

Mortgage Timing*

Ralph S.J. Koijen Otto Van Hemert Stijn Van Nieuwerburgh

September 5, 2008

Abstract

We study how the term structure of interest rates relates to mortgage choice, both at the household and the aggregate level. A simple utility framework of mortgage choice points to the long-term bond risk premium as theoretical determinant: when the bond risk premium is high, fixed-rate mortgage payments are high, making adjustable-rate mortgages more attractive. This long-term bond risk premium is markedly different from other term structure variables that have been proposed, including the yield spread and the long yield. We confirm empirically that the bulk of the time variation in both aggregate and loan-level mortgage choice can be explained by time variation in the bond risk premium. This is true whether bond risk premia are measured using forecasters' data, a VAR term structure model, or from a simple household decision rule based on adaptive expectations. This simple rule moves in lock-step with mortgage choice, lending credibility to a theory of strategic mortgage timing by households.

*First draft: November 15, 2006. Koijen: ralph.koijen@chicagogsb.edu; The University of Chicago Graduate School of Business; 5807 South Woodlawn Avenue Chicago, IL 60637. Van Hemert: ovanheme@stern.nyu.edu; Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012. Van Nieuwerburgh: svnieuw@stern.nyu.edu; Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; Tel: (212) 998-0673. Koijen is also associated with Netspar, Van Hemert with AQR, and Van Nieuwerburgh with the N.B.E.R. The authors would like to thank an anonymous referee, Yakov Amihud, Sandro Andrade, Andrew Ang, Jules van Binsbergen, Michael Brandt, Alon Brav, Markus Brunnermeier, John Campbell, Jennifer Carpenter, Michael Chernov, Albert Chun, João Cocco, John Cochrane, Thomas Davidoff, Joost Driessen, Gregory Duffee, Darrell Duffie, John Graham, Andrea Heuson, Vasso Ioannidou, Dwight Jaffee, Ron Kaniel, Anthony Lynch, Theo Nijman, Chris Mayer, Frank Nothaft, François Ortalo-Magné, Lasse Pedersen, Ludovic Phalippou, Amiyatosh Purnanandam, Adriano Rampini, Matthew Richardson, David Robinson, Anthony Saunders, Anjolein Schmeits, Walter Torous, Ross Valkanov, James Vickery, Annette Vissing-Jorgensen, Nancy Wallace, Bas Werker, Jeff Wurgler, Alex Ziegler, Stan Zin, and seminar participants at CMU, the University of Amsterdam, Princeton, USC, NYU, UC Berkeley, the St.-Louis Fed, Duke, Florida State, UMW, the Philadelphia Fed, the AREUEA Mid-Year Meeting in DC, the NYC real estate meeting, the Summer Real Estate Symposium in Big Sky, the Portfolio Theory conference in Toronto, the Asian Finance Association meeting in Chengdu, the Behavioral Finance conference in Singapore, the NBER Summer Institute Asset Pricing meeting in Cambridge, the CEPR Financial Markets conference in Gerzensee, and the EFA conference in Ljubljana for comments. The authors gratefully acknowledge financial support from the FDIC's Center for Financial Research and the Glucksman Institute.

This is a preprint version of the article. The final version may be found at
< <https://doi.org/10.1016/j.jfineco.2008.09.005> >.

One of the most important financial decisions any household has to make during its lifetime is whether to own a house and, if so, how to finance it. There are two broad categories of housing finance: adjustable-rate mortgages (ARMs) and fixed-rate mortgages (FRMs). The share of newly-originated mortgages that is of the ARM-type in the US economy shows a surprisingly large variation. It varies between 10% and 70% of all mortgages over our sample period from January 1985 to June 2006. In this paper, we seek to understand these fluctuations in the ARM share.

The main contribution of our paper is to understand the link from the term structure of interest rates to both individual and aggregate mortgage choice. While various term structure variables, such as the yield spread and the long-term yield (e.g., Campbell and Cocco (2003)), have been proposed before, the literature lacks a theory that predicts the precise link between the term structure and mortgage choice. A simple utility framework allows us to show that the long-term bond risk premium is the key determinant. This is the premium earned on investing long in a long-term bond and rolling over a short position in short-term bonds. The premium arises whenever the expectations hypothesis of the term structure of interest rates fails to hold, a fact for which there is abundant empirical evidence by now. We are the first to propose the bond risk premium as a predictor of mortgage choice and to document its strong predictive ability. We show that the long-term bond risk premium is conceptually and empirically very different from both the yield spread and the long yield. Because both of these variables are imperfect proxies for the long-term bond risk premium, they are imperfect predictors of mortgage choice.

What makes the bond risk premium a palatable determinant of observed household mortgage choice? Imagine a household that has to choose between an FRM and an ARM to finance its house purchase. With an FRM, mortgage payments are constant and linked to the long-term interest rate at the time of origination. With an ARM, matters are more complicated: future ARM payments will depend on future short-term interest rates not known at origination. We imagine that the household uses an average of short-term interest rates from the recent past in order to estimate future ARM payments. Under such expectations-formation rule, the difference between the long-term interest rate and the recent average of short-term interest rates is what the household would use to make the choice between the FRM and the ARM. Therefore, we label this difference the *household's decision rule*. The theoretical long-term bond risk premium that follows from our model is the -closely related- difference between the current long yield and the average expected future short yields over the contract period. The household decision rule is a proxy for the bond risk premium which arises when adaptive expectations are formed. Our motivation for this approximation is a suspicion that households may not have the required financial sophistication to solve complex investment problems (Campbell (2006)). The household decision rule is easy to compute, conceptually intuitive, and theoretically-founded.

