

Politecnico di Torino

Department of Management and Production Engineering

Master of Science in Engineering and Management

Change/Expected change in interest and mortgage rates and households' financial behavior

Supervisor:

Prof. Calcagno Riccardo

Candidate:

Citino Giovanni
328753

Academic Year 2025/2026

Abstract

The choice between a fixed-rate and an adjustable-rate mortgage determines how interest rate risk is allocated between the borrower and the lender. This study, organized by determinant and centered primarily on studies using the Survey of Consumer Finances, reviews what the empirical literature has established about the household characteristics and market conditions that shape this choice, assessing their finding against the predictions of the lifecycle and mortgage timing frameworks. The evidence provides consistent support for the role of relative contract prices, expected mobility and income volatility across four decades and several types of data. By contrast, the effect of liquidity constraints has consistently the opposite sign to that predicted by the theory, while findings on financial literacy conflict across studies and those on risk aversion are corroborated only in the American evidence. The determinants on which the theory performs best prove to be those observable without asking households to report on their own beliefs, while those that remain unsettled are precisely the ones requiring such reports. Interest rate expectations constitute the limiting case, since they have been recorded in the survey for over three decades and have never been entered into a model of contract choice.

Contents

Introduction 1

1. Theoretical framework and determinants of the choice between fixed and adjustable-rate mortgages 4

1.1 Mortgage choice as an intertemporal optimization problem under uncertainty 4

1.2 Campbell and Cocco (2003): the lifecycle benchmark 7

1.2.1 Model architecture 7

1.2.2 Simulation results and normative predictions 9

1.2.3 Critical limitations 10

1.2.4 Tension with empirical evidence 11

1.3 Koijen, Van Hemert, and Van Nieuwerburgh (2009): mortgage timing and time-varying bond risk premia 12

1.3.1 The bond risk premium as the key state variable 12

1.3.2 Empirical performance and the timing decision rule 14

1.3.3 Implications for the present thesis 16

1.3.4 Critical limitations 16

1.4 Further theoretical contributions 17

1.4.1 Extensions of the rational framework 18

1.4.2 Expectations and behavioral frictions 19

1.5 Synoptic table of determinants and testable predictions 20

1.5.1 Notes on the table 23

2. The U.S. macroeconomic context 25

2.1 Interest rates and mortgage rates, 2000 to 2024 25

2.2 The composition of the mortgage market and the adjustable-rate share 27

2.3 House prices, real income, and household indebtedness 30

2.4 The 2022 economic environment and household mortgage decisions 33

3. Determinants of mortgage choice: a critical review of the empirical evidence 35

3.1 Scope and method of the review 35

3.1.1 Organizing principle and delimitations 35

3.1.2 Search strategy and selection criteria 39

3.2 The Survey of Consumer Finances: structure and descriptive evidence on mortgage holders 41
3.2.1 The instrument 41
3.2.2 Descriptive evidence on adjustable and fixed rate borrowers 44
3.2.3 Summary of the descriptive evidence 45
3.3 Risk preferences and income risk 48
3.3.1 Evidence on stated risk aversion and income volatility 48
3.3.2 Measurement and its limits 50
3.4 Liquidity constraints and affordability 51
3.4.1 The empirical reversal 51
3.4.2 Explaining the contrasting evidence 52
3.4.3 How constraints are measured 53
3.4.4 Assessment 54
3.5 Mobility and planning horizon 55
3.5.1 The evidence 55
3.5.2 Measurement 56
3.5.3 The contract term 56
3.5.4 Assessment 57
3.6 Contract pricing and the ARM-FRM spread 57
3.6.1 The theoretical refinement 57
3.6.2 The evidence 58
3.6.3 Measurement 59
3.6.4 The contract term 60
3.6.5 Assessment 61
3.7 Interest rate expectations 61
3.7.1 Evidence from national markets 61
3.7.2 Evidence at household level 62
3.7.3 The gap in the American literature 63
3.7.4 Measurement 64
3.7.5 The relationship that the literature would predict 65
3.7.6 Assessment 66

3.8 Financial literacy and information gaps 66
3.8.1 What the estimates show 67
3.8.2 Bimodality 68
3.8.3 Survey measures 69
3.8.4 Knowledge of the contract itself 69
3.8.5 Assessment 71
3.9 Socioeconomic and demographic characteristics 71
3.9.1 Income 72
3.9.2 Wealth 74
3.9.3 Age 75
3.9.4 Other demographic characteristics 76
3.9.5 Assessment 77
3.10 The evidence assessed: theory and estimates compared 78
3.10.1 Three groups of rows 80
3.10.2 The pattern of empirical success 82
3.10.3 Toward the remaining chapters 83
4. Behavioral dimensions and limits of the rational benchmark 85
4.1 Rationale and scope of the chapter 85
4.2 The formation of interest rate expectations 86
4.2.1 The experience-based mechanism 87
4.2.2 Evidence from the mortgage market 87
4.2.3 Relation to the determinants reviewed in chapter 3 88
4.2.4 The mechanism applied to the 2022 environment 89
4.2.5 Limits of the account 90
4.3 Cognitive biases in the mortgage decision 90
4.3.1 Present bias and self-control 90
4.3.2 The distinction between product complexity and rate structure 91
4.3.3 Heuristics in the processing of macroeconomic information 92
4.3.4 What the literature has not established 93
4.4 The predictive adequacy of the rational benchmark 94
4.4.1 Where the framework holds 94

4.4.2 The compression of the decision horizon 94
4.4.3 Deviations in opposite directions 95
4.4.4 The aggregate puzzle 96
4.4.5 Implications for the interpretation of the estimates 96
4.4.6 Assessment 97
5. Synthesis 98
5.1 Established regularities 98
5.1.1 The pricing channel 98
5.1.2 The horizon channel 99
5.1.3 The risk channel 99
5.1.4 The income gradient 100
5.1.5 Established null results 100
5.1.6 Observability and measurements 101
5.2 Contested findings and the sources of divergence 101
5.2.1 Two forms of contestation 102
5.2.2 Divergence arising from measurement 103
5.2.3 Divergence arising from population heterogeneity 104
5.2.4 Divergence arising from regime and institutional setting 105
5.2.5 Requirements for further assessment 106
5.3 Gaps in the evidence base 106
5.3.1 Gaps that existing data could close 107
5.3.2 Gaps requiring improved measurement instrument 108
5.3.3 Gaps in the questions pursued 109
5.4 Implications for policy and further research 110
5.4.1 Disclosure and the limits of borrower comprehension 110
5.4.2 Financial education and uncertain returns 111
5.4.3 Market composition as an object of design 111
5.4.4 Priorities for further research 112
6. Conclusions 114
References 117

Introduction

Few financial decisions carry consequences as durable as the choice of a mortgage contract. For most households the mortgage is the largest liability they will ever assume, the payment it generates absorbs a substantial share of monthly income for decades, and the form of the contract determines who bears the risk that interest rates move. A household selecting a fixed-rate mortgage (FRM) transfers that risk to the lender and pays a premium for that protection, whereas the one who chooses an adjustable-rate mortgage (ARM) retains the risk and avoids paying the premium. The choice is typically made once, often under time constraints, and is rarely reconsidered except through refinancing. The economic significance of this decision extends beyond the individual household, as the aggregate composition of mortgage contracts affects the transmission of monetary policy. Housing debt constitutes the dominant component of household liabilities in the United States, and the proportion of that debt carrying a fixed rather than an adjustable rate determines how quickly a change in policy rates reaches household budgets. A market composed predominantly of fixed contracts insulates existing borrowers from tightening and, in doing so, attenuates the mechanism through which monetary policy is expected to operate. What appears at the individual level as a private matter of financial planning aggregates into a structural feature of the macroeconomy.

The two contracts have been analyzed within a common theoretical framework since Campbell and Cocco (2003) established that the choice is best understood as a problem in household risk management rather than as a comparison of expected costs. Their lifecycle model shows that neither contract is unambiguously safer, since the nominal fixed-rate mortgage exposes the household to variation in the real value of its debt while the adjustable-rate mortgage exposes it to short-term fluctuations in real payments, which are particularly problematic for households facing binding borrowing constraints. Koijen, Van Hemert, and Van Nieuwerburgh (2009) subsequently held household characteristics fixed and allowed the price of the fixed contract to vary, showing that the premium borrowers are willing to pay for payment certainty fluctuates with the state of the term structure and that the aggregate composition of new lending shifts with it. The two contributions address complementary questions, the one concerning which households select which contract at a given moment and the other concerning why the aggregate composition shifts over time. Against this theoretical apparatus, the empirical record is less settled than the volume of published work might suggest. Some determinants have been documented with

Artificial intelligence tools were used to support the drafting, editing, and linguistic refinement of selected parts of this chapter.

remarkable consistency across four decades, several types of data, and a range of national markets. Others have produced results that conflict across studies, sometimes for reasons that concern how a variable was constructed rather than how households behave. And a third group has scarcely been examined, including the determinant that both theoretical frameworks place closest to the center of the decision, namely what borrowers believe about the future course of interest rates. The principal survey used to study mortgage choice in the United States has recorded that belief for more than three decades, and no study has entered it into a model of the choice itself.

This configuration is the point of departure for the present work. The purpose is not to produce new estimates but to establish, through a critical review of the published evidence, what the literature has determined about the determinants of mortgage choice, where its findings diverge and why, and which questions remain open. A review of this kind is warranted for two reasons beyond the ordinary usefulness of consolidation. No dedicated systematic review of this literature appears to have been published, and recent contributions continue to describe the field as fragmented across countries and datasets, so that individual studies have accumulated without the periodic assessment that would reveal which of their disagreements are substantive and which arise from differences in construction. Moreover, the environment in which mortgage decisions are now taken differs sharply from those under which most of the accumulated evidence was produced, following the most rapid monetary tightening in four decades and a doubling of the fixed mortgage rate within a single year.

The review is organized around the Survey of Consumer Finances, which occupies a distinctive position among the sources available for studying this decision. Administrative loan records capture contract terms precisely but reveal little about the household that signed them, while panel surveys follow households over time but lack the balance sheet detail the decision requires. The Survey of Consumer Finances records, for the same household and within the same instrument, the type of contract held, the composition of assets and liabilities, the stated attitude towards financial risk, the experience of credit denial, the expectation of moving, and the expectation of where interest rates are heading. Evidence on the attitudinal and expectational determinants of mortgage choice in the United States rests almost entirely on this source, and organizing the review around it permits the determinants to be compared on a common evidential basis. Studies drawing on other data are included where they establish results the survey cannot, where they precede the survey-based literature historically, or where cross-country comparison distinguishes

findings that reflect household behavior from those that reflect national market structure. Three delimitations should be stated at the outset. The review treats the decision as a binary one between fixed and adjustable contracts, following both the convention of the literature and the distinction the survey records reliably, which sets aside the hybrid structures, introductory rates, and adjustment provisions that determine the risk a household actually bears. The analysis focuses on the demand side, taking the set of contracts offered by lenders as given, although the composition of the U.S. mortgage market is also shaped by lenders’ willingness to originate different contract types. The review is selective rather than exhaustive and is limited to studies available in English.

The argument proceeds as follows. Chapter 1 assembles the theoretical framework, presenting the two cardinal models and the complementary contributions that extend them, and closing with a synoptic table that records for each determinant the sign theory predicts and the mechanism which generates that sign. Chapter 2 documents the macroeconomic environment of the past two decades, tracing the movement of interest rates, the composition of mortgage originations, and the balance sheet position of American households, so that the setting in which recent decisions have been taken is established before the evidence on those decisions is examined. Chapter 3 conducts the review itself, organized by determinant rather than by study, and concludes by setting the estimated signs against the predicted ones in a comparative table. Chapter 4 turns to the behavioral evidence, examining how households form expectations about future rates and which cognitive mechanisms have been documented in this specific decision, and assesses how far a framework premised on informed optimization describes observed behavior. Chapter 5 synthesizes the material, distinguishing what the literature has established from what remains contested and from what it has not addressed, and considers the implications for policy and for further research. Chapter 6 concludes.

1. Theoretical framework and determinants of the choice between fixed and adjustable-rate mortgages

1.1 Mortgage choice as an intertemporal optimization problem under uncertainty

For most households, the mortgage contract on the primary residence is the largest and longest-duration financial obligation they will ever undertake. The decision between a fixed-rate mortgage (FRM) and an adjustable-rate mortgage (ARM) is therefore not a routine financial choice but an intertemporal commitment that conditions household consumption, savings, and risk exposure over horizons that frequently exceed two decades. Economic literature has accordingly framed the ARM/FRM decision as a constrained intertemporal optimization problem, in which the household allocates consumption, financial assets, and labor income across periods while servicing a long-duration liability whose real cost evolves stochastically. Four sources of uncertainty enter the household's optimization problem and condition the relative attractiveness of the two contractual forms. The first is interest rate uncertainty: short-term nominal rates fluctuate over the life of the mortgage, directly determining the periodic ARM payment and indirectly determining the present value of the FRM stream. The second is inflation uncertainty: because mortgage contracts in the United States are nominally denominated, unexpected changes in the price level alter the real burden of mortgage payments, eroding the real value of FRM payments in inflationary periods and leaving the real ARM payment relatively stable, since nominal ARM rates adjust to incorporate realized inflation through their indexation to short-term Treasury yields. The third is labor income uncertainty: shocks to household earnings, whether transitory (unemployment spells, illness) or permanent (career trajectories, occupational reallocations), affect the household's capacity to service the mortgage and to absorb payment volatility through savings. The fourth is tenure uncertainty: the household may move, by choice or necessity, before the mortgage matures, which shortens the effective horizon over which long-term rate insurance is valuable. These four sources of uncertainty are not independent: aggregate shocks correlate inflation and short-term rates with the macroeconomic conditions that drive labor income, and the household's tenure decision interacts with house prices and labor mobility. The contractual choice between ARM and FRM channels these correlated risks into two distinct categories of contractual risk. The first is what Campbell and Cocco (2003) define as risk to the real capital value of the mortgage: under the FRM, the nominal payment stream is

Artificial intelligence tools were used to support the drafting, editing, and linguistic refinement of selected parts of this chapter.

fixed, but its real present value fluctuates with realized inflation, exposing the borrower to long-horizon inflation risk; under the ARM, the nominal payment adjusts with short rates, leaving the real capital value relatively stable. The second is payment risk: the FRM delivers a constant nominal payment that is straightforward to budget against, whereas the payment variability associated with an ARM is particularly costly for households with binding borrowing constraints and limited liquid assets, as increases in interest rates can raise mortgage payments precisely when their ability to absorb additional financial pressure is limited. The two risks operate at different horizons and on different margins of household welfare: capital-value risk affects the household's lifetime resource position, while payment risk affects period-by-period consumption smoothing and the probability of liquidity-driven distress. A concept significantly tied to the two aforementioned risks is the tilt effect, which sits at the intersection of inflation and payment dynamics and predates the modern lifecycle literature. The tilt effect was articulated by Lessard and Modigliani (1975) in their contribution to the Federal Reserve Bank of Boston Conference Series volume on inflation and the housing market, where they observed that the standard nominal mortgage front-loads real payments in inflationary environments, because nominal payments incorporate an inflation premium that is paid early in real terms but recouped only as inflation erodes the principal in nominal terms. The result is that the standard FRM imposes its heaviest real burden on the household precisely in the early years of the contract, when the typical borrower is youngest, lowest in wealth, and most likely to be liquidity-constrained. The tilt effect explains why, even in the absence of payment risk, inflationary regimes generate household demand for alternative mortgage structures, and it remains conceptually relevant for interpreting the differential appeal of ARMs across different inflation regimes.

The interaction between payment risk and liquidity constraints is central to the welfare comparison between ARM and FRM. A household that can borrow freely against future income or run-down liquid assets without cost can absorb ARM payment volatility through standard consumption smoothing, and the welfare cost of payment risk is correspondingly small. A household that faces binding borrowing constraints, whether through formal credit denial, prohibitively expensive marginal credit, or institutional restrictions, cannot so easily absorb payment shocks; an upward repricing of an ARM may force a contemporaneous reduction in non-housing consumption, with welfare consequences that compound at the margin. This asymmetry between financially flexible and financially constrained households is the welfare channel through which the rational benchmark

generates its central prediction that risk-averse, liquidity-constrained borrowers should prefer the FRM despite its higher unconditional cost. The same channel, however, will be challenged in the empirical literature reviewed in chapter 3: a substantial body of evidence suggests that constrained borrowers in fact tilt toward the ARM, drawn by the lower initial payments that the contract typically offers. This tension between the normative prediction of the lifecycle benchmark and the pattern observed in survey data is the central question that the review of chapter 3 examines.

The remainder of chapter 1 examines how successive generations of theoretical models have organized these four sources of uncertainty and two contractual risks into formal frameworks delivering testable predictions on who should optimally choose ARM versus FRM. Section 1.2 presents the Campbell and Cocco (2003) lifecycle benchmark, which provides the normative reference point against which the empirical literature is conventionally evaluated. Section 1.3 introduces the Koijen, Van Hemert, and Van Nieuwerburgh (2009) extension, in which time-varying bond risk premia make the ARM/FRM choice depend on the current state of the term structure and on household expectations about future rates, the dimension most directly relevant to the title of the present thesis. Section 1.4 gathers a set of further theoretical contributions along two lines: one extends the rational account toward housing portfolio choice, household finance, and inflation hedging, while the other departs from the rational-expectations frame to consider behavioral and information frictions, namely experience-based expectation formation, financial literacy, and the finding that a non-trivial share of borrowers do not accurately know the terms of their own mortgage. Section 1.5 consolidates the chapter into a synoptic table of testable predictions that organizes the review conducted in chapter 3. A note of caution is warranted before proceeding. The framework introduced in this section assumes a household that fully understands the risk structure of each contract and that solves the optimization problem the theorist writes down on its behalf. This assumption will be challenged on two fronts. The behavioral and literacy literature reviewed in section 1.4 questions whether households possess the cognitive resources and informational inputs required to execute this optimization. The empirical evidence reviewed in chapter 3 documents systematic deviations from the predictions of the rational benchmark, deviations that are particularly concentrated among less-literate and more financially-constrained borrowers. The synoptic table of section 1.5 therefore records, where relevant, both the rational-benchmark prediction and the alternative behavioral prediction, an architecture that chapter 3 uses to set the estimated signs against both.

1.2 Campbell and Cocco (2003): the lifecycle benchmark

The canonical normative framework for the ARM/FRM choice is the lifecycle model of Campbell and Cocco (2003). The contribution of the paper is twofold. First, the paper provides the first fully articulated dynamic stochastic model of the ARM/FRM decision in which the household's choice is jointly determined with consumption, savings, and labor-income realizations over a finite horizon. Second, by simulating the model under empirically calibrated parameter values, the paper delivers a structured set of predictions about which household types should optimally select each contract, predictions that have organized the household-finance literature on mortgage choice for the subsequent two decades.

1.2.1 Model architecture

The Campbell and Cocco (2003) model embeds the mortgage choice within a finite-horizon lifecycle framework with the following structure. A household enters the model financing a house of given size, makes a once-and-for-all choice between an FRM and an ARM at origination, and thereafter chooses period-by-period consumption of non-housing goods to maximize expected lifetime utility under constant relative risk aversion. The household receives stochastic labor income with permanent and transitory components, faces stochastic short-term nominal interest rates and stochastic inflation, and is subject to a borrowing constraint that prevents it from capitalizing future labor income to smooth consumption. Savings are held entirely in riskless assets, so that all portfolio considerations other than the mortgage are set aside. The two contracts differ in how the periodic payment is determined, and the difference is most clearly seen in real terms. Under the fixed rate contract the nominal payment is set at origination so that the discounted stream of payments equals the loan principal, and the real payment obtained by deflating this nominal amount is

$$M_{jt}^{F} = \frac{(1-\lambda)P_{j1}^{H}H_{j}}{P_{t}\sum_{s=1}^{T}(1+Y_{T1}^{F})^{-s}} \quad , \quad (1)$$

where $\lambda$is the required down-payment ratio, $P_{j1}^{H}$the nominal price of the house at date one, $H_j$its size, so that $(1-\lambda)P_{j1}^{H}H_j$is the nominal principal borrowed, $P_t$the price level at date t, $Y_{T1}^{F}$the nominal rate on a fixed-rate contract of maturity Tset at origination, and $s$the

index of summation over the life of the loan.[1] Since the nominal payment is fixed at initiation, the real payment is inversely proportional to the price level and declines as prices rise. Under the adjustable rate contract the payment is recomputed each period on the principal then outstanding, and its real value is

$$M_{jt}^{A} = \frac{Y_{1t}^{A} D_{jt} + \Delta D_{j,t+1}}{P_t} \quad , \quad (2)$$

where $D_{jt}$is the nominal principal outstanding at date t, $\Delta D_{j,t+1}$the portion of the payment that reduces principal rather than servicing interest, and $Y_{1t}^{A}$the nominal rate on the contract, equal to the prevailing short rate plus a constant premium $\theta_A$that compensates the lender for default risk. The comparison of equations (1) and (2) yields the distinction on which the model turns, and it is a distinction between two forms of risk rather than between a risky and a safe contract. The nominal fixed-rate mortgage has a real capital value that is highly sensitive to inflation, so that the household holding it bears what the authors term wealth risk: the contract proves expensive when inflation remains stable and cheap when inflation rises, as it did during the 1960s and 1970s. The adjustable-rate mortgage removes this exposure, since its real capital value is largely unaffected by inflation, but replaces it with income risk, namely the short-term variability of the real payments required each period. When expected inflation and nominal rates increase, the nominal payment rises proportionally before the price level has adjusted, and the real payment rises with it. Variability of this kind would be immaterial to a household able to borrow against future income; it matters because the borrowing constraint binds precisely in the states where income and house prices are low and buffer-stock savings have been exhausted. Both contracts therefore transfer a risk to the household, and the choice between them is a choice of which risk to bear. The rate at which the fixed payment is discounted follows from the assumption the authors adopt about the term structure. Long-term nominal yields are taken to satisfy the log expectations hypothesis, so that the log yield on an n-period nominal bond is the expected average of successive one-period log yields plus a constant term premium,

$$y_{nt} = \frac{1}{n} \sum_{i=0}^{n-1} E_t \left[y_{1,t+i} \right] + \xi \quad , \quad (3)$$

[1] The summation index is denoted s rather than j as in the original, so as to avoid collision with the subscript identifying the household. All remaining notation follows Campbell and Cocco (2003).

where $y_{nt} = \log(1 + Y_{nt})$and ξdoes not vary with maturity. The rate charged on the fixed contract adds to this yield a further constant premium θ_F, compensating the lender for default risk and for the value of the prepayment option, while the adjustable contract carries no term premium at all, its rate resetting each period to the short rate plus θ_A. The fixed contract therefore costs the household a premium in exchange for insurance against wealth risk, and the adjustable contract forgoes both.

Equation (3) carries an implication that should be registered before the discussion proceeds, since it delimits what the model can address. If the expectations hypothesis holds with a constant premium, excess returns on long-term over short-term bonds are unpredictable, and a household cannot lower its expected borrowing cost by attempting to time the bond market. The authors state this consequence explicitly. It is the assumption that the framework examined in section 1.3 relaxes, and the two models are therefore not merely complementary but rest on incompatible premises about the predictability of the term structure. The model is calibrated to United States data. Inflation and interest rate parameters are estimated on annual data from 1962 to 1999, yielding mean annual inflation of 4.6 percent with a standard deviation of 3.9 percent and an autoregressive coefficient of 0.754, and a mean log real interest rate of 2.0 percent with a standard deviation of 2.2 percent. Income parameters are estimated from the Panel Study of Income Dynamics (PSID), with labor income decomposed into permanent and transitory components following the buffer-stock savings literature. The contractual parameters are set from the monthly interest rate survey of the Federal Housing Finance Board for the period 1986 to 2001, giving a fixed-rate premium $\theta_F$of 1.8 percent over the ten-year Treasury yield, an adjustable-rate premium $\theta_A$of 1.7 percent over the one-year yield, and a term premium $\xi$of 1.0 percent, with the required down-payment ratio λset at 20 percent. This calibration places the model within the postwar American inflation experience, including the high-inflation episode of the 1970s and early 1980s, a point revisited in section 1.2.4.

1.2.2 Simulation results and normative predictions

The main contribution of the paper consists in a set of welfare comparisons across different household profiles. The authors compute the certainty-equivalent consumption stream associated with each contract for households differing in risk aversion, mortgage size relative to income, income volatility, expected mobility, and the assumed cost of default. The pattern of results is summarized concisely in the paper's abstract: “numerical solution of a life-cycle model with borrowing constraints and income risk shows that an ARM is

generally attractive, but less so for a risk-averse household with a large mortgage, risky income, high default cost, or low moving probability". The normative implications follow from this central result. The FRM is welfare-dominant, optimally chosen by households that exhibit one or more of the following characteristics: high CRRA risk aversion, a large mortgage relative to financial assets and income, volatile labor income, a high marginal cost of default (whether through legal sanctions or non-pecuniary stigma), and a long expected tenure in the property such that the long-horizon insurance properties of the FRM are most valuable. Conversely, the ARM is welfare-dominant for households that are wealthy in financial assets, hold mortgages small relative to their resources, have stable labor income, and expect to move within a relatively short horizon. A subsidiary but consequential prediction concerns the liquidity-constraint channel. The welfare losses associated with the payment volatility of ARM arise mainly when households face borrowing constraints. In the absence of such constraints, households could smooth consumption over time and absorb temporary payment fluctuations more easily. Consequently, the model predicts that liquidity-constrained households should exhibit a stronger preference for FRM, as they are more exposed to the welfare costs generated by the payment uncertainty associated with ARM. This prediction is the cornerstone of the rational-benchmark view of mortgage choice and corresponds to row 2 of the synoptic table that closes the chapter.

The Campbell and Cocco (2003) framework also generates a macro-state contingency that will recur throughout the thesis. The FRM rate equals the expected average future short rate plus a term premium that compensates the lender for bearing duration risk. When the term premium is high, the FRM is expensively priced relative to the expected ARM cost path, and the ARM becomes more attractive even for risk-averse borrowers; when the term premium is low or negative, that is when the yield curve is flat or inverted, the FRM may even dominate on first-moment grounds, before any risk consideration enters.

1.2.3 Critical limitations

The Campbell and Cocco (2003) model is a benchmark, not a complete description of the household mortgage choice, and its limitations should be stated explicitly. Three deserve particular attention. First, the model abstracts from endogenous default. Mortgage default is a strategic option whose value depends on house prices, mortgage balances, and the legal regime governing recourse and recovery; abstracting from this option simplifies the

model considerably but precludes analysis of the channels through which contract type affects default risk. Subsequent work by the same authors, most notably Campbell and Cocco (2015) *Journal of Finance* article "A Model of Mortgage Default", has filled this gap, showing that contract-type effects on default can substantially modify the welfare comparison, an extension that falls outside the scope of the present review. Second, the model assumes complete markets in the sense that the household can freely trade financial assets at a single risk-free rate (subject to the borrowing constraint), and there is no scope for hedging the labor-income, inflation, or mortgage-rate risks through other financial instruments. Real households face substantial market incompleteness, particularly with respect to labor-income hedging, and incompleteness in markets outside the mortgage choice can in principle alter the welfare ranking of mortgage contracts. Third, the model is calibrated to U.S. postwar data, with substantial weight placed on the high-inflation episode of the 1970s and early 1980s. The structural parameters of inflation persistence, inflation volatility, and the term premium implied by this calibration are not necessarily representative of the post-2008 environment of low and anchored inflation expectations, nor of the post-2022 transition regime documented in chapter 2. This is not a flaw of the original paper but a constraint on its external validity: the normative predictions of the authors are predictions for households facing the inflation environment captured by the calibration, and households making mortgage choices in a different inflation regime may face a different welfare ranking.

1.2.4 Tension with empirical evidence

A final observation closes this section and motivates much of what follows.
The Campbell and Cocco (2003) prediction that liquidity-constrained households should prefer the FRM is the cleanest welfare result of the rational benchmark. Yet the empirical literature, and most prominently the SCF-based study of Coulibaly and Li (2009), who conclude that more financially constrained households are more inclined to choose ARMs, has consistently documented the opposite pattern. The same paper notes a parallel tension on the literacy dimension: highly educated households, who are rarely borrowing constrained, are more prone to opt for an ARM than households with a lower level of educational achievement, complicating any simple mapping between financial sophistication and contract choice. Brueckner and Follain (1988) similarly find that high earners holding a mortgage contract, who are also comparatively less likely to be constrained, tilt towards ARMs, replicating the same qualitative pattern in earlier loan-level data. This empirical-

theoretical mismatch is not a refutation of Campbell and Cocco (2003), but it may reflect either that constrained borrowers face a binding short-term affordability constraint that overrides longer-horizon welfare considerations (selecting the lower initial-payment ARM as the only contract they qualify for), or that the model omits an important dimension of the choice problem (such as the role of teaser-rate ARMs, liquidity-driven origination behavior by lenders, or borrower information frictions). The present thesis takes the position that the mismatch is informative rather than dispositive. Section 1.5 records the constraint coefficient as a *contested* row in the synoptic table. The other contested rows, as age, education/literacy, and the rational-versus-experience-based reading of expectations, will be developed in sections 1.3 and 1.4 and complete the structural framework against which the empirical literature is assessed in chapter 3.

