DISCUSSION PAPER SERIES

DP14934

THE LONG SHADOWS OF THE GREAT INFLATION: EVIDENCE FROM RESIDENTIAL MORTGAGES

Matthew J. Botsch and Ulrike M. Malmendier

FINANCIAL ECONOMICS

ISSN 0265-8003

THE LONG SHADOWS OF THE GREAT INFLATION: EVIDENCE FROM RESIDENTIAL MORTGAGES

Matthew J. Botsch and Ulrike M. Malmendier

Discussion Paper DP14934
Published 24 June 2020
Submitted 04 June 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

THE LONG SHADOWS OF THE GREAT INFLATION: EVIDENCE FROM RESIDENTIAL MORTGAGES

Abstract

A major puzzle in financial contracting is consumers' aversion to adjustable rates. In the mortgage market, the empirical mix of contracts (80% fixed-rate) is inconsistent with standard life-cycle consumption models. We argue that these choices reflect the longlasting effect of the Great Inflation, and have sizable welfare implications. First, we show that consumers who have experienced higher inflation expect higher future interest-rate increases, which explains their preference for fixed-rate financing. Next, we quantify the influence of personal inflation experiences on mortgage financing using linked data from the Census Bureau's Residential Finance Survey. We estimate a discrete-choice model over mortgage financing alternatives. The structural parameters indicate that one additional percentage point of experienced inflation increases a borrower's willingness to pay for a fixed-rate mortgage by 6 to 14 basis points, compared to the adjustable-rate alternative in a given origination year. This experience effect has a major impact on the product mix of FRMs versus ARMs: Nearly one in seven households would switch to an ARM if not for the longlasting effect of personal inflation experiences. Our simulations suggest that households who would otherwise have switched pay $8,000-$16,000 in year-2000, after-tax dollars for the embedded inflation protection of the FRM.

JEL Classification: D14, D83, D84, D91, E31, G41, G51

Keywords: Contract choice, household finance, Inflation expectations, Mortgage Choice, behavioral finance

Matthew J. Botsch - mbotsch@bowdoin.edu
Bowdoin College

Ulrike M. Malmendier - ulrike@econ.berkeley.edu
UC Berkeley, National Bureau of Economic Research and CEPR

Acknowledgements
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, Canyao Liu, Junjun Quan, and Jeffrey Zeidel for excellent research assistance.

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

Matthew J. Botsch and Ulrike Malmendier†

January 13, 2020

Abstract

A major puzzle in financial contracting is consumers' aversion to adjustable rates. In the mortgage market, the empirical mix of contracts (80% fixed-rate) is inconsistent with standard life-cycle consumption models. We argue that these choices reflect the longlasting effect of the Great Inflation, and have sizable welfare implications. First, we show that consumers who have experienced higher inflation expect higher future interest-rate increases, which explains their preference for fixed-rate financing. Next, we quantify the influence of personal inflation experiences on mortgage financing using linked data from the Census Bureau's *Residential Finance Survey*. We estimate a discrete-choice model over mortgage financing alternatives. The structural parameters indicate that one additional percentage point of experienced inflation increases a borrower's willingness to pay for a fixed-rate mortgage by 6 to 14 basis points, compared to the adjustable-rate alternative in a given origination year. This experience effect has a major impact on the product mix of FRMs versus ARMs: Nearly one in seven households would switch to an ARM if not for the longlasting effect of personal inflation experiences. Our simulations suggest that households who would otherwise have switched pay $8,000-$16,000 in year-2000, after-tax dollars for the embedded inflation protection of the FRM.

Keywords: Contract choice, household finance, inflation expectations, mortgage choice, behavioral finance.

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, 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

Whether to buy a home and how to finance the purchase is one of the biggest financial decisions for many households, with important consequences for lifetime savings and consumption. 67% of Americans live in their own home, and they typically take on significant leverage to finance the purchase.[1] The mortgage decision is complex: households have to consider loan-to-value ratios, contract durations, repayment schedules, and fixed versus adjustable rates. Choosing the wrong contract can have sizable financial consequences, and a large volume of prior research on financial contracting has investigated how households optimize their mortgage choice.[2]

But are households actually making optimal mortgage decisions? Especially since the 2008 financial crisis, policy makers and researchers have called this into question. An extensive literature documents mistakes in households' financial decision-making and how financial product design caters to their biases (Campbell et al. 2011). In the domain of mortgage choice, about half of all fixed-rate borrowers make refinancing mistakes, and their foregone savings exceed ten thousand dollars (Agarwal et al. 2013; Agarwal et al. 2015; Keys et al. 2016).

One major puzzle is the overwhelming preference of consumers for fixed-rate mortgages (FRMs). In the U.S., FRMs command an 80% market share, and households pay a premium of, on average, 170 basis points over equivalent-risk and -term adjustable-rate mortgages (ARMs). In other countries, consumer mortgage choices have developed very differently, with ARM-type contracts often constituting high market shares.[3]

The continued dominance of FRMs in the U.S. is puzzling since it is costly for consumers and hard to reconcile with standard consumption models. For example, in their life-cycle consumption model with borrowing constraints, Campbell and Cocco (2003, 2015) show that most households are predicted to prefer an ARM, particularly if they are younger and more mobile. Our own calculations below will confirm that far

[1] 2008-2015 average from U.S. Census Current Population Survey (Housing Vacancies & Homeownership).

[2] Demand-side determinants include borrower age, mobility, and risk preferences (e. g., Chambers et al. (2009); Campbell and Cocco (2003); Sa-Aadu and Sirmans (1995); Brueckner (1992)), while supply-side features include term structure, relative prices, secondary-market liquidity, and mechanism design such as the use of discount points (e. g., Badarinza et al. (2018); Moench et al. (2010); Koijen et al. (2009); Stanton and Wallace (1998)). See also Brueckner and Follain (1988) and Dhillon et al. (1987) among the earlier literature.

[3] Countries with *primarily* variable mortgage rates include Australia, Belgium, Chile, Estonia, Finland, Greece, Hungary, Ireland, Israel, Korea, Luxembourg, Mexico, Norway, Poland, Portugal, Slovakia, Slovenia, Spain, Switzerland, and Turkey; cf. Andrews et al. (2011) and Andrews and Caldera Sanchez (2011).

more households choose FRMs than the standard economic model predicts, especially the Baby Boomers' generation in the wake of the Great Inflation: They should have taken out 1 million fewer FRMs in the late 1980s, and half a million fewer in the late 1990s. The costs of these deviations are large. Given expected refinancing behavior and mobility, Baby Boomers overpayed more than $14 billion on their fixed-rate mortgages in the late 1980s, and almost $9 billion in the late 1990s.

How can we explain the choices of U.S. households? Prior papers have explored a vast array of market-specific factors (e. g., payment structure, interest-rate deductability, and rental-market regulation) and demographic determinants (e. g., life-cycle stages, age, fertility, household size, and mobility), surveyed in Campbell (2013). In this paper, we explore a novel channel: the long-lasting effect of high-inflation experiences during the Great Inflation in the 1970s and early 1980s. We show that, after accounting for other known determinants of mortgage choice, personal inflation experiences have significant predictive power for interest-rate expectations and mortgage choices. Our structural estimates imply that one additional percentage point of experienced inflation increases a borrower's willingness to pay for a fixed-rate mortgage by 6 to 12 basis points in a given origination year. As a result, households who would otherwise have switched to an ARM pay between $8,000 and $16,000 in year-2000, after-tax dollars for their experience-driven choice of an FRM over their expected tenure in the house.

Our approach builds on a notion frequently discussed among practitioners, namely, that the Great Inflation cast "long shadows," which continued to affect beliefs and fears of consumers for decades to come. An early literature on mortgage financing from the time of the Great Inflation first proposed that inflation expectations might distort housing decisions, including Kearl (1979), Baesel and Biger (1980), and Alm and Follain (1982). The recent macro-finance literature has formalized and tested the notion that personal experiences of macro-finance outcomes, such as the high inflation of the 1970s, have a lasting impact on individuals' beliefs and attitudes. Much of the evidence demonstrates that personal stock-market experiences have a longlasting effect on beliefs about future stock returns and on stock-market participation, cf. Kaustia and Knüpfer (2008), Malmendier and Nagel (2011), Strahilevitz 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 empirical 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]

To make these models and evidence on experience effects operational for analyzing the contractual mix and composition of mortgage markets, as well as to provide structural estimates of their magnitude, impact on the market structure, and welfare implications, we formalize the concept of experience-based learning. As in prior work, we allow households to overweight their lifetime experiences relative to the optimal Bayesian scheme with roughly linearly declining weights going into the past. For Bayesian agents, the effect of personal experiences on beliefs is predicted to be zero, after accounting for all other standard-model determinants. For experience-based learners, higher experienced inflation predicts higher expected price increases in the future.

Experience-based overweighting of lifetime inflation has two main implications for mortgage markets. First, different experiences drive a wedge between individual valuations of fixed-rate and adjustable-rate liabilities. Those with high inflation experiences overvalue fixed-rate mortgages (relative to adjustable-rate contracts) since they estimate the present value of future repayments to be lower, in real terms, than individuals with low inflation experiences. Higher valuation, in turn, implies that such individuals are willing to pay more for a fixed-rate mortgage, relative to adjustable-rate instruments.

Second, the experience-effect hypothesis also implies that individuals with shorter lifetime histories overweight recent experiences more than those with longer histories. Hence, younger borrowers are predicted to respond more strongly to recent inflation experiences. Consider, for example, young borrowers coming of age during the 1970s. In the 1980s, these cohorts had recently experienced a period of high inflation, and they had no personal memory of low inflation. The testable prediction of experience-based learning is that, in the 1980s, younger cohorts should be more likely to choose fixed-rate mortgages than older cohorts. In the late 1990s and early 2000s, instead, then-young mortgagors, who came of age after the Volcker Fed tamed inflation, should behave more like the cohorts that came of age prior to the Great Inflation. As the raw data on FRM versus ARM choices in Figure 1 indicates, this is exactly the case. (We will discuss the data and figure in more detail below.)

Turning to the empirical implementation, we first utilize the Survey of Consumer Finances (SCF) to show that inflation experiences affect interest-rate expectations: Those with higher inflation experiences are more likely to expect interest rates to rise.

[4] Past inflation also has the potential to explain differences in the homeownership rates across European countries (Malmendier and Steiny 2016), and to influence even FOMC members in their inflation forecasts and voting behavior (Malmendier et al. 2018).