This simple rule is highly effective at choosing the right mortgage at the right time. Section 1

shows that it has a correlation of 81% with the observed ARM share in the aggregate time series. We also use a new, nation-wide, loan-level data set that allows us to link the household decision rule to several hundred thousand individual mortgage choices. We find that it alone classifies 70% of mortgage loans correctly. The marginal impact of the household decision rule is essentially unaffected once we control for loan-level characteristics and geographic variables. In fact, the rule is an economically more significant predictor of individual mortgage choice than various individual-specific measures of financial constraints. The loan-level data reiterate the problem with the yield spread and the long yield as predictors of mortgage choice.

Section 2 presents our model; its novel feature is allowing for time variation in bond risk premia. The model is kept deliberately simple, as in Campbell (2006), and strips out some of the rich life-cycle dynamics modeled elsewhere.¹ It models risk-averse households who trade off the expected payments on an FRM and an ARM contract with the risk of these payments. The ARM payments are subject to real interest rate risk, while the presence of inflation uncertainty makes the real FRM payments risky. The model generates an intuitive risk-return trade-off for mortgage choice: the ARM contract is more desirable the higher the nominal bond risk premium, the lower the variability of the real rate, and the higher the variability of expected inflation. We explicitly aggregate the mortgage choice across households that are heterogeneous in risk preferences. Time variation in the aggregate ARM share is then caused by time variation in the bond risk premium. The mean and dispersion parameters of the cross-sectional distribution of risk aversion map one-to-one into the average ARM share and its sensitivity to the bond risk premium, respectively. The model also helps us understand the problem with the yield spread and long yield as predictors of mortgage choice. The yield spread is a noisy proxy for the long-term bond risk premium because average expected future short rates could differ from the current short rate. This creates an errors-in-variables problem in the regression of the ARM share on the yield spread. The problem is so severe in the data that the yield spread is effectively uninformative about the future ARM share. Intuitively, the yield spread fails to take into account that future ARM payments will adjust whenever the short rate changes. A similar, though empirically less pronounced, errors-in-variables problem occurs for the long yield.

In Section 3, we bring the theory to the data, and regress the ARM share on the nominal bond risk premium. We first show formally that the household decision rule arises as a measure of the bond risk premium when expectations of future nominal short rates are computed with an adaptive expectations scheme. This provides the theoretical underpinning for the empirical success of the household decision rule in predicting mortgage choice. The simple proxy for the bond risk premium explains about 70% of the variation in the ARM share. We also explore more academically conventional ways of measuring expected future short rates: based on Blue Chip

¹For instance, Campbell and Cocco (2003), Cocco (2005), Yao and Zhang (2005), and Van Hemert (2007).

forecasters' data and based on a vector auto-regression model of the term structure. These two forward-looking bond risk premia measures generate the same quantitative sensitivity of the ARM share: a one standard deviation increase in the bond risk premium leads to an 8% increase in the ARM share. This is a large economic effect given the average ARM share of 28%.

While the forward-looking measures of the bond risk premium deliver similar results to the household decision rule over the full sample, their performance diverges in the last ten years of the sample. This is mostly due to the increase in the ARM share in 2003-04, which is predicted correctly by the simple rule, but not by the other two forward-looking measures of the bond risk premium. Section 4 explains this divergence. Part of the explanation lies in product innovation in the ARM mortgage segment. But most of the divergence is due to large forecast errors in future short rates in this episode. This motivates us to consider the *inflation* risk premium component of the nominal risk premium, for which any forecast error that is common to nominal and real rates cancels out. We construct the inflation risk premium using real yield (TIPS) data and either Blue Chip forecasters' data or a VAR model for inflation expectations, and show that both measures have a strong positive correlation with the ARM share and deliver a similar economic effect.

In Section 5, we investigate the robustness of our results. First, we analyze the impact of the prepayment option, typically embedded in US FRM contracts, on the utility difference between the ARM and FRM. We show that the prepayment option reduces the exposures to the underlying risk factors. However, it continues to hold that the higher the bond risk premium, the more attractive is the ARM. The presence of the option does not materially alter the results. Second, we investigate to what extent aggregate variation in the ARM share is driven by variation in compositional shifts between U.S. states or income groups, rather than by common variation induced by the bond risk premium. We find no evidence for such composition effects. The bond risk premium continues to be a robust predictor in state-level or income group-level fixed-effect regressions of the ARM share. Third, we study the role of financial constraints using mostly loan-level data. The loan-level data allow us to investigate the importance of the loan-to-value ratio, the credit score, and the loan balance for the relative desirability of the ARM. While they are statistically significant predictors of mortgage choice, they do not add much to the explanatory power of the bond risk premium, nor significantly reduce it. In the context of financial constraints, we also investigate the role of short investment horizons as captured by a high rate of impatience or a high moving probability in a dynamic version of our model. When households are so impatient or have such high moving probability that they only care about the first mortgage payment, the yield spread fully captures the FRM-ARM tradeoff. For realistic values for moving rates or rates of time preference, the bond risk premium is the relevant determinant. Fourth, we study the incentives that mortgage lenders may have to recommend one type of mortgage versus the other and ask whether they predict the ARM share. We study the FRM-ARM rate spread, purged from a Treasury spread,