1.3 Koijen, Van Hemert, and Van Nieuwerburgh (2009): mortgage timing and time-varying bond risk premia

The Campbell and Cocco (2003) lifecycle model treats household characteristics, such as risk aversion, mortgage size, income volatility, and expected mobility, as the primary determinants of the optimal ARM/FRM choice.
The contribution of Koijen, Van Hemert, and Van Nieuwerburgh (2009) complements this household-characteristics view with a *macroeconomic-state* view: the optimal contract choice also depends, in a principled way, on the current state of the term structure of interest rates, and in particular on the bond risk premium embedded in long-term yields. The paper thereby shifts the analytical focus from who should choose ARM to when households should choose ARM and provides a theoretical foundation for the empirical observation that aggregate ARM origination shares move strongly with term-structure conditions.

1.3.1 The bond risk premium as the key state variable

The starting point of Koijen, Van Hemert, and Van Nieuwerburgh (2009) is a decomposition of the long-term nominal interest rate. Under the zero-profit condition for competitive mortgage lenders, the fixed rate on a long-term mortgage (FRM) is approximately equal to the long-term nominal bond yield. This rate can be decomposed into the average short rate expected to prevail over the life of the contract plus a long-term nominal bond risk premium that compensates the lender for bearing interest rate risk (arising from real interest rate and inflation fluctuations). Conversely, the adjustable rate (ARM) is

approximately equal to the current nominal short rate and embeds no such premium. For a mortgage contract maturing in τ years, this long-term nominal bond risk premium is defined as the difference between the observed nominal yield on a τ-year bond and the average of the future one-year short yields expected over its life:

$$\phi_t(\tau) = y_t(\tau) - \frac{1}{\tau}\sum_{j=1}^{\tau} E_t\left[y_{t+j-1}(1)\right] \quad , \quad (4)$$

where $y_t(\tau)$denotes the nominal yield at date t on a bond maturing in τ years, and y_{t+j-1} is the one-year nominal yield at date $t + j - 1$, with $E_t[\cdot]$ denoting the expectation operator conditional on the information available at date t. The expectations hypothesis would set this quantity to zero or to a constant, but the authors decline to impose it, as a substantial body of evidence indicates that it fails in the data. Two consequences follow. First, under the model's approximations, the sum of expected nominal payments on the fixed-rate contract exceeds that on the adjustable-rate contract by the bond risk premium multiplied by the maturity of the loan, $\tau\ \phi_t(\tau)$. Second, because the bond risk premium varies systematically with the state of the economy, the relative attractiveness and cost of the two mortgage contracts fluctuate dynamically over time. The utility framework then delivers a condition under which the adjustable contract is preferred. Within the model's structure, a borrowing-constrained borrower j with risk aversion γ_j and mortgage balance B prefers the adjustable-rate mortgage at date t when the nominal bond risk premium exceeds the risk-return tradeoff between the two contracts. For a two-period horizon, this condition is expressed as:

$$\phi_t(2) > \frac{\gamma_j}{2} B\left(\sigma_y^2 - \sigma_x^2\right) \quad , \quad (5)$$

where σ_y^2 denotes the variability of the real interest rate (which drives ARM risk) and σ_x^2 denotes expected inflation risk (which drives FRM risk). In the more general multi-period setting, a household prefers the ARM when the certainty-equivalent consumption differential between the two contracts is positive, which is dynamically driven by the fluctuations in the nominal bond risk premium $\phi_t(\tau)$. This framework makes the relationship to the existing literature explicit. While Campbell and Cocco (2003) hold the bond risk premium constant and let household characteristics determine which contract dominates, Koijen, Van Hemert, and Van Nieuwerburgh (2009) hold the cross-sectional distribution of households fixed and let the bond risk premium vary over time. Consequently, the

same borrower may satisfy the selection condition in one period and fail it in another, showing that the two contributions address complementary sides of the same mortgage-choice inequality. Equation (4) also identifies what earlier work had been measuring. Studies including Brueckner and Follain (1988), Dhillon, Shilling and Sirmans (1987) and Vickery (2007) documented that adjustable-rate origination responds to the spread between long and short yields, but that spread is not the premium. Subtracting the current nominal short rate, $y_t(1)$, from both sides of the definition of the nominal bond risk premium given in equation (4) yields:

$$y_t(\tau) - y_t(1) = \phi_t(\tau) + \left(\frac{1}{\tau} \sum_{j=1}^{\tau} E_t \left[y_{t+j-1}(1) \right] - y_t(1) \right) \quad , \quad (6)$$

so that the yield spread combines the nominal bond risk premium with the gap between expected future short rates and the current short rate. Only the first term (the nominal bond risk premium) bears on the mortgage choice, since the second term represents expected interest rate movements that are priced into both contracts alike, leaving the borrower indifferent. In the authors' estimates, these two terms move in opposite directions, showing a negative correlation of approximately -0.57 for the five-year maturity. This negative covariance makes the yield spread a highly contaminated proxy for the actual bond risk premium, introducing a severe errors-in-variables problem that accounts for its extremely weak explanatory power in predicting the ARM share.

1.3.2 Empirical performance and the timing decision rule

The empirical contribution of the paper is to show that this mechanism accounts for a substantial share of the observed variation in mortgage choice over time. The authors measure the premium in three ways, drawing on professional forecasters, on a vector autoregressive model of the term structure, and on a simple decision rule available to households themselves, and each measurement supports the same conclusion at both the aggregate and the loan level. The third is of particular interest for the present review.
The authors show that the nominal bond risk premium can be successfully approximated without the sophisticated econometric apparatus required by their formal estimation methods. Replacing the conditional expectation of future short-term rates in equation (4) with a simple average of the nominal short rates observed over the recent past yields the household decision rule:

$$k_t(\rho;\tau) = y_t(\tau) - \frac{1}{\rho}\sum_{u=o}^{\rho-1} y_{t-u}(1) \qquad (7)$$

where ρ is the length of the backward-looking window and τ represents the maturity of the loan. This intuitive rule of thumb requires only the current long-term yield, a historical record of recent short-term rates, and basic arithmetic, yet it nests two prominent measures previously used in the literature. Setting $\rho = 1$ recovers the conventional yield spread, which is the theoretically optimal proxy for mortgage choice only under the restrictive assumption that short-term rates follow a random walk. Conversely, letting the window grow without bound ($\rho \to \infty$) recovers the long-term yield net of the unconditional mean of the short rate.

For intermediate look-back horizons, the correlation with the observed ARM share is remarkably high, peaking at approximately 81% when the window covers between two and four years of past data (specifically, three years for Treasury yields and two years for mortgage rates). The measure performs robustly whether computed from Treasury yields, as in the baseline, or directly from mortgage market rates. This close comovement strongly supports the hypothesis that households engage in strategic timing when choosing their mortgage contracts, suggesting that close-to-rational financial decision-making requires significantly less cognitive and mathematical sophistication than standard rational expectations frameworks imply. A crucial empirical refinement further clarifies the economic drivers behind this behavior. The nominal bond risk premium is the sum of a component compensating for inflation risk and a component compensating for real interest rate risk, which are negatively correlated in the authors' estimates. Interestingly, when decomposed, these two risk components explain mortgage choice better than their aggregate sum. In the United States, where fixed-rate mortgages traditionally carry very long (30-year) maturities, the inflation risk premium component alone accounts for more than 60% of the variation in the ARM share, while the real rate risk premium contributes very little. This pattern is reversed in the United Kingdom, where fixed-rate contracts have much shorter maturities (usually 1 to 3 years), making households highly sensitive to short-term real rate fluctuations rather than long-term inflation risk. This cross-country contrast indicates that households do not respond to the aggregate nominal premium blindly; rather, they selectively respond to the specific risk factor (inflation vs. real rate) that the prevailing contract structure of their mortgage market exposes them to.

1.3.3 Implications for the present thesis

Three implications of the framework proposed by Koijen, Van Hemert, and Van Nieuwerburgh (2009) shape the way the empirical literature is read in the chapters that follow. The relative price of the two contracts at the moment of origination emerges as a market-level determinant that no account of mortgage choice can omit. The theoretical object is the bond risk premium, and its observable counterpart in the literature is the spread between the fixed and adjustable rates prevailing when the contract was signed, reported for the United States by the Primary Mortgage Market Survey of Freddie Mac. This determinant is recorded as row 10 of the synoptic table presented in section 1.5, and section 3.6 examines what the empirical work has established about it.

Household expectations about the future path of interest rates acquire the status of a determinant in their own right, and this implication bears most directly on the question from which the present work takes its title. Under the rational reading of the framework the relevant object is the household's subjective expectation of future short rates; under the adaptive rule of equation (7) it is the household's extrapolation from rates observed in the recent past. On both interpretations, what a household believes about where rates are heading is theoretically prior to its choice of contract rather than incidental to it. The Survey of Consumer Finances records this belief directly, and section 3.7 considers how far the literature has made use of it. The framework also makes explicit a dependence on the macroeconomic state that the lifecycle model leaves implicit. Since the premium varies over time while household characteristics do not, the contract that is optimal for a given household in one period need not be optimal for the same household in another, and estimates obtained under one configuration of the term structure need not carry over to a different one. This consideration is based on the interpretation of a literature whose findings were obtained across several decades and several distinct rate environments, and it is revisited in chapter 5. It also motivates the attention devoted in chapter 2 to the environment of the early 2020s, in which the Federal Reserve conducted the sharpest tightening in four decades and both the premium and the fixed mortgage rate moved substantially within a matter of months.

1.3.4 Critical limitations

Two limitations of the Koijen, Van Hemert, and Van Nieuwerburgh (2009) framework deserve explicit mention.

Firstly, the paper's decision-rule analysis assumes that households behave as if they observe the bond risk premium and act on it, whether through sophisticated forecasting or through simple adaptive rules. Whether households are actually able to follow such a strategy remains an open empirical question. This issue is particularly relevant for less financially literate households, since Bucks and Pence (2008) show that such category of borrowers is more likely to not fully understand even the basic characteristics of their own mortgage contracts. Chapter 3 returns to this question when examining the evidence on financial literacy and mortgage choice. Second, the framework abstracts from the household-level idiosyncratic factors (income volatility, mobility, risk aversion) that organize the Campbell and Cocco (2003) analysis. The two papers are therefore complementary rather than substitutes.

1.4 Further theoretical contributions

Three limitations of the framework proposed by Koijen, Van Hemert, and Van Nieuwerburgh (2009) deserve explicit mention. The first concerns the behavioral requirement the analysis imposes. The decision-rule argument assumes that households act as if they observed the bond risk premium, whether by forecasting it or by applying the adaptive approximation of equation (7). Whether households are in fact able to do so remains an open empirical question, and one that bears particularly on the less financially literate, since Bucks and Pence (2008) find that borrowers in that group frequently fail to grasp even the basic terms of the contracts they hold. Chapter 3 returns to this question when examining the evidence on financial literacy and mortgage choice. The second concerns the structure of the model from which the theoretical condition is derived. The framework is deliberately austere, and its assumptions differ markedly from those of the lifecycle model discussed in section 1.2. The household borrows the full value of the property, makes interest payments only, is constrained to consume whatever income remains after servicing the mortgage, and receives an income stream uncorrelated with aggregate conditions. The authors adopt these simplifications explicitly, in order to isolate the role of time-varying premia, and the resulting apparatus is a device for motivating the empirical analysis rather than a description of how the decision is taken. The third follows from the second. Because the model holds the household fixed, it abstracts from the idiosyncratic characteristics that organize the Campbell and Cocco (2003) analysis, namely income volatility, expected mobility and risk aversion, and it has correspondingly little to say about which households choose which contract at a given moment. The two contributions

are therefore complementary rather than competing: one accounts for the cross-sectional distribution of contract choice at a point in time, the other for its movement over time, and the empirical literature reviewed in chapter 3 draws on both.

1.4.1 Extensions of the rational framework

Campbell (2006) argues that households face constraints absent from textbook models, including fixed costs, uninsurable income risk, borrowing limits, and contracts that are not neutral with respect to inflation, each of which bears directly on the adjustable versus fixed decision. It also documents a pattern that runs through the present thesis, namely that while most households invest competently, a minority make significant mistakes, and that this minority is concentrated among the poorer and less educated. This observation supplies the conceptual bridge between the rational benchmark and the behavioral considerations that follow, since it implies that models incorporating education and financial literacy measures capture heterogeneity in decision-making quality, rather than serving solely as controls for potential confounding factors. Cocco (2005) embeds the mortgage decision in a lifecycle portfolio model with house-price risk and shows that housing investment crowds out stockholding, an effect that is larger for households with low financial net worth. Two implications follow for contract choice. The housing position consumes the household's risk-bearing capacity, so that households heavily exposed to housing relative to their financial wealth have little residual capacity to absorb the payment volatility of an adjustable contract, which reinforces from a portfolio angle the lifecycle prediction that low-wealth households should prefer the fixed rate. At the same time, the mortgage amplifies housing exposure through leverage, so that the contractual form of that leverage governs how rate shocks pass into consumption. Net worth, leverage, and balance-sheet composition are therefore theoretically motivated determinants, and the evidence on each is examined in section 3.9.

The institutional dimension of mortgage choice is examined by Miles (2004), who investigates why long-term fixed-rate mortgages remained relatively uncommon among borrowers in the United Kingdom despite their potential welfare advantages. The review argues that this outcome is driven less by household preferences and more by the institutional features of mortgage funding and pricing. In particular, factors such as lenders' willingness to bear long-term interest rate risk and the degree of development of securitization markets play a central role in shaping the availability and attractiveness of fixed-rate contracts. The review also characterizes fixed-rate mortgages as instruments that

provide protection against inflation risk, while suggesting that households may undervalue these long-term benefits because they place greater emphasis on short-term affordability considerations.

Complementary evidence is provided by Vickery (2007), whose analysis of a large loan-level dataset shows that borrowers respond strongly to differences in the relative pricing of mortgage contracts. His findings indicate that the share of adjustable-rate mortgages closely follows movements in the spread between adjustable-rate and fixed-rate mortgage rates. When considered alongside the mortgage timing framework developed by Koijen, Van Hemert, and Van Nieuwerburgh (2009), this evidence identifies the mortgage rate spread as one of the most consistently supported determinants of mortgage contract choice in the literature and provides a strong justification for its inclusion in the empirical model.

1.4.2 Expectations and behavioral frictions

The rational account assumes that households form expectations consistently with the public information set and understand the products they choose. A body of evidence suggests that both assumptions fail for a meaningful share of borrowers. Malmendier and Nagel (2016) show that individuals outweigh the inflation realised during their own lifetimes when forecasting future inflation, with younger individuals updating more strongly because recent experience accounts for a larger share of their history. The authors apply the mechanism to borrowing decisions, including mortgage choice, and the implication for the 2022 setting is direct. The households making active mortgage decisions in 2022 had formed their financial expectations overwhelmingly within the post-Volcker low-rate regime, and under experience-based learning they would have read the inflation surge of that year as a transitory deviation rather than a regime change, which could leave them more willing than the rational benchmark predicts to accept an adjustable contract in anticipation of a return to low rates. Lusardi and Mitchell (2014) treat financial knowledge as a form of human capital and document that financial literacy is low, unevenly distributed, and predictive of financial outcomes across many domains. For mortgage choice the relevance is conceptual rather than mechanical, since the sign of the literacy coefficient is theoretically ambiguous. More literate households may recognize when an adjustable contract is appropriately priced and choose it for that reason, or less literate households may be drawn into adjustable contracts by low initial payments without grasping the future risk, in which case literate households would avoid the product. The two readings

deliver opposite signs, and the empirical answer depends on which mechanism dominates. The two strands meet in the work of Bucks and Pence (2008) that compares mortgage characteristics, as reported by households in the SCF, with those reported by lenders in administrative data. Most borrowers report the broad terms of their mortgages accurately, including the adjustable versus fixed distinction, which is reassuring for the dependent variable used throughout this literature. Adjustable-rate borrowers, however, are markedly less aware of the specific features governing how their rates can change, and those least informed about these features tend to be the lower-income households most exposed to financial strain should rates rise. The finding does not overturn the rational benchmark, since the broad choice on which the benchmark operates is well understood, but it marks an empirical boundary on how far that benchmark can be pushed, and it lends weight to the behavioral and literacy considerations that the contested rows of the synoptic table record.

1.5 Synoptic table of determinants and testable predictions

The five preceding sections have built the theoretical apparatus required to address the central question of this thesis: which household characteristics and which macroeconomic state variables predict the choice between an adjustable-rate and a fixed-rate mortgage, and how do these predictions stand against the empirical record? The present section consolidates the chapter into a single structured device, Table 1, that records for each candidate determinant the predicted sign on the probability of choosing ARM, the source paper or paradigm from which the prediction is drawn, and the underlying economic mechanism. Where the rational-benchmark framework of sections 1.2 to 1.4 and the behavioral or empirical evidence of sections 1.3 and 1.4 produce competing predictions, the table records both signs and flags the row as a contested determinant. The table organizes eleven candidate determinants into three categories. The first category, household preferences and beliefs, captures the household's risk attitudes, expectations about future macroeconomic conditions, and planned mobility horizon. The second category, household economic position, captures income, net worth, age, education, and the related dimension of income volatility. The third category, mortgage and market characteristics, captures the loan-level features and the macro-level pricing environment in which the choice is made. Three columns operationalize the synthesis. The predicted-sign column reports the direction in which an increase in the determinant is expected to move the probability of choosing ARM. The source column attributes the prediction to a specific paper or class of

papers from sections 1.2 to 1.5. The mechanism column states, in one or two clauses, the economic channel through which the determinant operates. Rows marked as contested record two predictions, each attributed to the competing theoretical frameworks.

Table 1. Determinants of ARM choice, with theoretical predictions and empirical benchmarks

#	Determinant	Predicted sign on P(ARM)	Source paper or paradigm	Economic mechanism
1	Risk aversion	Negative	Campbell and Cocco (2003); Coulibaly and Li (2009)	Risk-averse households value the payment-stability insurance provided by the FRM and accept the term premium as the cost of that insurance. Coulibaly and Li (2009) confirm empirically on SCF data that more risk-averse borrowers prefer fixed-rate mortgages.
2	Credit and liquidity constraints (affordability)	Contested. Negative under rational benchmark; positive under empirical evidence.	Campbell and Cocco (2003) vs Coulibaly and Li (2009)	Rational benchmark: constrained households bear the highest welfare cost of ARM payment volatility and should therefore prefer FRM as insurance. Empirical pattern: constrained borrowers tilt toward ARM, drawn by the lower initial payments that affordability considerations make decisive. This is the central contested determinant of the thesis.
3	Expected mobility (short tenure)	Positive	Campbell and Cocco (2003); Coulibaly and Li (2009); Vickery (2007)	Households with short expected tenure do not bear long-horizon interest-rate risk because they expect to prepay before the FRM's insurance becomes valuable. The term premium embedded in the FRM rate is therefore an unattractive cost. Empirically robust across studies.

#	Determinant	Predicted sign on P(ARM)	Source paper or paradigm	Economic mechanism
4	Expected direction of interest rates	Contested. Negative if rising rates expected; positive if falling rates expected, under rational benchmark. Sign under experience-based view depends on cohort.	Koijen, Van Hemert, and Van Nieuwerburgh (2009); Malmendier and Nagel (2016)	Rational interpretation: households who expect rising rates lock in current FRM rates, while those expecting falling rates prefer ARM to capture future declines. Experience-based interpretation: households with formative experience in low-rate regimes extrapolate forward and may hold ARM in the expectation of mean reversion regardless of current macro signals. Central variable of the thesis title.
5	Household income	Positive	Brueckner and Follain (1988); Coulibaly and Li (2009); Bacon and Moffatt (2012)	Higher-income households can self-insure against ARM payment volatility through buffer-stock financial wealth and accept the lower expected cost of the ARM in exchange. Multiple empirical studies converge on the positive sign, contradicting the simpler intuition that high-income households should be most insurable against payment shocks through the FRM.
6	Net worth (financial buffer stock)	Positive	Campbell and Cocco (2003); Cocco (2005)	Wealthy households can absorb ARM payment volatility from financial reserves without compromising non-housing consumption. The FRM term premium is therefore an unattractive cost of redundant insurance for households with substantial buffer-stock wealth.
7	Age of household head	Contested. Negative under lifecycle view; positive under experience-based view depending on cohort.	Campbell and Cocco (2003) vs Malmendier and Nagel (2016)	Lifecycle channel: older households have shorter remaining horizons over which to bear interest-rate risk and may prefer ARM. Experience-based channel: older households who lived through high-inflation episodes hold higher inflation expectations and value FRM insurance more. The empirical sign reflects which channel dominates in the sample.

#	Determinant	Predicted sign on P(ARM)	Source paper or paradigm	Economic mechanism
8	Education and financial literacy	Contested. Positive under rational interpretation; negative under protection interpretation.	Lusardi and Mitchell (2014); Campbell (2006); Bucks and Pence (2008)	Rational interpretation: more literate households understand ARM risks and select the product when it is appropriately priced for their circumstances. Protection interpretation: less literate households are funneled into ARM by lower initial payments without understanding the future payment risk. The empirical sign discriminates between the two views.
9	Mortgage horizon (contract term)	Positive	Koijen, Van Hemert, and Van Nieuwerburgh (2009)	Longer FRM duration concentrates more of the bond risk premium in the FRM rate, making the relative pricing of the ARM more attractive at the margin. The effect operates through the same pricing channel that drives row 10.
10	Current ARM–FRM spread at origination	Positive	Koijen, Van Hemert, and Van Nieuwerburgh (2009); Vickery (2007)	When the term premium is high, the FRM rate is expensively priced relative to the expected ARM cost path, and households substitute toward the ARM. This is the market-level pricing variable most strongly supported by empirical evidence at both household and aggregate levels.
11	Labor income volatility	Negative	Campbell and Cocco (2003); Coulibaly and Li (2009)	Volatile labor income amplifies the welfare cost of ARM payment shocks because adverse income realizations and adverse rate realizations may coincide, producing joint distress. The FRM provides hedging against this joint adverse outcome. Coulibaly and Li (2009) confirm the negative sign empirically on SCF data.

1.5.1 Notes on the table

Six rows of Table 1, i.e. 1, 3, 6, 9, 10, 11, deliver the same predicted sign across rational-benchmark, timing, and empirical frameworks; they function in the empirical analysis as validation checks, since large deviations on these well-established regularities would signal either a sample anomaly or a regime change. The analytical backbone of chapter 5 lies

instead in rows 2, 4, 7, and 8, where rational-benchmark and behavioral or empirical predictions diverge. Row 5 (income) warrants a separate observation: although the positive sign appears to contradict the textbook reading of Campbell and Cocco (2003), the empirical literature converges on it through the buffer-stock channel by which higher-income households self-insure against payment shocks and treat the FRM term premium as redundant insurance, a result independently confirmed by Coulibaly and Li (2009) and by Brueckner and Follain (1988). Bacon and Moffatt (2012) reach a compatible conclusion on British data through a different route, inferring from the choices of some 280,000 borrowers that risk aversion declines with income, which points to preferences rather than to buffer-stock self-insurance as the operative channel. The contestation of row 4, the variable from which the thesis takes its title, differs in nature from that of rows 2, 7, and 8: in a sample dominated by households formed in the post-Volcker low-rate regime, as is the case for the SCF 2022 mortgaged subsample, the rational interpretation of Koijen, Van Hemert, and Van Nieuwerburgh (2009) and the experience-based interpretation of Malmendier and Nagel (2016) may converge on the same predicted sign, so that the aim of the corrisponding section in chapter 5 is not to discriminate the two frameworks but to document the coefficient's magnitude and interpret it against both, treating a null result as itself substantive.

2. The U.S. macroeconomic context

2.1 Interest rates and mortgage rates, 2000 to 2024

The choice a household makes between a fixed-rate and an adjustable-rate mortgage cannot be understood apart from the interest rate environment in which the decision is taken. Because that environment changed profoundly over the two decades leading up to 2022 and changed with unusual speed in the final stretch of that period, it is worth setting out the trajectory of both the policy rate and the mortgage rate before turning, in the chapters that follow, to the household characteristics that the literature has identified as determinants of the choice. This section traces the path of the federal funds rate and of the thirty-year fixed mortgage rate from the early 2000s to the middle of the present decade, distinguishing three broad phases that differ not only in the level of rates but in the volatility and direction that borrowers had to contend with. The first phase covers the years from 2000 to the eve of the financial crisis. After the collapse of the technology-stock bubble and the recession of 2001, the Federal Reserve lowered its policy rate through a sequence of cuts that brought the federal funds target to one percent by the middle of 2003, a level that at the time seemed remarkably low by the standards of the preceding decades. The subsequent recovery prompted a gradual reversal, and between 2004 and 2006 the Fed raised the target in seventeen consecutive quarter-point steps to a peak of 5.25 percent. The thirty-year fixed mortgage rate over this stretch moved within a comparatively narrow band, opening the decade around eight percent and settling into the upper five to six percent range for much of the mid-2000s. For the purposes of the present study, what matters about this phase is that it was a period of moderate and cyclically normal rate movements, the kind of environment in which the adjustable-rate product retained a meaningful presence in the market and in which the empirical literature reviewed in chapter 3 was largely produced. The second phase opens with the financial crisis of 2008 and extends through the pandemic period to early 2022. In response to the crisis the Federal Reserve cut its policy rate almost to zero by the end of 2008 and then held it there, in what became the longest stretch of near-zero rates in its history, from December 2008 until December 2015. The first tentative increase in December 2015 was followed by a slow normalization over the following years, but the arrival of the pandemic in early 2020 sent the policy rate back to the zero lower bound, where it remained from March 2020 until March 2022. Alongside the near-zero policy rate, the Federal Reserve conducted large-

Artificial intelligence tools were used to support the drafting, editing, and linguistic refinement of selected parts of this chapter.

scale purchases of Treasury and mortgage-backed securities, and these purchases pressed long-term borrowing costs down further still. The thirty-year fixed mortgage rate, which had already drifted below four percent during the 2010s, reached its lowest level in the entire history of the Freddie Mac survey in December 2020, when it averaged 2.68 percentage points over that month. The salient feature of this long second phase, and the one that bears most directly on the analysis of household choice in 2022, is that an entire cohort of borrowers formed its expectations and its sense of what constituted a normal mortgage rate during more than a decade of exceptionally and persistently cheap credit. The third phase is the abrupt reversal of 2022 and 2023. As inflation accelerated through 2021 and into 2022, reaching a year-over-year peak of 9.10 percent in June 2022, the Federal Reserve abandoned its earlier characterization of the price pressures as transitory and moved to restrictive policy with exceptional speed. Between March 2022 and July 2023 the federal funds target surged from near zero to a range of 5.25 to 5.50 percent, the fastest tightening cycle since the Volcker disinflation of the early 1980s. Mortgage rates followed the broad direction of policy, though they are tied more closely to the ten-year Treasury yield than to the policy rate itself. The thirty-year fixed rate began 2022 at 3.45 percent average, climbed through the spring and summer, and reached a peak of around seven percent in October 2022, before touching still higher levels during 2023. In the space of less than a year the cost of a thirty-year fixed mortgage had doubled, as displayed in Figure 1. It is against this backdrop that the households covered by the 2022 wave of the survey of consumer finances made their borrowing decisions, and the significance of this timing is discussed further in the sections that follow.

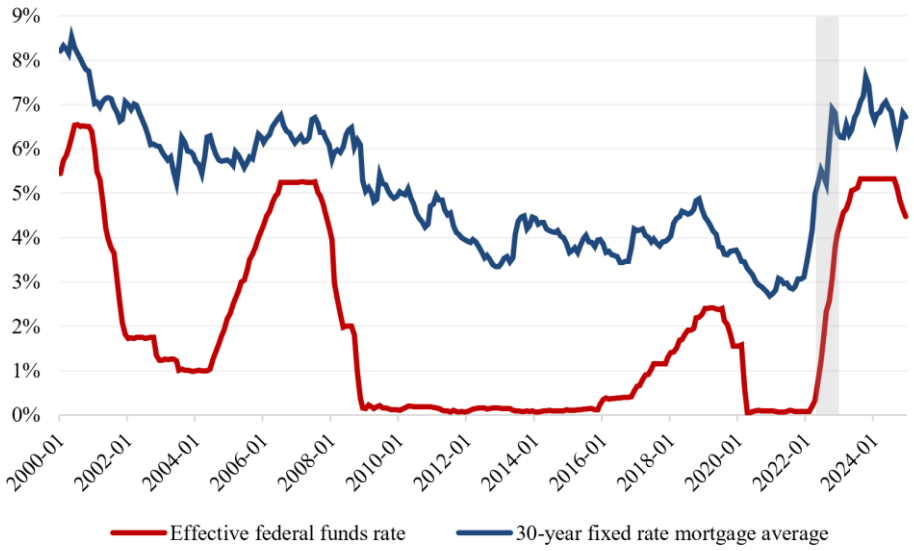


Figure 1.

Effective federal funds rate and 30-year fixed mortgage rate, January 2000 to December 2024.

Monthly observations of the effective federal funds rate (FRED series: FEDFUNDS) and the 30-year fixed mortgage rate (FRED series: MORTGAGE30US, aggregated to monthly averages). The shaded area denotes the 2022 field period of the Survey of Consumer Finances (SCF). Source: Federal Reserve Bank of St. Louis (FRED), using data from the Board of Governors of the Federal Reserve System and Freddie Mac.

Two features of this trajectory deserve emphasis because they frame the interpretation of household behavior in the period under study. The first is the sheer magnitude and speed of the 2022 reversal, which compressed into a few months a movement in borrowing costs that in earlier cycles would have unfolded over a much longer span, leaving households little time to adjust their expectations or their plans. The second is the contrast between the long preceding decade of cheap and stable credit and the sudden shift to expensive and volatile credit, a contrast that made the direction of future rates a genuinely open and salient question for borrowers in a way it had not been for many years. Households approaching a mortgage decision in 2022 did so having spent much of their adult financial lives in an environment of falling or historically low rates, and the rapid rise confronted them with conditions that had no recent precedent in their own experience.