Figure 1: FRM Share and Experienced Inflation by Age Group

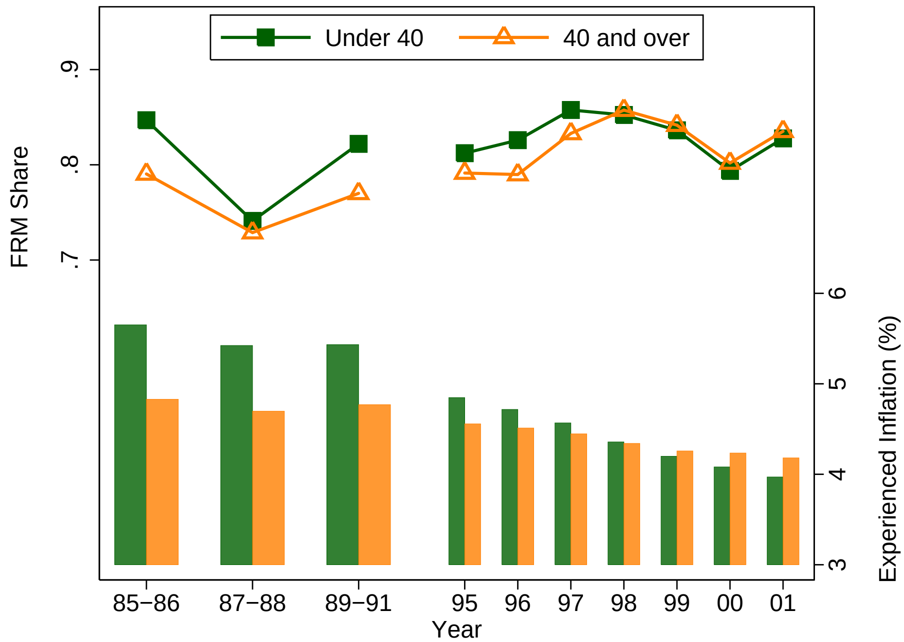

Notes. Data from the 1991 & 2001 RFS and BLS CPI. The 1991 RFS reports origination years in two- or three-year intervals.

We then test the implications for mortgage choice. Here, 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. The data are unique in that they survey both the household and the mortgage servicer, providing detailed and high-quality access to both demographic information and mortgage contract terms in the two cross-sections.

We estimate the structural parameters of a discrete-choice model over mortgage financing alternatives. Estimating this model is challenging for two reasons. 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 that are observable to every household. 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 region-year. In Step 2, we estimate the fixed- or variable-rate mortgage terms for each household 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 seniority of the mortgage. We implement the semiparametric Newey (2009) series estimator to correct for selection bias which arises since these equations are estimated 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 of both 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 approximately 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 lifetime inflation experiences. 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 current inflation as well as the entire history of all past inflation realizations (common to all market participants) at the time of origination. 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 experienced inflation, compared to other individuals in the same origination year.

For robustness, we re-estimate the reduced-form, binomial mortgage-choice model on the SCF data from 1989-2013. The SCF is conducted at a higher frequency, but has a smaller overall sample, fewer contract terms, and does not identify the geographic location. Despite these differences, the point estimates from the SCF data are remarkably similar to those from the RFS, providing a strong corroboration of our results.

To 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 experienced-inflation 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.

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. (Follain 1990 provides an overview of the earlier literature.) Mortgage choice is also a core element of the growing field of household finance, especially since the financial crisis (Guiso and Sodini 2013, Green and Wachter 2005, Mayer et al. 2009).

Our analysis builds on the earlier work of Case and Shiller (1988) and Shiller (1999, 2005) and adds to the literature on the role of (non-standard) belief formation in explaining mortgage choice and homeownership. For example, Koijen et al. (2009) argue that aggregate U.S. mortgage choice is well-explained by an adaptive-expectations "rule of thumb" where households use only the most recent three years of yield curve data.[5] 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. Extrapolative expectations formation is also a candidate to explain the house price boom and bust of the mid-2000s, as discussed in the survey of Glaeser and Nathanson (2015).[6] Our paper differs from these previous studies in introducing prior lifetime experiences as a novel determinant of households' beliefs and mortgage decisions. 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 our cost estimates suggest potentially significant welfare consequences.

More broadly, prior research has provided insights into the implications of behavioral factors for mortgage contract design and regulation. For example, Gottlieb and Zhang

[5] Badarinza et al. (2018) find that this rule performs less well in an international context.

[6] 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).

(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 later. As discussed above, most related is the work on inflation experiences. Relative to the evidence of experience-based inflation expectations in Malmendier and Nagel (2016), we add new evidence that this relationship translates into subjective interest-rate beliefs, and hence has the potential to affect the relative choice of fixed- versus variable-rate instruments. 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. We overcome these difficulties using the RFS. Most importantly, our paper is the first to provide structural estimates of the magnitude of experience effects on financial-contract choices and their payoff consequences, which suggest substantial welfare implications. The results aim to be a first stepping stone toward more complete welfare estimations.

2 Institutional Background and Framework

2.1 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.

Following the onset of the S&L crisis, the Garn-St. Germain Depository Institutions Act of 1982 allowed banks to originate adjustable-rate mortgages. 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, 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 then variable, and "interest-only" mortgages, under which no principal is paid in early periods to keep initial payments low. Most of our analysis will focus on the dominant contract types, FRMs and ARMs, with some comparison to mortgages with balloon payments.

The Census Bureau's RFS data on outstanding residential mortgages reveals the persistent dominance of FRMs, at around 80% market share (cf. Figure 1 above). 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). A major determinant of the variation in FRM market share (from 73% in 1987-88 to 86% in 1998) is the relative cost of FRMs versus ARMs. The average annual rate differential reported in the PMMS fluctuated between a low of 34 basis points (in 2009) and a high of 302 basis points (in 1994). In the years also covered by the RFS, the correlation between the FRM share and the FRM-ARM rate spread in the PMMS is -0.49.

2.2 Inflation Experiences and Belief Formation

The key hypothesis in this paper is that, to better understand the puzzling asset composition of the mortgage market, and the significant cost consumers incur as a result, it is helpful to consider the long-lasting consequences of the Great Inflation. To make this hypothesis testable, we build on the theoretical framework and empirical evidence on experience effects. 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 similarly: while the most recent years receive the highest weight, experiences earlier in life still carry significant weight. We thus calculate experienced inflation $\pi^e_{s,t}$ in year t for individuals belonging to the cohort born in year s as:

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

This formula places the highest weight on the most recent observation, zero weight on the year of birth, and connects these endpoints linearly.

We obtain the CPI-U from BLS for 1913-2013, and 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_k as the log change in the annual average

Figure 2: Actual and Hypothetical Inflation

Panel A: Inflation **Panel B: Inflation Experiences**

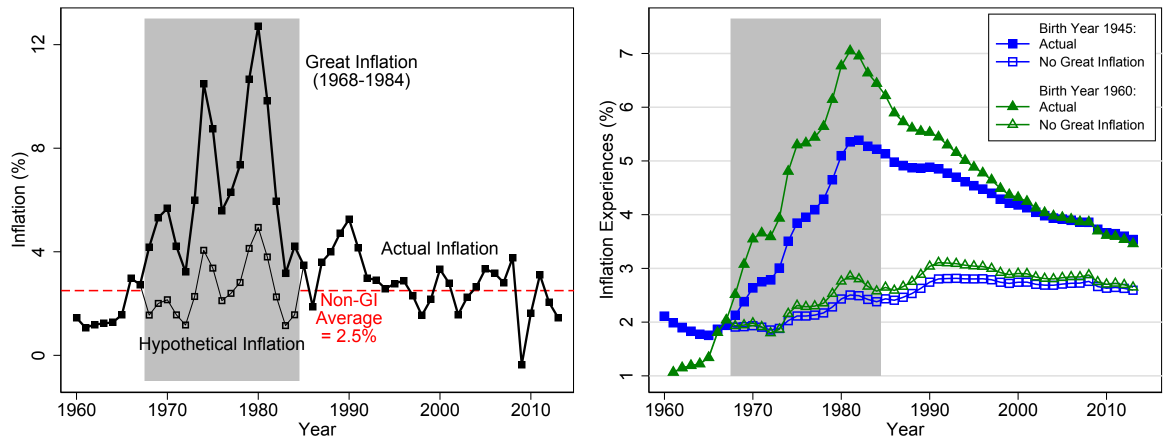

Notes. Data from BLS CPI-U. In Panel A, Actual Inflation (solid boxes) is the annual log change in CPI, and Hypothetical Inflation (hollow boxes) is a location-scale transformation of actual 1968-84 inflation (mean 6.68%, S.D. 2.84%) to the 1960-2013 (excluding 1968-84) mean of 2.55% and S.D. of 1.13%: $\pi_t^H := (\pi_t - 0.0668)/0.0284 \times 0.0113 + 0.0255$. 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.

level of the price index between years $k-1$ and k. We then calculate experienced inflation $\pi_{s,t}^e$ in year t for individuals belonging to the cohort born in year s using (1).

Figure 2 illustrates the effect of this weighting formula on inflation for two representative households, an "older" household from the 1945 birth-year cohort, and a "younger" household from the 1960 birth-year cohort. The top panel (2a) plots annual CPI-U inflation rates from 1960 to 2013 with the solid line (filled squares); the time of the "Great Inflation" is shaded in grey.[7] The lower line (hollow squares) indicates a hypothetical alternative inflation path if the Great Inflation had not occurred, using a location-scale transformation of actual inflation to the No-Great-Inflation mean of 2.5% and S.D. of 1.1%. In the bottom graph (2b), we use these actual and hypothetical inflation paths to calculate the corresponding lifetime weighted-average inflation experiences, separately for the "young" and "old" cohorts.

Figure 2 provides two main insights. First, young borrowers are particularly affected by inflation shocks, since they have the shortest personal histories of inflation experiences. Even under the hypothetical "No Great Inflation" scenario, the lifetime average of the younger cohort increases by 30 basis points more than that of the older cohort

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

following the second oil crisis in 1979. In reality the lifetime average of the younger cohort increases by significantly more, 200 basis points. However, by the late 1990s their personal inflation experiences 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 Data

Our main source of individual-level data on mortgage financing and demographics is the Residential Finance Survey (RFS), which the Census Bureau conducted the year after every decennial Census between 1950 and 2000. 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.[8] 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. Total household income in the Census year is imputed back to the origination year by the peak-to-peak log growth rate in U.S. nominal median household income

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

over 1980-2001 from CPS Historical Table H-6 (approximately 4.14% annually).

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, 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 A.

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).

We supplement the RFS with data from 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 mortgages with an 80% LTV and a 30-year term. These baseline rates are charged to high-quality borrowers. We take annual averages of the weekly 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.

Borrower attributes are summarized by mortgage product choice in Table 1. 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 (above the conforming loan limit). There is no significant age difference between FRM and ARM borrowers, contrary to the prediction of standard theory. Lifetime inflation experiences (defined below) are actually 5 bp lower for the typical FRM borrower than for the typical ARM borrower (4.74% versus 4.79%). However, this simple comparison pools across all origination years and ignores time-series variation in the relative cost of the two products. As we will be seen below, individuals who have experienced higher inflation within an origination year are more likely to choose an FRM.

We also use data from the Survey of Consumer Finances (SCF), conducted triennially by the Federal Reserve Board. The SCF has the advantage of being conducted at a higher frequency than the RFS. An important limitation of the SCF 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). Our identification strategy relies on the inclusion of year fixed effects, so the lack of within-survey geographic variation prevents us from estimating some parameters of interest. The SCF also includes a less extensive list of mortgage contract characteristics than the RFS. For example, we only have information about refi and non-conventional status for the first mortgage on the primary residences, not for junior mortgages or for mortgages on second homes. In addition, the SCF did not ask about first-time homeowner status until 2007. On the other hand, the SCF includes a more complete picture of the overall household balance sheet, notably allowing us to control for household net worth.

4 Interest Rate Expectations

Experience-based learning generates several predictions for the beliefs and choices of mortgage borrowers. First, we consider beliefs about future interest rates, which are a key determinant of the ARM-versus-FRM choice. Since personal experiences of inflation are known to affect consumers' beliefs about future inflation (Malmendier and Nagel (2016)), then, by the Fisher equation, $i = r + \mathbb{E}\pi$, they should also affect beliefs about future interest rates. Specifically, overweighting lifetime experiences of inflation should lead individuals coming of age during periods of high inflation to expect not only higher inflation but also higher nominal interest rates in the future.

