6	6.4

Notes. The table reports expected additional interest paid by switching households, allowing for heterogeneity in the probability of moving based on head of household's age or discount points paid. All dollar amounts are in constant year-2000 units. Positive values indicate that the FRM is more expensive than the ARM. Welfare-relevant treatment effect, PDV calculations, and refinancing scenarios same as in Table 8. Column (1) reproduces the estimation from the final column of Table 8 for comparison. In columns (2)-(5), discount points paid at time of origination are used to calculate the time to break even, τ^*, for each household, assuming an 8% nominal discount rate. In columns (2) and (3), the time of moving events $\tau \sim$ Negative Exponential (λ) distribution, with λ picked to fit τ^* to the mean and median of the distribution for each household. In columns (4) and (5), the time of moving events $\tau \sim$ Weibull $(\lambda, 0.7)$ distribution, so the hazard rate of moving is decreasing over time, with λ picked to fit τ^* to the mean and median of the distribution for each household. Average median tenure is calculated as the median tenure for each household, then averaged over all switching households.