2.2 The composition of the mortgage market and the adjustable-rate share

Having traced the movement of rates, it is natural to ask how borrowers responded, and the most direct evidence lies in the changing weight of adjustable-rate mortgages within the flow of new lending. The share of originations that carry an adjustable rate is a useful barometer of household behavior precisely because it aggregates thousands of individual

decisions into a single series that can be read against the interest rate environment described in the previous section. What that series reveals is a market whose appetite for adjustable products has swung dramatically over the past two decades, and whose recent movements bear directly on the interpretation of household choices in 2022. Before turning to the numbers, it is worth recalling the mechanism that links the adjustable-rate share to the rate environment, since it was set out in theoretical terms in chapter 1. Demand for adjustable products tends to rise when the gap between short-term and long-term interest rates widens, because a steeper term structure makes the initial payment on an adjustable loan cheaper relative to the fixed alternative, offering borrowers a larger saving in exchange for accepting future rate risk. This relationship, which underlies the pricing logic of Koijen, Van Hemert, and Van Nieuwerburgh (2009) and the loan-level evidence of Vickery (2007), means that the adjustable share should not be read as a free-standing measure of borrower preference but as a response to the relative price of the two contracts at each moment. The trajectory of the adjustable share can be divided along much the same lines as the rate history itself. During the housing boom of the mid-2000s, adjustable products were a substantial presence in the market. According to the Mortgage Bankers Association's weekly applications survey, the share of adjustable-rate mortgage applications fluctuated considerably during the 1990s and the first half of the 2000s, eventually reaching a record 36.6% in March 2005. This was the era of teaser-rate structures and of the exotic adjustable products that would later feature prominently on accounts of the crisis, and it was also, not coincidentally, the period during which much of the empirical literature on mortgage choice was written. Following the 2008 global financial crisis, the share of adjustable-rate mortgage applications declined sharply and remained at historically low levels throughout the 2010s, as the combination of ultra-low fixed rates and the removal of the riskier adjustable products from the market left borrowers with little reason to accept rate risk. Since the attractiveness of the ARM depends on the ARM-FRM spread, as shown by Vickery (2007), the potential cost savings associated with choosing an ARM are generally limited when fixed-rate mortgage rates are already very low. As a result, households have little incentive to accept the additional uncertainty associated with future interest rate adjustments. This pattern was further reinforced by the regulatory reforms introduced after the global financial crisis, which restricted the availability of the loosely underwritten adjustable-rate mortgage products that had contributed to their

widespread use during the mid-2000s.[2] Consequently, for more than a decade, adjustable-rate mortgages represented only a small share of the U.S. mortgage market.

The reversal of 2022 is the development most relevant to the present study. As fixed rates climbed through that year, the relative attractiveness of the adjustable product returned, and the adjustable share rose sharply from its depressed base. The Urban Institute's housing finance chartbook, drawing on the Mortgage Bankers Association data, records that the adjustable share of applications increased from 3.1 percent in the week ending 7 January 2022 to 12.8 percent by the week ending 14 October 2022. In relative terms this was a striking movement, roughly a fourfold increase over the course of a few months, even though in absolute terms the share remained far below its 2005 peak and modest by the standards of the pre-crisis market. The share then eased as fixed rates fluctuated, falling to 5.9 percent by late July 2023 before rising again.

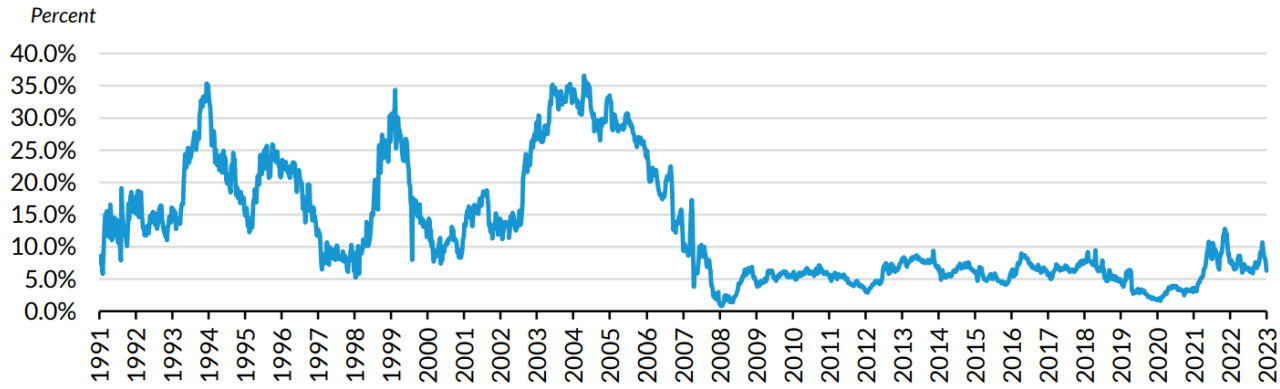


Figure 2.

Adjustable-rate share of mortgage applications, 1991 to 2023.

Source: Mortgage Bankers Association, Weekly Mortgage Applications Survey, as reported in Urban Institute, Housing Finance Policy Center, Housing Finance at a Glance: A Monthly Chartbook, December 2023. The series includes both purchase and refinance applications and is updated through 8 December 2023.

The trajectory just described is shown in Figure 2, which plots the adjustable-rate share of applications from the early 1990s through 2023 and makes visible at a glance the mid-2000s peak, the prolonged post-crisis trough, and the revival of 2022. One qualification should be noted for transparency. Freddie Mac discontinued the publication of its

[2] For the restriction of loosely underwritten adjustable-rate products, see the Ability-to-Repay and Qualified Mortgage Standards Under the Truth in Lending Act (Regulation Z), issued by the Consumer Financial Protection Bureau to implement sections 1411, 1412, and 1414 of the Dodd-Frank Act. Final Rule, 78 Fed. Reg. 6408 (January 30, 2013), effective January 10, 2014.

adjustable-rate series in late 2022, so a continuous adjustable-rate quotation matching the full span of the fixed-rate series is no longer available from that source. The application-based share reported by the Mortgage Bankers Association, on which the figure relies, is therefore the more reliable basis for documenting the trends described here. Beyond documenting the trend, these movements in the adjustable share invite a further observation, because they bear on how the composition of adjustable-rate borrowers in 2022 should be understood. When the adjustable share is low, as it was for most of the preceding decade, the households that opt for an adjustable product tend to form a self-selected group, often wealthier or more financially sophisticated borrowers able to exploit favorable pricing or comfortable bearing rate risk. When the share climbs sharply in response to a jump in fixed rates, the pool widens to include households drawn to the product chiefly by its lower initial payment, which for some may be the only means of meeting an affordability constraint in a high-rate environment. The adjustable-rate population of 2022 is therefore unlikely to mirror that of earlier periods, a consideration that connects to the discussion of affordability and financial constraints reviewed in chapter 3 and that colors any reading of who chose adjustable products in that year.

2.3 House prices, real income, and household indebtedness

The mortgage decision is not taken in isolation from the rest of the household balance sheet, and the years surrounding the 2022 survey were marked by unusually large movements in the assets and liabilities that frame that decision. To understand the position of households approaching a mortgage in 2022, it helps to look beyond the rates themselves to three related dimensions of the economic environment: the value of the homes being financed, the real purchasing power of the incomes servicing the loans, and the weight of debt already resting on household balance sheets. Each of these moved in ways that bear on how affordability and financial capacity should be understood in the period under study. The most visible of these movements was the surge in house prices during the pandemic. According to Board of Governors of the Federal Reserve System (2022), the S&P Case-Shiller national home price index rose by approximately 37 percent in nominal terms in the two years between March 2020 and March 2022, an increase without recent precedent over so short a span. For much of 2020 and 2021 this rise in prices was offset, from the buyer's point of view, by the record-low mortgage rates described earlier in this chapter, so that the monthly cost of financing a purchase remained manageable even as the sticker price climbed. The offset did not last. When mortgage rates began their ascent

in early 2022, the combination of elevated prices and rising rates pushed the required monthly payment on a newly originated purchase mortgage to roughly 50 percent above its pre-pandemic level, according to the same source.

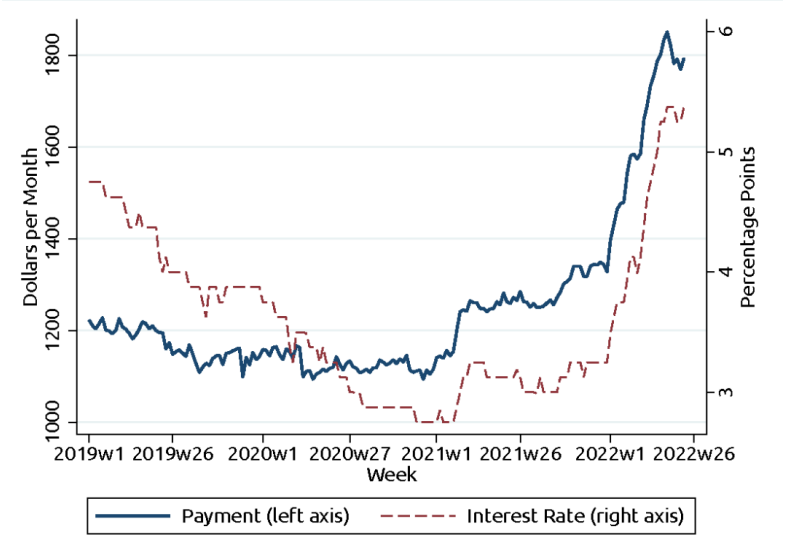


Figure 3.

Median monthly mortgage payment and median interest rate on home purchase loans, 2019 to 2022.

Source: Optimal Blue LLC, Optimal Blue Mortgage Price Data, as reported in Board of Governors of the Federal Reserve System, Declining affordability and home purchase borrowing by lower income households, FEDS Notes, July 8, 2022. The solid line reports the median monthly principal and interest payment, and the dashed line the median interest rate, on mortgages for the purchase of owner-occupied homes, grouped by the week in which the borrower locked the rate.

The household contemplating a home purchase in 2022 thus faced a double burden that had not existed in either of the two preceding years, when at least one of the two components of cost had been favorable. This shift has an important consequence for the way household wealth interacts with mortgage choice, and the consequence cuts in two directions that are worth distinguishing carefully. On one side, the appreciation of house prices meant that existing homeowners entered 2022 with substantially more housing equity than they had held a few years earlier, which enlarged the cushion available to those already on the property ladder. On the other side, that same appreciation raised the sums that new buyers had to borrow and the loan-to-value positions they had to accept, tightening the financial constraints facing those attempting to enter the market. The point is that rising house values do not affect all households in the same way, and the balance-sheet strength implied by higher home equity accrues chiefly to those who already own, not to those seeking to buy. A note of caution follows from this, and it echoes a distinction drawn elsewhere in this work. The wealth that the buffer-stock logic identifies as relevant

for bearing the payment volatility of an adjustable loan is liquid financial wealth, not the illiquid equity locked in a house, so the surge in home values inflates measured net worth without necessarily expanding the resources a household can draw on to absorb a rise in its monthly payment. The behavior of real income over the same period reinforced the pressure on affordability. Nominal wage growth was robust through 2021 and into 2022, but it was outpaced for a time by the sharp acceleration of inflation, which peaked at 9.1 percent in the middle of 2022 as noted in section 2.1. The result was that the real purchasing power of household income came under strain precisely when mortgage costs were rising, so that the resources available to meet higher payments were being squeezed from two sides at once. Even where nominal income figures appeared healthy, the erosion of their real value meant that households servicing or seeking a mortgage in 2022 were doing so under tighter real conditions than the headline wage numbers alone would suggest.
The third dimension, the overall weight of household debt, provides the backdrop against which these pressures played out. Aggregate household debt in the United States reached 17.06 trillion dollars by the end of the second quarter of 2023, of which housing debt accounted for 12.01 trillion, or roughly 70 percent of the total. Mortgage debt thus remained by far the dominant component of what American households owe, as it has been throughout the modern period. Yet the aggregate figures conceal a feature of the 2022 landscape that is central to interpreting it. A large share of existing mortgage holders had locked in fixed rates at the exceptionally low levels available during 2020 and 2021, and these households carried a debt-service burden far lighter than current rates would imply. This lock-in effect, whereby homeowners with cheap fixed-rate loans have little incentive to move or refinance once rates rise, created a divide between the favorable position of those who had borrowed earlier and the far less favorable position of those entering the market in 2022. It is a feature without close parallel in the earlier periods studied by the mortgage-choice literature, and it forms part of the distinctive character of the moment at which the households in the survey of consumer finances were observed.
Taken together, these three dimensions describe an environment that was, in an important sense, internally divided. Households approaching a mortgage in 2022 faced simultaneously high asset values, high nominal incomes, eroded real incomes, historically low debt-service burdens for those already locked into cheap loans, and sharply higher borrowing costs for those taking on new debt. The standard determinants of mortgage choice pointed in different directions for different segments of the population, and no single summary of the macroeconomic backdrop captures the situation of every household

within it. This heterogeneity is worth keeping in view, because it means that the aggregate picture, however striking, translates into the circumstances of individual households only through the varied positions they occupied on the balance sheet.

2.4 The 2022 economic environment and household mortgage decisions

The three previous sections have described the interest rate environment, the structure of the U.S. mortgage market, and the balance sheet position of American households around the time of the 2022 survey. Considered separately, each provides a consistent picture of one aspect of the economic environment. Taken together, however, they reveal a setting characterized by important tensions. It is therefore useful to bring these elements together into a single picture before turning, in the following chapters, to the evidence on how households choose between fixed-rate and adjustable-rate mortgages.

The most striking aspect of the 2022 landscape is that its various components pointed in different directions at once. Interest rates had risen at a speed unmatched in four decades, making new borrowing far more expensive than it had been a year earlier. House prices stood near historic highs, having climbed roughly 37 percent in the two years to early 2022, which enlarged the equity of existing owners while raising the sums new buyers had to finance. Nominal incomes were growing, yet inflation had eroded their real value at the very moment mortgage costs were climbing. Aggregate debt service burdens remained low by historical standards, but this apparent comfort concealed a sharp divide between households that had locked in cheap fixed rates during 2020 and 2021 and those confronting the market afresh in 2022. No single summary captures a situation in which high asset values, high nominal incomes, eroded real incomes, light debt service for some, and steep borrowing costs for others all coexisted. The environment was, in a real sense, several environments at once, experienced differently by households depending on when they had borrowed and where they stood on the balance sheet. This heterogeneity matters for the way the rest of this work approaches the mortgage decision, because it means that the aggregate picture translates into the circumstances of individual households only through the varied positions they occupied. A household that had bought and locked a low fixed rate in 2021 faced the 2022 environment from a position of relative safety, insulated from the rate rise and sitting on enlarged equity. A household entering the market in 2022 faced the same environment as a wall of higher costs, and for such a household the adjustable mortgage, with its lower initial payment, could take on an appeal it had lacked throughout the preceding decade of cheap fixed credit. The point

developed in section 2.2, that the composition of adjustable-rate borrowers shifts with the rate environment, follows directly from this. When the environment is internally divided, the households that select each contract type are sorted not only by their preferences but by their position relative to the rate cycle.

There is one dimension of the 2022 setting that deserves particular emphasis, because it connects the macroeconomic backdrop to the behavioral questions that later chapters take up. For more than a decade, the direction of interest rates had been a settled matter. Rates were low, and they showed little sign of moving decisively, so that a household approaching a mortgage had little reason to agonize over where rates might head. The events of 2022 changed this abruptly. Rates were not only high but visibly rising, and their future path had become genuinely uncertain, with informed observers divided over whether tightening would continue or reverse. For the first time in years, a household choosing a mortgage in 2022 had to form a view about the future of interest rates, and that view could plausibly tip the decision one way or the other. A borrower expecting rates to keep climbing had reason to lock in a fixed rate before it rose further, while a borrower expecting a return to the low rates of the previous decade might prefer an adjustable contract and wait for its payments to fall. The salience of interest rate expectations, dormant through the long years of cheap and stable credit, returned with force in 2022. This observation is the natural bridge from the context to the substance of the review that follows. The cohort of households making mortgage decisions in 2022 had formed its sense of what constitutes a normal interest rate during more than a decade of historically low rates, a fact that connects directly to the experience-based account of expectation formation introduced in section 1.4. How households form expectations about rates, how well they understand the contracts they choose, and how these subjective and behavioral factors interact with the objective determinants of the decision are questions that the macroeconomic backdrop raises but cannot answer. The chapters that follow turn to the body of research that has addressed them, examining what the literature has established about the determinants of mortgage choice, how households form and act on their expectations, and where the accumulated evidence remains incomplete. The environment described in this chapter is the stage on which that research has been conducted, and keeping its internal tensions in view will help in reading what the literature has, and has not, been able to conclude.

3. Determinants of mortgage choice: a critical review of the empirical evidence

3.1 Scope and method of the review

3.1.1 Organizing principle and delimitations

The two preceding chapters have laid the groundwork for what follows. Chapter 1 assembled the theoretical vocabulary of the mortgage decision, setting out the models that treat the choice between a fixed and an adjustable contract as a problem in household risk management and closing with a table of determinants and the signs that theory predicts for each. Chapter 2 described the market in which that decision is taken, tracing two decades of movements in interest rates, in the share of adjustable lending, and in the balance sheet position of American households. The present chapter addresses the remaining question, namely what the empirical literature has established about those determinants. The organization of the material is not a neutral decision, since it conditions how the evidence on each determinant can be assessed. Reviews of empirical literature are commonly built study by study, proceeding chronologically or grouping contributions by the dataset they employ. This practice preserves the internal logic of each contribution and documents the chronological development of the field, but it disperses the evidence on any single determinant across multiple separate treatments. Establishing what is known about liquidity constraints, for instance, would require reassembling findings distributed throughout the chapter. Since the purpose here is precisely to establish what is known about each determinant, that fragmentation would work against the aim of the review. The material is therefore organized by determinant rather than by study. Each of the sections that follow takes up one of the factors that theory identifies as relevant to the mortgage decision and asks what the empirical work has established about it, drawing on however many studies bear on the question. The organizing principle is not invented for the occasion but inherited from the theoretical apparatus: the sections correspond, with some grouping where the underlying mechanism is shared, to the rows of Table 1 in section 1.5. Risk aversion and income volatility are treated together because both concern exposure to risk, one from the side of preferences and the other from the side of objective circumstances. Contract term and the price differential between the two products are likewise treated together, since both operate through the pricing channel that the timing framework identifies. The remaining determinants receive a section apiece. This arrangement makes

Artificial intelligence tools were used to support the drafting, editing, and linguistic refinement of selected parts of this chapter.

the correspondence between theoretical prediction and empirical result explicit in each section, rather than presenting the contributions in sequence.

The backbone of the review consists of studies based on the Survey of Consumer Finances. This choice reflects a feature of the data rather than a preference for any particular research tradition. Administrative loan-level records, which have supplied much of the evidence on mortgage pricing, capture the terms of the contract with great precision but reveal almost nothing about the household that signed it beyond what the lender recorded at origination. Panel surveys follow households over time but lack the balance sheet detail that the mortgage decision requires. The Survey of Consumer Finances occupies a distinctive position because it records, for the same household and in the same instrument, the type of contract held, the composition of assets and debts, the stated attitude towards financial risk, the experience of being refused credit or of expecting to be refused, the expectation of moving, and the expectation of where interest rates are heading. No other nationally representative American source brings these dimensions together, and the last three are unobservable in administrative data. Evidence on risk attitudes, credit constraints and expectations as determinants of mortgage choice in the United States therefore rests almost entirely on this survey. The study by Coulibaly and Li (2009) is referred to throughout this chapter as the reference study, and the logit model it estimates as the reference specification. The designation reflects the standing of the work in the literature rather than an assessment of its merits relative to other contributions. It constitutes the most comprehensive econometric treatment of the choice between fixed and adjustable contracts conducted on the Survey of Consumer Finances, and a substantial proportion of the estimates discussed in the sections that follow derives from it.

The specification models the probability that household $i$selects an adjustable-rate mortgage as

$$\Pr(\mathrm{ARM}_i = 1 \mid x_i) = \frac{\exp(x'_i\beta)}{1 + \exp(x'_i\beta)} \quad , \quad (8)$$

where the dependent variable takes the value one if the household holds an adjustable-rate contract and zero if it holds a fixed-rate one. The vector $\boldsymbol{x}_i$collects the covariates entering the reference specification, which comprise pricing variables, indicators of affordability and financial stress, measures of expected mobility and stated risk attitude, and a set of demographic controls. The estimated coefficients reported in the sections that follow are the elements of $\boldsymbol{\beta}$. They record the effect of each covariate on the log odds of selecting an

adjustable rather than a fixed contract, and not the marginal effect on the probability itself. Studies drawing on other sources are also included, and they serve three distinct purposes. A first group is included on grounds of precedence, since the econometric analysis of mortgage choice originated in the 1980s with work on single-lender and market-level data, and the results established then, particularly on borrower responsiveness to relative contract prices, remain part of the accepted evidence. Others are included because they observe with greater precision determinants that survey data measure only roughly, contract pricing being the clearest case. A third group is included for the sake of comparison across countries, since the institutional arrangements that shape mortgage markets differ considerably between the United States and Europe, and studies conducted elsewhere help to distinguish findings that reflect household behavior from findings that reflect national market structure. Work of this kind, whether on euro area data or on individual national surveys, will be brought in wherever it sharpens the interpretation of the American evidence.

The criteria governing the selection of studies require explicit statement, since undisclosed selection would compromise the interpretation of the conclusions. Priority was given to work whose dependent variable is the mortgage contract itself, that is to study modelling the choice between fixed and adjustable rates rather than studies in which mortgage type appears incidentally. Both peer-reviewed articles and working papers from central banks and research institutions were considered, the latter because a good deal of the most relevant work in this field circulates in that form and because some of it, including the study that has become the standard reference for the American evidence, appeared as a Federal Reserve working paper before publication. The review is confined to work available in English, which excludes national literatures that may contain relevant findings. No claim to exhaustiveness is made. The field is comparatively small, and the studies discussed include those most frequently cited and most directly pertinent, but the review constitutes a selective survey rather than a systematic census. A delimitation of a different order concerns the object of the review rather than its coverage. The literature examined here treats mortgage choice as a binary decision between a fixed and an adjustable contract, and the present chapter follows that convention, since it is the distinction that the Survey of Consumer Finances records reliably and the one on which the theoretical predictions of chapter 1 are formulated. The convention nonetheless simplifies a considerably richer contractual menu. Sa-Aadu and Sirmans (1995) emphasize precisely this point in their analysis of differentiated mortgage contracts, showing that the terms on

which borrowers and products are matched extend well beyond the form of interest rate. Hybrid contracts fixed for an initial period and adjustable thereafter, introductory rates set below the fully indexed level, and differing provisions on the frequency of resets and on the caps limiting them all bear on the risk a household actually assumes. Two borrowers recorded identically as holders of an adjustable contract may therefore face materially different exposures. The binary framing is retained here for reasons of measurement and comparability, but the results reviewed in the sections that follow should be read as concerning the broad form of the contract rather than the full structure of the obligation. A further restriction concerns the side of the market examined. The studies reviewed here model the choice as one made by the household, taking as given the contracts available to it. Supply is not neutral in this respect, and the composition of the American market is determined in part by the willingness of lenders to originate each type of contract. Fuster and Vickery (2015) establish the point by exploiting a regulatory discontinuity in access to securitization together with the freeze in private securitization liquidity of 2007 and find the fixed-rate share sharply lower where mortgages are difficult to securitize, which they attribute to lenders' reluctance to retain the prepayment and interest rate risk those contracts embed. That literature is not reviewed here, since its object is the behavior of intermediaries rather than of households, but the estimates assembled in this chapter should be read as conditional on a menu whose composition it helps to explain. A final observation concerns the state of the field rather than the scope of this review. No dedicated systematic review of the determinants of mortgage choice appears to have been published, and recent contributions continue to characterize the literature as fragmented across countries and datasets, with only a handful of studies attempting comparison across national contexts. This absence motivates the present review, but it also entails that no established synthesis exists against which the organization adopted here can be validated.

The chapter is organized as follows. Section 3.2 introduces the Survey of Consumer Finances and presents descriptive evidence from published sources on the characteristics of households holding adjustable-rate and fixed-rate mortgages, so that the size and composition of the two groups are established before their determinants are examined. Sections 3.3 to 3.9 consider the determinants in turn, including risk preferences and income risk, liquidity constraints and affordability, mobility and the planning horizon, contract pricing, interest rate expectations, financial literacy and information gaps, and the socioeconomic and demographic characteristics that appear repeatedly in the literature. Section 3.10 brings these findings together and compares them with the theoretical predictions from

which they were derived, in a table reporting for each determinant the sign predicted by theory, the sign prevailing in the estimates, the evidence supporting the relationship, and the principal qualification attaching to it, and it considers what the resulting pattern of agreement and disagreement indicates about the state of the field. The broader implications, concerning where the evidence converges, where it remains divided, and what has yet to be addressed, are taken up in chapter 5.

3.1.2 Search strategy and selection criteria

The procedure by which the literature was identified is set out here so that its coverage can be assessed and, if required, reproduced. Searches were conducted between December 2025 and August 2026. Three types of sources were consulted. Scopus was used for peer-reviewed journals in economics and finance; Google Scholar for forward and backward citation tracking; and the SSRN repository for working papers, which in this field frequently circulate for several years before publication and occasionally contain material absent from the published version. Scopus was preferred to the other multidisciplinary databases because its coverage of the relevant journals overlaps extensively with theirs. SSRN was preferred to the search interfaces of individual series because its Economics Research Network indexes, within a single query, the discussion paper series of the Board of Governors of the Federal Reserve System, those of the individual Federal Reserve banks, and those of several European central banks, which between them account for the unpublished work relevant to this review. Search terms were organized in three blocks, with terms within a block treated as alternatives and blocks combined as joint requirements. The first block identified the object of the decision, through the terms adjustable-rate mortgage, fixed-rate mortgage, mortgage choice, mortgage contract choice and mortgage type. The second identified the determinants examined in this chapter, through household risk management, interest rate expectations, financial literacy, liquidity constraints, mortgage timing and bond risk premium. The third identified the principal data source, through Survey of Consumer Finances and its abbreviation. Searches paired the first block with each of the other two in turn, and were run on title, abstract and keyword fields, with no restriction on year of publication. The repository interface handles queries of this length less reliably, and the searches conducted there were correspondingly simplified, retaining the first block in full and reducing the second to its two most productive terms. Records were retained for full assessment when the mortgage contract was itself the object of analysis, that is when the study modelled or described the choice between

fixed and adjustable rates. Theoretical contributions were retained on a separate criterion, namely that they supply a framework within which the empirical results can be interpreted, and these are discussed in chapter 1 rather than in the present chapter. Peer-reviewed articles, working papers issued by central banks and research institutes, and analytical publications of statistical agencies were all admitted, the last category on the grounds that descriptive tabulations of the Survey of Consumer Finances are frequently published in that form and are not otherwise available. Three exclusion rules were applied. Studies were excluded when mortgage type appeared only as a control variable in an analysis directed at another question; when the object was default, prepayment or refinancing rather than the initial choice of contract; and when the work was not available in English. A criterion of relevance governed the treatment of studies conducted outside the United States. Such work was retained where it established results that the American data cannot yield, where it preceded the survey-based literature historically, or where comparison across countries separates findings that reflect household behavior from those that reflect national market structure, as set out in section 3.1.1. Studies satisfying none of these conditions were not retained even when methodologically sound. Two additions were made outside the database searches. Backward citation tracking from the studies retained identified the loan-level literature of the 1980s, which predates the indexing coverage of the databases consulted. Forward and backward citation tracking from the two cardinal theoretical contributions identified both the framework literature discussed in chapter 1, part of which predates database coverage or falls outside the search terms used here, and more recent empirical work, including studies published after the initial searches were conducted. Table 2 reports the outcome of the procedure for the literature used throughout this work, and not for the empirical review alone.

Table 2. Identification and selection of the literature reviewed

Stage	Records
Records identified through database searching	112
Records identified through repository searching	106
Duplicates removed	53
Records screened on title and abstract	165

Stage	Records
Records excluded at first screening	87
Full texts assessed for eligibility	78
Full texts excluded on eligibility criteria	63
Studies added through citation tracking (+)	19
Total works included	34

Note. The counts refer to scholarly works, of which twenty-five are the empirical studies reviewed in chapter 3 and nine the theoretical contributions discussed in chapter 1. Institutional data sources and regulatory documents are cited throughout but were not identified through the search procedure and are excluded from the table. The theoretical contributions were identified principally through citation tracking, several of them predating the coverage of the databases consulted.

3.2 The Survey of Consumer Finances: structure and descriptive evidence on mortgage holders

This section describes the instrument from which the literature reviewed in this chapter draws most of its evidence and reports the descriptive characteristics of adjustable and fixed rate borrowers documented in published analyses. Both elements are preliminary to the analysis of determinants, since the conclusions of the studies depend on features of the survey design and the descriptive magnitudes establish the scale of the phenomenon under examination.