To test whether lifetime experiences of inflation affect households' views on future interest rates, we use SCF question: "Five years from now, do you think interest rates will be higher, lower, or about the same as today?" In each survey wave, we calculate, separately for each cohort, the net fraction of respondents expecting interest rates to rise as the fraction answering "higher" minus the fraction answering "lower." We then relate the net fraction expecting rising rates to their lifetime experience of inflation.

Figure 3 illustrates the resulting relationship. For the graphical representation, we group cohorts above and below the sample median of age, representing "older" and "younger" cohorts. (We also estimate the relationship in a regression framework where we use each cohort-year lifetime experience separately.) Figure 3 plots the deviation

Figure 3: Interest-Rate Expectations and Inflation Experiences

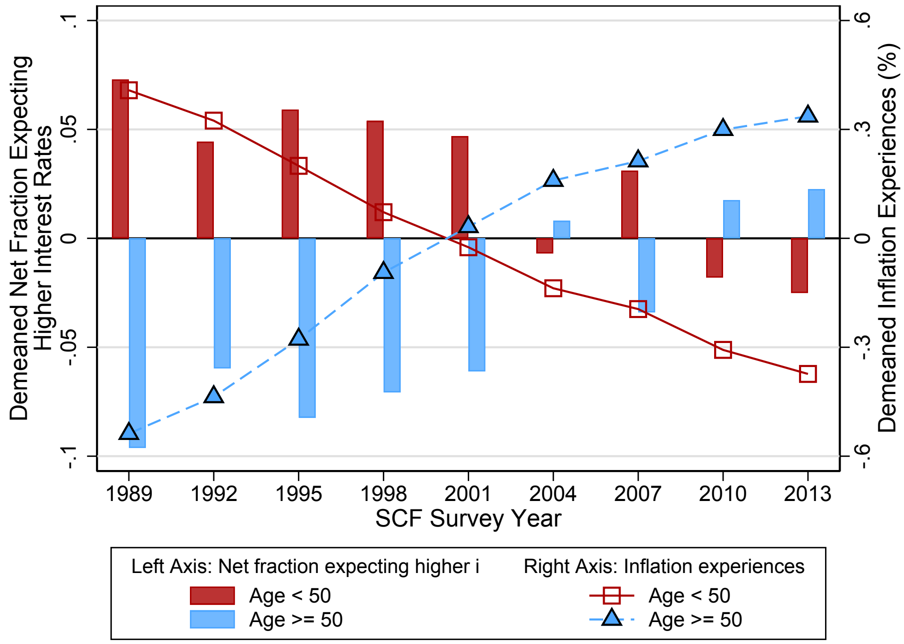

Notes. Values shown are cohort deviations from survey-year mean (average across implicates).

of each group's response from the overall survey-year mean. We see that, in the early SCF years (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. This relationship reverses in the mid-2000s, as the memory of the Great Inflation is fading (i. e., weighted less) and as the households who have experienced the Great Inflation age and become older households. At that time new, younger households who put relatively less weight on the Great Inflation enter the sample and have less positive expectations. The timing of this reversal in interest-rate expectations coincides almost exactly with cross-sectional differences in survey respondents' lifetime inflation experiences, calculated using (1). The link between lifetime experiences of inflation and interest-rate expectations lends plausibility to the hypothesized relationship between experiences and mortgage choice.

5 Mortgage Choice

Experience-based learning and its effect on inflation expectations as well as interest-rate expectations has two implications for mortgage choice: (i) Higher lifetime experiences of inflation induce a preference for FRMs. Experience-based learners with higher inflation experiences, who expect higher future nominal rates, 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. (ii) Younger individuals with shorter lifetime histories so far overweight recent experiences more than those with longer histories, and hence, younger borrowers respond more strongly to recent inflation realizations.

Figure 1 in the introduction illustrated that these 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 had experienced higher 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, and so did their mortgage choices. In our main analysis, we test for this pattern formally, in a rich econometric framework. We then quantify the magnitude and economic cost of experience effects on mortgage financing.

5.1 Estimation Methodology

We provide an overview of our methodology here, and describe it in full detail in Appendix C. Suppose that each household n derives (indirect) utility

$$U_{n,i} = \beta_{0,i,y} + \beta_{R,i} Rate_{n,i} + \beta_{\pi,i}\pi_n^e + \beta_{Inc,i} Income_n + f_i(Age_n) + \varepsilon_{n,i} \tag{2}$$

when choosing alternative i from a menu of J alternatives, $i \in \{FRM,\ ARM,\ Balloon\}$. A household lives in Census region r and chooses a mortgage once, in year y (unless they take a junior mortgage), so we omit time subscripts for notational simplicity. 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).[9] 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 $\varepsilon_{n,i}$ accounts for any unobservable factors affecting mortgage choice. The explanatory variable of interest is borrower n's lifetime experience of inflation at the time of the choice situation, π_n^e.

[9] 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.

Note that (2) includes alternative-specific year fixed effects $\beta_{0,i,y}$, which 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 thus 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. Normalizing $\beta_{\pi,ARM} = 0$, the experience-effect hypothesis implies $\beta_{\pi,FRM} > 0$, while the standard rational framework predicts $\beta_{\pi,FRM} = 0$. Alternative i is chosen by household n if

$$D_{n,i} \coloneqq \mathbb{I}\{U_{n,i} > U_{n,j} \quad \forall j \neq i\} \tag{3}$$

equals 1. This could be estimated by standard discrete-choice methods such as 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:

$$Rate_{n,i} = \gamma_{0,i} + \gamma_{R,i} PMMSRate_{y,r,i} + z_n' \gamma_i + v_{n,i}. \tag{4}$$

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 (4) 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).[10] Plugging (4) into (2), we obtain a *reduced-form* choice model that we can estimate:

$$U_{n,i} = \tilde{\beta}_{0,i,t} + \tilde{\beta}_{R,i} PMMSRate_{y,r,i} + \beta_{\pi,i}\pi_n^e + \tilde{\beta}_{Inc,i} Income_n + f_i(Age_n) + \tilde{z}_n' \tilde{\gamma}_i + \tilde{\varepsilon}_{n,i}. \tag{5}$$

We place tildes on coefficients and variables that represent different objects in (5) than in (2). 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 alter-

[10] 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.

natives. Moreover, since model (4) does not include inflation experiences, the reduced-form model (5) consistently estimates the structural coefficient $\beta_{\pi,i}$.

Our three-step estimation procedure is as follows. First, we estimate the reduced-form choice model (5), where households' decisions depend on region- and time-specific average FRM and ARM rates from the PMMS. In the second step, we estimate two mortgage pricing equations (4), 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 C) 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. Since the SPSC control function absorbs the intercept, we follow the suggestion of Heckman (1990) and estimate the intercept of (4) 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 (2) using the household-level menu of predicted prices from the second step.

5.2 Choice Model Estimates

We first estimate the reduced-form multinomial logit model in equation (5) using both the RFS and the SCF data. We present the RFS estimates in the main text, and show the replication with the SCF data in Appendix D.

The sample consists of all borrowers aged 25 to 74 at origination for whom all covariates are available. All specifications include alternative-specific year fixed effects. 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 experienced inflation. So, a positive coefficient indicates higher relative utility of, and probability of choosing, an FRM versus the baseline of an ARM.

Table 2 presents the estimation results. In our baseline specification (column 1), we restrict the coefficients on the FRM rate and the ARM initial rate from the PMMS to be the same (i. e., households only pay attention to the FRM−ARM rate spread). The negative coefficient estimate of $\hat{\hat{\beta}}_R = -0.483$ indicates that individuals derive less utility from, and are less likely to choose, the FRM when the FRM survey rate is higher. Turning to the variable of interest, we estimate a significant, positive coefficient of 0.220 for experienced 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 for a fixed-rate mortgage if their lifetime inflation experience were 1 pp higher. We take the total derivative of utility for alternative i and set it equal to zero (cf. Train 2009, ch. 3): $dU_{n,i} = \beta_R \partial Rate_{n,i} + \beta_{\pi,i} \partial \pi^e_n = 0$. This generates the following increase in individuals' willingness to pay for an FRM:

$$WTP := \left. \frac{\partial Rate_{n,i}}{\partial \pi^e_n} \right|_{dU_{n,i}=0} = -\frac{\beta_{\pi,i}}{\beta_R}. \tag{6}$$

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 2-5 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 the baseline in 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 4: Actual and Counterfactual FRM Shares

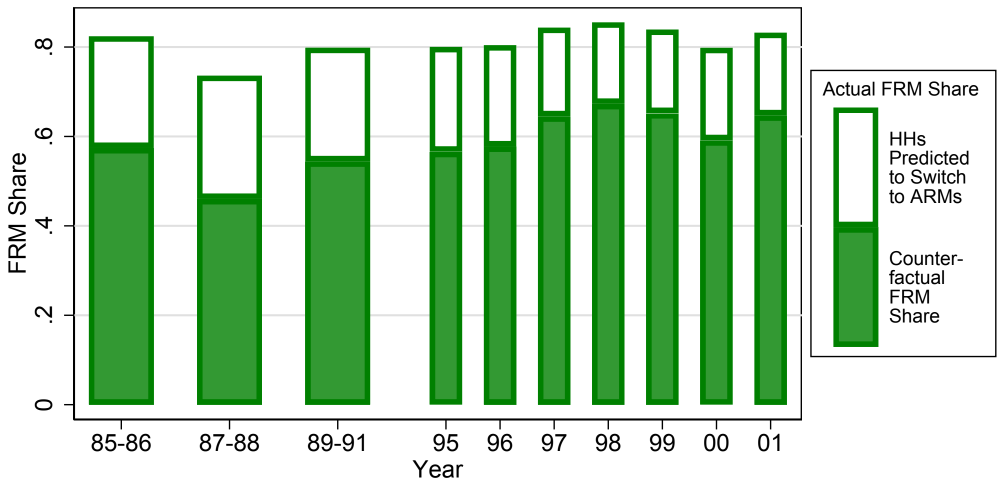

Notes. Data from the 1991 and 2001 RFS. The 1991 RFS reports origination years in two- or three-year intervals. The counterfactual FRM share is based on estimates from Table 2 column 4, with coefficient on experienced inflation set to zero.

To visualize the economic impact of the experience effect on aggregate mortgage choice behavior, Figure 4 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 2, 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 at the time and all information common to all households (including the full history of past inflation), it is sensible to compare the product shares with and without the experienced-inflation effect. The latter captures the prediction of the standard-model mortgage choice determinants.

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 experienced inflation diminishes as memory of the Great Inflation recedes (but 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 the choice of mortgage contract.

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 3 shows censored LAD estimates of the pricing equation (4) for $i \in \{FRM,\ ARM\}$. We use all of the exogenous explanatory variables from Table 2 in the first-step selection model in (5), 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.[11]

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.7) in Appendix C) for $1 \leq K \leq 4$, on all possible leave-one-out subsamples, for both the FRM and ARM rate equations.[12] 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. The coefficient on the PMMS survey rate is closer to unity in the selection-corrected model (97 vs. 84

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

[12] 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.

bps). We also estimate that married couples receive a bigger discount ("joint owners" coefficients of -18 vs. -4 bps without selection correction), possibly reflecting wealth effects; that mortgages in rural counties are more expensive than we would detect without selection correction (+33 vs. +12 bps); and that the jump in loan prices for jumbo loans whose dollar amounts exceed the conforming loan limit (CLL) and cannot be purchased by Fannie Mae and Freddie Mac is almost twice as big (+68 vs. +36 bps).