3.2.1 The instrument

The Survey of Consumer Finances is conducted every three years by the Federal Reserve Board, in cooperation with the Department of the Treasury, and is carried out by NORC at the University of Chicago. It is the principal source of information on the balance sheets of American households, and it has been fielded in its modern form since 1989. Its sampling design departs from standard practice in ways that condition its use. The survey combines a standard geographically based random sample, which secures representativeness of the broad population, with a supplemental sample drawn from statistical records derived from tax data, which deliberately oversamples wealthy households. This second

component exists because wealth in the United States is highly concentrated, and a purely random sample would capture too few affluent families to describe the upper tail of the distribution with any precision. The two samples are combined through a system of weights that restores representativeness, and any descriptive statistics computed from the survey must apply these weights to be meaningful. Two further features of the survey condition the work built on it. To protect the confidentiality of respondents and to deal with the fact that people do not always know or report every figure on their balance sheet, the Federal Reserve releases the data with missing values filled in through multiple imputation, producing five separate copies of the dataset that differ only in the imputed entries. Proper analysis uses all five and combines the results in a way that reflects the uncertainty introduced by imputation. In addition, the reliability of estimates is assessed using replicate weights that account for the complexity of the sampling design. These procedures are not discussed further here, but they account for the methodological requirements that empirical work on the survey must satisfy and for the unreliability of tabulations computed without them. The relevance of the survey to the study of mortgage choice derives less from its sampling design than from the breadth of its content. For a single household observed at a single moment, the survey records the type of mortgage held, whether fixed or adjustable, together with the full composition of the household's assets and liabilities, its income, its demographic characteristics, its self-assessed willingness to take financial risk, its experience of having been denied credit or of having refrained from applying for fear of denial, its expectation of moving house, and its expectation of the future course of interest rates and the economy. No administrative dataset contains the attitudinal and expectational items, and no other nationally representative American survey combines them with balance sheet detail of comparable depth. This combination is the reason the survey has anchored the literature reviewed in this chapter, and it is also the reason that certain determinants of mortgage choice, particularly those involving attitudes and expectations, have been studied in the United States almost exclusively through this instrument. The value of the Survey of Consumer Finances as a source for analyzing mortgage choice ultimately depends on the reliability with which it records the mortgage contract selected by each household. Figure 4 assesses this issue by comparing the share of ARM originations derived from the SCF with the corresponding series published by the Mortgage Bankers Association (MBA). Although both series measure the prevalence of adjustable-rate mortgages, they are constructed differently. The SCF series is based on mortgages originated within the two years preceding each survey interview and therefore reflects completed

mortgage originations. By contrast, the MBA series is based on mortgage applications submitted to lenders and therefore includes applications that do not ultimately result in a completed loan.

Despite these methodological differences, the two series display a remarkably similar cyclical pattern. Both identify peaks in 1994 and 2000 and troughs in 1998 and 2001, and they move in the same direction in most years. The main exceptions occur in 1996, when the SCF series declines while the MBA index increases, and in 1999, when the MBA index rises sharply whereas the SCF series remains broadly unchanged. The relationship between the levels of the two series also changes over time. The SCF reports a higher share of adjustable-rate mortgages until 1995, after which the MBA index becomes consistently larger, with the difference widening from 1999 onwards. Overall, the close correspondence in the timing of cyclical movements provides reassuring evidence that mortgage contract type is measured reliably in the SCF and that the survey captures the main fluctuations in the aggregate mortgage market. The remaining differences in levels are consistent with the different populations represented by the two series, since one measures mortgage applications and the other completed originations. They may also partly reflect the sampling variability associated with the relatively small number of recent mortgage originations observed in each SCF wave.

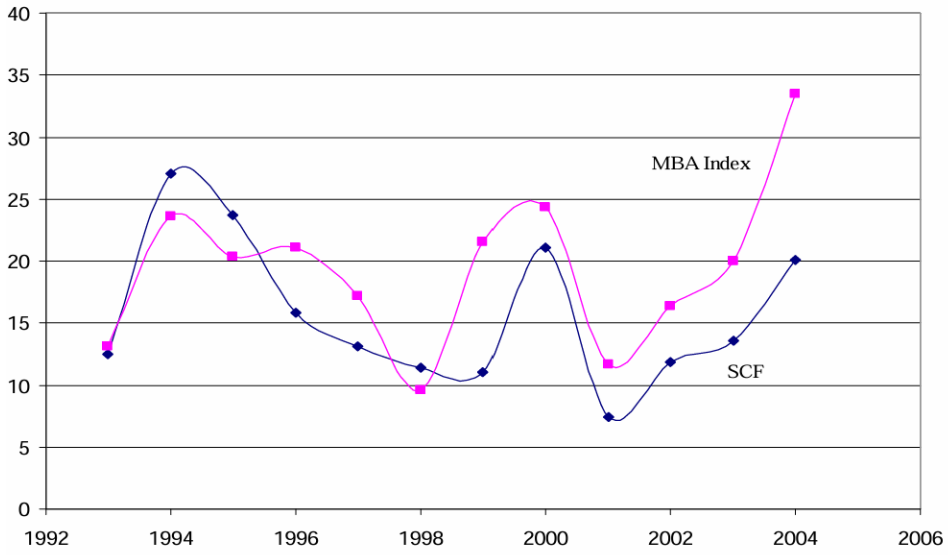


Figure 4.

ARM share in the Survey of Consumer Finances and in the Mortgage Bankers Association index, 1993 to 2004.

Source: Coulibaly and Li (2009), figure 1.

3.2.2 Descriptive evidence on adjustable and fixed rate borrowers

Published analyzes of recent waves yield a consistent descriptive picture. The most detailed of these, produced by Chiang and Dueholm (2024) using the 2019 wave, document that the adjustable-rate mortgage is very much a minority product. Adjustable contracts represented a small share of outstanding mortgages, and the households that held them were younger, earned higher incomes, and carried larger initial loan balances and larger outstanding balances than the households that held fixed-rate contracts. The concentration of adjustable borrowing among higher earners is pronounced and does not move in a simple straight line across the income distribution. In the lowest income decile, about 6.5 percent of mortgaged households held an adjustable contract; this share falls to around 3.9 percent in the fifth decile before rising steadily to 18.8 percent in the highest decile. The overall pattern is one in which adjustable borrowing is most common at the top of the income distribution and least common in the middle, with a modest elevation at the very bottom. This gradient is relevant to the interpretation of the income and wealth determinants examined later in the chapter. The observation that adjustable borrowers are on average younger is consistent with the mobility channel, since younger households are more likely to expect to move, while the observation that they earn more and borrow more

speaks to the buffer stock and affordability mechanisms that recur throughout the literature. These descriptive associations do not establish causal relationships, but they provide an empirical reference point for the analysis that follows. One methodological caution accompanies these figures. The authors Chiang and Dueholm (2024) deliberately based their tabulations on the 2019 wave rather than on the more recent 2022 wave, on the grounds that the 2019 data are not entangled with the disruptions of the pandemic period. This choice is defensible for their purposes, but it means that the most careful published descriptive breakdown of adjustable and fixed borrowers predates the rate environment that is of central interest to the present work.

Tabulations cross-classifying contract type against the attitudinal and expectational variables are available for earlier waves, principally through the descriptive sections of the studies reviewed in this chapter, but none appears to have been published for the 2022 wave. The gap is more striking because the relevant items are not missing from the survey: households' expectations about the future direction of interest rates have been elicited since 1989, yet no study has brought them to bear on the 2022 data. The wave conducted during the sharpest tightening in four decades has served for broad descriptive portraits of family finances, but not, so far, for a focused analysis of how contract choice relates to the attitudes and expectations that the survey uniquely records. This absence is taken up in chapter 5.

3.2.3 Summary of the descriptive evidence

The table below collects the principal descriptive statistics discussed in this section. It reports only characteristics of the population of mortgage holders, drawn from published tabulations, and indicates for each figure the wave to which it refers and the source from which it is taken. It contains nothing about the results of the econometric studies reviewed later in the chapter, which are summarized separately in the comparative table of section 3.10. The two tables serve distinct functions, the former documenting the characteristics of contract holders and the latter summarizing the findings of the econometric literature.

Table 3. Determinants of mortgage choice: measurement and descriptive evidence in the Survey of Consumer Finances

Determinant	Measure in the SCF	Published descriptive evidence	SCF wave	Source
Contract type (dependent variable)	Whether the rate can change over the life of the loan	ARM held by 14.6 percent of 2,887 mortgage-holding households	1995–2004	Coulibaly and Li (2009)
Contract type	Whether the rate can change over the life of the loan	ARM share of originations: 17.5 percent in 1990, 23.9 percent in 2005, 11.9 percent in 2007	1992–2007	Johnson and Li (2014)
Contract type	Whether the rate can change over the life of the loan	ARM held by about 8 percent of mortgaged households	2019	Chiang and Dueholm (2024)
Risk aversion (row 1)	Four-category item on willingness to take financial risk, fielded since 1983; eleven-point scale added in 2016	40.8 percent unwilling to take any risk, 38.6 percent average, 16.4 percent above average, 4.2 percent substantial, all families	2016	Kim, Hanna and Ying (2021)
Risk aversion (row 1)	Four-category item on willingness to take financial risk, fielded since 1983; eleven-point scale added in 2016	69.6 percent of mortgage holders classified as risk-averse	1995–2004	Coulibaly and Li (2009)
Credit constraints (row 2)	Indicators of financial stress and of the terms on which credit was obtained	Overspending reported by 12.4 percent; subprime mortgages held by 11.5 percent; median non-housing debt to income ratio 12.4	1995–2004	Coulibaly and Li (2009)
Expected mobility (row 3)	Self-assessed probability, on a scale from 0 to 100, of still living at the current address two years hence; rescaled by the authors into a probability of moving. Available from 1995	Mean probability of moving within two years 15.2 percent	1995–2004	Coulibaly and Li (2009)

Determinant	Measure in the SCF	Published descriptive evidence	SCF wave	Source
Interest rate expectations (row 4)	Whether interest rates in five years' time are expected to be higher, lower or about the same as today (item X302), with a parallel item on expected economic conditions	Responses tabulated by cohort from 1989 to 2013, showing systematic divergence between older and younger households; the item does not enter as a covariate in the contract-choice specifications reviewed here	1989-2013	Botsch and Malmendier (2023)
Household income (row 5)	Total family income in the year preceding the interview	Median 70,300 dollars, mean 141,900 dollars, all families	2007-2022	Aladangady et al. (2023)
Household income (row 5)	Total family income in the year preceding the interview	ARM share by income decile: 6.5 percent lowest, 3.9 percent fifth, 18.8 percent highest	2019	Chiang and Dueholm (2024)
Net worth (row 6)	Total assets minus total liabilities	Median net worth up 37 percent from 2019 to 2022 for homeowners	2007-2022	Aladangady et al. (2023)
Age (row 7)	Age of the reference person	Mean 43.5 years among mortgage holders	1995–2004	Coulibaly and Li (2009)
Age (row 7)	Age of the reference person	Median 32.1 years among ARM holders against 50.3 among FRM holders	2019	Chiang and Dueholm (2024)
Education and financial sophistication (row 8)	Educational attainment; no dedicated literacy module	College educated 48.5 percent of mortgage holders	1995–2004	Coulibaly and Li (2009)
Contract term (row 9)	Indicator of whether the mortgage matures in fifteen years or less; term in years also recorded	Mean term 24.1 years among mortgage holders	1995–2004	Coulibaly and Li (2009)

Determinant	Measure in the SCF	Published descriptive evidence	SCF wave	Source
Income volatility (row 11)	Not surveyed directly. Standard deviation of the residuals of a Mincer income equation estimated by occupation, expressed relative to household real income	Mean value of the index 92.8	1995–2004	Coulibaly and Li (2009)

Note. The spread between fixed and adjustable rates, recorded as row 10 of Table 1, is a market-level quantity identical across households observed in the same period and therefore has no descriptive counterpart at household level; the corresponding series appear in chapter 2. The wave column records the release from which each figure is drawn. Entries from different waves rest on different samples and definitions and are not directly comparable. The 2022 wave is the reference throughout this work, and figures from earlier waves appear where no published tabulation of the 2022 data exists for the variable concerned. No published tabulation cross-classifies contract type against risk attitudes, credit constraints or expectations for the 2022 wave, a gap discussed in chapter 5.

3.3 Risk preferences and income risk

Two of the determinants listed in Table 1 concern the household's exposure to risk: the coefficient of risk aversion recorded in row 1 and the volatility of labor income recorded in row 11. They are examined jointly here because they operate through a common mechanism. The lifecycle framework treats the fixed-rate contract as an insurance instrument whose premium is the term spread embedded in its rate, and the value of that insurance depends both on how strongly a household dislikes variability in its mortgage payment and on how likely it is that an unfavorable payment realization will coincide with an unfavorable income realization. Risk aversion governs the first quantity, income volatility the second. Both are predicted to reduce the probability of selecting an adjustable contract.

3.3.1 Evidence on stated risk aversion and income volatility

The reference result for the United States comes from Coulibaly and Li (2009), whose logit estimates on pooled waves of the Survey of Consumer Finances indicate that attitudes toward financial risk substantially influence the choice of contract, with more risk-averse borrowers favoring fixed-rate mortgages. The sign accords with the theoretical prediction, and the estimate is obtained while controlling for pricing variables, affordability and mobility, so the association does not appear to be a proxy for those channels. A

secondary finding of the same study qualifies the picture in a way that has received less attention than it deserves. When the sample is partitioned by stated risk attitude, the structure of the decision differs across groups: among less risk-averse borrowers, the choice responds less to pricing variables and to income volatility, and affordability measures lose statistical significance altogether. The lifecycle model does not predict this pattern, since it treats the welfare cost of payment variability as increasing smoothly in risk aversion rather than as governing whether other determinants operate at all. The result suggests that risk aversion may function less as a continuous parameter shifting the probability of an adjustable contract and more as a condition determining which considerations enter the decision in the first place. Evidence from outside the United States requires careful reading. Ehrmann and Ziegelmeyer (2017), analyzing euro area households through the Euro system Household Finance and Consumption Survey, report general support for the predictions of Campbell and Cocco (2003), but the support concerns the income risk channel rather than stated risk attitude. Their survey contains a direct question on willingness to take financial risk, which they exclude from the benchmark specification because it is unavailable for France and introduce in a robustness test covering the remaining countries. The estimated coefficient is close to zero and statistically insignificant. The euro area evidence therefore corroborates the framework on the objective dimension of risk while failing to reproduce the effect of the stated attitude, which places it alongside the Italian result rather than alongside the American one on this particular determinant. The convergence is not universal, however, and one dissenting result merits attention. Paiella and Pozzolo (2007) examine Italian household data from the Bank of Italy’s Survey of Household Income and Wealth and find limited evidence that household characteristics related to non-mortgage risk exposure or risk preferences systematically affect mortgage contract choice. Their results instead point to the relative cost of the two mortgage types and the presence of liquidity constraints as the main factors associated with the decision. Their conclusion runs directly against the prediction of row 1 and against the American evidence. Whether the divergence reflects a genuine difference in behavior, a difference in the institutional setting, or a difference in how risk aversion is measured across the two surveys cannot be settled from the published results, and the question is revisited in chapter 5. The negative association between income volatility and adjustable-rate borrowing is among the more stable results in the literature. Coulibaly and Li (2009) obtain the predicted negative coefficient using a proxy that they construct rather than observe, since the survey does not elicit perceived income risk. They estimate a Mincer income equation by

occupation and take the dispersion of its residuals, expressed relative to household real income, as an approximation of the riskiness of future earnings. The sign accords with the proposition that households facing uncertain earnings place a higher value on payment certainty. The euro area evidence reproduces it on a different instrument and a different population, and no published study appears to reverse it. Row 11 therefore stands as one of the better corroborated entries in Table 1, although the corroboration concerns a constructed measure rather than a directly observed one, a qualification developed below.

3.3.2 Measurement and its limits

The empirical support for both determinants rests on a narrow measurement base, and the limitations warrant explicit statement. The risk attitude variable used in almost all American work is the item that has appeared in the Survey of Consumer Finances since 1983, which asks respondents to select among four statements describing their willingness to take substantial, above average, average, or no financial risk. Three difficulties follow from its construction. The question is framed in terms of saving and investment decisions rather than borrowing, and there is no assurance that willingness to accept variance in an investment portfolio corresponds to willingness to accept variance in a mortgage payment, since risk attitudes are known to vary by domain. The distribution of responses is also heavily skewed, with only around four percent of respondents selecting the highest category, which leads most analysts to collapse the four levels into two or three and thereby to discard variation and to make the estimated coefficient dependent on the collapsing rule adopted. Evidence on the reliability of the item, moreover, remains limited relative to the volume of research that relies on it. From the 2016 wave onward the survey has fielded an additional eleven-point scale alongside the original question, which improves granularity but complicates comparison with the earlier waves on which the mortgage choice literature was built. The three difficulties just noted concern how the item is built. A fourth concerns what it measures. Cenzon and Szabó (2024) find that households who experienced high and volatile inflation were less inclined towards accepting financial risk, which indicates that the stated attitude is shaped in part by the macroeconomic environment the respondent has lived through. If risk tolerance is itself in part a function of experience, then a specification containing the attitude but no measure of that experience attributes to preference what may originate in circumstance, and the coefficient on risk aversion absorbs part of the channel examined in section 4.2.

The income volatility measure raises a difficulty of a different order. The survey contains

no question on perceived income risk, so the variable used in the literature is imputed rather than reported: it derives from the dispersion of residuals in an earnings equation estimated by occupation. Every household in a given occupation therefore receives the same value, irrespective of whether its own earnings have been stable or erratic. What the coefficient captures is the association between contract choice and the risk typical of an occupational category, which is informative but is not the household-level income uncertainty that the theory specifies.

A final consideration bears on the interval between the moment of choice and the moment of observation. Household characteristics in the survey are recorded at the date of the interview, whereas the contract was selected at origination, and in a cross-section the two may lie years apart. Coulibaly and Li (2009) address this directly by restricting their sample to mortgages originated within two years of the interview, on the reasoning that characteristics reported at the interview will then approximate those prevailing when the contract was signed. The restriction reduces the sample from 3,636 recent mortgages to 2,887 after further exclusions, and it removes the concern for the reference study at the cost of confining the analysis to a subset of borrowers who had transacted recently. Caution applies more forcefully to work that does not impose an equivalent restriction, and to determinants whose value is revised frequently, a point developed in section 3.7.

3.4 Liquidity constraints and affordability

Row 2 of Table 1 records the determinant on which theory and evidence diverge most sharply. The lifecycle framework treats the borrowing constraint as the mechanism that gives payment volatility its welfare cost: a household unable to smooth consumption by drawing on wealth or credit suffers most when its mortgage payment rises and should therefore be the most willing to pay the term premium embedded in the fixed rate. The prediction is that financially constrained households select fixed-rate contracts. The empirical literature reports the opposite.

3.4.1 The empirical reversal

Coulibaly and Li (2009) find that affordability and financial stress play an important role in the decision for the general population, though they qualify this by noting these factors are not significant for less risk-averse households. Their logit estimates give the result substance: households that reported spending more than they earned in the previous year are 27 percent more likely to hold an adjustable contract, and the ratio of the mortgage

balance to income enters positive and significant, which the authors read as evidence of affordability concerns. The ratio of non-housing debt to income is positive but not significant. A further indication points the same way: mortgages of fifteen years or less are markedly more likely to be fixed rate, and the authors observe that this admits an affordability reading, since a household able to discharge the debt over a shorter term is less exposed to liquidity pressure and correspondingly readier to accept the fixed contract.

The largest coefficient in the specification, by a wide margin, attaches to an indicator of subprime status, whose interpretation is taken up below. The corroborating evidence comes from Johnson and Li (2014), who address the question directly and combine the Survey of Consumer Finances with the Consumer Expenditure Survey in order to observe both balance sheets and spending behavior. Their conclusion is that adjustable-rate borrowers do indeed appear to be borrowing constrained. What gives their result its interest is the form the constraint takes. The demographic and financial characteristics of the two groups of borrowers are broadly similar, and the differences lie instead in how they use credit and in the attitudes they express toward borrowing. The constraint, in other words, is not visible in the standard balance sheet aggregates but in behavior, which suggests that measures based on income and assets alone will understate it.

Evidence from outside the United States points the same way. Paiella and Pozzolo (2007), working with the Bank of Italy household survey, find that the choice depends critically on the relative price of the two contracts and on whether the household is liquidity constrained, with constrained households drawn to the adjustable contract by its lower initial installment. Their result is the more striking because in the same estimation the proxies for risk aversion prove irrelevant, so that in the Italian data the constraint channel operates while the risk channel does not.

3.4.2 Explaining the contrasting evidence

Two lines of interpretation have been offered, and they are not mutually exclusive. The simpler one turns on the horizon over which the constraint binds. The lifecycle model asks which contract leaves the household better off across the life of the loan, and answers that the constrained household should buy insurance. A household facing a binding affordability constraint at origination is answering a different question, namely which contract it can afford to service in the coming months. The adjustable contract offers the lower initial payment, and for a household near the limit of what it can qualify for, that difference may determine whether the purchase is feasible at all. The theoretical

prediction and the observed behavior are therefore not strictly contradictory: they concern different decision problems, and the empirical result indicates which of the two dominates when they conflict. The second line qualifies the picture more substantially. Smith, Finke and Huston (2011) show that adjustable-rate borrowing in the survey is bimodal with respect to financial sophistication, drawing in both the least and the most sophisticated households but for opposite reasons. The less sophisticated are attracted by the low initial payment under conditions of income constraint; the more sophisticated respond to a wide spread between the two contract rates. A single coefficient on a constraint indicator therefore aggregates two distinct populations, and its magnitude depends on their relative weight in the sample. The implication is that the constraint coefficient is not a stable structural parameter but a composite whose value may shift with the composition of adjustable-rate borrowers, which is itself a function of the rate environment. Section 3.8 returns to the sophistication dimension, and chapter 5 takes up the consequences for interpreting estimates obtained under different regimes.

3.4.3 How constraints are measured

The evidence rests on proxies whose construction deserves scrutiny. The measures entering the reference the specification estimated by Coulibaly and Li (2009) are indirect: an indicator of reported overspending, an indicator of whether the mortgage is subprime, and ratios of mortgage and non-mortgage debt to income. Each captures something related to financial pressure, but none identifies a constraint in the sense the theory intends, that is an inability to borrow at the prevailing terms. The subprime indicator raises the sharpest difficulty, and it is the variable carrying the largest estimated effect, with an odds ratio of 15.51. The survey does not record whether a mortgage is subprime, and the authors approximate it by classifying as subprime any loan whose interest rate exceeds that on a Treasury security of comparable maturity by more than three percentage points, a criterion that identifies 11.5 percent of the sample. Two features of this construction bear on how the coefficient should be read. The variable is defined on the contract rate, which is itself the outcome of the choice being explained, so a loan carrying a high rate is classified as subprime and adjustable contracts carried higher initial rates in part by design. The authors also note that more than two thirds of subprime mortgages are issued as adjustable contracts, which they cite as the reason for including the control. The coefficient therefore records to a considerable extent the composition of the subprime segment rather than a behavioral response by constrained households, and its magnitude should not be

read as the effect of financial constraint on contract choice. A further observation concerns what the survey offers and the literature does not use. The Survey of Consumer Finances asks directly whether a household has been refused credit and whether it has refrained from applying in the expectation of refusal, and these items are the standard instruments for identifying constrained households in the wider consumer finance literature. They do not appear among the covariates of the studies reviewed here, which rely instead on the indirect measures described above. The pattern parallels the treatment of expectations discussed in section 3.7: the survey records the variable that theory calls for, and the empirical literature works with something else.

3.4.4 Assessment

Row 2 stands as the clearest case in which the empirical record contradicts the normative prediction, and the contradiction is robust across datasets, countries and specifications. The interpretation of that record is nonetheless less settled than the sign alone suggests. The constraint appears to operate through a short-horizon affordability channel rather than the long-horizon insurance channel that the model specifies; its estimated magnitude rests partly on a proxy that conflates constraint with contract type; and the underlying population is heterogeneous in ways that a single coefficient conceals. The finding that constrained households choose adjustable contracts can be regarded as established. What remains unresolved is the reason. One consideration bearing on the reason concerns the environment in which the constraint is measured. The proxies employed in this literature identify households under financial pressure, but the economic content of that pressure is not invariant across rate regimes. Under conditions of elevated housing wealth and rising borrowing costs, of the kind documented in chapter 2 for the early 2020s, a household may hold substantial equity in its property while facing difficulty in meeting a monthly installment that has risen sharply. Such a household is constrained in cash flow while being comfortably placed in assets, a configuration that the standard measures do not distinguish from constraint arising out of low wealth. The distinction matters for the interpretation of row 2, since the two situations imply different motives for selecting the adjustable contract, the one reflecting an inability to sustain the higher installment of the fixed alternative and the other a temporary mismatch between illiquid wealth and current obligations. It also indicates why a coefficient estimated under one configuration of house prices and rates need not carry the same meaning under another, a point developed in section 3.10 and revisited in chapter 5.

3.5 Mobility and planning horizon

Rows 3 and 9 of Table 1 concern the length of time over which the household expects to hold the contract. The mechanism is the most transparent in the whole framework. The fixed-rate mortgage insures against movements in interest rates over the life of the loan, and the term premium is the price of that insurance. A household that expects to sell the property and discharge the mortgage within a few years will never collect on the policy it has paid for, and should therefore prefer the adjustable contract, whose lower initial payments accrue precisely over the period it intends to remain. Expected mobility should raise the probability of choosing an adjustable rate; a longer intended horizon should lower it.

3.5.1 The evidence

The determinant is among the best supported in the literature, and it is one of the few for which the Survey of Consumer Finances provides a direct measure rather than a proxy. Respondents are asked to state, on a scale from zero to one hundred, the chance that they will still be living at their current address two years hence. Coulibaly and Li (2009) rescale this into a probability of moving and obtain a mean of 15.2 percent among mortgage holders. The item was introduced in 1995, and its absence from earlier releases is the reason their analysis is confined to four waves. Their logit estimates return a positive coefficient for mobility, which is statistically significant in the pooled sample. Specifically, a one standard deviation increases in the probability of moving increases the likelihood of choosing an ARM by 10 percent (odds ratio of 1.10). This effect, while significant, is relatively small compared to other factors like the mortgage spread or subprime status. When the sample is split by stated risk attitude, mobility retains its significance among less risk-averse borrowers but loses it among the more risk-averse. The mobility channel therefore appears to operate selectively rather than as a general feature of the decision, which is consistent with the broader pattern noted in section 3.3, where risk attitude conditions whether other determinants enter the choice at all. Indirect corroboration comes from the age composition of the two borrower groups. Chiang and Dueholm (2024) report a median age of 32.1 years among adjustable-rate holders against 50.3 among fixed-rate holders, a gap of eighteen years. Since residential mobility is concentrated in the early stages of the life cycle, the age gradient is what the mobility channel would predict, although age also proxies for wealth, income trajectory and experience, so the inference is

suggestive rather than conclusive. The determinant is examined in its own right in section 3.9. The earlier loan-level literature reached compatible conclusions with a different instrument. Dhillon, Shilling and Sirmans (1987), working with data from a single lender, used length of tenure at the current address as their proxy for mobility prospects. The construct is not the same: tenure to date is a realized quantity looking backwards, whereas the survey item is a subjective expectation looking forward. Those two measures of this kind should point in the same direction lends the finding a robustness that a single operationalization would not provide.

3.5.2 Measurement

Three difficulties qualify the evidence, and the first is specific to this determinant. The horizon of the survey item does not match the horizon of the theory. Respondents are asked about the next two years, whereas the insurance value of the fixed-rate contract accrues over a term of twenty or thirty years. A household expecting to remain in place for two years may be planning to stay for thirty, and the item cannot distinguish between the two. The measure identifies imminent movers with reasonable accuracy but says little about the planning horizon that the model specifies, and the modest size of the estimated coefficient may partly reflect this mismatch rather than a weak underlying relationship. Stated intentions are also an imperfect guide to behavior. The variable records what households expect at the moment of interview, and expectations of moving are revised frequently in response to employment, family circumstances and the housing market itself. No study in this literature validates the item against realized moves, though the two-year window that Coulibaly and Li (2009) impose on the age of the mortgage leaves limited scope for intentions to be revised between origination and interview.

3.5.3 The contract term

Row 9 is estimated, though indirectly. The reference specification includes a dummy for mortgages maturing in fifteen years or less, which enters negative and highly significant: short-term loans are considerably more likely to be fixed rate, and by implication longer terms are associated with adjustable contracts, as the prediction requires. The descriptive statistics record a mean term of 24.1 years among mortgage holders, with almost no difference between risk-attitude groups. The authors offer two readings of the coefficient. Many lenders extend adjustable contracts only on thirty-year terms while offering shorter maturities on fixed-rate loans, so the association may record the structure of supply rather

than the preferences of borrowers. This reading finds independent support in Fuster and Vickery (2015), who show that the availability of fixed-rate contracts responds to conditions in the securitization market that are unrelated to what borrowers want. Alternatively, a household able to discharge its debt over fifteen years faces less liquidity pressure and is correspondingly readier to accept the fixed contract, which places the result in the affordability channel of section 3.4. The two readings are not separable in these data, and the row should be counted as corroborated in sign but ambiguous in mechanism.

3.5.4 Assessment

Expected mobility is one of the determinants on which theory and evidence agree, and the agreement holds across survey and loan-level data and across two different operationalization. The qualification is that the agreement concerns the sign rather than the magnitude. The measure available in the survey addresses a horizon far shorter than the one the theory concerns, it records intentions that are known to be revisable, and it is recorded at a date that may lie years after the decision it is used to explain. Each of these considerations pushes the estimated coefficient toward zero, which makes the finding of a positive and significant effect more rather than less credible, but it also means that the literature offers little basis for judging how large the effect actually is.

3.6 Contract pricing and the ARM-FRM spread

Row 10 of Table 1 records the determinant with the strongest empirical support in the literature. The mechanism requires no behavioral assumption beyond a preference for paying less: when the initial rate on an adjustable contract falls well below the fixed rate available at the same moment, the immediate saving grows, and the adjustable product attracts a larger share of borrowers. What makes the determinant theoretically interesting is not the direction of the response, which is scarcely in doubt, but the question of what borrowers are responding to.

3.6.1 The theoretical refinement

The observed difference between fixed and adjustable rates confounds two components. Part of it reflects the market's expectation that short rates will rise or fall over the life of the loan; the remainder is compensation to the lender for bearing duration risk, that is the term premium introduced in section 1.2.

Koijen, Van Hemert, and Van Nieuwerburgh (2009) argue that only the second component

should drive the choice. A spread arising purely from expected rate increases offers no genuine saving, since the adjustable payment will rise as those increases materialize, whereas a spread arising from a high term premium represents a real transfer that the borrower forgoes by choosing the fixed contract. Their empirical work supports the distinction: variation in the adjustable share tracks a measure of the bond risk premium more closely than it tracks the raw differential between quoted rates. The distinction matters for reading the evidence that follows, because the empirical literature generally works with the observed differential rather than with the decomposed premium. Estimated coefficients on the spread therefore capture the two components jointly, and a positive coefficient is consistent with rational timing of the premium but also with a simpler responsiveness to the initial payment.

3.6.2 The evidence

Vickery (2007) provides complementary evidence using administrative loan-level data. His analysis shows that borrowers are highly responsive to differences in the relative prices of mortgage contracts and that the aggregate share of adjustable-rate mortgages closely follows changes in the interest rate spread between adjustable-rate and fixed-rate mortgages. Because the study relies on administrative records, in which mortgage contract characteristics are directly observed rather than self-reported, its findings are not affected by the measurement issues that may arise in survey-based data. The survey-based evidence leads to similar conclusions. Coulibaly and Li (2009) find that the FRM–ARM interest rate spread has a positive and statistically significant effect on mortgage contract choice at the 5 percent significance level. They also identify the spread as one of the most important pricing variables influencing borrowers' decisions. Their results further reveal an interesting source of heterogeneity. The estimated coefficient on the spread is larger for more risk-averse households (0.80) than for less risk-averse households (0.40), indicating that borrowers who generally prefer fixed-rate mortgages are also more responsive to changes in their relative cost. One possible explanation is that highly risk-averse households evaluate the trade-off between payment stability and borrowing costs more carefully. In contrast, households that are less concerned about payment volatility may base their mortgage choice on factors other than relative pricing. The result predates the survey literature. Dhillon, Shilling and Sirmans (1987), working with single-lender data in the 1980s, identified the rate differential as a principal determinant of the choice, and Brueckner and Follain (1988) documented the same relationship at market level in accounting

for the rise and subsequent decline of adjustable lending during that decade. That a finding first established on lender records in the 1980s should survive re-estimation on national survey data two decades later, across a period spanning several distinct rate regimes, places it on firmer ground than any other entry in Table 1. A qualification comes from the sophistication literature. Smith, Finke and Huston (2011) report that responsiveness to the spread is concentrated among more financially sophisticated households, which exploit a wide differential when one is available, while less sophisticated households enter adjustable contracts for reasons connected to affordability rather than to relative price. The two heterogeneity results are not in conflict, since sophistication and risk aversion are distinct attributes, but taken together they indicate that the aggregate spread coefficient averages over populations responding to quite different considerations. A related result concerns the circumstances in which the contract is chosen. The specification estimated by Coulibaly and Li (2009) distinguishes households refinancing an existing mortgage from those originating a new one, and refinancing enters with a negative and significant coefficient: households that refinance are more likely to move to a fixed rate. The authors attribute this to the motive for refinancing itself, which is commonly to lock in an attractive rate before it rises. The finding is worth recording because the samples used throughout this literature contain both purchasers and refinancers, a distinction that is otherwise left implicit. It also indicates that the pricing channel operates differently at the two moments: a purchaser weighs the differential against the payment it can sustain, whereas a refinancer is acting on a judgment about the level of rates and where they are heading, which brings the decision closer to the expectational channel of section 3.7 than to the affordability considerations of section 3.4.

3.6.3 Measurement

Because this variable is determined outside the household, it is not subject to the reporting errors that may affect the attitudinal measures discussed previously. However, its use introduces a different set of empirical challenges. Assigning a spread to a household requires knowing when the contract was originated, since the relevant price is the one prevailing at that date rather than at the interview. Where origination year is recorded the assignment is feasible, though it introduces error whenever rates moved sharply within a year, as they did in 2022. The construction of the differential is also less straightforward than it appears. Quoted adjustable rates refer to a specific product, and the composition of adjustable lending has changed considerably, from the one-year adjustable contract that

dominated earlier decades to the hybrid structures that prevail today. A differential computed against a one-year adjustable rate and one computed against a five-year hybrid measures different things. The problem has become more acute since Freddie Mac discontinued publication of its adjustable-rate series at the end of 2022, so that no continuous quoted series now spans the period in which the fixed-rate series is available. One feature of the reference specification warrants attention. Alongside the spread, Coulibaly and Li (2009) include a yield curve variable, which has a negative and statistically significant coefficient. The interpretation of this result depends on the precise definition of the variable. Under a conventional interpretation, a steeper yield curve would generally be associated with a larger difference between fixed-rate and adjustable-rate mortgage rates, thereby increasing the attractiveness of adjustable-rate mortgages. The estimated negative coefficient therefore suggests that the yield curve may capture a different dimension of the interest rate environment than the mortgage rate spread. One possible interpretation is that the spread reflects the immediate cost advantage of adjustable-rate mortgages, whereas the yield curve captures expectations about future interest rate movements. This interpretation is broadly consistent with the decomposition proposed by Koijen, Van Hemert, and Van Nieuwerburgh (2009). However, because Coulibaly and Li (2009) do not explicitly provide this interpretation or discuss the exact construction of the variable, this explanation should be regarded as a plausible hypothesis rather than an established conclusion.

3.6.4 The contract term

Row 9, the maturity of the loan, belongs to the same pricing channel, since a longer term concentrates more duration risk in the fixed rate and widens the differential the borrower must pay. The evidence discussed in section 3.5 is consistent with this: shorter maturities are markedly more likely to carry a fixed rate. The interpretation is nonetheless complicated by the supply side, since adjustable contracts have generally been offered on thirty-year terms alone, and the evidence of Fuster and Vickery (2015) indicates that the availability of each contract type responds to conditions in the securitization market independently of borrower demand. The observed association may therefore reflect the menu available rather than a response to duration pricing.

3.6.5 Assessment

The spread is the one determinant on which theory, survey evidence and administrative evidence converge without qualification as to sign. Two reservations attach to the interpretation rather than to the finding. What the estimated coefficients measure is the responsiveness to the observed differential, not to the term premium that theory identifies as the relevant quantity, and no study reviewed here decomposes the two on household data. The response is moreover not uniform, varying with risk attitude and with financial sophistication in ways that suggest different segments of the market are reacting to different features of the same price. The proposition that borrowers respond to relative contract prices can be regarded as settled; the proposition that they do so in the manner the timing framework describes remains supported at market level and untested at household level.

3.7 Interest rate expectations

Row 4 of Table 1 is the determinant that occupies a distinctive position in the existing literature. It is theoretically important, empirically supported when examined at the level of aggregate mortgage markets, but still largely unexplored when considering individual household decisions in the United States. The underlying mechanism is straightforward. A borrower who expects interest rates to increase has an incentive to choose a fixed-rate mortgage in order to lock in current payments before future increases occur. Conversely, a borrower who anticipates declining rates may prefer an adjustable-rate mortgage, allowing future payments to adjust downward. These expectations influence mortgage choice independently of the current interest rate differential because households facing the same mortgage prices may make different decisions if they hold different views about future rate movements. As discussed in section 1.3, the strongest theoretical formulation of this mechanism assumes that households are able to distinguish between the component of the mortgage spread associated with expected future interest rate changes and the component reflecting compensation for bearing duration risk, responding only to the latter.

3.7.1 Evidence from national markets

The most important contribution in this area is provided by Badarinza, Campbell and Ramadorai (2018), who construct a nine-country panel dataset combining mortgage origination volumes and contract interest rates to examine how current and expected future interest rates influence the share of adjustable-rate mortgages among new loans. To address the potential endogeneity of expectations, the authors employ instrumental variable

techniques and evaluate the results under alternative assumptions about expectation formation, including rational expectations, adaptive expectations, and survey-based measures of household expectations. Three main findings emerge, and their joint interpretation provides a more complete understanding of household mortgage choice. First, short-term expectations play an important role: expected changes in adjustable mortgage rates over the following year significantly affect the composition of new mortgage lending, with this relationship becoming stronger in the post-2001 period. By contrast, expectations over longer horizons appear to have a more limited effect, as three-year rate forecasts have only a weak influence on mortgage contract choice. At the same time, the current mortgage rate spread remains an important explanatory factor even after controlling for both short- and long-term expectations. The authors interpret this pattern as evidence that households consider both immediate borrowing costs and longer-term mortgage affordability, although their decisions appear to place greater weight on near-term developments. This interpretation is consistent with the view that households do not fully optimize over the entire life of a long-term mortgage contract. A borrower following a strict lifecycle optimization framework would be expected to assign at least as much importance to longer-horizon interest rate expectations as to short-term forecasts. The stronger influence of one-year expectations therefore suggests that mortgage decisions are more closely tied to present and near-future financial conditions than predicted by a fully forward-looking framework. This result is consistent with the affordability mechanism discussed in section 3.4 and is further reflected in the interpretation developed in chapter 5.

3.7.2 Evidence at household level

Direct household evidence is available, although not for the United States. Devine, McCarthy and O'Toole (2023) conducted a specifically designed survey among approximately two thousand Irish mortgage holders, collecting information on households' expectations about future interest rates and estimating the choice between fixed-rate and adjustable-rate mortgages using a probit model. Their results show that expectations represent a major determinant of mortgage choice, operating alongside current price conditions rather than simply reflecting them. The study provides evidence that the expectations channel can be directly identified at the household level when survey instruments are specifically designed to capture it. Evidence of a different kind comes from a considerably larger sample. Bacon and Moffatt (2012) model the choice between variable and fixed

rates for approximately 280,000 British borrowers observed between 1992 and 2001 in the Survey of Mortgage Lenders, estimating an ordered probit in which expectations about future rates enter alongside pricing variables and borrower characteristics. Fixing proves more likely when rates are expected to rise, and demand for fixed contracts falls when the level of rates is high, which the authors read as reflecting an expectation of subsequent decline. The expectations entering their model are not reported by borrowers but constructed by the authors from market data, since the underlying source is a survey of lenders rather than of households. The finding therefore corroborates the channel on a sample an order of magnitude larger than the Irish one, while leaving untouched the question of whether households act on the expectations they themselves report. Ilan and Mugerman (2025) examine the relationship from a complementary perspective by analyzing how expectations interact with financial competence in the choice between indexed and non-indexed mortgage contracts. Their findings indicate that households with lower levels of financial literacy tend to rely more heavily on currently observed inflation rather than on expectations of future inflation. This result suggests that the expectation and financial literacy channels discussed in section 3.8 are closely interconnected: the formation of expectations depends not only on the information available to households but also on their ability to process and interpret that information. Chapter 4 develops this argument further by examining the role of past experiences in shaping expectation formation.

3.7.3 The gap in the American literature

The Survey of Consumer Finances collects information on households' expectations regarding future interest rates by asking respondents whether they expect rates five years ahead to be higher, lower, or approximately unchanged relative to current levels. It also includes related questions about expectations for overall economic conditions over one-year and five-year horizons. These variables are therefore available within the same nationally representative dataset that records mortgage contract characteristics, household balance sheets, and measures of risk attitudes. The SCF item in question has been used, but not in the manner the theory would require. Botsch and Malmendier (2023) tabulate responses by cohort from 1989 to 2013 and document that older and younger households hold systematically different views of where rates are heading, which establishes both that the variable carries information and that it can be linked to contract choice on American data. What has not been done is to include it among the covariates of a model of the

decision itself. This statement rests on the studies identified in the course of this review, and no claim of exhaustiveness attaches to it.

The specification estimated by Coulibaly and Li (2009) includes seventeen covariates capturing pricing conditions, affordability, mobility expectations, attitudes, and demographic characteristics, but none directly measures households' own expectations about future interest rates. Instead, the interest rate environment is represented only through market-based variables, such as the mortgage rate spread and the yield curve.

Similarly, the financial sophistication literature discussed in section 3.8 does not exploit these expectation variables. As a result, a determinant that theoretical models place at the center of mortgage choice, and that the SCF has measured for several decades, has not yet been systematically examined using U.S. household-level data. Evidence from other surveys suggests that incorporating expectations into mortgage choice analysis is empirically feasible. A section of the European Central Bank's Economic Bulletin reports households' expectations about future interest rates separately for borrowers holding fixed-rate and adjustable-rate mortgages in the euro area. The analysis shows that, from mid-2022 onwards, an increasing share of adjustable-rate borrowers expected interest rates to decline. Although the evidence is primarily descriptive and captures preferences as well as expectations, it demonstrates that linking mortgage contract characteristics with households' views about future interest rate developments is possible in survey data.

3.7.4 Measurement

Part of this omission may be explained by the characteristics of the expectation variable itself, which merit careful examination rather than being taken for granted. The most important limitation concerns the time horizon over which expectations are measured. The SCF asks respondents whether they expect interest rates to be higher, lower, or about the same five years into the future. However, the evidence presented by Badarinza, Campbell, and Ramadorai (2018) suggests that households' mortgage decisions are influenced primarily by short-term expectations. Their analysis shows that one-year interest rate expectations have a significant effect on mortgage choice, whereas expectations over a three-year horizon exert only a limited influence. Consequently, a variable measuring expectations five years ahead may not correspond to the time horizon that households actually consider when selecting a mortgage contract. Under these circumstances, an insignificant estimated coefficient would be difficult to interpret, as it could reflect either the absence of an expectations channel or a mismatch between the horizon captured by the survey

question and the horizon that is relevant for household decision-making. The expectation variable also has important measurement limitations. It is qualitative in nature, asking respondents only whether they expect interest rates to increase, decrease, or remain unchanged, without capturing the expected magnitude of the change. Consequently, households anticipating a very small increase in interest rates are classified in the same category as those expecting a much larger increase, even though these expectations may have very different implications for mortgage choice. This distinction is particularly relevant because theoretical models emphasize the size of the expected interest rate movement relative to the cost of choosing a fixed-rate mortgage. In addition, the survey does not measure respondents' confidence in their expectations, making it impossible to distinguish between firmly held beliefs and highly uncertain forecasts. The timing issue discussed in section 3.3 is particularly important for this variable. Households continuously update their expectations about future interest rates as new economic information becomes available. As a result, the expectation reported in the 2022 survey may differ substantially from the expectation that influenced a mortgage decision made several years earlier, such as in 2015. Although this timing mismatch introduces measurement error for all explanatory variables considered in this chapter, its consequences are likely to be especially severe for interest rate expectations, because they are inherently forward-looking and subject to frequent revision. Therefore, the expectation variable is likely to provide meaningful information only for households that originated or refinanced their mortgages close to the survey date. This consideration suggests that restricting the analysis by mortgage origination or refinancing year is more appropriate than simply including origination year as a control variable.

3.7.5 The relationship that the literature would predict

Since the item in question has not been deployed, its expected role must be inferred from studies that measure the same underlying construct in other ways. The evidence reviewed above leads to four predictions that make this gap empirically testable. First, the coefficient should be negative: households expecting higher interest rates should be less likely to hold an adjustable-rate mortgage, as found by Devine, McCarthy and O'Toole (2023) for Ireland and by Bacon and Moffatt (2012) for the United Kingdom. Second, the effect is likely to be modest and may not be statistically significant, given the limitations related to the forecast horizon and the coarse measurement discussed in section 3.7.4. Third, expectations should complement rather than replace pricing variables, since Badarinza,

Campbell and Ramadorai (2018) show that the contemporaneous rate spread retains explanatory power even after expectations are taken into account. Finally, the effect should vary with financial competence. Devine, McCarthy and O'Toole (2023) find that education and financial knowledge increase the likelihood of forming accurate expectations, while Ilan and Mugerman (2025) show that less financially literate borrowers tend to base their expectations on observed inflation rather than on its expected path. This suggests that an interaction between the expectation measure and financial competence should produce a stronger effect among households toward the upper end of the competence distribution. The fourth prediction is also the one that has not been tested in the existing literature. It would help distinguish an expectations channel operating through beliefs about future rates from a mechanism in which expectations are shaped by the inflation experience that Botsch and Malmendier (2023) identify as an important source of those beliefs.

3.7.6 Assessment

The evidence reviewed above supports a role for expectations in mortgage choice, subject to the horizon qualification of section 3.7.4 and to the predictions stated in section 3.7.5. The Survey of Consumer Finances records and describes households' expectations about future interest rates, but none of the studies reviewed in this thesis incorporates this information into a model of contract choice. The existing literature does not explain this omission, making it unclear whether it reflects the measurement limitations of the expectation variable or simply the research priorities that have shaped the field. Consequently, the evidence supporting the role of expectations rests on aggregate cross-country analyses, on a purpose-built household survey conducted in a single European country, and on a large British sample in which expectations are imputed rather than reported, with the American contribution confined to descriptive and structural work on inflation experience.

3.8 Financial literacy and information gaps

Row 8 of Table 1 is contested for a different reason than row 2. In the case of row 2, the disagreement arises because theoretical predictions and empirical evidence point in opposite directions. For row 8, however, the uncertainty originates within the theoretical literature itself, which does not provide a clear prediction regarding the expected effect. Two alternative interpretations can be identified, and both are theoretically plausible. The first argues that greater financial competence enables households to evaluate the trade-off between fixed-rate and adjustable-rate mortgages more effectively. Financially sophisticated

borrowers are better able to compare the initial cost savings offered by an adjustable-rate mortgage with the additional interest rate risk it involves. They are also more likely to recognize situations in which the premium charged for a fixed-rate mortgage is unusually high and to adjust their mortgage choice accordingly. Under this interpretation, financially sophisticated households would be expected to choose adjustable-rate mortgages more frequently when market conditions make them relatively attractive. The second interpretation reaches the opposite conclusion. From this perspective, adjustable-rate mortgages are viewed as relatively complex financial products whose risks are not immediately apparent from their lower initial payments. Financial competence therefore serves as a form of protection against underestimating these risks. Households with lower levels of financial knowledge may focus primarily on the initial monthly payment and overlook the possibility of future payment increases, whereas more financially sophisticated borrowers are better able to understand the implications of future interest rate resets. Under this interpretation, higher financial competence would be associated with a lower probability of choosing an adjustable-rate mortgage.

3.8.1 What the estimates show

The analysis conducted by Coulibaly and Li (2009) highlights education as an important determinant of mortgage choice. In their reference specification, households with a college degree are substantially more likely to choose an ARM than households without a college education. In the pooled sample, the estimated effect of education is among the largest predictors of mortgage choice, following only market conditions, such as mortgage pricing variables, and loan characteristics, such as subprime status. The relationship between education and ARM choice is particularly strong among households with lower levels of risk aversion. For this group, higher educational attainment is associated with a 77 percent increase in the probability of selecting an adjustable-rate mortgage. By contrast, among more risk-averse households, the corresponding increase is smaller, although still economically significant, at approximately 41 percent. These results indicate that the positive relationship between education and the choice of adjustable-rate mortgages is more pronounced among households that are less concerned about the uncertainty associated with future payment adjustments. Taken at face value, these findings appear to support the interpretation that greater financial competence encourages households to select adjustable-rate mortgages when they are financially advantageous. More educated borrowers may be better able to evaluate the trade-off between lower initial payments and

future interest rate risk, leading them to choose the contract that best fits market conditions. However, this interpretation cannot be considered conclusive. Further analysis is required to determine whether education captures genuine differences in financial decision-making ability or simply broader socioeconomic heterogeneity among borrowers.

3.8.2 Bimodality

Smith, Finke and Huston (2011) construct a scale of financial sophistication from survey items covering formal education, the respondent's comprehension of the interview, and search behavior around borrowing decisions, following the factor-analytic approach developed for the literacy literature. Their result is that adjustable-rate borrowing draws on both tails of that scale. The least sophisticated households hold adjustable contracts when they are income constrained, attracted by the lower initial payment in the manner described in section 3.4. The most sophisticated hold them when the differential between the two contract rates is wide, exploiting the pricing opportunity discussed in section 3.6. The two groups arrive at the same contract from opposite starting points and for unrelated reasons. The consequence for row 8 is that neither reading is wrong and neither is complete. A linear coefficient estimated across the whole distribution averages two mechanisms of opposite sign, and its value depends on the relative weight of the two populations in the sample. This weight is not a fixed property of the mortgage market: it varies with the spread, which determines how much is there for sophisticated households to exploit, and with affordability conditions, which determine how many constrained households are pushed toward the adjustable product. A coefficient estimated in one rate environment therefore may not be directly applicable to another, a point that chapter 5 develops as a broader challenge of the mortgage choice literature. Taken together, the existing evidence also highlights the importance of interactions between household characteristics. As discussed in section 3.6, more risk-averse households respond more strongly to changes in the mortgage rate spread. By contrast, the results reviewed in this section show that the positive relationship between education and the choice of adjustable-rate mortgages is stronger among less risk-averse households. These findings indicate that risk preferences and financial sophistication interact with other determinants of mortgage choice in ways that cannot be fully captured by a single-equation specification. Although the literature has documented these interaction effects, it has not yet developed an integrated framework that explains them jointly.

3.8.3 Survey measures

The Survey of Consumer Finances does not include a dedicated financial literacy module and therefore does not contain the standard measures commonly used in the financial literacy literature. As a result, the studies reviewed in this thesis rely either on educational attainment as a proxy for financial knowledge or on financial sophistication indices constructed from survey questions that were originally designed for other purposes. Huston, Finke, and Smith (2012) describe one such approach and show that incorporating a financial sophistication measure improves the explanatory power of standard demographic variables. This finding suggests that education alone captures only part of the variation associated with households' financial decision-making abilities. Both approaches, however, have important limitations. Educational attainment reflects years of formal schooling rather than actual financial knowledge, and the relationship between the two is far from perfect. Similarly, financial sophistication indices constructed from existing survey items depend on the specific variables included and on the statistical method used to combine them. Consequently, results obtained using one measure of financial sophistication may not be directly comparable with those obtained using another. Moreover, neither approach directly measures whether households fully understood the specific mortgage contract they ultimately selected. Despite these measurement challenges, the underlying theoretical framework is well established. Lusardi and Mitchell (2014) show that financial literacy is generally low, varies substantially across the population, and is an important predictor of financial behavior in many different contexts. Similarly, Campbell (2006) demonstrates that although most households make sound financial decisions, a minority consistently make important mistakes, particularly among those with lower levels of income and education. The main limitation of the mortgage choice literature is therefore not the absence of a theoretical role for financial literacy, but the lack of a direct and reliable measure of this concept.

3.8.4 Knowledge of the contract itself

A complementary strand of the literature examines whether borrowers accurately understand the mortgage contracts they have chosen. Bucks and Pence (2008) address this question by comparing mortgage characteristics reported by households with corresponding information obtained from lenders' administrative records. Their analysis combines data covering approximately 80 percent of outstanding U.S. mortgages at the end of 2001

with information from the Residential Finance Survey, restricting both datasets to first-lien mortgages. Their findings provide a nuanced picture of borrowers' mortgage knowledge. Households generally report the main characteristics of their mortgages accurately. In particular, the distributions of mortgage type, monthly payment, and loan maturity reported in the survey closely match those recorded by lenders, and only a small proportion of respondents indicate that they do not know these basic features. These results suggest that the distinction between fixed-rate and adjustable-rate mortgages is measured reliably in survey data, providing reassurance for the empirical studies reviewed in the present work. However, this accuracy does not extend to the contractual features that determine the future cost of adjustable-rate mortgages. Borrowers holding adjustable-rate contracts are substantially less informed about the size of potential interest rate adjustments, the frequency with which interest rates can reset, and the contractual caps that limit those adjustments. Importantly, the lowest levels of understanding are observed among lower-income households, who are also the borrowers most vulnerable to increases in interest rates.

This evidence suggests that while households generally know the broad characteristics of their mortgage contracts, many have only a limited understanding of the features that determine the risks associated with adjustable-rate borrowing. These findings have two important implications. First, they provide strong support for the reliability of the dependent variable used throughout the mortgage choice literature. The distinction between fixed-rate and adjustable-rate mortgages appears to be measured accurately, which is particularly important given its central role in empirical analyses. Second, the measurement error affecting variables related to informed decision-making is unlikely to be random. Instead, it is concentrated among households with lower levels of education and income, which are also the groups whose financial literacy and sophistication the empirical models aim to capture. This type of non-classical measurement error tends to bias estimated coefficients toward zero, suggesting that the reported effects of financial literacy and education are more likely to underestimate than overestimate their true influence on mortgage choice. The findings of Bucks and Pence (2008) also provide additional insight into the interpretation of the positive coefficient on education reported in the literature. If borrowers with lower levels of education are less likely to understand the adjustment features of adjustable-rate mortgages, then the observed relationship between education and the choice of an adjustable-rate contract may reflect not only differences in mortgage preferences but also differences in borrowers' understanding of the contracts they hold.

Consequently, the available data cannot determine whether more educated households actively choose adjustable-rate mortgages because they evaluate them as more advantageous, or whether less educated households are simply less aware of the characteristics and risks of their mortgage contracts.

3.8.5 Assessment

Row 8 remains an unresolved issue, and the underlying disagreement cannot be addressed simply by estimating the same empirical specification with additional data. Although the positive relationship between education and adjustable-rate mortgage choice appears robust in the reference study, the evidence on financial sophistication indicates that this average effect combines the behavior of different groups of households that respond to adjustable-rate mortgages for fundamentally different reasons. Moreover, the measurement foundation for this determinant is weaker than for most other variables examined in this chapter. The SCF does not provide a direct measure of financial literacy, the available proxies are constructed indirectly rather than specifically designed to capture financial competence, and the only study that examines borrowers' actual understanding of mortgage contracts finds that knowledge of adjustable-rate mortgage features is systematically lower among the households whose decision-making behavior these variables are intended to explain.

The most reliable conclusion emerging from this literature is therefore the existence of heterogeneous responses across the financial sophistication distribution. This finding suggests that the analysis should move beyond the search for a single average coefficient and instead adopt empirical specifications that allow the relationship between financial competence and mortgage choice to vary across different groups of households.

Evidence obtained outside the United States with a purpose-designed literacy measure bears directly on this ambiguity and is examined in section 4.3.

3.9 Socioeconomic and demographic characteristics

Rows 5, 6, and 7 of Table 1 refer to income, wealth, and age, respectively. These variables are included in virtually all empirical specifications in the mortgage choice literature and are typically treated as standard control variables. However, this characterization understates their economic relevance. Income is associated with one of the more unexpected findings in the literature, as its relationship with mortgage choice does not always align with conventional theoretical predictions. Wealth, despite its importance in determining

households' ability to absorb payment risk, is not directly measured in the reference study. Age, meanwhile, reveals a notable divergence between descriptive patterns and conditional regression results, highlighting a broader challenge that emerges throughout this chapter: the relationship between household characteristics and mortgage choice can change substantially once other determinants are considered.

3.9.1 Income

Row 5 examines the relationship between household income and mortgage contract choice. From a theoretical perspective, the expected direction of this relationship is not straightforward. Fixed-rate mortgages generally have higher initial borrowing costs because they include compensation for long-term interest rate risk, whereas adjustable-rate mortgages typically offer lower initial monthly payments. From an affordability perspective, this difference suggests that households with more limited financial resources should be more likely to choose adjustable-rate mortgages in order to reduce their initial borrowing costs. This prediction is consistent with the liquidity constraint mechanism discussed in section 3.4, according to which financially constrained households are more sensitive to initial payment affordability. Extending this reasoning to household income implies that the likelihood of choosing an adjustable-rate mortgage should decrease as income increases.

The empirical evidence, however, paints a different picture. Studies have consistently found that higher-income households are more likely to choose adjustable-rate mortgages, a result that has remained remarkably stable for nearly four decades. Brueckner and Follain (1988) were among the first to document a positive relationship between household income and the use of adjustable-rate mortgages during the expansion and subsequent decline of the U.S. adjustable-rate mortgage market in the 1980s. More recent survey-based evidence confirms the persistence of this pattern.

The relationship is also more nuanced than a simple positive income gradient. Chiang and Dueholm (2024) show that the share of adjustable-rate mortgages is 6.5 percent among mortgaged households in the lowest income decile, falls to 3.9 percent in the fifth decile, and then rises sharply to 18.8 percent among households in the highest decile. Figure 5 displays the full distribution across deciles and makes the resulting shape apparent. The share declines gradually from the first decile to the fifth, recovers moderately through the sixth and seventh, and increases steeply only across the three highest deciles. Adjustable-

rate borrowing is therefore comparatively rare through the middle of the income distribution, markedly more frequent at the top, and modestly elevated at the bottom.