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 3 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.6) 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. As hypothesized, the data suggest 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 ($p = 0.78$). 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[13] 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 4 presents the estimates of (2), where the dependent variable indicates that the household chose an FRM. We use predicted interest rates from the pricing equations (Table 3) 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 (4). 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. The experience-effect estimates, instead, are robust with and without selection correction.

Columns 1 and 2 include only the FRM and the initial ARM teaser rate predictions 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 3, altering it from an (insignificant) estimate of 0.2 bps to a highly

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

significant estimate of -114 bps, we explore whether non-conventional status has an additional effect on mortgage choice, above and beyond its structural impact on mortgage prices. We test this by including non-conventional status as an additional explanatory variable in columns 5 and 6. This specification generates "correct," negative demand elasticities both with and without the selection correction, suggesting that future ARM resets do play a role in households' mortgage contract decision.

Turning to the variable of interest, we estimate consistent experience effects across all specifications. Higher levels of inflation experiences are associated with a greater probability of choosing 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 re-estimate the reduced-form, binomial mortgage choice model on 1989-2013 SCF data. While this survey has smaller sample size, fewer contract terms, and does not identify the geographic location of the borrower, it is conducted at a higher frequency and may suffer from less survivorship bias. All details are in Appendix D. As shown in Table A.2, the point estimates are remarkably similar to the RFS results. The replication in such different data 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 E 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. First, we test the validity of the logit choice model using the specification test of Horowitz and Härdle (1994). The 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 F.1. The result leads us to reject the logit model. 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 F.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 F.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 ($i = r + \mathbb{E}\pi$).[14] 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 much of the variation in lifetime experiences comes from the 1970s, when nominal interest rates and inflation rose together.[15] So, rather than running a horse race between the two, we investigate whether

[14] 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).

[15] Clarida et al. (2000) find a breakpoint in monetary policy in 1979: the pre-Volcker Fed was

this alternative specification generates similar results. In Appendix G, we re-estimate our reduced-form mortgage choice model, replacing $\pi^e_{n,t}$ with $i^e_{n,t}$. 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 motivating evidence in Section 4 that individuals coming of age during the Great Inflation expected not only higher inflation, but also higher nominal interest rates (cf. Figure 3), and that this experience significantly affected their valuation of fixed- versus variable-rate debt contracts.

5.3 Mortgage Amount Estimates

Our analysis so far has estimated the extent to which individuals with higher inflation experiences are more likely to choose an FRM. In this section, we show that they also originate and hold larger balances of fixed-rate liabilities. We follow the methodology laid out in Malmendier and Nagel (2016), who explore this question using the SCF.

For each birth-year cohort in the RFS, we calculate the average, per-capita amount of fixed- and adjustable-rate mortgage holdings and estimate:

$$log(\overline{Balance}_{c,t}) = \alpha_t + \alpha_{Age} + \beta_\pi \pi^e_{c,t} + \gamma log(\overline{HHInc}_{c,t}) + u_{c,t}, \tag{7}$$

where c indexes a birth-year cohort and t is the year of observation. Since we observe each cohort in multiple years, we can separately control for year and age fixed-effects. Unlike Malmendier and Nagel (2016), the RFS does not report household wealth.

Table 5 shows the estimation results. Columns 1-4 of the table calculate per-capita mortgage balances as of the RFS survey year ($t = 1991$ or 2001), among all homeowners from age 25 to 74 surveyed in the RFS, including individuals with zero mortgage balances. We provide separate estimates for the remaining balance of all mortgages (columns 1-2) and the remaining balances of mortgages originated recently, in the two years prior to the survey (columns 3-4). Columns 5-6 calculate per-capita mortgage originations among members of the cohort who originated a mortgage during the six years prior to the survey ($t \in \{1985\text{-}91, 1995\text{-}2001\}$).

All six columns tell a similar story. Cohorts with higher inflation experiences both hold higher amounts of fixed-rate mortgage balances in the survey year, and originate larger fixed-rate loans in the years prior to the survey year. However, we find no cor-

"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.

relation between inflation experiences and ARM balances or loan amounts, conditional upon the other control variables in equation (7). These results are reassuringly similar to the results reported in Malmendier and Nagel (2016), despite coming from a different survey covering different years and sampling from a different population (homeowners rather than consumers).

6 Financial Costs and Welfare Implications

Our evidence on mortgage choices and mortgage balances 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 experience effects are for consumers: Whether experience effects induce a welfare loss *ex post*, depends on the interest rate dynamics, and whether they induce a welfare loss *ex ante* depends on the hypothetical alternative choices and their expected costs and benefits.

In this section, we provide estimates of the financial costs of experience effects in 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 the financial costs of experience effects, we need to (1) identify whose choice is affected, and (2) calculate whether their experience-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 experienced inflation. The relevant subset are the "switchers:" households who chose an FRM only because experienced 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

$$h_n = \Pr(D_n = 1 | \beta_\pi = \beta_\pi) - \Pr(D_n = 1 | \beta_\pi = 0), \tag{8}$$

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 2 or 4 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 3 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. 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 3, 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 3, 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).

The scenarios are summarized below, where each scenario makes progressively greater adjustments for risk characteristics, at the cost of increasing sensitivity to our modeling assumptions. However, since many ARM contracts have caps on annual or lifetime interest-rate adjustments (while our simulated ARMs are not capped), all three scenarios overstate the amount of interest-rate risk in an ARM and underestimate the potential savings from choosing an ARM over an FRM. Hence, this simulation assumption biases against the hypothesis that experienced inflation is welfare-reducing and generates conservative estimates.

Scenario:	**1**	**2**	**3**
FRM Rate	Freddie Mac PMMS	Risk-adj. (CLAD)	Risk-adj. (CLAD)
ARM Initial Rate	Freddie Mac PMMS	Risk-adj. (CLAD)	Risk-adj. (CLAD)
ARM Margin	1-year T-bill + 2.75	Seniority-adj. (CLAD)	Risk-adj. (OLOGIT)
Choice probabilities	Table 2, col. 5	Table 4, col. 2	Table 4, col. 6

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 simulated *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 ΔY_n across all households, using their switching probabilities as weights:

$$\widehat{WRTE} := \frac{1}{N} \sum_{n=1}^{N} \Delta \hat{y}_n \cdot \left(\frac{\hat{h}_n}{\sum_n \hat{h}_n} \right). \tag{9}$$

We show in Appendix C.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.[16] 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,

[16] 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.

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 experienced 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 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 3 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.[17] 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 gives us a set of probabilities describing, t years after origination, the probability that the household holds a mortgage last (re-

[17] 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).

)financed s years after origination, $0 \le s \le t$. We use these probabilities to calculate the household's expected FRM payments across the entire distribution of possible time-t interest rates. See Appendix C.3 for further details.

Illustration. Table 6, Panel A, illustrates these calculations for one of our sample households. "Potential FRM Rates" are the rates that we estimate this household would receive if it refinanced into a new FRM that year based on Scenario 3. The "Optimal Threshold" for refinancing reflects the household's current outstanding mortgage balance and time remaining to maturity according to ADL's square-root rule (equation (A.13)). At the beginning of year 4, mortgage rates had fallen by 100 basis points since year 1 (= 10.42 – 11.42). This exceeds the optimal refinancing threshold of 96 basis points, so the household should refinance. However, we estimate that there is only a 29% chance that the household will do so. This reflects both types of refinancing errors discussed earlier – the household might have already refinanced in year 2 or 3, or the household might ignore the incentive and wait an extra year.

Figure 5: FRM Rate for Mortgage ID 500

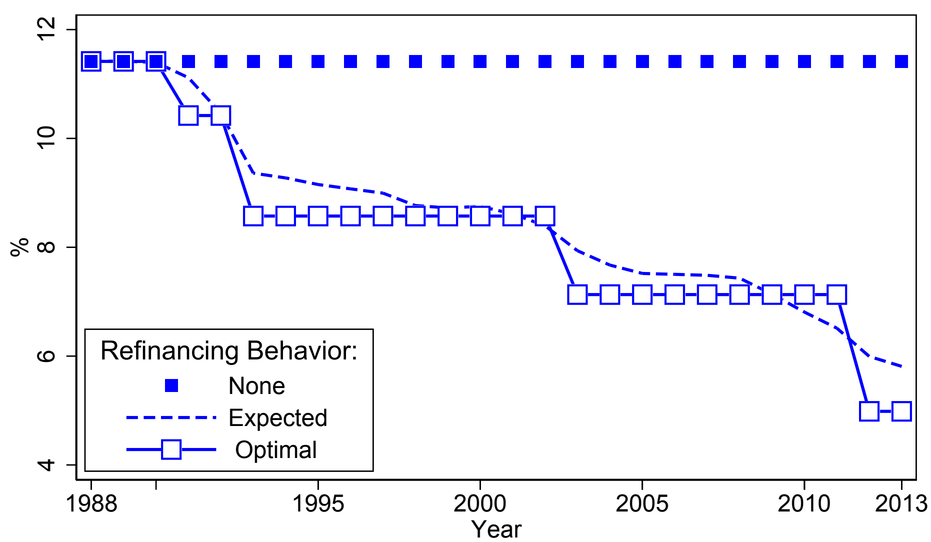

Notes. Date from the 1991 RFS. Refinancing behavior is described in notes to Table 6. The expected FRM rate for a household n is integrated over all potential time-t rates using the time-t state probability distribution: $E[i_{n,t}] = \int_{1 \le s \le t} i_{n,s} P_n(S_t = s) ds$.

Figure 5 shows the paths of FRM rates for this household under the three cases graphically. Optimal refinancing occurs in years 4, 6, 16, and 25. "Expected" rates, integrating over all potential rates using the time-t state probability distribution, track "optimal" rates with a lag, indicating the importance of delayed refinancing behavior.

Panel B of Table 6 illustrates mortgage payments for the same household under the three sets of assumptions about refinancing. Under "Expected Refinancing," the present value of interest payments from 1988 to 2012 is $200,000 using an 8% nominal discount rate. Assuming that the household faces a marginal tax rate of 25%, the present value of its interest deductions are $50,000. The present value of refinancing costs (which are not tax deductible) is $4,600, generating a bottom-line cost of $154,000.

The present value of FRM payments is highest if never refinanced, and lowest if optimally refinanced. Regardless of refinancing behavior, though, an ARM would have been cheaper over this time horizon. Under optimal behavior, the FRM is over $27,000 more expensive. This figure rises to $32,000 if the household refinances as expected, and could be as high as $54,000 if the household never refinances. The WRTE is the average value of this difference across all households, weighted by the households' switching probabilities.

Simulation Results. Turning to the full simulation results, we show the cost estimates for all switching households in Table 7. 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.

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 repeat 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 H 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 7, 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 $3,5000-$7,000 in PDV terms.[18] 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.[19] 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

[18]Using an estimate of 6-12 bp per percentage point of lifetime inflation experiences for an average-sized mortgage, this amounts to $60-$120 per year, or $700-$1,400 over 30 years (discounting at 8%) per pp of lifetime inflation experiences, times 4.75 pp.