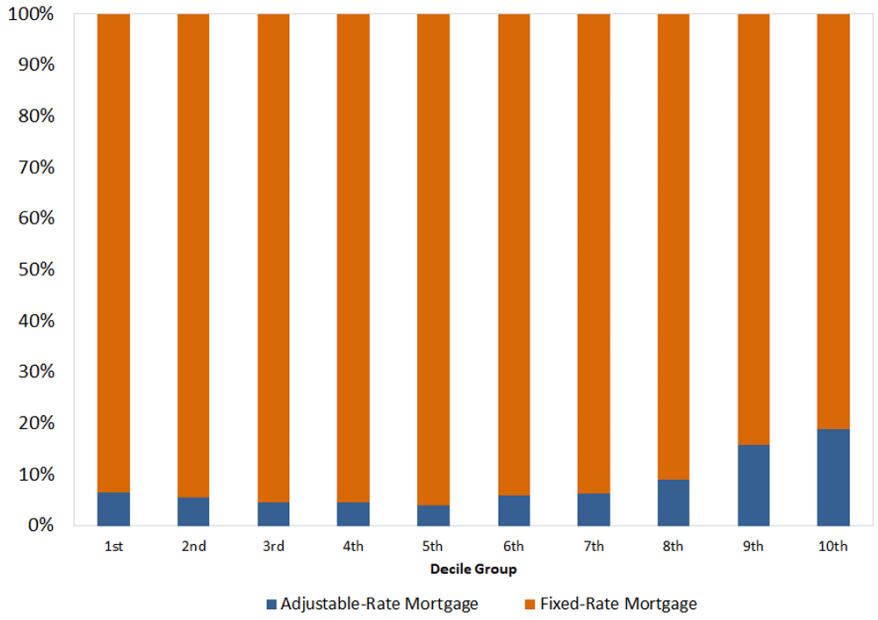


Figure 5.

Distribution of mortgage contract type by household income decile, 2019.

Source: Chiang and Dueholm (2024), based on the 2019 Survey of Consumer Finances. Deciles are ordered from lowest to highest income.

This pattern closely resembles the bimodal distribution discussed in section 3.8. The use of adjustable-rate mortgages is concentrated at both ends of the income distribution, although for different reasons. Among lower-income households, the choice of an adjustable-rate mortgage is primarily driven by affordability constraints, as the lower initial payments reduce the immediate financial burden. Among higher-income households, by contrast, the decision appears to reflect a more strategic response to relative mortgage pricing, with borrowers taking advantage of the lower cost of adjustable-rate contracts when market conditions make them attractive. Although these two groups select the same type of mortgage, they do so for fundamentally different reasons. Notably, the two-group structure emerges independently from two distinct strands of the literature, one based on financial sophistication measures and the other on income distributions. The consistency of these findings strengthens the evidence that mortgage choice reflects heterogeneous decision-making across different segments of the population.

The explanation commonly proposed for the behavior of higher-income households is based on the buffer stock mechanism. Households with substantial financial resources are better able to absorb temporary increases in mortgage payments without significantly reducing their consumption. Consequently, the insurance provided by a fixed-rate mortgage against future interest rate increases is less valuable for these borrowers, as they are capable of managing payment shocks using their own financial resources. This reduces the incentive to pay the additional cost associated with a fixed-rate mortgage and increases the attractiveness of an adjustable-rate contract.

The evidence from the reference study, however, should be interpreted with some caution. Household income is not included in the empirical model as a standalone explanatory variable. Instead, it appears only as the denominator of two leverage measures: the mortgage balance-to-income ratio and the non-housing debt-to-income ratio. Of these, only the mortgage balance-to-income ratio has a positive and statistically significant effect on mortgage choice. Consequently, the estimated model identifies the influence of household leverage relative to income rather than the independent effect of income itself. This distinction is important because the descriptive evidence presented earlier and the regression results capture different aspects of household financial conditions and therefore should not be interpreted as equivalent.

3.9.2 Wealth

Row 6 derives from the same buffer-stock logic and from the portfolio framework of Cocco (2005), in which the housing position absorbs the household's capacity to bear risk and crowds out other risky holdings, the effect being strongest among households with low financial net worth. A household heavily committed to housing relative to its financial assets retains little residual capacity to withstand payment variability and is predicted to prefer the fixed contract. Within the buffer stock framework, the relevant concept is liquid financial wealth rather than total household wealth. The mechanism relies on the availability of financial resources that can be used to absorb increases in mortgage payments following an interest rate adjustment. Home equity does not serve this function because it cannot be accessed without selling the property or borrowing against it. Moreover, its value is particularly difficult to realize during periods of stress in the housing market. A similar distinction was discussed in section 2.3 in relation to the sharp increase in house prices after 2020. Although rising house values substantially increased households measured net worth, they did not necessarily expand the liquid financial resources

available to cope with higher mortgage payments. As a result, total net worth, which is the wealth measure most readily available in the Survey of Consumer Finances, combines liquid and illiquid assets and may therefore overstate the financial buffer available to households whose wealth is largely tied to their primary residence. This is the case for a large proportion of homeowners with moderate levels of wealth. The reference study does not include household net worth among its explanatory variables. Instead, it incorporates an indicator of stock ownership, which has a positive and statistically significant coefficient of 0.39 and is observed for 28.3 percent of the sample. However, the authors explicitly interpret this variable as a proxy for financial sophistication rather than as a measure of liquid financial wealth, motivating its inclusion by the finding that many households do not recall the more complex terms of their adjustable contracts. Under the author's own interpretation, the variable belongs to the discussion of section 3.8. Row 6 has nonetheless been tested, and outside the United States. Ehrmann and Ziegelmeyer (2017) include liquid financial wealth measured in quintiles of the national distribution, defined as the sum of current and savings accounts, mutual funds, bonds, shares and managed accounts net of uncollateralized debt. This is the construct the mechanism specifies, and the result is null: none of the four quintile indicators attains significance in the benchmark specification, and only one does so in a single robustness variant, with a negative sign. Two qualifications attach to this finding. Wealth is recorded at the date of the survey rather than at origination, and the authors introduce the variable to capture the affordability of the fixed contract rather than the capacity to absorb payment shocks, so a null result admits either reading. The row therefore moves from untested to tested without result, which is a stronger statement than the absence of evidence but a weaker one than corroboration.

3.9.3 Age

Among the determinants examined in this chapter, age displays the greatest difference between descriptive evidence and regression-based results. The descriptive statistics reveal a large gap in the average age of borrowers across mortgage types. The median age of households holding ARM contracts is 32.1 years, compared with 50.3 years for households holding FRMs, a difference of approximately eighteen years. However, once other household characteristics are considered, the estimated effect of age becomes very small. In the reference specification, the age coefficient is reported to be 0.01. Although statistically significant, its economic magnitude is negligible, indicating that age has only a

limited direct effect on the probability of choosing an ARM after controlling for other relevant factors, such as mobility expectations, leverage, stock ownership, and education. These two findings are not contradictory because they capture different aspects of the relationship between age and mortgage choice. The descriptive comparison reflects all the characteristics that typically distinguish younger households from older ones. Younger borrowers are generally more likely to relocate, are at an earlier stage of their earnings profile, have accumulated less wealth, and are more likely to have originated their mortgage recently. In the regression analysis, these characteristics are captured directly by the corresponding explanatory variables. Consequently, the coefficient on age measures only the variation that remains after accounting for these observable factors. The fact that this residual effect is very close to zero suggests that age influences mortgage choice primarily through these observable channels rather than through an independent effect. This result also helps explain the theoretical ambiguity associated with age discussed in Table 1. The lifecycle framework and the experience-based approach generate different predictions regarding the effect of age on mortgage choice. According to the experience-based perspective, older households may place greater value on the payment stability provided by fixed-rate mortgages because they have experienced periods of high inflation and volatile interest rates. However, if this mechanism operates through households' expectations about future economic conditions, it cannot be identified in a model that does not include measures of those expectations. Consequently, the small estimated coefficient on age should not be interpreted as evidence against the experience-based hypothesis. Instead, it reflects the inability of the reference specification to capture this channel directly. This issue is examined in chapter 4, where the formation of expectations is analyzed explicitly, and is closely related to the omission of expectation variables discussed in section 3.7.

3.9.4 Other demographic characteristics

The remaining demographic variables play a relatively limited role in explaining mortgage choice. In the reference specification, marital status, household size, and race are all statistically insignificant. This result is obtained from a sample in which 76.8 percent of respondents are married and 82.0 percent are white. The authors note that the lack of statistically significant effects for these standard demographic characteristics is consistent with the findings of several earlier studies. The result for race, however, deserves careful interpretation. The reference model includes an indicator for subprime borrowing, which has the largest estimated effect among all explanatory variables. Previous research has

consistently shown that subprime lending in the United States has been disproportionately concentrated among minority borrowers. Avery et al. (2010), in a Federal Reserve Bulletin, report that black and Hispanic households experienced the highest incidence of high-priced home purchase loans during the peak of the subprime lending boom. Similarly, Mayer and Pence (2008) document that subprime mortgage originations were more common in ZIP codes with larger black and Hispanic populations. This pattern is observed from the late 1990s onward and therefore covers much of the period represented by the SCF waves used in the reference study. Moreover, the measure of subprime status employed by Coulibaly and Li (2009) closely resembles the approach adopted in these studies, as it is based on the spread between the mortgage interest rate and a market benchmark. Consequently, the absence of a statistically significant coefficient on race should not be interpreted as evidence that racial differences in mortgage choice do not exist. Rather, it suggests that much of this variation is already captured by the subprime borrowing variable included in the model. In other words, the relationship between race and mortgage choice may operate indirectly through differences in access to, or participation in, the subprime mortgage market.

A second consideration also calls for caution when interpreting this result. Because 82 percent of the sample consists of white households, fewer than 600 observations are available to estimate the coefficient for minority groups. The lack of statistical significance may therefore reflect limited statistical power rather than the absence of a meaningful relationship. Finally, the available evidence does not clarify whether the observed relationship between minority status and subprime borrowing reflects differences in household preferences, differences in credit access, or discriminatory lending practices. This distinction is important because it has different implications for interpreting mortgage choice. However, the studies reviewed in this chapter do not address this issue directly. More broadly, the literature on household mortgage choice and the literature on racial disparities in mortgage lending have largely developed independently, and their limited integration represents one of the research gaps identified in chapter 5.

3.9.5 Assessment

The positive income gradient ranks among the better established findings in this literature, holding across four decades, several datasets, and both descriptive and conditional evidence. Its interpretation nonetheless remains incomplete, since the shape of the descriptive relationship indicates the presence of two distinct populations rather than one,

and the reference specification identifies the effect of leverage relative to income rather than that of income itself.

The evidence on wealth is the least satisfactory of the three. The mechanism specified by the theory concerns liquid reserves, whereas the measure available in the survey aggregates liquid and illiquid holdings, and the reference study omits the variable from its specification, the single candidate proxy having been introduced by the authors for a different purpose.

The evidence on age is indicative of a difficulty common to this literature. A substantial descriptive difference is reduced to a negligible magnitude once controls are introduced, a result that is informative as to the channels through which age operates but that leaves the theoretical ambiguity unresolved, since the mechanism underlying the competing prediction is not observed in any of the specifications reviewed.

3.10 The evidence assessed: theory and estimates compared

The preceding sections have examined each determinant against the studies bearing upon it. The present section sets those results alongside the predictions from which they were derived and considers what the pattern of agreement and disagreement indicates about the state of the field. Table 4 performs this comparison. Its rows correspond to those of Table 1, so that each theoretical prediction appears alongside the estimate obtained for it, and the rows on which the two diverge become apparent without further argument. Two columns are added to the comparison itself. The first records the evidence supporting each estimate, since a sign replicated on survey, administrative and cross-country data carries different weight from one obtained within a single specification. The second records the principal qualification attaching to the result, which in several instances governs how the sign is to be interpreted. The table does not report the descriptive characteristics of the samples, which appear in Table 3, nor does it reproduce the economic mechanisms set out in Table 1.

Table 4. Theoretical predictions and empirical findings on the determinants of mortgage choice

Determinant	Expected sign	Estimated sign	Evidence base	Principal qualification
Risk aversion	Negative	Negative	SCF	Not reproduced on Italian data; item elicited with reference to saving and investment rather than borrowing and shown to be shaped by the respondent's own inflation history
Liquidity constraints	Negative	Positive	SCF; Italian SHIW; SCF combined with expenditure data	Reversal robust across countries; largest coefficient attaches to subprime status, a variable defined on the contract rate itself
Expected mobility	Positive	Positive	SCF; loan-level data from the 1980s	Survey horizon of two years against a contract term of thirty; significant only among less risk-averse borrowers
Interest rate expectations	Negative if increases expected	The sign differs across institutional and geografical settings (negative in U.S.)	Nine-country panel; purpose-built Irish survey; U.S. and euro area evidence on experience effects	The SCF item is described but does not enter contract-choice specifications; experience raises fixed-rate demand in the United States and lowers it in the euro area, where prepayment penalties alter the risk the contract carries
Household income	Positive	Positive	Loan-level data from the 1980s; SCF descriptive tabulations	Relationship non-monotonic across deciles in the American data; the American reference specification identifies leverage relative to income, while the euro area estimate places the top quintile twelve percentage points above the bottom
Net worth	Positive	Null	Euro area HFCS	Liquid financial wealth, the construct the mechanism specifies, enters the euro area specification in quintiles and is insignificant throughout; absent from the American reference specification

Determinant	Expected sign	Estimated sign	Evidence base	Principal qualification
Age	Ambiguous	Positive but negligible	SCF	Descriptive difference of eighteen years absorbed by other covariates; experience channel unobservable without expectations
Education and literacy	Ambiguous	Positive	SCF; SCF sophistication scales; UK survey with direct literacy measure	Effect bimodal across the sophistication distribution; survey contains no literacy module
Contract term	Positive	Positive	SCF	Estimated through a dummy for maturities of fifteen years or less; result attributable in part to lender supply
ARM-FRM spread	Positive	Positive	SCF; loan-level; market-level; nine-country panel	Estimates capture the observed differential rather than the term premium specified by the theory
Income volatility	Negative	Negative	SCF; euro area HFCS	Measure imputed by occupation rather than observed; in the euro area only the unemployment proxy is significant, the remaining proxies showing no effect

Note. Rows correspond to those of Table 1. Expected signs are the predictions derived in chapter 1; estimated signs are those prevailing in the studies reviewed in this chapter, which comprise American, European and cross-country evidence. Descriptive characteristics of the populations concerned appear in Table 3, which is confined to the Survey of Consumer Finances. A null entry in the estimated sign column indicates that the determinant has been examined without a significant relationship being found and is distinct from an entry recording that no estimate exists.

3.10.1 Three groups of rows

The determinants reported in Table 4 can be grouped into three broad categories. This classification provides a clearer picture of the current state of the literature than considering each determinant in isolation. The first category includes determinants for which the empirical evidence is both consistent and well established. Risk aversion, expected mobility, ARM-FRM spread, mortgage maturity, and income volatility all exhibit the signs predicted by theory, and their effects have been confirmed across multiple datasets and

empirical specifications. Among these variables, the interest rate differential receives the strongest empirical support. Its influence on mortgage choice has been documented consistently using survey data, loan-level administrative records, aggregate market data, and cross-country evidence over a period of more than forty years. These findings represent the most robust empirical regularities in the mortgage choice literature and provide important benchmarks that any theoretical or empirical model should be able to explain.

The second category consists of determinants for which the empirical evidence either conflicts with theoretical predictions or remains inconclusive. Liquidity constraints belong to the first case. While the lifecycle framework predicts that financially constrained households should prefer fixed-rate mortgages because they provide insurance against future payment increases, empirical studies consistently find the opposite relationship across different countries, datasets, and model specifications. This pattern is sufficiently robust to suggest that it cannot be explained by the characteristics of any single study.

Age falls into the second case. Its estimated effect is consistently close to zero and precisely estimated, but this finding does not resolve the theoretical ambiguity surrounding its role. The competing theoretical explanation based on experience-driven expectations cannot be adequately tested because the empirical models reviewed in this chapter do not directly observe households' expectations about future interest rates.

The remaining determinants have received only limited empirical attention or have not been examined directly. On one hand, the effect of interest rate expectations has been estimated on American household data only indirectly, through the inflation experience that shapes them, and the expectation item recorded in the survey has not itself entered a model of contract choice. On the other hand, household net worth is not included among the explanatory variables in the reference specification. Income is considered only indirectly through measures of leverage relative to income rather than in levels in the reference specification. Education is included in the empirical models, but its interpretation remains uncertain. Evidence from the financial sophistication literature indicates that the estimated education coefficient combines the behavior of different groups of households that choose adjustable-rate mortgages for different reasons. Consequently, this coefficient should not be interpreted as representing a single, homogeneous effect of education on mortgage choice.

3.10.2 The pattern of empirical success

When considered together, the three groups of determinants reveal two broader patterns that are less apparent when each variable is examined individually.

The determinants for which the evidence is strongest are generally those that can be measured objectively or with relatively little measurement error. The interest rate spread is directly observed from market data, income volatility is derived from external occupational information, the mortgage term is recorded in administrative loan records, and expected mobility is measured through a survey question referring to a concrete and easily verifiable circumstance. Risk aversion is the main exception to this pattern because it is based on self-reported information. Notably, it is also the only well-supported determinant for which the international evidence is not entirely consistent.

By contrast, the determinants for which the evidence remains mixed or incomplete are those that depend on households' subjective beliefs, knowledge, or perceptions. Variables such as interest rate expectations, financial literacy, and borrowing constraints cannot be observed directly. In many cases, the survey either does not measure these concepts, measures them only imperfectly, or includes relevant information that has not been fully exploited in the existing literature. As a result, researchers have often relied on indirect proxies or have omitted these determinants altogether.

This pattern suggests that the strength of the empirical evidence largely reflects the quality of the available data rather than the validity of the underlying theoretical mechanisms. Determinants that can be measured directly have generally received consistent empirical support, whereas those involving subjective judgments or latent characteristics remain more difficult to evaluate because of measurement limitations. This observation should not be interpreted as a criticism of the individual studies reviewed in this chapter, all of which were necessarily constrained by the data available at the time. Instead, it highlights a broader limitation of the literature. Many of the remaining unanswered questions arise not because the theoretical mechanisms are implausible, but because they are difficult to measure accurately with existing survey data.

A second pattern emerging from the literature concerns the pervasive presence of heterogeneous effects across households. Whenever researchers have examined whether the impact of a determinant varies across borrower groups, they have consistently found evidence that it does. The responsiveness to the FRM-ARM spread is stronger among more risk-averse households, whereas the positive association between education and

adjustable-rate borrowing is more pronounced among less risk-averse borrowers. Similarly, expected mobility affects mortgage choice primarily within this latter group, while both financial sophistication and income display non-linear relationships with adjustable-rate borrowing rather than a single monotonic pattern. Although these findings are generally presented as secondary results, together they carry an important methodological implication. The effect of several determinants is not constant across the population but depends on observable household characteristics. As a result, coefficients estimated from pooled single-equation models should be interpreted as average effects across heterogeneous groups rather than as structural parameters that apply uniformly to all borrowers. The estimated coefficients therefore reflect not only the underlying behavioral relationship but also the composition of the sample on which they are estimated. This consideration is particularly relevant when comparing studies conducted in different periods or under different interest rate environments, where the composition of mortgage borrowers may itself change.

3.10.3 Toward the remaining chapters

This assessment naturally leads to two questions, each of which is addressed in the following chapters. The first concerns the determinants for which theory and empirical evidence do not fully align. These variables share an important common feature: they are closely related to households' beliefs, expectations, and understanding of mortgage products. The positive coefficient on liquidity constraints suggests that mortgage decisions may place greater weight on short-term affordability than on the long-term consumption smoothing predicted by the lifecycle framework. This interpretation is consistent with the international evidence showing that households respond primarily to short-term interest rate expectations, as well as with the bimodal relationship between financial sophistication and ARM choice. Together, these findings point to the importance of behavioral factors that are not fully captured by the standard rational benchmark. Chapter 4 examines these issues by focusing on expectation formation and on the extent to which observed household behavior departs from the predictions of fully rational decision-making models.

The second question concerns the stability of the estimated relationships reported in Table 4. Existing evidence suggests that the characteristics of households choosing adjustable-rate mortgages change over time as interest rate conditions evolve. As discussed in chapter 2, and further supported by the evidence on heterogeneity across financial

sophistication groups, the composition of adjustable-rate borrowers is not constant across different market environments. Consequently, coefficients estimated during one interest rate regime may not remain valid under another. Chapter 5 addresses this issue by examining the extent to which the determinants identified in the existing literature remain stable in the 2022 environment. It also discusses the principal gaps highlighted throughout this review, with particular attention to the absence of household-level evidence on the role of interest rate expectations, which remains one of the most important unresolved questions in the mortgage choice literature.

4. Behavioral dimensions and limits of the rational benchmark

4.1 Rationale and scope of the chapter

Chapter 3 tested each determinant of mortgage choice against the predictions of the theoretical framework set out in chapter 1, and while several determinants showed an estimated sign consistent with theory, others did not, and these divergent cases turn out to share a common feature. This chapter looks more closely at those cases and reviews the behavioral explanations that have been proposed to account for them. Three findings from chapter 3 are particularly relevant here. The coefficient on liquidity constraints carries the opposite sign to what the lifecycle model predicts, and since this reversal appears consistently across several countries and datasets, it cannot be put down to a single specification. A similar pattern emerges from the cross-country evidence on interest rate expectations: households respond strongly to forecasts one year ahead but only weakly to forecasts three years out, a result that sits uneasily with a framework meant to capture the cost of a thirty-year obligation. Finally, the relationship between financial sophistication and adjustable-rate borrowing is bimodal, in a way that the least and the most sophisticated households end up choosing the same contract, though for strongly different reasons.
What these three findings have in common is that they concern how the decision is reached rather than the incentives households face. The first two suggest that households give more weight to the immediate installment than the theory assumes, while the third suggests that it is precisely the ability to assess the contract that determines the final choice. The pattern for which household financial decisions depart systematically from normative benchmarks, is the founding premise of the field defined by Campbell (2006), and two aspects of that approach apply directly to mortgage decisions. On one hand, households face constraints that standard portfolio theory sets aside, such as fixed costs of participation, uninsurable labor income risk, borrowing limits, and contracts whose real value depends on inflation. On the other hand, errors are not evenly spread across the population: most households manage their finances competently, but a minority make costly mistakes, and these mistakes tend to be concentrated among poorer and less educated households. Because of this, such errors end up correlated with the very characteristics used as explanatory variables in mortgage choice models, which limits how far they can be treated as random.
The relation between this chapter and the preceding one requires clarification, since both

Artificial intelligence tools were used to support the drafting, editing, and linguistic refinement of selected parts of this chapter.

deal with expectations and financial competence, but approach them in different ways. Chapter 3 treated these attributes as variables within estimated models, asking what coefficient the literature has obtained for each and how reliably each has been measured. This chapter instead treats them as processes, asking how households form a view about future interest rates, what determines whether that view is well founded, and how beliefs interact with an understanding of the contract on offer. A coefficient on an expectation variable simply records an association between a stated belief and an observed choice, but it says nothing about how that belief was formed, whether it responds to available information, or why it varies systematically with age and circumstance. The rational framework retains explanatory value, as the corroborated determinants from chapter 3 show, and the evidence reviewed below qualifies that framework rather than replacing it. The individual determinants are not revisited, since doing so would simply repeat material already covered. The broader question this chapter addresses is how far a framework based on informed optimization can describe the decisions households actually make, and where the departures from it become systematic enough to warrant separate treatment.

The chapter is organized into three parts. Section 4.2 looks at how expectations are formed, drawing on the experience-based account proposed by Malmendier and Nagel (2016), in which the macroeconomic conditions a household has lived through shape its forecasts in ways that vary predictably by age and cohort. Section 4.3 reviews the evidence on cognitive biases documented specifically in the context of mortgage decisions, distinguishing findings established in this setting from the broader behavioral finance literature. Section 4.4 then assesses how well the rational benchmark predicts observed behavior and considers what the accumulated evidence implies for the estimates assembled in chapter 3.

4.2 The formation of interest rate expectations

The theoretical framework presented in chapter 1 treats expectations about future interest rates as an input into the household's decision-making process, without specifying how these expectations are formed. The models implicitly assume that households exposed to the same public information generate identical forecasts. This section considers an alternative explanation, according to which this assumption does not hold systematically, and reviews the evidence showing that the proposed mechanism is relevant in the mortgage market.

4.2.1 The experience-based mechanism

Malmendier and Nagel (2016) argue that individuals form expectations about inflation by overweighting the realizations they have observed during their own lifetimes. The rule they estimate assigns weight to past inflation according to how recently it occurred and how large a share of the individual's lifetime it represents, so that a given realization counts for more in the forecast of a young person than in that of an older one. Younger individuals consequently revise their expectations more strongly in response to new data, since recent observations constitute a greater proportion of their accumulated experience. The model is estimated on fifty-seven years of microdata from the Reuters and University of Michigan Survey of Consumers, and the age-dependent updating pattern it predicts is recovered in those data. Two implications follow for the study of household decisions. Individuals of different ages, observing identical current conditions, hold measurably different forecasts, and the disagreement is largest in periods of volatile inflation such as the 1970s. The dispersion of expectations across households should not be interpreted as random noise, but as a predictable outcome of differences in the macroeconomic experiences of each cohort. Malmendier and Nagel (2016) extend the framework to household borrowing and lending, including the choice of mortgage contract, and the application is developed in detail in the work discussed below.

4.2.2 Evidence from the mortgage market

Botsch and Malmendier (2023) examine the choice between fixed and adjustable contracts directly, taking as their point of departure a fact that the review of chapter 3 leaves unexplained. Fixed-rate contracts account for roughly eighty percent of the American market, a proportion that standard lifecycle models do not predict, particularly given that the fixed rate has been priced on average 170 basis points above the comparable adjustable alternative.

Their argument proceeds in two stages. Households that have experienced higher inflation over their lifetimes expect larger future increases in interest rates, which follows from the Malmendier and Nagel (2016) rule applied to the term structure. Because they anticipate such increases these households place greater value on locking in the payment and are correspondingly more likely to choose the fixed contract. The empirical implementation estimates a structural discrete-choice model over mortgage financing alternatives, using the Residential Finance Survey conducted by the Census Bureau, which links household

characteristics to contract terms reported by servicers, and supplementing it with the Survey of Consumer Finances and with mortgage rate surveys. The estimated magnitudes indicate that the mechanism is economically substantial: an additional percentage point of experienced inflation raises a borrower's willingness to pay for a fixed-rate contract by between 7 and 17 basis points in a given origination year. Aggregating across the population, the authors calculate that close to one in six households would have selected an adjustable contract in the absence of the experience effect. The welfare consequences follow from the price of the insurance so acquired, meaning that households that would otherwise have switched contract are estimated to pay between 8,000 and 19,000 dollars, measured in year-2000 after-tax terms, for the inflation protection embedded in the fixed contract.

The most relevant aspect for the present review is how concentrated these effects are by cohort. The excess demand for fixed-rate contracts is attributed mainly to households whose formative years coincided with the Great Inflation[3], and the authors estimate that this cohort's aggregate overpayment amounted to 12.6 billion dollars during the 1980s and a further 9.1 billion during the 1990s. Evidence from a second market bears on the generality of the mechanism and complicates it. Cenzon and Szabó (2024), examining household data from nine euro area countries, find that for a one log-point increase in lifetime experienced inflation the associated decrease in the odds of holding a fixed rate contract is computed at 71 percent. The sign is the reverse of that obtained on American data, where the same experience raises the demand for fixed contracts. The authors attribute the reversal to the prepayment penalties prevalent in euro area markets, which prevent a borrower from exiting a fixed contract without cost and thereby heighten the inflation risk it carries, whereas prepayment is effectively free in the United States and the fixed contract there operates as a one-sided option, protecting against increases while permitting refinancing if rates fall. Experience therefore appears to shape the choice in both markets, while the direction in which it operates depends on the contractual structure prevailing in each.

4.2.3 Relation to the determinants reviewed in chapter 3

Three connections to the preceding chapter are worth noting.

First, this account offers a possible mechanism behind the ambiguity recorded in row 7 of

[3] The defining macroeconomic period of the second half of the 20th century lasted from 1965 to 1982.

Table 1. Section 3.9 reported that the estimated coefficient on age is precise and close to zero once other characteristics are controlled for, and it noted that this result cannot distinguish between the lifecycle channel and the experience channel, since the latter operates through expectations that the specifications in question do not observe. Botsch and Malmendier (2023), by contrast, do observe these expectations, and find the experience channel to be substantial. The near-zero age coefficient found in the survey literature is therefore consistent with an experience effect that those specifications are simply not equipped to detect.

Second, this account also has implications for how the expectation variable discussed in section 3.7 should be interpreted. Under the experience-based rule, a household's stated forecast is not an independent input to the decision, but is itself shaped by the household's macroeconomic history. Estimating the effect of expectations on contract choice without accounting for how those expectations are formed therefore risks attributing to belief something that actually originates in cohort membership.

Finally, a qualification concerning data should be noted. The authors draw on the expectation item recorded in the Survey of Consumer Finances, which they tabulate by cohort to establish that older and younger households hold systematically different views of the future path of rates, but the estimation of contract choice itself rests on the Residential Finance Survey, which links household characteristics to contract terms reported by servicers. The gap identified in section 3.7 is therefore narrowed rather than closed. The expectation item has been shown to carry information on American data, but what remains absent is a model of the choice between the two contracts in which that item enters as a covariate.

4.2.4 The mechanism applied to the 2022 environment

The framework permits an inference about the period examined in chapter 2, and the inference is less straightforward than it may appear. Households originating or refinancing mortgages in 2022 had, for the most part, accumulated their financial experience within the low and stable inflation regime that prevailed from the mid-1980s onward. Applied mechanically, the experience rule implies that such households would hold subdued expectations of future rate increases and would therefore attach comparatively little value to the protection afforded by the fixed contract, which is consistent with the recovery in the adjustable share documented in chapter 2.