[19]Cf. investopedia.com/articles/pf/06/payingforpoints.asp or 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.

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 H, 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 purchase 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 adjustment in Appendix H, 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 historical inflation subsequent to each mortgagor's origination date. 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).[20] We use Scenario 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 (9) that we change is the subsequent interest payments ΔY_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

[20] Since ARMs were not introduced in the U.S. until 1982, we impute the Freddie Mac PMMS initial ARM rate for 1971 and 1981 assuming that it would have taken its average value over the 1-year constant-maturity Treasury rate of 1.5 percentage points.

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 assess whether the cost we estimate is a reasonable price for households wishing to reduce their exposure to inflation risk, we turn to a wider range of possibilities and simulate 100 different inflation environments.

In each simulation, we draw 30 years of inflation and nominal mortgage rates. We simulate the economic environment using the following processes. 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:

$$y_t^T = \frac{1}{T}\sum_{s=1}^{T} \mathbb{E}_t y_{t+s-1}^1 + \theta^T, \tag{10}$$

where $\mathbb{E}_t y_{t+s}^1 = \rho + \phi^s(\pi_t - \mu) + \mu$. ARM rates equal the one-year nominal bond rate plus a term premium: the ARM teaser 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, each simulation has two independent sources of variation: the sequences of inflation rates and the sequences of one-year real interest rates. All other variables are derived by exact, linear relationships. Table 8 gives the values and sources for all the simulation parameters.

Figure 6 plots the counterfactual WRTEs against average inflation in each of the 100 simulations. As expected, we observe a strong inverse correlation between 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 simulations with average inflation exceeding about 5.5%, the expected WRTE becomes negative, indicating that the FRM is *ex post* cheaper. For the large majority of scenarios (82% of simulations), however, average inflation is below 5.5% and the expected the choice of FRM implies a welfare loss for the switching households.

The scatter plot suggests that FRMs were not unusually expensive given actual

Figure 6: Average Inflation and E[WRTE] in 100 Simulations

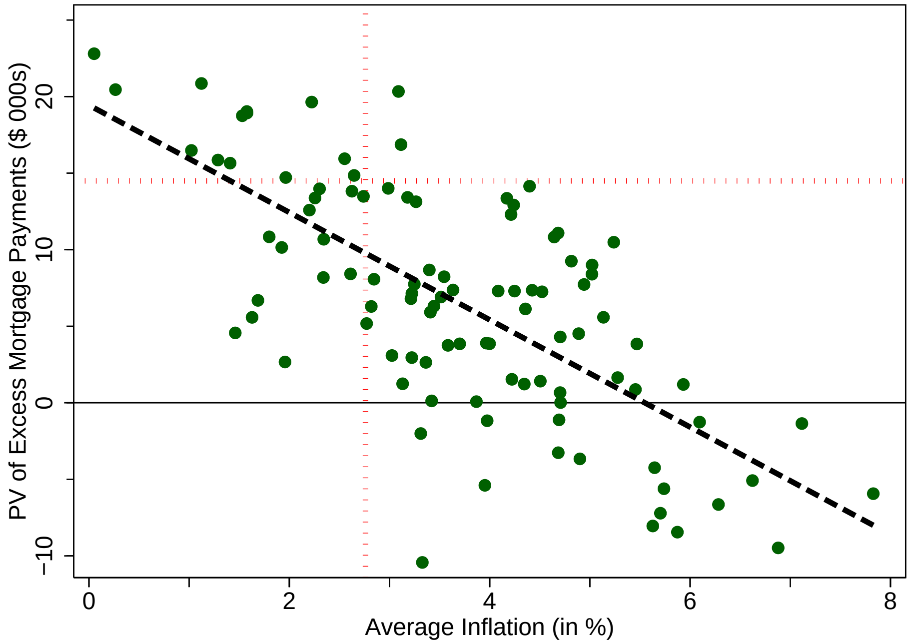

Notes. Each point represents one simulation. Calculations based on Scenario 3 estimates and age-based mobility. Thick dashed line is OLS regression line across 100 simulations. Horizontal and vertical crossing lines indicate average values of inflation over 1986-2013 and E[WRTE] from the RFS.

subsequent economic conditions in the 1990s and 2000s. In fact, the simulations assume that inflation reverts to a long-run mean of 3.8%, based on an historical average that includes the Great Inflation. However, U.S. inflation averaged only 2.8% over 1986-2013. In simulations with similarly low values of average inflation, we predict that the WRTE would normally be $9,718 (s.e. 632). The bottom line is that while the low-inflation experience of the 1990s was particularly 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 overweighting their personal inflation experiences.

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 2, 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 9). 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 mistakes 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 12 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 Gen-X generation of households who were first-time homeowners then.

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.

Andrews, D. and A. Caldera Sanchez (2011). Drivers of homeownership rates in selected OECD countries. *OECD Economics Department Working Papers* (849).

Andrews, D., A. Caldera Sanchez, and A. Johansson (2011). Housing Markets und Structural Policies in OECD Countries. *OECD Economics Department Working Papers* (836).

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.

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.

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. (1992). Borrower Mobility, Self-Selection, and the Relative Prices of Fixed- and Adjustable-Rate Mortgages. *Journal of Financial Intermediation 2*(4), 401–421.

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.

Campbell, J. Y., H. E. Jackson, B. C. Madrian, and P. Tufano (2011). Consumer Financial Protection. *Journal of Economic Perspectives 25*(1), 91–114.

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.

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.

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.

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.

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.

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.

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. 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.

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.

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

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 Fi-*

nance 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.

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.

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.

Moench, E., J. Vickery, and D. Aragon (2010). Why is the market share of adjustable-rate mortgages so low? *Current Issues in Economics and Finance, Federal Reserve Bank of New York 16*(8), 1–11.

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

Nathanson, C. G. and E. Zwick (2018, 12). Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market. *The Journal of Finance 73*(6), 2587–2633.

Newey, W. K. (2009, January). Two-Step Series Estimation of Sample Selection Models. *Econometrics Journal 12*, S217–S229.

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.

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.

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.

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

Table 1: 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*
KHN decision rule	0.34	0.45	0.38	-0.12*
Borrower Characteristics				
Primary owner age	41.4	41.8	42.8	-0.4
Experienced inflation (%)	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. The KHN (2009) decision rule is the difference between the five-year Treasury yield and a three-year moving average of the one-year Treasury yield. All other variable definitions are in Appendix A. * $p < 0.05$.

Table 2: 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)
Experienced inflation (%)	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					
Experienced inflation (%)	-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
Sociodemographic 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 $\leq$ 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 3: 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 rate (%)	84.21***	96.61***	77.28***	86.71***	-11.83***	-6.513**
	(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 4: 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)
Experienced inflation (%)	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
$-\beta_{\pi,\text{FRM}} / \beta_{\text{Rate, FRM}}$	-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 3 and 2, 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 5: Inflation Experiences and FRM Balances

	(1)	(2)	(3)	(4)	(5)	(6)
	All Balances		Recent Balances		Loan Amount	
	Log(FRM)	*Log(ARM)*	*Log(FRM)*	*Log(ARM)*	*Log(FRM)*	*Log(ARM)*
Experienced	0.352***	0.0951	0.740***	0.284	0.304**	-0.159
Inflation (%)	(0.05)	(0.13)	(0.15)	(0.24)	(0.14)	(0.19)
Log(Household	0.523*	1.807*	0.084	-1.117	0.502***	0.607***
Income)	(0.26)	(0.90)	(0.61)	(1.40)	(0.07)	(0.22)
Survey Year FE	YES	YES	YES	YES		
Origination Year FE					YES	YES
Age FE	YES	YES	YES	YES	YES	YES
As of	Survey Year		Survey Year		Origination Year	
Sample	Homeowners		Homeowners		Mortgagors	
Observations	100	100	100	97	490	408
R-squared	0.989	0.889	0.960	0.860	0.573	0.284

Notes. The table reports OLS regressions of log per-capita mortgage amounts on birth-year cohorts' inflation experiences as of the survey or origination year. Each observation is a cohort-year average, in columns 1-4 over the sample of all homeowners with a head of household between ages of 25 and 74 as of the survey year, and in columns 5-6 over the sample of all mortgagors with a head of household between 25 and 74 as of the origination year. The dependent variable is the natural log of the per-capita mortgage holdings, defined as the outstanding mortgage balance as of the survey year in columns 1-2; as the outstanding balance of mortgages financed $\leq$ 2 years prior to survey year in columns 3-4; and as the original loan amount of mortgages financed $\leq$ 6 years prior to survey year in columns 5-6, in all cases aggregating senior and junior liens at the household level. Log transformations are applied to per capita amounts, dropping three cohort-years with zero balances in column 4. All amounts are deflated to constant year-2000 dollars using CPI-U. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 6: Refinancing Behavior and Interest Payment Calculations for a Sample Household

Panel A: Refinancing Behavior

Year	Potential FRM Rates	Optimal Threshold (pp)	Probability of refinancing (%)	FRM Rate, Optimal Refi	ARM Rate
1	11.42	0.95	0.0	11.42	7.70
2	11.45	0.95	13.0	11.42	11.58
3	11.27	0.95	15.2	11.42	10.92
4	10.42	0.96	28.7	10.42	8.90
5	9.62	0.96	40.8	10.42	6.93
6	8.57	0.97	50.2	8.57	6.47
7	9.58	0.97	14.1	8.57	8.35
8	9.14	0.98	17.2	8.57	8.99
9	9.04	0.99	15.9	8.57	8.55
10	8.86	1.00	16.6	8.57	8.67
⋮	⋮	⋮	⋮	⋮	⋮

Panel B: Present Value of Interest Payments

	FRM ($)			ARM ($)
	No Refi	Expected Refi	Optimal Refi	
PDV	**235,498**	**199,637**	**193,659**	**163,074**
- Int. Deduct.	-58,874	-49,909	-48,415	-40,768
+ Refi Cost	0	4,633	3,895	0
Total	**176,623**	**154,361**	**149,139**	**122,305**

Notes. The table illustrates how refinancing behavior affects the present value of interest payments under different contract types for a sample household (mortgage ID 500 in our data). All dollar figures are in constant year-2000 units. This joint-owner household took out a loan for $200,751 in 1988 on a property in the Midwest Census region. The head of household was 30 years old, and the household reported total income of $166,972. **Panel A** illustrates the first ten years of interest rates. Potential FRM rates are predicted based on Freddie Mac PMMS rate in region-year and HH risk characteristics (Table 3, column 2). The optimal threshold for refinancing is calculated using the Agarwal et al. (2013) square-root rule. The conditional probability of refinancing from previous interest rate i_0 to current rate i is based on Andersen et al. (2015) Table 8 column 1: $Pr(\text{refi}|i_0) = \Phi\left(-1.921 + e^{1.033}(i_0 - i - OT)\right)$, converted from monthly to annual horizon. The column labeled "Probability of Refinancing" reports the unconditional probability of refinancing into a time-t FRM, given all possible previous interest rates i_0: $\sum_{s<t} Pr(\text{refi}|i_0 = i_s) \times Pr(i_0 = i_s)$. The ARM rate in year 1 is predicted from the Freddie Mac PMMS rate in the origination region-year and HH risk characteristics (Table 3, column 4). Years 2-30 are 1-year US Treasury + HH's predicted margin, estimated by ordered logit (Table 3, column 6). **Panel B:** In the "No Refi" column, the household holds the initial FRM until maturity. In the "Expected Refi" column, the household is assumed to refinance probabilistically, according to the probit function of the interest rate differential estimated in Andersen et al. (2015). (The timing of principal repayment is the same as in Optimal Refi scenario.) In the "Optimal Refi" column, the household refinances deterministically whenever $i_0 - i > OT$. PDV calculations assume a nominal discount rate of 8% / year ($r = .04$, $\pi = .04$). The mortgage interest deduction is calculated assuming a 25% marginal tax rate. Refinancing costs $2,000 and is not tax-deductible.

Table 7: 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 2 col. 5, scenario 2 = Table 4 col. 2, scenario 3 = Table 4 col. 6). PDV calculations assume a nominal discount rate of 8% / year (r = .04, π = .04). "No Refi," "Expected Refi," and "Optimal Refi" are defined as in Table 6. 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 8: 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

Notes. Inflation follows an AR(1) process with normally-distributed innovations. Real one-year interest rate innovations are independent, normally-distributed, and mutually-independent with the inflation innovations. One-year nominal interest rates equal real rates plus inflation. Long-term rates are calculated via the expectations hypothesis. PMMS data begin in 1971 for FRMs, and 1984 and 1987 for ARMs.

Table 9: Aggregate Cost of the Great Inflation

(1) Survey Year - Cohort	(2) % Switching HHs	(3) E[WRTE] per switching HH ($)	(4) # of switching HHs (1000s)	(5) Total Cost ($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 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. Dollar figures are in constant year-2000 units. For purposes of this table, 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. The % of Switching HHs in column (2) is the predicted change in the FRM product share per survey year-cohort, under the hypothetical inflation experiences if the Great Inflation had not occurred, i.e., under the location-scale transformation of actual 1968-84 inflation to $\mu = 2.5\%$, $\sigma = 1.1\%$ as shown in Figure 2. The Expected WRTE in column (3) is calculated assuming expected refinancing behavior and tenure given the head of household's age. Calculations use Scenario 3 interest rates. Our calculation of the number of switching households in column (4) assumes that every sample household represents 2,599 population households in 1991 and 3,655 population households in 2001. The Total Cost in column (5) is the product of the E[WRTE] and number of switching households: (5) = (3) $\times$ (4).

Online Appendix

A Variable Definitions

Variable	Units	Definition	Source
Experienced Inflation	%	Weighted average inflation (log change in annual average CPI-U) over 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
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.	Census Bureau RFS
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.	Census Bureau RFS
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.	Census Bureau RFS
Joint owners	$\{0, 1\}$	=1 if number of property owners exceeds one.	Census Bureau RFS
Rural county	$\{0, 1\}$	=1 if property is located outside of an MSA.	Census Bureau RFS
Junior mortgage	$\{0, 1\}$	=1 for second or third mortgage on a property.	Census Bureau RFS
Non-conventional	$\{0, 1\}$	=1 if mortgage is FHA-, VA-, or FmHA/RHS-insured or guaranteed.	Census Bureau RFS

Variable	Units	Definition	Source
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.	Census Bureau RFS
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.	Census Bureau RFS
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.	Census Bureau RFS; Fannie Mae
Jumbo loan	$\{0, 1\}$	=1 if Loan / CLL > 1.	Census Bureau RFS; Fannie Mae
Points paid	% or bps	Discount points paid as interest at inception of first mortgage, excluding loan origination and non-interest fees.	Census Bureau RFS

B Dating the Great Inflation

Figure A.1

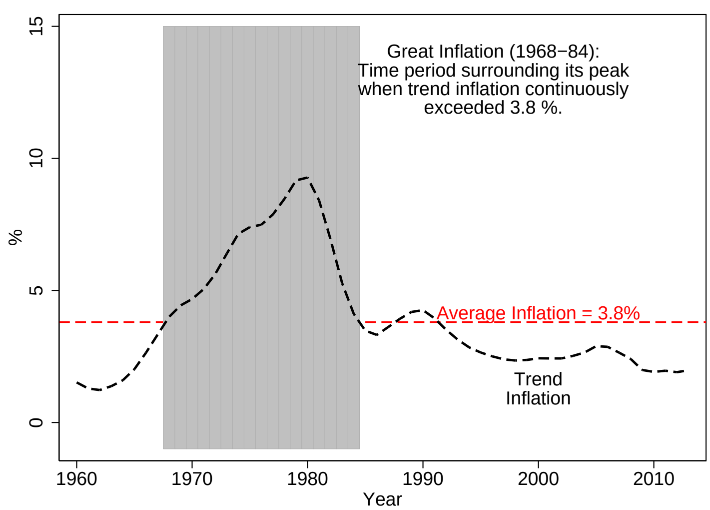

We determine the dates for the Great Inflation in a data-driven manner, proposed by Scrimgeour (2008). We first extract the trend component of BLS CPI-U log annual inflation using a triangular moving-average filter:

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

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.

C Methodology in Detail

C.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 (5)).

2. Predict FRM and ARM mortgage rates at the household level, correcting for selection bias (equation (4)).

3. Estimate a structural model of mortgage choice using individual-level predicted mortgage rates (equation (2)).

We begin by assuming that a household n derives utility $U_{n,i} = x'_{n,i}\beta_i + \varepsilon_{n,i}$ 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 $\varepsilon_{n,i}$. 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

$$\begin{aligned} D_{n,i} &\coloneqq \mathbb{I}\{U_{n,i} > U_{n,j} \quad \forall j \neq i\} \\ &= \mathbb{I}\{\varepsilon_{n,j} - \varepsilon_{n,i} < x'_{n,i}\beta_i - x'_{n,j}\beta_j \quad \forall j \neq i\} \end{aligned} \tag{A.2}$$

equals 1.[21] Marley (cited by Luce and Suppes (1965)) and McFadden (1974) show necessary and sufficient conditions on the distribution of the unobserved utility components ε_{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 scale).[22] The definition of $D_{n,i}$ implies that,

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

[22] 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 ε's to $\pi^2/6$ before estimating the coefficients.

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 ((2) in the paper):

$$U_{n,i} = \beta_{0,i,y} + \beta_{R,i} Rate_{n,i} + \beta_{\pi,i}\pi^e_n + \beta_{Inc,i} Income_n + f_i(Age_n) + \varepsilon_{n,i}, \tag{A.3}$$

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 would use the subset of households n choosing alternative i to estimate the following equation ((4) in the paper) for all J all alternatives:

$$\begin{aligned} Rate_{n,i} &= \gamma_{0,i} + Z'_{n,i}\Gamma_{n,i} + v_{n,i} \\ &= \gamma_{0,i} + \gamma_{R,i} PMMSRate_{y,r,i} + z'_n\gamma_i + v_{n,i}. \end{aligned} \tag{A.4}$$

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 $v_{n,i}$ captures all remaining, unobserved factors that affect the interest rate for alternative i being offered to household n.

The goal of estimating equation (A.4) 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: $\mathbb{E}\left[v_{n,i}|Z_{n,i}, D_{n,i}=1\right] = f(Z_{n,i}) < 0$. 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 $\mathbb{E}\left[\operatorname{sgn}(v_{n,i})|Z_{n,i}\right] = 0$, rather than the usual conditional mean restriction $\mathbb{E}\left[v_{n,i}|Z_{n,i}\right] = 0$, 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.4).

Although our coefficient estimates from (A.4) do not provide us directly with predicted rates, we can plug them into (A.3) and obtain a *reduced-form* choice model that we can estimate ((5) in the paper):

$$\begin{aligned} U_{n,i} &= \tilde{x}'_{n,i}\tilde{\beta} + \tilde{\varepsilon}_{n,i} \\ &= \tilde{\beta}_{0,i,t} + \tilde{\beta}_{R,i} PMMSRate_{y,r,i} + \beta_{\pi,i}\pi^e_n + \tilde{\beta}_{Inc,i} Income_n + f_i(Age_n) + \tilde{z}'_n\tilde{\gamma}_i + \tilde{\varepsilon}_{n,i}. \end{aligned} \tag{A.5}$$

We place tildes on coefficients and variables that represent different objects than in equation (A.3). For example, the coefficient on the PMMS rate in equation (A.5) is the structural coefficient from equation (A.3), scaled by the partial correlation between household interest rates and PMMS rates from equation (A.4): $\tilde{\beta}_{R,i} \coloneqq \beta_{R,i}\gamma_{R,i}$. We write $\tilde{z}_n$ to represent the subset of variables in z_n from equation (A.4) that do not appear directly in (A.3) (e.g., excluding household income). The pricing errors from (A.4), v_{ni}, are absorbed into the unobserved component of latent utility: $\tilde{\varepsilon}_{n,i} \coloneqq \varepsilon_{n,i} + \beta_{R,i}v_{n,i}$.

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.4), 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.5) first, (A.4) second, and (A.3) third. Model (A.5) 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.4) semiparametrically. Specifically, let $\tilde{\eta}_{n,i,j} \coloneqq \tilde{x}'_{n,i}\tilde{\beta}_i - \tilde{x}'_{n,j}\tilde{\beta}_j$ 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.4) as

$$\begin{aligned} v_{n,i} &= \mathbb{E}\left[v_{n,i} \mid Z_{n,i}, D_{n,i}=1\right] + w_{n,i} = \mathbb{E}\left[v_{n,i} \mid Z_{n,i}, \tilde{\varepsilon}_{n,j} - \tilde{\varepsilon}_{n,i} < \tilde{\eta}_{n,i,j}\ \forall j \neq i\ \right] + w_{n,i} \\ &= g(\tilde{\eta}_{n,i,1}, \ldots, \tilde{\eta}_{n,i,J}) + w_{n,i}, \end{aligned} \tag{A.6}$$

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 (4), we obtain

$$Rate_{n,i} \approx \gamma_{R,i} PMMSRate_{y,r,i} + z'_n\gamma_i + \sum_{k=0}^{K} \tau_k \cdot p(\tilde{\eta}_{n,i,j})^k + w_{n,i}. \tag{A.7}$$

In the special case where $K = 1$ and $p(\cdot)$ is the inverse of Mill's ratio, equation (A.7) 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.7) drops the intercept $\gamma_{0,i}$ from (A.4) 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.3). As mentioned, the pricing equation intercept $\gamma_{0,i}$ is not identified in the two-step series estimator (A.7). 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.