The rule nonetheless weighs recent observations most heavily, and 2022 delivered the

sharpest inflation and the fastest tightening in four decades. The same mechanism that produces subdued expectations from a long low-inflation history also produces a substantial upward revision when a shock of that magnitude occurs, and the revision is predicted to be largest among the young, for whom the episode represents the greatest proportion of accumulated experience. The two forces operate in opposite directions, and their relative strength cannot be settled from the theory alone. The observation is advanced here as an implication of the framework rather than as an established result, since no study has yet estimated experience effects on contract choice for that period.

4.2.5 Limits of the account

The framework accounts for part of the phenomenon rather than the whole of it, and it operates alongside the pricing, affordability and mobility channels documented in chapter 3 rather than taking their place. Its evidential base has broadened with the addition of euro area evidence, which corroborates the mechanism on an independent market, but the corroboration comes with a reversal of sign that no purely behavioral account explains. What the two studies establish jointly is that inflation experience shapes contract choice; what determines the direction is the institutional setting, and neither study derives that dependence from a common framework. The mechanism should therefore be regarded as documented rather than as understood.

4.3 Cognitive biases in the mortgage decision

The behavioral finance literature has identified a wide range of systematic deviations from normative decision-making. However, applying this broad set of behavioral biases to the mortgage market without specific evidence would be inappropriate. This section therefore focuses only on biases for which there is direct evidence in the mortgage context, based either on surveys of actual borrowers or on observed mortgage choices. Biases that are well documented in other areas of finance but have not been specifically examined in relation to mortgage choice are mentioned where relevant but are not assumed to apply to mortgage decisions without supporting evidence.

4.3.1 Present bias and self-control

The most direct evidence is provided by Gathergood and Weber (2017), who developed a specific set of questions for a representative sample of mortgage holders in the United Kingdom. The survey measures financial literacy on a five-point scale and identifies

present bias through a self-reported measure of impulsive buying behavior. Since both characteristics are measured for the same individuals, the study can examine their relationship directly, rather than relying on educational attainment as a proxy for financial literacy. The study focuses mainly on the choice of what the industry refers to as alternative mortgage products, namely contracts with back-loaded payment structures, with interest-only mortgages being the main example. Both lower financial literacy and stronger present bias are associated with a higher probability of holding this type of contract. The effect of financial literacy is particularly large, given that a one-point decrease, on the five-point scale, increases the probability of holding an alternative mortgage product by approximately 53 percent. Instrumental variable estimates further suggest that this relationship is not driven by reverse causality. Present bias also remains strongly associated with the choice after controlling for patience, which the authors interpret as evidence that the mechanism is related to self-control rather than to conventional differences in time preferences. Households that place a relatively high weight on the immediate mortgage payment compared with future obligations may therefore find back-loaded contracts more attractive, even when their total cost is higher.

4.3.2 The distinction between product complexity and rate structure

A finding from the same study is directly relevant to one of the determinants discussed in chapter 3, although it points in a direction that remained unresolved in section 3.8. Financially literate individuals are more likely to choose an adjustable-rate mortgage, which the authors explain by arguing that they are better able to recognize the term premium embedded in the fixed rate and therefore less willing to pay for it.

Two aspects of this result make it particularly relevant. First, the finding is based on a direct measure of financial literacy rather than on educational attainment or a constructed measure of financial sophistication. It therefore addresses the measurement concern raised in section 3.8 regarding the evidence from the United States. Second, the result is also consistent with the positive coefficient reported by Coulibaly and Li (2009) using the Survey of Consumer Finances, despite being obtained in a different national mortgage market. This provides evidence that the relationship between financial literacy and adjustable-rate mortgage choice may reflect a broader pattern rather than an isolated finding from a single study.

The study also shows that the two behavioral mechanisms affect different aspects of mortgage choice. Present bias and low financial literacy increase the likelihood of choosing

back-loaded mortgage products, whereas financial literacy is associated with a greater likelihood of choosing an adjustable-rate contract. These findings concern different features of the mortgage contract and should therefore be considered separately. Treating them as part of the same decision would overlook the fact that the behavioral characteristics associated with these choices operate in opposite directions. This distinction is also consistent with the caveat introduced in section 3.1 regarding the binary framework adopted in this review. Alternative mortgage products represent an additional dimension of the contractual choice that is not captured by the fixed-rate versus adjustable-rate distinction, and the behavioral literature suggests that different characteristics may influence household decisions along this broader dimension.

4.3.3 Heuristics in the processing of macroeconomic information

A second line of evidence concerns not household preferences, but the way households process the information on which their expectations rest. Ilan and Mugerman (2025) examine this issue using a specific feature of the Israeli mortgage market, where borrowers can choose whether to index their mortgage to the consumer price index. The index is published in the middle of each month and refers to inflation observed in the previous month. As a result, borrowers who make their decision during the period between two publications base their choice on an inflation figure that no longer reflects the most recent conditions.

The authors find that observed inflation affects the indexation decision beyond what would be expected based on expected inflation alone. This effect is stronger among borrowers from lower socioeconomic status (SES) groups, which the authors employ as a proxy for lower financial literacy, who tend to place greater weight on the inflation rate currently observed rather than on the rate they expect in the future. The authors interpret this result in terms of the availability and representativeness heuristics, whereby recently observed information receives disproportionate weight in the decision-making process.

An important feature of the analysis is that it addresses the possibility that this behavior could still represent a rational response to uncertainty. The effect remains present during periods of low inflation, when uncertainty about future price movements is relatively limited and the precautionary motive for choosing indexation should therefore be weaker. The persistence of the effect under these conditions provides stronger support for a behavioral interpretation and makes it less likely that the observed pattern can be explained solely by a rational response to future inflation risk. The relationship between this result

and the experience-based explanation discussed in section 4.2 is one of complementarity rather than competition. Malmendier and Nagel (2016) focus on how individuals weigh inflation outcomes accumulated over several decades, whereas the heuristic mechanism discussed here gives greater importance to the most recent observation. Both mechanisms imply that household expectations may differ from the forecasts supported by available information, and both suggest that these deviations vary systematically across households. However, they operate over different time horizons and can therefore coexist rather than represent competing explanations.

Earlier work by Mugerman, Ofir and Wiener (2016) showed that borrowers deviate from the choices predicted by normative models in ways that are consistent with psychological explanations of decision-making under risk. More recent studies build on this evidence by moving beyond the identification of deviations from standard models and examining the specific behavioral mechanisms that may explain them.

4.3.4 What the literature has not established

A complete assessment of this literature also requires acknowledging what remains unexplored, since several behavioral biases that could plausibly influence mortgage decisions have not yet been directly examined in this context. For example, inertia and inattention have been widely documented in the refinancing literature, where households often fail to exercise valuable options. However, refinancing decisions concern whether to modify an existing mortgage contract, rather than which type of contract to choose at the time of origination. Therefore, evidence from refinancing behavior cannot be directly extended to initial mortgage choice without additional assumptions that have not been established. Similarly, loss aversion, overconfidence, and framing effects are frequently discussed as potentially relevant factors in mortgage decisions. Nevertheless, their role in this specific choice has not been demonstrated by the studies reviewed in this section. The theoretical plausibility of these mechanisms should therefore be distinguished from empirical evidence showing that they actually influence mortgage choice.

A further limitation concerns the geographical scope of the available evidence. The studies discussed above focus mainly on the United Kingdom and Israel, neither of which shares the institutional characteristics of the U.S. mortgage market, particularly the widespread availability of long-term fixed-rate contracts supported by a developed secondary market. Whether the same behavioral mechanisms operate with similar intensity in markets where the dominant mortgage product differs remains an open empirical question.

4.4 The predictive adequacy of the rational benchmark

The two previous sections have analyzed how households form their expectations and which systematic biases have been identified in their mortgage choices. The remaining question is to what extent the framework developed in chapter 1 is able to explain the decisions observed in practice, and which aspects require further refinement. The answer differs across the various dimensions of the decision, and the areas where the framework succeeds or fails provide useful insights into the factors shaping household mortgage choices.

4.4.1 Where the framework holds

The determinants confirmed in section 3.10 are those for which the rational framework provides accurate predictions. Households respond to the pricing difference between fixed and adjustable contracts, and this relationship has been consistently documented across survey data, administrative records, and international evidence over several decades. Expected mobility increases the likelihood of choosing an adjustable-rate contract, in line with the insurance mechanism proposed by the theory. Conversely, greater labor income volatility reduces this likelihood, while higher stated risk aversion also leads households away from adjustable contracts, both in accordance with the predictions of the lifecycle model. These findings retain particular significance. Contract pricing and the expected duration of the mortgage are the two main factors identified by the theoretical framework as driving household choice, and the empirical evidence supports both. An evaluation of the framework that ignored these results would therefore provide an incomplete representation of the evidence. The behavioral mechanisms discussed in this chapter should be interpreted as refinements that improve the explanatory power of the model, instead of a replacement for a framework that already captures a substantial part of observed mortgage choices.

4.4.2 The compression of the decision horizon

The determinants for which the framework does not provide a satisfactory explanation share a common feature that becomes apparent only when they are considered together. Each of these findings was initially reported as an isolated anomaly within the study in which it was identified.

Section 3.4 showed that liquidity-constrained households are more likely to choose adjustable-rate contracts, which is contrary to the prediction that financial constraints should

increase the demand for payment insurance. Section 3.7 showed that the cross-country evidence presented by Badarinza, Campbell and Ramadorai (2018) indicates that households respond more strongly to interest rate forecasts for the following year than to forecasts three years ahead. The evidence reviewed in the present chapter provides a third finding pointing in the same direction: present bias increases the probability of choosing contracts with deferred payments, independently of measured patience.

Three literatures, working on different data and addressing different questions, converge on a single conclusion. Households weigh the near-term installment more heavily than a framework concerned with the discounted cost of a thirty-year obligation would imply. The lifecycle model treats the mortgage as an instrument for smoothing consumption over the full life of the loan, whereas the empirical evidence suggests that household decisions are influenced more strongly by conditions closer to the time at which the mortgage is chosen. This finding does not require a general rejection of the theoretical framework, but rather a specific modification. The framework's basic account of which factors households should consider remains relevant; what needs to be reconsidered is the time horizon over which these factors are assumed to be evaluated.

4.4.3 Deviations in opposite directions

A second observation emerges when the behavioral findings are considered together, and it makes it difficult to describe household behavior as a single deviation from the benchmark model. The explanation proposed by Botsch and Malmendier (2023) accounts for an excess demand for fixed-rate mortgages. Households whose past inflation experience leads them to expect higher future interest rates are willing to pay for protection that they would not choose based solely on a neutral assessment of its cost. The authors estimate that almost one household in six holds a fixed-rate contract for this reason. The mechanisms discussed in sections 3.4 and 4.3 point in the opposite direction. Liquidity-constrained households and households with stronger present bias are more likely to choose adjustable-rate or back-loaded contracts, even when a full assessment of their costs would not necessarily justify this choice. The behavioral literature therefore does not suggest that households systematically deviate from the benchmark in the same direction. Instead, different groups of households appear to deviate for different reasons and in opposite directions. This result is closely related to the bimodal pattern documented in section 3.8, and the two findings support each other. Adjustable-rate borrowing is chosen by households at both ends of the financial sophistication distribution, while the behavioral

deviations from the benchmark also occur across different segments of the population and lead households toward different types of contracts. This has an important implication for empirical analysis. A coefficient estimated for the population as a whole combines these opposing effects, and its size depends not only on the strength of any individual mechanism but also on the relative importance of the different groups in the sample. The heterogeneity documented in section 3.10 is consistent with this interpretation and helps explain why estimates may differ across samples and periods even when the underlying behavioral mechanisms remain the same.

4.4.4 The aggregate puzzle

The clearest evidence that the framework underestimates demand for fixed-rate contracts comes from the composition of the U.S. mortgage market. Approximately 80 percent of outstanding mortgages have a fixed interest rate, even though fixed-rate contracts have historically been priced, on average, around 170 basis points above comparable adjustable-rate alternatives. Standard lifecycle models are unable to generate a market share of this magnitude at such a price difference.

The experience-based account addresses part of the gap. Households formed under the Great Inflation attach an inflated value to the protection embedded in the fixed contract, and the welfare cost of that valuation has been estimated between 8,000 and 19,000 dollars for the households whose choice is driven by this mechanism. The size of this distortion is therefore economically meaningful rather than negligible. However, this explanation does not fully account for the observed prevalence of fixed-rate borrowing. Institutional factors, such as those discussed by Miles (2004) in the context of the United Kingdom, also influence the mortgage market by affecting the structure of mortgage funding and lenders' willingness to bear interest rate risk. These supply-side considerations interact with household preferences in determining which contracts are available and at what price. Since none of the studies reviewed here separate the aggregate market share of fixed-rate mortgages into its demand and supply components, the extent to which the behavioral literature explains the overall pattern remains necessarily limited.

4.4.5 Implications for the interpretation of the estimates

Three consequences follow from the interpretation of the results reviewed in chapter 3. First, the estimated relationships should be understood as describing household behavior within a specific institutional and macroeconomic environment rather than as identifying

structural parameters that remain constant over time. If the relative size of the two groups of households that deviate from the benchmark changes with market conditions, and if the relative price of fixed and adjustable contracts affects which households choose each type of product, then estimates obtained in one period may have limited applicability to other periods. Second, the relationship between measurement quality and empirical performance discussed in section 3.10 becomes even more relevant in light of the behavioral evidence. The characteristics that are most difficult to observe, such as household expectations and the ability to evaluate mortgage contracts, are precisely those through which behavioral deviations arise. Improving their measurement would therefore allow researchers to capture relationships that current empirical specifications are unable to represent, rather than simply producing more precise estimates of existing ones. Finally, the unequal distribution of prediction errors highlighted by Campbell (2006) remains an important consideration. Deviations from the benchmark model are concentrated among households with specific characteristics, and these same characteristics are typically included as explanatory variables in empirical models. A specification that treats these variables as controls while assigning the behavioral deviations they predict to the error term may therefore be incorrectly specified in a way that affects the interpretation of the results.

4.4.6 Assessment

The rational framework provides a satisfactory explanation of how households respond to the relative price of the two contracts and to the expected length of time for which they will hold the mortgage. However, it provides a less complete explanation of the time horizon over which households evaluate their decisions, of how the expectations underlying their choices are formed, and of how the ability to assess mortgage alternatives differs across households.

The evidence therefore supports specific adjustments to the framework instead of a complete revision of its foundations. Incorporating a shorter effective decision horizon, allowing expectations to be shaped by personal experience rather than by common information alone, and recognizing that decision-making processes may vary with financial sophistication would help explain the documented deviations while preserving the elements of the model that successfully account for the main empirical relationships. Whether such an extended framework would improve predictions outside the samples used for estimation remains an open question. The literature reviewed in this chapter does not provide a definitive answer, and this issue is further examined in chapter 5.

5. Synthesis

5.1 Established regularities

The two previous chapters have examined the determinants of mortgage choice individually, together with the behavioral evidence that qualifies some of these relationships. This chapter brings these findings together, starting with the relationships that can be considered well established in the literature. A clear criterion is needed before reaching such a conclusion, since a finding reported in a single study, however precise, is not necessarily a result on which the literature has reached a consensus. The criterion adopted here has three components. A finding is considered established when it has been replicated using different types of data, including household surveys, administrative loan records, and aggregate series; when it remains valid across periods with different interest rate conditions; and when it does not depend on a specific model specification or on a particular way of measuring the variable. Findings that satisfy all three conditions are treated below as regularities, while those that satisfy only some of them are reported with the relevant qualifications.

5.1.1 The pricing channel

The relationship between mortgage contract choice and the relative price of the two products meets these criteria more clearly than any other determinant examined in this study. The relationship was documented using single-lender data from the 1980s by Dhillon, Shilling, and Sirmans (1987), at the market level during the same period by Brueckner and Follain (1988), using national survey data by Coulibaly and Li (2009), using a large administrative dataset by Vickery (2007), and across nine countries by Badarinza, Campbell, and Ramadorai (2018). The theoretical framework developed by Koijen, Van Hemert, and Van Nieuwerburgh (2009) further clarifies what borrowers are responding to by distinguishing the bond risk premium from the observed mortgage rate spread, but it does not change the predicted direction of the relationship. Across four decades, several types of data, and different institutional settings, the evidence consistently points in the same direction. Two qualifications should nevertheless be considered. First, household-level estimates measure the response to the observed difference between mortgage rates rather than directly to the premium identified by the theoretical framework as the relevant factor, and none of the studies reviewed here separates these two components using micro-level data. Second, the strength of the response varies across households. Coulibaly

Artificial intelligence tools were used to support the drafting, editing, and linguistic refinement of selected parts of this chapter.

and Li (2009) find a stronger response among more risk-averse households, while Smith, Finke, and Huston (2011) find that the relationship is more pronounced among financially sophisticated borrowers. The regularity therefore concerns the average response in the population, while the underlying relationship varies across households in ways that are not fully captured by the aggregate coefficient.

5.1.2 The horizon channel

The finding that households expecting to move are more likely to choose adjustable-rate contracts is the second robust relationship identified in the literature. It is supported by two measures that differ in how they are constructed. Coulibaly and Li (2009) use a direct survey question on the probability of remaining at the current address, which has been available since 1995, while earlier loan-level studies use length of residence as a backward-looking proxy for expected mobility. The fact that a subjective expectation and an observed household characteristic point in the same direction provides stronger support than replication based on a single measure. The descriptive evidence reported by Chiang and Dueholm (2024) is also consistent with this relationship, with a median age of 32.1 years among adjustable-rate mortgage holders compared with 50.3 years among fixed-rate mortgage holders. Mortgage maturity appears to belong to the same mechanism and shows a consistent relationship with contract choice, as shorter maturities are substantially more likely to be associated with fixed-rate mortgages. However, this finding does not fully satisfy the third part of the criterion. As discussed in section 3.6, the relationship may partly reflect the structure of mortgage supply, since adjustable-rate contracts have generally been offered with thirty-year maturities. The available data do not allow these supply-side and demand-side explanations to be clearly separated.

5.1.3 The risk channel

Volatile labor income reduces the likelihood of choosing an adjustable-rate contract, as predicted by the lifecycle framework. This relationship is reported by Coulibaly and Li (2009) using U.S. data and is also found by Ehrmann and Ziegelmeyer (2017) using survey data from the euro area, with no published study reviewed here reporting the opposite result. However, the evidence is based on a constructed measure rather than on a directly observed perception of income risk. The surveys do not ask households to assess their own income uncertainty, so the measure is instead inferred from the dispersion of earnings within occupational categories. The evidence for stated risk aversion is less

consistent and should therefore be presented with appropriate qualification. The negative relationship is found in U.S. data and replicated in the euro area survey, while Paiella and Pozzolo (2007) find no significant effect of their risk-attitude measures in the Italian data. In their analysis, mortgage choice is instead mainly related to the relative price of the contracts and to liquidity constraints. The available evidence does not allow it to be determined whether this difference reflects actual behavioral differences, institutional characteristics, or differences in how risk attitudes are measured. Risk aversion is therefore better described as a supported determinant in the U.S. and euro area evidence rather than as a fully established relationship.

5.1.4 The income gradient

The positive relationship between income and adjustable-rate borrowing has also been documented at different points over the past four decades, including in the studies of Brueckner and Follain (1988) and Coulibaly and Li (2009), as well as in the descriptive evidence reported by Chiang and Dueholm (2024). Its persistence is notable because it contrasts with the affordability mechanism discussed above, according to which financial constraints can push households toward adjustable-rate contracts. The explanation proposed in section 3.9, namely that adjustable-rate borrowing is concentrated at both ends of the income distribution for different reasons, is consistent with the non-monotonic relationship observed in more recent descriptive evidence. The regularity was until recently confined to the descriptive gradient, since income enters the American reference specification only through leverage ratios. The euro area evidence supplies the conditional counterpart: Ehrmann and Ziegelmeyer (2017) estimate income in quintiles of the national distribution and find households in the top quintile twelve percentage points more likely to hold an adjustable contract than those in the bottom, an effect significant at the one percent level and the largest among the household characteristics they examine. The determinant is therefore established in both the unconditional and the conditional form, though the two rest on different populations.

5.1.5 Established null results

Some of the findings that can be considered well established are negative, and this is also worth reporting. Marital status, family size, and race do not appear to predict mortgage contract choice once pricing, affordability, and planning horizon are taken into account. The authors of the reference study note that these results are consistent with earlier

evidence. As discussed in section 3.9, the absence of a significant effect of race does not imply that racial differences in mortgage contract choice do not exist. Rather, it suggests that these differences may operate through the other channels already included in the model. This distinction is important, but the reviewed literature has not examined it further.

5.1.6 Observability and measurements

When considered together, the findings that meet the criterion share a feature that was already identified in section 3.10 and is further supported by the behavioural evidence discussed in chapter 4. They mainly concern characteristics that can be observed without requiring households to report their own beliefs or understanding. The interest rate spread is a market price, the mortgage term is recorded in loan data, mobility is measured through a question referring to a concrete and observable outcome, and income volatility is estimated using occupational information. Stated risk aversion is the main exception, and it is also the determinant for which the international evidence is least consistent. A final qualification applies to all these findings. What the literature has established with some consistency is the direction of the relationships, rather than their magnitude. The estimates reviewed in chapter 3 are based on different samples, use variables constructed in different ways, and report coefficients on scales that are not directly comparable. Moreover, none of the studies considered here attempts to reconcile the size of the estimated effects across different contributions. The regularities identified in this chapter should therefore be interpreted primarily as evidence of the direction in which each determinant affects the probability of choosing an adjustable-rate contract, rather than as precise estimates of the size of these effects.

5.2 Contested findings and the sources of divergence

A finding may fail to be established for two different reasons. It may not have been studied often enough to support a clear conclusion, in which case the appropriate description is a lack of evidence. Alternatively, it may have been examined repeatedly, but the results may point in different directions, meaning that the available evidence does not provide a clear answer. This section focuses on the second case, while the first is discussed in section 5.3. Three determinants fall into this category, and a fourth sits at its edge. They do not fail in the same way, however, and the distinction is drawn before the sources of disagreement are examined. The estimated effect of liquidity constraints has the opposite sign

to that predicted by the lifecycle framework, and this is not an isolated finding. Coulibaly and Li (2009) report the same relationship using pooled waves of the U.S. survey, Johnson and Li (2014) obtain similar evidence by combining survey information with consumption data, and Paiella and Pozzolo (2007) find the same pattern using Italian data. In these settings, financially constrained households are more likely to choose adjustable-rate contracts because of their lower initial payments. The empirical relationship is therefore relatively consistent, while the unresolved issue concerns how it can be reconciled with the theoretical framework, which predicts the opposite effect. Age presents a different type of disagreement. The unconditional difference between the two groups of borrowers is substantial, with a median age of 32.1 years among adjustable-rate mortgage holders compared with 50.3 years among fixed-rate mortgage holders. Once other characteristics are controlled for, however, the estimated coefficient is only 0.01 and has little economic significance. Neither result is necessarily inconsistent with the data. The unresolved issue is instead which theoretical mechanism better explains the observed age pattern, namely the lifecycle mechanism or the experience-based mechanism discussed in the behavioral literature.

The relationship between financial literacy, education, and mortgage choice is also contested, both in terms of its sign and its interpretation. The coefficient in the reference specification is positive and equal to 0.40, while Gathergood and Weber (2017) find a positive relationship using British data and a direct measure of financial literacy rather than a proxy. However, Smith, Finke, and Huston (2011) show that adjustable-rate borrowing is concentrated at both ends of the financial sophistication distribution. A single coefficient may therefore combine different mechanisms that operate in opposite directions, making its interpretation more difficult. Stated risk aversion is at the boundary of this category, as discussed in section 5.1. The evidence from the United States and the euro area points in the same direction, while the Italian evidence does not. The determinant can therefore be considered supported in some settings, but the available evidence is not sufficiently consistent to treat it as an established regularity.

5.2.1 Two forms of contestation

The determinants discussed above are not contested in the same way, and distinguishing between two types of disagreement helps clarify what the remainder of this section can and cannot explain. The effect of liquidity constraints is not disputed among the studies that estimate it. Coulibaly and Li (2009) find a positive relationship using pooled U.S.

survey waves, Johnson and Li (2014) obtain similar evidence using consumption data, and Paiella and Pozzolo (2007) report the same pattern for Italy. Three studies, covering two countries and using three different types of data, therefore point in the same direction. What is in question is the theoretical prediction. The lifecycle framework implies that constrained households should value payment insurance more strongly and should therefore prefer fixed-rate contracts. The empirical result instead points towards adjustable-rate borrowing. A case of this kind calls for a modification of the theoretical framework rather than a comparison between conflicting empirical findings. As discussed in chapter 4, the evidence suggests that the relevant amendment is to place greater weight on the short-term affordability horizon in the mortgage decision. The effects of risk aversion and financial literacy are contested in a different and stricter sense, because the empirical studies do not reach the same conclusions. The negative relationship between risk aversion and adjustable-rate borrowing is found in the U.S. and euro area evidence but not in the Italian data. The coefficient on financial literacy is positive in the reference specification and in the British evidence reported by Gathergood and Weber (2017), while Smith, Finke, and Huston (2011) show that the overall relationship can combine mechanisms that operate in opposite directions. These disagreements cannot be resolved by modifying the theoretical framework, because the source of the disagreement lies in the empirical evidence rather than in the theoretical prediction. The age coefficient does not fit neatly into either category. The large unconditional difference between adjustable- and fixed-rate borrowers and the small conditional estimate are both valid and answer different questions, as discussed in section 3.9.3. The unresolved issue is therefore not which estimate is correct, but which of the two theoretical mechanisms, lifecycle considerations or experience-based expectations, provides the better explanation for the observed age pattern. The three sources examined below bear principally on the second category, in which the disagreement lies among the findings themselves. Measurement difficulties also affect the first, since a result may be unanimous and still be hard to interpret, and the discussion that follows notes where this is so.

5.2.2 Divergence arising from measurement

A second source of disagreement arises from the way the determinants are measured. This problem is more serious because it cannot be resolved simply by clarifying what each variable is intended to capture. Liquidity constraints provide the clearest example. The variable with by far the largest coefficient in the reference specification is an indicator of

subprime status. However, the survey does not directly record whether a mortgage is subprime, for this reason Coulibaly and Li (2009) constructed a measure that classifies loans as subprime when their interest rate is more than three percentage points above a comparable Treasury yield. The measure is therefore based on the mortgage rate, which is itself related to the contract choice being explained. The authors also note that more than two thirds of subprime mortgages were issued as adjustable-rate contracts. As a result, it is difficult to determine how much of the estimated effect reflects liquidity constraints and how much reflects the specific characteristics of the subprime mortgage segment. Similar measurement issues affect the other contested determinants. Income volatility is inferred from occupational categories, which means that households in the same occupation receive the same measure regardless of their individual earnings history. Financial literacy has no dedicated module in the survey, and the studies work either with educational attainment or with scales constructed from items collected for other purposes. Section 3.8 showed that estimates based on one measure of the attribute do not necessarily replicate when a different measure is used. This suggests that part of the disagreement in this area may be due to differences in measurement rather than differences in household behavior. The comparison with Gathergood and Weber (2017) is particularly informative. They use a tested financial literacy measure on a British sample and obtain an effect in the same direction as the U.S. estimates, suggesting that the use of constructed proxies may reduce the precision with which the relationship is identified. The findings of Bucks and Pence (2008) add another dimension to these measurement problems. Since knowledge of mortgage contract terms is positively related to income and education, measurement errors in variables associated with informed mortgage choice are unlikely to be randomly distributed across households. They are more likely to affect less affluent and less educated households. This type of measurement error can attenuate estimated coefficients, suggesting that the contested relationships involving financial literacy and liquidity constraints may be underestimated rather than overstated.

5.2.3 Divergence arising from population heterogeneity

The third source of disagreement is also the one that appears most frequently in this literature, but it has received the least systematic attention. Whenever researchers have examined differences across households, they have found meaningful heterogeneity. Responsiveness to the interest rate spread is stronger among more risk-averse households, the relationship between education and contract choice is stronger among less risk-averse

households, and expected mobility affects contract choice mainly among less risk-averse households. Meanwhile, adjustable-rate borrowing is concentrated at both ends of the financial sophistication distribution and shows a non-monotonic relationship with income. The finding of bimodality is particularly important because it changes the interpretation of estimated coefficients. If less sophisticated households choose adjustable-rate contracts because of affordability constraints, while more sophisticated households choose them to benefit from more favorable pricing, the estimated coefficient on sophistication combines two mechanisms that operate in opposite directions. Its value will therefore depend partly on the relative size of these groups in the sample, rather than only on the strength of either mechanism. As a result, two studies could obtain different coefficients even if the underlying behavior were the same, simply because their samples contained different proportions of the two groups. The behavioral evidence reviewed in chapter 4 supports and extends this interpretation. Botsch and Malmendier (2023) explain the relatively high demand for fixed-rate mortgages among households whose inflation experience leads them to place an unusually high value on interest rate protection, while the mechanisms related to liquidity constraints and present bias push households towards adjustable or back-loaded contracts. Deviations from the rational benchmark therefore do not appear to affect all households in the same way. Instead, they are concentrated in different segments of the population and can lead to choices in opposite directions, which is consistent with the bimodal pattern found in the empirical evidence.