C.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 (3), the value of D_n depends on the difference in latent utility in equation (2) between the alternatives: the FRM is chosen if the difference in observed components of latent utility exceed the difference in unobserved components, $-(\varepsilon_{n,1} - \varepsilon_{n,0}) < x'_{n,1}\beta_1 - x'_{n,0}\beta_0$. Observed latent utility may include alternative characteristics, such as prices, as well as household characteristics, and experienced inflation. The coefficients in Table 4 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

$$D_n = \int A_n(\beta_\pi) D_n(b_\pi) db_\pi, \tag{A.8}$$

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.[23]

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

$$\mathbb{E}[Y_{n,1} - Y_{n,0} | D_n(\beta_\pi) = 1, D_n(0) = 0], \tag{A.9}$$

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 $Y_{n,1} - Y_{n,0}$ 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.9), because we do not observe households' counterfactual choices $D_n(0)$. However, Bayes' rule lets us rewrite (A.9) as

$$\begin{aligned}
\mathbb{E}\left[\Delta Y_n | D_n(\beta_\pi) = 1, D_n(0) = 0\right] &= \int \Delta y \cdot f(\Delta y | D_n(\beta_\pi) = 1, D_n(0) = 0) d\Delta y \\
&= \frac{\int \Delta y \cdot h(D_n(\beta_\pi) = 1, D_n(0) = 0 | \Delta y) f(\Delta y) d\Delta y}{g(D_n(\beta_\pi) = 1, D_n(0) = 0)}.
\end{aligned} \tag{A.10}$$

The first line of equation (A.10) 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.10) 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:

$$h(D_n(\beta_\pi) = 1, D_n(0) = 0 | \Delta y) = \Pr(D_n = 1 | b_\pi = \beta_\pi, \Delta y) - \Pr(D_n = 1 | b_\pi = 0, \Delta y). \tag{A.11}$$

[23] 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 $\pi^e > 0$.

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 2 or 4. 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 3.

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):

$$\begin{aligned}\widehat{WRTE} &:= \widehat{\mathbb{E}}\left[Y_{n,1} - Y_{n,0} | D_n(\beta_\pi) = 1, D_n(0) = 0\right] \\ &= \sum_{n=1}^{N} \Delta\hat{y}_n \cdot \left\{ \frac{\widehat{\Pr}(D_n(\hat{\beta}_\pi) = 1 | \Delta\hat{y}_n) - \widehat{\Pr}(D_n(0) = 1 | \Delta\hat{y}_n)}{\sum_n \left(\widehat{\Pr}(D_n(\hat{\beta}_\pi) = 1 | \Delta\hat{y}_n) - \widehat{\Pr}(D_n(0) = 1 | \Delta\hat{y}_n)\right)} \right\}, \end{aligned} \tag{A.12}$$

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.[24]

C.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 to the optimal threshold:

$$OT_{n,t} \approx -\sqrt{\frac{\sigma\kappa}{M_{n,t}(1-\tau)}\sqrt{2(\rho + \lambda_{n,t})}}, \tag{A.13}$$

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,

[24] 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.

28%.) The mortgage prepayment process parameterized by $\lambda_{n,t}$ is derived from three exogenous sources of principal repayment:

$$\lambda_{n,t} = \mu + \frac{i_n}{\exp(i_n(T-t)) - 1} + \pi \tag{A.14}$$

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

$$\Pr(\mathit{Refi}_{n,y,m}|i_0) = \Phi\left(-1.921 + \exp(-1.033) \times (OT_{n,y} - (i_{n,y} - i_0))\right), \tag{A.15}$$

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.[25] 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 previous 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

$$P_n(S_t = t|S_{t-1} = s) \coloneqq \Pr(\mathit{Refi}_{n,\mathit{OrigYr}_n+t}|i_0 = i_{n,\mathit{OrigYr}_n+s}) \cdot \mathbb{I}\{s < t\}. \tag{A.16}$$

$S_t \in \{0, 1, 2, ..., t\}$ 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.[26]

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

[26] 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.

D Mortgage Choice using the SCF

We replicate the reduced-form mortgage-choice model of equation (5) from Table 2 using SCF data. The advantage of the SCF data over the RFS is the availability of more survey waves; the disadvantage is the lack of mortgage-specific data on interest rates and on some of the other controls we employ in the RFS analysis. For example, the SCF provides information on "percentage points paid" and about "refinance" and "non-conventional mortgage" also for mortgages on the secondary homes, neither of which is available in the SCF. In addition, the SCF provides records whether the respondents are first-time home owners only from the 2007 survey on. The lack of these controls, as well as the significantly smaller sample size of the SCF, might reduce the precision of our estimates relative to the RFS.

For the replication analysis, we pool the full public data from every (triennial) survey wave from 1989-2013, merged with the summary extract public data sets to obtain each household's net worth. We keep information on the primary mortgage and any secondary mortgages on the principal residences, plus the first two reported mortgages on any secondary residences. Each mortgage enters as a separate observation (choice situation) in our analysis. Both income and net worth are deflated to constant year-2013 dollars using the CPI-U-RS from BLS. We exclude a small number of households reporting negative values for income or net worth. We include mortgages that were originated up to two years before the survey year for all the surveys except 1989; we extend the origination year back to 1985 for the 1989 SCF in order to match the time period with the RFS. (The 1986 survey differs in design from the later surveys, so is not directly comparable.)

Table A.2 reports estimates of the reduced-form binomial mortgage-choice model in (5), including as many of the same controls as possible from Table 2, column 5. All point estimates use SCF sample weights, and we adjust the standard errors for multiple imputation using the standard Rubin (1987) formulas. The first column reports binomial logit coefficients for the full 1985-2013 time period, and the second column restricts to the same origination years as the RFS: 1985-1991 and 1995-2001. Note that we cannot include the Freddie Mac PMMS rate indices since we do not observe borrowers' geographic locations, and the remaining time-series variation in the national rate indices is absorbed by the year fixed effects.

Despite the differences in controls and sample size, the estimate of the within-origination year effect of lifetime inflation experiences is remarkably similar. The point estimate for the logit index coefficient is between 0.25 and 0.27 (compared to 0.19 in Table 2, column 5). The precision of the estimates is lower, likely due to the much smaller sample sizes of the SCF surveys. For example, the RFS has 60% more observations (14,337/8,929 = 1.606) than column 1, so we would expect the SCF standard error to be about 27% larger ($\sqrt{1.606} \approx 1.27$). The actual standard error in column 1 is 36% larger than in the RFS table (= 0.133/0.098 - 1), and statistically significant at the 10% level in column 1. When we further shrink the sample to match the sample period of the RFS sample, 1985-1991 and 1995-2001, the sample to about one fifth (only 3,161 observations), the coefficient remains again very stable, at 0.273, though the estimate is noisily estimated. (We would expect the SCF standard error to be 113% larger, and estimate it to be 150% larger.)

Table A.2: Logit Model of Mortgage Choice using the SCF

	Origination year coverage	
	1985-2013	**1985-1991; 1995-2001**
Experienced inflation in %	0.246*	0.273
	(0.133)	(0.250)
Log(Normal Income)	0.077	0.294**
	(0.078)	(0.144)
Log(Net worth)	-0.023	-0.077
	(0.039)	(0.059)
Joint owners? (=1 if married)	0.110	0.006
	(0.100)	(0.158)
Junior mortgage dummy	-0.201	0.307
	(0.126)	(0.198)
Nonconventional dummy	0.427***	0.345*
	(0.125)	(0.200)
Loan-to-CLL ratio	-0.175***	-0.151**
	(0.050)	(0.074)
Jumbo dummy	-0.759***	-0.647***
	(0.140)	(0.235)
Observations (per imputation)	8,929	3,161
Pseudo R2	0.079	0.073
Age & age^2 controls	YES	YES
Origination year FE	YES	YES

Notes. The table reports binomial logit coefficients, adjusted for multiple imputations. The dependent variable is an indicator equal to 1 for FRMs, and 0 for ARMs. Each observation is a mortgage. Balloon mortgages are excluded. All sample mortgages are originated between 1985 and 2013, using the most recent wave of the Survey of Consumer Finances (administered in 1989, 1992, ..., 2013). Column (1) includes the mortgages originated ≤ 2 years ago for SCF 1992-2013 and originated after 1985 for SCF 1989. Column (2) further restricts origination to occur in 1985-1991 or 1995-2001, making it the same as in the RFS analysis. Both regressions use SCF "revised consistent" sampling weights (variable X42001). The weights are scaled so that each survey wave receives equal weight. We adjust for multiple imputations using the Rubin (1987) methodology: (i) Point estimates are the average of coefficients, estimated separately within each imputation; and (ii) the multiple-imputation variance-covariance matrix V is $V = U + (1 + 1/M) \times B$, where U is the average within-imputation VCV matrix, B is the between-imputation VCV matrix, and M is the number of imputations ($M = 5$ in the SCF). Income and net worth are adjusted for inflation (2013 \$). We use the normal income concept (variable X7326) starting in 1995, and total household income (variable X5729) in 1989 and 1992. We drop observations with negative values before taking logs. The numbers of observations refer to the numbers in a single imputation. Pseudo R^2 is the average of the Logit pseudo R^2 on each of the five imputations. Robust standard errors, adjusted for multiple imputations, in parentheses. *** p<0.01, ** p<0.05, * p<0.1

E Robustness Check: Supply-Side Constraints

Table A.3: 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
Socidemographic 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

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.3, we re-estimate the reduced-form binomial choice model separately on the 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 2, column 5). We conclude that supply-side constraints in the mortgage lending process are not driving our results.

F Robustness Check: Alternate Estimation Methods

F.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 (2), 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,

$$H_0 : E[D_n|x_n] = F(x_n'\beta) \text{ versus } H_1 : E[D_n|x_n] = G(x_n'\beta), \tag{A.17}$$

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

$$HH \coloneqq h^{1/2} \sum_{n=1}^{N} \omega_n \cdot (D_n - F(x_n'\hat{\beta})) \cdot (\hat{F}(x_n'\hat{\beta}) - F(x_n'\hat{\beta})). \tag{A.18}$$

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

$$V = 2 \int K(u)^2 du \cdot \int \omega(z)^2 \sigma^4(z) dz. \tag{A.19}$$

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

under H_0 is

$$\hat{V} = 2\int K(u)^2 du \cdot \frac{1}{N}\sum_{n=1}^{N} \omega_n^2 \frac{[(F(x_n'\hat{\beta}))(1-F(x_n'\hat{\beta}))]^2}{\hat{f}(x_n\hat{\beta})}. \tag{A.20}$$

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 $\sigma_n^2 = \mathrm{Var}(D_n|x_n'\hat{\beta})$ 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:

$$K^{(4)}(u) \coloneqq \frac{15}{8}\left(1 - \frac{7}{3}u^2\right) \times K^{(2)}(u) \times \mathbb{I}\{|u| \leq 1\}, \tag{A.21}$$

$$\text{where } K^{(2)}(u) \coloneqq \frac{3}{4}\left(1 - u^2\right) \times \mathbb{I}\{|u| \leq 1\}. \tag{A.22}$$

$K^{(2)}$ 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.2 illustrates the bandwidth selection procedure. The index $x'\hat{\beta}$ is calculated using the reduced-form, binomial logit choice model coefficients reported in Table 2, 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.3 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 ($h \in \{0.25, 0.45, 0.65\}$).

Figure A.2: Cross-validation Function

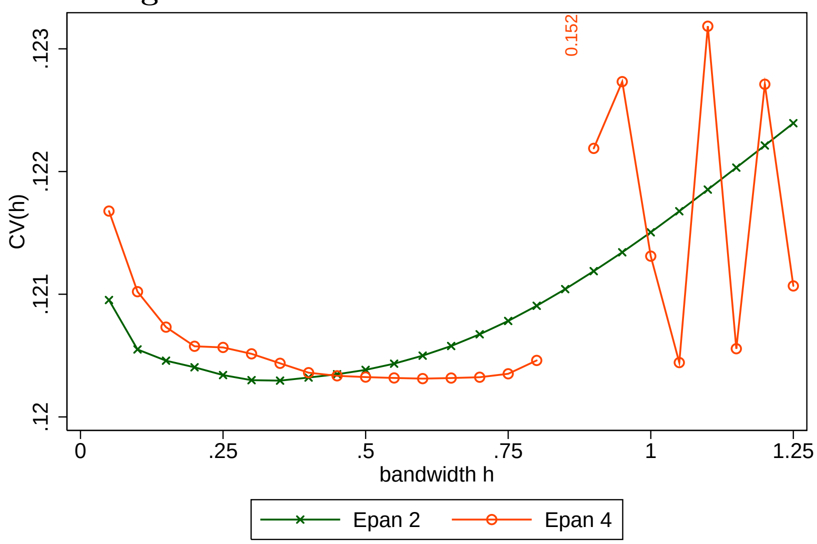

Notes. The figure shows the cross-validation function for the reduced-form binomial mortgage choice model (5), with the Epanechnikov kernel ($r = 2$) and its fourth-order analogue ($r = 4$).