5.2.4 Divergence arising from regime and institutional setting

The composition of households choosing each type of contract is not constant over time. When the ARM share is low, the households choosing such type of contract form a relatively small and self-selected group. When the share increases sharply because fixed rates become more expensive, the group becomes broader and includes households attracted by the lower initial payment. If the relative importance of these two groups changes with the interest rate environment, coefficients estimated in one period may not apply to another. This is particularly relevant because the studies reviewed in chapter 3 were conducted under substantially different market conditions. Institutional differences can create similar effects across countries. Koijen, Van Hemert, and Van Nieuwerburgh (2009) find that the inflation risk premium is particularly important in the United States, where fixed-rate mortgages generally have long maturities, while the real interest rate premium is more important in the United Kingdom, where rates are typically fixed for only one to three

years. Borrowers in each market therefore respond to the type of risk associated with the prevailing mortgage structure. Whether the different findings on risk aversion between the Italian and U.S. evidence can be explained by similar institutional differences has not yet been examined. A more recent instance concerns the experience channel. Botsch and Malmendier (2023) find that inflation experience raises the demand for fixed-rate contracts in the United States, while Cenzon and Szabó (2024) find that it lowers the probability of holding one across nine euro area countries, a reversal they attribute to prepayment penalties that make exit from a fixed contract costly. Two studies employing the same mechanism on comparable data reach opposite conclusions because the contracts available to their respective populations differ in a feature neither study treats as a variable.

5.2.5 Requirements for further assessment

The four sources of disagreement identified above cannot be resolved simply by estimating the same specifications again using the same variables. To clarify what each result captures, descriptive and conditional evidence should be reported together rather than treating one as a substitute for the other. Addressing the measurement problems would require variables specifically designed to capture the relevant concepts, particularly a dedicated financial literacy module and a direct measure of perceived income risk, neither of which is available in the Survey of Consumer Finances. Accounting for heterogeneity would require models that allow the relevant mechanisms to differ across households rather than imposing a single coefficient on the entire population. Finally, determining whether the estimated relationships change across market conditions would require comparable models to be estimated consistently across survey waves. No study reviewed here has carried out this type of analysis. These issues are therefore considered among the main priorities for future research in section 5.4.

5.3 Gaps in the evidence base

Section 5.2 examined questions that have been studied in the literature but remain unresolved. This section instead considers questions that have not yet been examined. The distinction is important because an issue that remains open due to conflicting evidence requires a different response from one that remains open simply because it has not yet been investigated. The research gaps identified below are grouped according to the main obstacle to addressing them. Some remain unresolved even though the necessary data are

already available, meaning that further analysis is the main requirement. Others concern concepts that existing data sources do not measure directly, so addressing them would require changes to the survey or the collection of new information. A third group is more diverse in nature but has a common origin: these are questions that the literature has not pursued because they lie between different research areas, because the necessary identification strategy has not been developed, or because established research practices have left certain aspects of the decision unobserved. This classification is more informative than simply listing the gaps, because it distinguishes between questions that could be answered through further analysis of existing evidence and those that require new data or research designs.

5.3.1 Gaps that existing data could close

The 2022 wave of the Survey of Consumer Finances has not yet been used for an econometric analysis of mortgage contract choice to the time of this review. It has mainly been used to describe household finances, while the most detailed published comparison of adjustable- and fixed-rate borrowers, by Chiang and Dueholm (2024), deliberately relies on the 2019 wave because it avoids the potential distortions associated with the pandemic period. This choice is reasonable for their purpose, but it leaves the wave conducted during the strongest monetary tightening in four decades largely unexplored in relation to the question considered in this review. The available data allow this analysis to be conducted. The expectations variable included in the survey have also not been used directly in a model of mortgage contract choice. This omission should be interpreted carefully, because the variable itself has been studied. Botsch and Malmendier (2023) compare responses across cohorts between 1989 and 2013 and show that younger and older households have systematically different expectations about the future direction of interest rates. This demonstrates both that the measure contains relevant information and that it can be linked to U.S. household-level data. Using Irish data, Devine, McCarthy, and O'Toole (2023) also show that expectations can be incorporated into models of mortgage choice when they are directly elicited for this purpose. What remains to be done is to combine these approaches by including the expectation measure already available in the Survey of Consumer Finances as an explanatory variable in a model of mortgage contract choice. Given the theoretical importance of expectations, discussed in section 1.3 and reviewed in section 3.7, this is arguably the most important omission identified in this section. The stability of the estimated relationships across different interest rate

environments has also not been tested. The studies reviewed in chapter 3 use data from the 1990s, the period surrounding the financial crisis, and the years before the pandemic, while section 5.2 identified several reasons why the relevant relationships may change with market conditions. Estimating the same specification across different survey waves, while keeping the variable definitions and model structure consistent, would make it possible to test whether the coefficients are stable over time. The necessary data are available for all Survey of Consumer Finances waves since 1995.

5.3.2 Gaps requiring improved measurement instrument

A second group of gaps cannot be addressed using the data currently available because the concepts identified by the theoretical framework are not directly measured by the existing surveys. Section 5.2 identified three main cases: the absence of a dedicated financial literacy module, the measurement of income risk through occupational categories rather than at the household level, and the lack of a decomposition of the bond risk premium using household-level data. The requirements for addressing these gaps differ, and it is therefore useful to consider them separately. Financial literacy is the most straightforward case, since validated measures already exist and have been applied to mortgage holders in other settings, as shown by Gathergood and Weber (2017) using British data. Including a similar module in an existing survey would mainly require changes to the questionnaire rather than a new methodological approach. This contrasts with more complex behavioral constructs, for which measurement remains a more difficult methodological issue. Present bias and self-control, for example, can be elicited through experimental or survey-based measures, as in Gathergood and Weber (2017), but there is less agreement on how these concepts should be measured consistently in large-scale household surveys and across different settings. More generally, the empirical measurement of complex cognitive biases and time preferences in real-world financial decisions remains less standardized than the measurement of financial literacy. This distinction is relevant because a lack of measurement does not necessarily indicate a lack of theoretical importance: in some cases, the main limitation is that the available surveys were not designed to capture the relevant behavioral construct. Measuring perceived income risk is more difficult because it requires households to report their expectations about the distribution of future income rather than an observable characteristic. No widely established measure for collecting this information in household surveys has yet emerged. Decomposing the bond risk premium at the household level is more demanding still, as it would

require linking each mortgage to the term structure prevailing when the loan was originated and then separating the expected future path of interest rates from the compensation required for bearing duration risk. Although this decomposition has been performed at the aggregate market level, it has not yet been carried out for individual mortgage contracts.

5.3.3 Gaps in the questions pursued

Some gaps in the literature cannot be explained by a lack of data or suitable measurement tools, because they concern questions that have not yet been clearly formulated. The absence of a significant effect of race provides a useful example. Section 3.9 showed that racial differences in mortgage contract choice disappear once subprime status is taken into account and noted that this may indicate that differences of this type operate through the channels already included in the model rather than that they do not exist. Whether these channels reflect differences in household preferences or differences in access to credit is an important distinction, but the studies reviewed here do not address it. Research on mortgage choice and research on disparities in mortgage lending have mostly been conducted as separate strands of literature, and their interaction has not been systematically examined. The relationship between mortgage maturity and interest rate structure raises a similar issue. Shorter maturities are substantially more likely to be associated with fixed-rate contracts, but adjustable-rate mortgages have generally been offered with thirty-year maturities only. The observed relationship may reflect the range of products offered by lenders rather than borrower preferences. Separating supply-side effects from demand-side preferences would require variation in the contracts offered by lenders that are independent of borrower demand, and no study reviewed here has attempted to identify such variation. The behavioral evidence is also more limited than the frequent use of behavioral explanations in discussions of mortgage markets might suggest. Chapter 4 showed that present bias and salience-based heuristics affect mortgage contract choice, but other mechanisms often discussed in the mortgage literature, such as loss aversion, overconfidence, and framing effects, have not been directly examined in this context. Inertia has been extensively documented in refinancing decisions, but refinancing concerns whether households act on an existing mortgage rather than which contract they choose at origination. The evidence from refinancing therefore cannot be directly applied to initial contract choice without further justification. Moreover, the behavioral evidence reviewed here comes mainly from the United Kingdom and Israel, whose mortgage markets differ from the U.S. market in important institutional respects. It therefore remains unclear

whether the same mechanisms operate with similar strength when the standard mortgage product is different. Finally, the literature has generally treated mortgage choice as a binary decision between fixed- and adjustable-rate contracts, and the present review has followed this convention for the reasons discussed in section 3.1. This approach leaves aside hybrid contracts, introductory rates below the fully indexed rate, and features such as interest rate caps and reset frequencies that determine the actual risk faced by borrowers with adjustable-rate mortgages. Section 4.3 showed that different behavioral characteristics may influence choices along these additional dimensions.

5.4 Implications for policy and further research

The three preceding sections established what the literature has settled, where it divides, and what it has not addressed. This one considers what follows for those who design the rules under which mortgages are sold, and for those who will take up the questions left open. A caution governs the first of these. A review of empirical findings supports policy inference only within limits, since the studies examined here estimate associations on observational data and were not designed to evaluate interventions. What follows is accordingly framed as implications of the evidence rather than as recommendations, and the distinction is not merely rhetorical, since several of the findings most relevant to policy are among those which section 5.2 identified as contested.

5.4.1 Disclosure and the limits of borrower comprehension

The finding of Bucks and Pence (2008) bears more directly on regulation than any other result reviewed in this work. Borrowers report the broad terms of their mortgages accurately, including whether the rate can change, but those holding adjustable contracts are markedly less aware of the magnitude of possible adjustments, of how often resets occur, and of the caps that limit them. Those least informed about these features are concentrated among lower-income households, who are also the most exposed should rates rise. The implication is that disclosure requirements which emphasize the distinction between the two contract types address a matter that borrowers already grasp, while the features determining what an adjustable contract will actually cost are precisely those the evidence identifies as poorly understood. The regulatory response that followed the financial crisis in the United States, comprising the ability-to-repay and qualified mortgage standards issued under the Dodd-Frank Act, operates through underwriting requirements rather than through comprehension, and section 2.2 noted its contribution to the contraction of the

products most associated with payment shocks. Whether comprehension itself improved is not established by any study reviewed here.

5.4.2 Financial education and uncertain returns

The case for financial education appears straightforward given the association between literacy and mortgage outcomes documented by Lusardi and Mitchell (2014), but the evidence reviewed in chapters 3 and 4 complicates it in three respects. The bimodality documented by Smith, Finke, and Huston (2011) implies that adjustable-rate borrowing draws on households at both extremes of the sophistication distribution, so that raising literacy would not move the population uniformly toward either contract. Gathergood and Weber (2017) point the same way from a different angle, finding that more financially literate borrowers are more likely to select the adjustable contract, declining to pay the premium embedded in the fixed rate. In this case, greater financial literacy would increase rather than reduce the use of a product that is often considered riskier. Their second finding qualifies the case further. Present bias raises the probability of selecting contracts that defer payment, and it does so independently of measured patience, which the authors interpret as evidence of a self-control problem rather than of a lack of financial knowledge. Education can improve what households know about a contract, but it does not necessarily change how they behave when the immediate payment receives greater weight than the eventual cost. The shorter decision horizon documented in section 4.4 also suggests that this behavioral margin may be relevant for a substantial share of households. Botsch and Malmendier (2023) introduce a complication of a different order. If expectations are shaped by the inflation a household has lived through, then the beliefs underlying mortgage choice cannot be easily changed through financial education, since past experience cannot simply be replaced by new information. Their estimates indicate that close to one household in six holds a fixed contract for this reason alone, at a cost of between eight and nineteen thousand dollars measured in year-2000 after-tax terms. None of these argues against financial education. It indicates that its effects on this particular decision are unlikely to be uniform in direction, and that a program evaluated on the assumption that greater literacy shifts borrowers toward fixed contracts may find the opposite.

5.4.3 Market composition as an object of design

The implications considered so far concern what borrowers understand. A separate group concerns what they are offered, and here the evidence indicates that the composition of a

national mortgage market is determined only in part by household preference. The pattern documented in section 3.5, whereby shorter maturities are markedly more likely to carry a fixed rate, illustrates the difficulty at contract level, since adjustable contracts have generally been offered on thirty-year terms alone. Fuster and Vickery (2015) establish the point at market level, finding the fixed-rate share sharply lower among mortgages that cannot readily be securitized, an effect they identify through a regulatory discontinuity and through the collapse of private securitization liquidity in 2007, and which they attribute to lenders' reluctance to retain the prepayment and interest rate risk embedded in long-term fixed contracts. The dominance of the fixed-rate instrument in the United States is therefore sustained by the arrangements that allow lenders to transfer that risk, and any policy altering those arrangements would alter the composition of the market irrespective of what households would otherwise prefer. The international evidence is consistent with this reading. Ehrmann and Ziegelmeyer (2017) document that the adjustable share varies substantially across euro area countries in ways that track national market structures rather than differences in what households want, and Koijen, Van Hemert, and Van Nieuwerburgh (2009) find that the risk premium households respond to differs between the United States and the United Kingdom according to the maturities their markets offer. The premise on which the Miles (2004) review proceeded in the United Kingdom, namely that the prevailing contract structure is an object of institutional design, receives support from both. One consequence runs in the opposite direction and concerns an instrument of policy rather than its object. De Stefani and Mano (2025) find, on a panel of some thirty-five countries, that a larger stock of outstanding fixed-rate contracts weakens the effect of monetary policy on the wider economy, since households shielded by a fixed rate do not experience rising rates immediately, while tightening in turn reduces the share of newly originated fixed contracts. Measures that affect the balance between the two contracts therefore have consequences extending well beyond the households whose choices they affect, and the regime-dependence discussed in section 5.2 acquires a further dimension, since the environment shaping household choice is itself partly shaped by the accumulated stock of choices already made.

5.4.4 Priorities for further research

The gaps set out in section 5.3 differ in the feasibility and potential contribution of addressing them, and the ordering below reflects the expected informational gain relative to the resources required.

Estimating a model of contract choice on the 2022 wave of the Survey of Consumer Finances, with the expectation item entered as a covariate, would address the two most consequential absences at once. The data exist, the variable is recorded, and the sample was fielded during a rate environment no household-level study has examined. The design would need to restrict attention to mortgages originated close to the interview date, following the practice adopted by Coulibaly and Li (2009), since an expectation recorded years after a contract was signed carries little information about the belief that informed it. Harmonized estimation across survey waves, with the specification and variable construction held constant, would allow testing whether the estimated coefficients vary systematically across interest-rate regimes. Section 5.2 identified several reasons to expect such regime dependence, making this question directly relevant to the interpretation and comparability of the estimates accumulated in the literature. The required data are available for all Survey of Consumer Finances waves since 1995. Specifications permitting the mechanism to differ across the sophistication distribution would address the difficulty the bimodality result poses, since estimating a single coefficient where two mechanisms of opposite sign operate produces a parameter whose value depends on sample composition. Interaction terms between constraint indicators and literacy proxies are the natural device, subject to the limitations of the proxies available. Adding a literacy module and a direct measure of perceived income risk to the survey would resolve at the source the measurement problems that section 5.2 identified as a principal cause of divergence. This is the most demanding of the priorities the least within researchers' direct control, but it is also the one with the greatest potential to expand the set of questions that the literature can address rather than merely refine existing estimates. Two further directions warrant mention without being ranked alongside the preceding four, since what stands in their way differs in nature. Examining the intersection between contract choice and disparities in mortgage lending requires the integration of two literatures that have developed in parallel, and the main challenge lies in the conceptual framing rather than in data availability. Extending the analysis beyond the binary distinction, to the hybrid structures and adjustment features that determine the risk actually borne, requires abandoning a convention adopted throughout this field for reasons of measurement, and the behavioral evidence reviewed in section 4.3 indicates that different traits govern choices along those additional dimensions.

6. Conclusions

This work sets out to assess what the empirical literature has learned about how households choose between fixed-rate and adjustable-rate mortgages, and where that literature remains incomplete. The question is not merely academic. The mortgage is the largest financial obligation most households ever assume, and the two contracts distribute risk between borrower and lender in ways that carry consequences for consumption, for financial distress, and for the transmission of monetary policy. The theoretical apparatus assembled in chapter 1 offered two complementary perspectives on this decision. The lifecycle model of Campbell and Cocco (2003) treats the choice as one between two forms of risk, since the nominal fixed contract carries a real capital value sensitive to inflation while the adjustable contract transfers the variability of short-term payments to a household whose borrowing constraint binds precisely when income is low. Koijen, Van Hemert, and Van Nieuwerburgh (2009) hold the household fixed and let the price of the fixed contract vary, showing that the premium the borrower pays for payment certainty fluctuates with the state of the term structure. The first framework therefore helps to explain which households choose which contract at a given moment, whilst the second explains why the aggregate composition of mortgage contracts changes over time. Both were shown to rest on assumptions the subsequent chapters examined against evidence. The review conducted in chapter 3 established that the literature has settled a smaller set of questions than the volume of published work might suggest. Households respond to the relative price of the two contracts, and this finding holds across four decades, across survey, administrative, and aggregate data, and across several national markets. Those expecting to move choose the adjustable contract more often, and volatile labor income pushes in the opposite direction. A positive association between income and adjustable borrowing has been documented at intervals since the 1980s. Beyond these, the record is less secure than the frequency of citation implies. What proved most informative was not the list of established determinants itself, but the characteristic they have in common. Each concerns a variable that can be observed without requiring households to report on their own beliefs or understanding. The price spread is a market variable, the mortgage term is recorded in the loan data, and expected mobility is based on a question with a clear and concrete reference. The determinants that remain contested are those requiring subjective report, namely the experience of credit constraint, the capacity to evaluate a contract, and expectations about the future path of rates. The distribution of empirical

Artificial intelligence tools were used to support the drafting, editing, and linguistic refinement of selected parts of this chapter.

success across this literature tracks the availability of reliable measurement rather than the plausibility of the underlying theory, and this observation, developed in section 3.10, is the conclusion of the present work that seems to me most worth carrying forward. Chapter 4 examined the behavioral evidence and found that it qualifies the rational framework in a specific way rather than displacing it. Three results obtained on different data converge on a single amendment. The constraint coefficient reverses the predicted sign, households respond to interest rate forecasts one year ahead far more strongly than to forecasts three years ahead, and present bias increases the probability of selecting contracts that defer payment independently of measured patience. Each suggests that the decision is anchored considerably nearer to the present than a framework concerned with the discounted cost of a thirty-year obligation would imply. The framework therefore remains useful for identifying the factors households should consider, but the time horizon over which these factors are evaluated requires revision. A second observation from that chapter complicates any simple characterization of household error. The experience-based account explains an excess of fixed-rate borrowing among households whose inflation history leads them to overvalue protection, while the constraint and present-bias mechanisms push in the opposite direction. Departures from the rational benchmark are not common to the population but concentrated in distinct segments moving in opposite directions, which is the same structure the bimodality in financial sophistication displays. A coefficient estimated across the whole distribution averages these opposing deviations, and its magnitude depends on their relative weight in the sample rather than on the strength of any single mechanism. Several limitations of this review should be acknowledged. The survey is selective rather than exhaustive, confined to work available in English, and organized around a dataset whose distinctive content determined which determinants could be examined at all. The classification of disagreements proposed in section 5.2, and the research priorities proposed in section 5.4, depend on the literature included in the review. No established synthesis of this literature exists against which the organization adopted here could be validated, a circumstance already noted in section 3.1. The review also followed the convention of treating mortgage choice as binary, setting aside the hybrid structures, introductory rates, and adjustment provisions that determine the risk a borrower actually bears, and the behavioral evidence suggests that different traits govern choices along those additional dimensions. The most consequential absence identified in this work concerns the determinant reflected in its title. The Survey of Consumer Finances has recorded households' expectations about the future direction of interest rates

since the late 1980s, and Botsch and Malmendier (2023) draw on that series to document systematic divergence between cohorts. No study has entered the measure as a covariate in a model of contract choice. The wave conducted in 2022, during the sharpest monetary tightening in four decades, has served for descriptive portraits of family finances but not for an analysis of how contract choice relates to the attitudes and expectations the survey uniquely records. The relevant data and variables are therefore already available, placing this question among the gaps that could be addressed using existing evidence rather than requiring new measurement instruments. The persistence of this gap in a literature that spans four decades is itself informative about how the field has developed. Mortgage choice has been studied predominantly as a response to prices and to observable household characteristics, while less attention has been given to the beliefs about an uncertain future that may also shape the decision. The theoretical frameworks reviewed in chapter 1 place those beliefs at the center of the choice, but the empirical literature has not yet fully incorporated them into the analysis. This provides a clear direction for future research and, at the same time, highlights the main unresolved question identified by this review.

References

Aladangady, A., Bricker, J., Chang, A. C., Goodman, S., Krimmel, J., Moore, K. B., Reber, S., Henriques Volz, A., & Windle, R. (2023). *Changes in U.S. family finances from 2019 to 2022: Evidence from the Survey of Consumer Finances*. Board of Governors of the Federal Reserve System. https://doi.org/10.17016/8799

Avery, R. B., Bhutta, N., Brevoort, K. P., Canner, G. B. & Gibbs, C. N. (2010). The 2008 HMDA data: The mortgage market during a turbulent year. *Federal Reserve Bulletin, 96*, A169–A211. https://doi.org/10.17016/bulletin.2010.96-4

Bacon, P. M., & Moffatt, P. G. (2012). Mortgage choice as a natural field experiment on choice under risk. *Journal of Money, Credit and Banking, 44*(7), 1401–1426. https://doi.org/10.1111/j.1538-4616.2012.00537.x

Badarinza, C., Campbell, J. Y., & Ramadorai, T. (2018). What calls to ARMs? International evidence on interest rates and the choice of adjustable-rate mortgages. *Management Science, 64*(5), 2275–2288. https://doi.org/10.1287/mnsc.2016.2629

Board of Governors of the Federal Reserve System (US). (2026). *Federal funds effective rate* (FEDFUNDS) [Data set]. FRED, Federal Reserve Bank of St. Louis. Retrieved 14 August 2026, from https://fred.stlouisfed.org/series/FEDFUNDS

Board of Governors of the Federal Reserve System. (2022, July 8). *Declining affordability and home purchase borrowing by lower income households*. FEDS Notes. https://doi.org/10.17016/2380-7172.3160

Board of Governors of the Federal Reserve System. (2023). *Codebook for 2022 Survey of Consumer Finances*. https://www.federalreserve.gov/econres/files/codebk2022.txt

Botsch, M. J., & Malmendier, U. (2023). *The long shadows of the Great Inflation: Evidence from residential mortgages* [Working paper]. SSRN. https://doi.org/10.2139/ssrn.3888762

Brueckner, J. K., & Follain, J. R. (1988). The rise and fall of the ARM: An econometric analysis of mortgage choice. *The Review of Economics and Statistics, 70*(1), 93–102. https://doi.org/10.2307/1928154

Bucks, B., & Pence, K. (2008). Do borrowers know their mortgage terms? *Journal of Urban Economics, 64*(2), 218–233. https://doi.org/10.2139/ssrn.6551978

Campbell, J. Y. (2006). Household finance. *The Journal of Finance, 61*(4), 1553–1604. https://doi.org/10.1111/j.1540-6261.2006.00883.x

Campbell, J. Y., & Cocco, J. F. (2003). Household risk management and optimal mortgage choice. *The Quarterly Journal of Economics, 118*(4), 1449–1494. https://doi.org/10.2139/ssrn.301325

Campbell, J. Y., & Cocco, J. F. (2015). A model of mortgage default. *The Journal of Finance, 70*(4), 1495–1554. https://doi.org/10.1111/jofi.12252

Cenzon, J., & Szabó, B. E. (2024). Mortgage choice and inflation experiences in the Eurozone. *Journal of Monetary Economics, 147*, Article 103611. https://doi.org/10.1016/j.jmoneco.2024.103611

Chiang, Y.-T., & Dueholm, M. (2024, February 6). *Which households prefer ARMs vs. fixed-rate mortgages?* Federal Reserve Bank of St. Louis, On the Economy. https://www.stlouisfed.org/on-the-economy/2024/feb/which-households-prefer-arms-fixed-rate-mortgages

Cocco, J. F. (2005). Portfolio choice in the presence of housing. *The Review of Financial Studies, 18*(2), 535–567. https://doi.org/10.2139/ssrn.258428

Consumer Financial Protection Bureau. (2013, January 30). Ability-to-repay and qualified mortgage standards under the Truth in Lending Act (Regulation Z). Final rule. *Federal Register, 78*, 6408.

Coulibaly, B., & Li, G. (2009). Choice of mortgage contracts: Evidence from the Survey of Consumer Finances. *Real Estate Economics, 37*(4), 659–673. https://doi.org/10.2139/ssrn.1034286

De Stefani, A., & Mano, R. C. (2025). *Long-term debt and short-term rates: Fixed-rate mortgages and monetary transmission* (IMF Working Paper No. 2025/024). International Monetary Fund. https://doi.org/10.5089/9798400296086.001

Devine, K., McCarthy, Y., & O'Toole, C. (2023). The role of borrower expectations in mortgage choice. *Journal of Behavioral and Experimental Economics, 107*, Article 102104. https://doi.org/10.2139/ssrn.4113816

Dhillon, U. S., Shilling, J. D., & Sirmans, C. F. (1987). Choosing between fixed and adjustable rate mortgages: A note. *Journal of Money, Credit and Banking, 19*(2), 260–267. https://doi.org/10.2307/1992281

Dodd-Frank Wall Street Reform and Consumer Protection Act, Pub. L. No. 111-203, 124 Stat. 1376 (2010).

Ehrmann, M., & Ziegelmeyer, M. (2017). Mortgage choice in the euro area: Macroeconomic determinants and the effect of monetary policy on debt burdens. *Journal of Money, Credit and Banking, 49*(2–3), 469–494. https://doi.org/10.1111/jmcb.12386

European Central Bank. (2023). Consumers' interest rate expectations in a monetary policy tightening cycle. *Economic Bulletin*, (5), Box 2.

Federal Reserve Bank of New York, Center for Microeconomic Data. (2023). *Quarterly report on household debt and credit, second quarter 2023*. Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/interactives/household-credit/data/pdf/HHDC_2023Q2

Freddie Mac. (2026). *30-year fixed rate mortgage average in the United States* (MORTGAGE30US) [Data set]. FRED, Federal Reserve Bank of St. Louis. Retrieved 14 August 2026, from https://fred.stlouisfed.org/series/MORTGAGE30US

Fuster, A., & Vickery, J. (2015). Securitization and the fixed-rate mortgage. *The Review of Financial Studies, 28*(1), 176–211. https://doi.org/10.1093/rfs/hhu060

Gathergood, J., & Weber, J. (2017). Financial literacy, present bias and alternative mortgage products. *Journal of Banking & Finance, 78*, 58–83. https://doi.org/10.1016/j.jbankfin.2017.01.022

Huston, S. J., Finke, M. S., & Smith, H. (2012). A financial sophistication proxy for the Survey of Consumer Finances. *Applied Economics Letters, 19*(13), 1275–1278. https://doi.org/10.1080/13504851.2011.619485

Ilan, Y., & Mugerman, Y. (2025). Misguided mortgage choices. *Journal of Behavioral and Experimental Finance, 47*, Article 101077. https://doi.org/10.1016/j.jbef.2025.101077

Johnson, K. W., & Li, G. (2014). Are adjustable-rate mortgage borrowers borrowing constrained? *Real Estate Economics, 42*(2), 457–471. https://doi.org/10.1111/1540-6229.12033

Kim, K. T., Hanna, S. D., & Ying, D. (2021). The risk tolerance measure in the 2016 Survey of Consumer Finances: New, but is it improved? *Journal of Financial Counseling and Planning, 32*(1), 86–103. https://doi.org/10.1891/JFCP-19-00022

Koijen, R. S. J., Van Hemert, O., & Van Nieuwerburgh, S. (2009). Mortgage timing. *Journal of Financial Economics, 93*(2), 292–324. https://doi.org/10.1016/j.jfineco.2008.09.005

Lessard, D. R., & Modigliani, F. (1975). Inflation and the housing market: Problems and potential solutions. In *New mortgage designs for stable housing in an inflationary environment* (Conference Series No. 14, pp. 13–45). Federal Reserve Bank of Boston.

Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature, 52*(1), 5–44. https://doi.org/10.1257/jel.52.1.5

Malmendier, U., & Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics, 131*(1), 53–87. https://doi.org/10.1093/qje/qjv037

Mayer, C., & Pence, K. (2008). *Subprime mortgages: What, where, and to whom?* (NBER Working Paper No. 14083). National Bureau of Economic Research. https://doi.org/10.3386/w14083

Miles, D. (2004). The UK Mortgage Market: Taking a Longer-Term View. Final report and recommendations.

Mugerman, Y., Ofir, M., & Wiener, Z. (2016). How do homeowners choose between fixed and adjustable rate mortgages? *Quarterly Journal of Finance, 6*(4), Article 1650012. https://doi.org/10.1142/s2010139216500130

Paiella, M., & Pozzolo, A. F. (2007). Choosing between fixed- and adjustable-rate mortgages. In S. Agarwal & B. W. Ambrose (Eds.), *Household credit usage: Personal debt and mortgages* (pp. 219–236). Palgrave Macmillan. https://doi.org/10.1057/9780230608917_13

Sa-Aadu, J., & Sirmans, C. F. (1995). Differentiated contracts, heterogeneous borrowers, and the mortgage choice decision. *Journal of Money, Credit and Banking, 27*(2), 498–510. https://doi.org/10.2307/2077880

Smith, H. L., Finke, M. S., & Huston, S. J. (2011). The impact of financial sophistication on adjustable-rate mortgage ownership. *Journal of Financial Counseling and Planning, 22*(2), 3–15.

Urban Institute, Housing Finance Policy Center. (2023). *Housing finance at a glance: A monthly chartbook, December 2023*. Urban Institute. https://www.urban.org/research/publication/housing-finance-glance-monthly-chartbook-december-2023

Vickery, J. (2007). *Interest rates and consumer choice in the residential mortgage market* [Working paper]. Federal Reserve Bank of New York.