Figure A.3: Parametric and Non-parametric Estimates of F

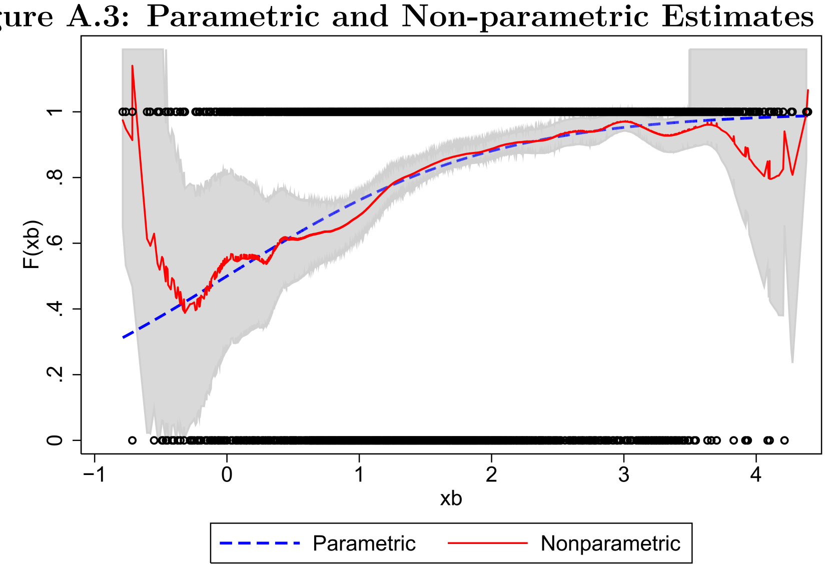

Notes. The figure shows the parametric and nonparametric estimates of the link function $\Pr(D_n = 1 \mid x_n'\hat{\beta})$ conditional on the reduced-form choice model logit coefficients estimated in Table 2, column 5. The nonparametric estimator $\hat{F}$ is calculated using bandwidth $h = 0.35$ and $K^{(4)}$. Shaded area is a uniform 2-SE confidence interval for the nonparametric estimator, constructed using the Bonferonni correction for multiple testing.

F.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" 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:

$$g^*(e) = P(e)\varphi(e) = \left(\sum_{r=0}^{R} H_r(e)\right)\varphi(e). \tag{A.23}$$

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.4 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 2, 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.4: 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
Sociodemographic 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 $\leq$ 6 years ago. All three columns rescale the coefficients so $b_{\text{FRMRate}} = -1$ 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

F.3 Three-Step Estimation by Fully Parametric Methods

We report estimates of the structural choice model using fully-parametric predicted values in Table A.6. 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.5: 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 2, 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.6: 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

G Learning from Nominal Interest-Rate Experiences

We have shown that individuals with personal exposure to high inflation expect not only higher inflation in the future, but also higher nominal interest rates (cf. Figure 3). Hence, individuals could be forming their expectations about future interest rates and making their mortgage choices based on their lifetime experiences of past interest rates, rather than inflation. As discussed in Section 5.2, this approach is not theoretically distinct since expected inflation and nominal interest rates are related by the Fisher equation, $i = r + \pi$. Nevertheless, we investigate whether, utilizing lifetime experiences of nominal interest rates, we obtain similar estimates on experience effects on mortgage choice behavior.

We re-estimate our reduced-form mortgage choice model, replacing $\pi^e_{n,t}$ with $i^e_{n,t}$, using short-term (90-day) Treasury-bill rates as well as long-term (10-year) Treasury rates. For short-term nominal rates, we obtain historical 90-day T-bill rates between 1926-2001 from the CRSP US Treasuries and Inflation Indexes database and use the average annual return. For long-term nominal rates, we splice the U.S. government "long-term bond yield," reported in the Historical Statistics of the United States (series Cj1192) over 1919-1961, to the 10-year constant maturity Treasury yield that begins in 1962 (from the Fed's H.15 Selected Interest Rate release). As before we weight historical nominal interest rates using weights that put the highest weight on the current year and linearly decline to zero in the year that the decision-maker was born. Neither of the two historical nominal interest rate series go back to 1915, the earliest birth year in our dataset. We drop these few early years when data are unavailable in the construction of the interest-rate experience measure and then re-normalize the weights. The resulting short-term and long-term interest-rate experience measures are highly correlated with lifetime inflation experiences ($\rho = 0.81$ and 0.69, respectively).

Estimation results are in Table A.7. The direction and magnitude are similar albeit statistically weaker. Model-selection criteria point to inflation beliefs as the preferable independent variable, followed by the short-term interest-rate-experiences model, albeit by a small margin (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; and the maximized log-likelihood is −5692.90 for the baseline model, followed by −5692.96 for the short-term interest rate model and −5693.21 for the long-term interest rate model).

Table A.7: Learning from Nominal Interest Rates

	(1)	(2)
Freddie Mac PMMS FRM index rate (%)	-3.590***	-3.588***
	(0.816)	(0.816)
Freddie Mac PMMS ARM index rate (%)	0.845***	0.845***
	(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
Socidemographic 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 2, 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 2, col. 5. Robust standard errors in parentheses. * $p<0.10$, ** $p<0.05$ and *** $p<0.01$.

H 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.

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

Figure A.4: Age and Mobility

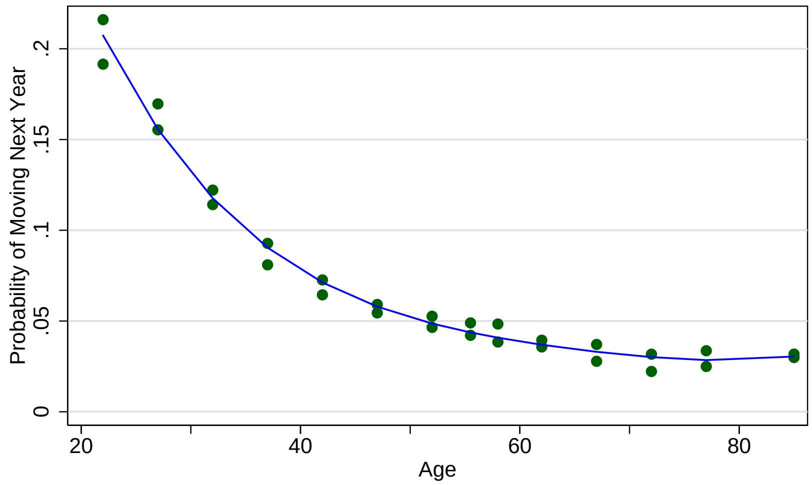

Notes. The data source is the CPS ASEC from 2005 and 2010. Fitted values are calculated using fourth-order polynomial function of age.

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.4. 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:

$$\begin{aligned}\widehat{\text{MoveProb}}^{1y}(age) &= \underset{(0.077)}{0.696} - \underset{(0.007)}{0.0355} \times age + \underset{(0.0002)}{0.000752} \times age^2 \\ &\quad - \underset{(2.80\cdot 10^{-6})}{7.40 \cdot 10^{-6}} \times age^3 + \underset{(1.30\cdot 10^{-8})}{2.80 \cdot 10^{-8}} \times age^4 \qquad \text{(A.24)}\end{aligned}$$

(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:

$$\text{StayProb}(a, T) = \prod_{s=0}^{T-1} \left(1 - \widehat{\text{MoveProb}}_{1y}(a + s)\right). \qquad \text{(A.25)}$$

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 description of the average schedule: $r(p) = r_0 - 0.0027p + 0.0002p^2$. This is the same order of magnitude that Brueckner (1994) finds using National Association of Realtors data for the early 1990s. In our data, 16.5 percent of households pay discount points, with a median of 2 points paid.

With some additional assumptions, we can use this information to construct a distribution of moving probabilities for each household. First, we assume risk-neutral rational agents without liquidity constraints, an assumption that is standard in prior literature but likely to significantly underestimated implied tenure rates, as discussed in Section 6.2. We then calculate the break-even horizon τ^* for each household in our dataset, assuming that the monthly interest savings are discounted at an annual nominal rate of 8%. We consider two plausible parametrizations of moving-time distributions. If we assume that the moving events arrive according to a Poisson process, then the distribution of moving times τ is negative exponential with intensity parameter $\lambda = 1/E[\tau] = 1/\tau^*$. This distribution implies that the hazard rate of moving is stationary. Alternately, individuals form an attachment to their communities over time (Dynarski 1985, Quigley 1987), and the hazard rate of moving might decrease with time. To model this, we let moving times follow a Weibull distribution with shape parameter $\alpha = 0.7$.[27] paid 3 discount points, implying a break-even horizon of $\tau^* = 6$ years, 9 months.

Table A.8 reports, separately for each scenario of interest rate calculation, the expected additional interest payments made by switching households, over different distributions of moving times for different refinancing scenarios. The first column reproduces the last column of Table 7 for reference. Starting from the bottom of each panel we see the significantly lower estimates of median tenure relative to our previous age-based calculations. Instead of an estimated average of 12.5 years based on household age (in column 1), we now obtain tenure estimates of 4.7-4.9 years across switching households based on discount points paid under the negative exponential distribution (column 2), and of 3.6 under the Weibull distribution (column 4). This large discrepancy in occupancy times reflects that most households do not pay any discount points. If households purchase fewer 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.

Ignoring these concerns, we estimate the present discounted values of excess interest payments to be somewhat lower using the points paid methodologies; for example, column 2 reports a WRTE of \$9,106 under Scenario 3 interest rates, expected refinancing behavior, and Poisson moving events, roughly \$5,000 less than our age-based estimate.

If, instead, we want to allow for the possibility that individuals with a given expected tenure choose less than the number of points predicted for risk-neutral rational individuals, our estimates of occupancy time are too short. As a simple remedy, we fit each household's intensity parameter to the median rather than the mean of the distribution: $F_\lambda^{-1}(\tau^*) = 0.5$. This raises the average median time of occupancy to 6.4 years, and reduces the gap between the dollar costs estimated under the two methodologies, as shown in column (3) for the negative exponential distribution and in column (5) for the Weibull distribution. For example, column 5

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

reports a WRTE of \$11,629 under Scenario 3, expected refinancing behavior, and a decreasing hazard rate of moving, approximately \$3,000 less than our age-based mobility estimate.

Table A.8: 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 aditional 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 amoutns 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 7. Column (1) reproduces the estimation from the final column of Table 7 for comparison. In columns (2)-(5), discount points paid at time of origination are used to calculate the time to breakeven, 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 or median of 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 or median of distribution for each household. Average median tenure is calculated as the median tenure for each household, then averaged over all switching households.