a prepayment risk premium, and a default risk premium, as a direct measure of incentives. We find that it does not forecast the ARM share. We also study various indirect channels, but end up concluding that supply-side effects are unlikely to drive mortgage choice. Fifth, we investigate to what extent the bond risk premium measure picks up time variation in bond return volatility, which the households may mistake for bond risk. We find that it does not. Instead, adaptive expectations measures of the bond risk premium are associated with future excess bond returns. Sixth, the same household decision rule that predicts purchase-money mortgage choice also predicts individual and aggregate mortgage refinancing activity. Furthermore, when the household decision rule is high, indicating a high bond risk premium, households tend to invest more in bonds. Both findings suggest a wider usefulness of the household decision rule. Seventh, we discuss the robustness of the statistical inference, and conduct a bootstrap exercise to calculate standard errors. Finally, we discuss liquidity issues in the TIPS markets and how they may affect our results on the inflation risk premium. We conclude that the bond risk premium is a robust determinant of mortgage choice.

Our findings resonate with recent work in the portfolio literature by Campbell, Chan, and Viceira (2003), Sangvinatsos and Wachter (2005), Brandt and Santa-Clara (2006), and Koijen, Nijman, and Werker (2007). This literature emphasizes that forming portfolios that take into account time-varying risk premia can substantially improve performance for long-term investors.² Because the mortgage is a key component of the typical household's portfolio, and because an ARM exposes that portfolio to different interest rate risk than an FRM, choosing the wrong mortgage may have adverse welfare consequences (Campbell and Cocco (2003) and Van Hemert (2007)). In contrast to these studies, our exercise suggests that mortgage choice is an important financial decision where the use of bond risk premia is not only valuable from a normative point of view. Time variation in risk premia is also important from a positive point of view, to explain observed variation in mortgage choice both at the aggregate and at the household level.

Finally, our paper also relates to the corporate finance literature on the timing of capital structure decisions. The firm's problem of maturity choice of debt is akin to the household's choice between an ARM and an FRM. Baker, Greenwood, and Wurgler (2003) show that firms are able to time bond markets. The maturity of debt decreases in periods of high bond risk premia.³ Our findings suggest that households also have the ability to incorporate information on bond risk premia in their long-term financing decision.

²Campbell and Viceira (2001) and Brennan and Xia (2002) derive the optimal portfolio strategy for long-term investors in the presence of stochastic real interest rates and inflation, but assume risk premia to be constant.

³See Butler, Grullon, and Weston (2006) and Baker, Taliaferro, and Wurgler (2006) for a recent discussion. In ongoing work, Greenwood and Vayanos (2007) study the the relationship between government bond supply and excess bond returns.

1 A Simple Story for Household Mortgage Choice

We imagine a household that is choosing between a standard fixed-rate and a standard adjustable-rate mortgage contract. On the FRM contract, it will pay a fixed, long-term interest rate while the rate on the ARM contract will reset periodically depending on the short-term interest rate. The household knows the current long-term interest rate, but lacks a sophisticated model for predicting future short-term interest rates. Instead, it naively forms an average of the short rate over the recent past as a proxy of what it expects to pay on the ARM.⁴ The relative attractiveness of the ARM contract is the difference between the current long rate and the average short rate over the recent past. We label this difference at time t the *household decision rule* κ_t .

Figure 1 displays the time series of the share of newly-originated mortgages that is of the ARM type (solid line, left axis) alongside the household decision rule $\kappa_t(3, 5)$ (dashed line, right axis). The latter is formed using the 5-year Treasury bond yield (indicated by the second argument) and the 1-year Treasury bill yield averaged over the past three years (indicated by the first argument). The ARM share is from the Federal Housing Financing Board, the standard source in the literature. Appendix A discusses the data in more detail and compares it to other available series. The figure documents a striking co-movement between the ARM share and the decision rule; their correlation is 81%. In Section 3 below, we present similar evidence from a regression analysis.

[Figure 1 about here.]

Figure 2 shows that this high correlation not only holds when the household decision rule is formed using Treasury interest rates (left panel), but also using mortgage interest rates (right panel). In both panels the household decision rule κ has the strongest association with the ARM share (highest bar) for intermediate values of the horizon over which average short rates are computed. The correlation is hump-shaped in the look-back horizon.

[Figure 2 about here.]

We not only find such high correlation between the household decision rule and the ARM share in aggregate time series data, but also in individual loan-level data. We explore a new data set which contains information on 911,000 loans from a large mortgage trustee for mortgage-backed security special purpose vehicles. The loans were issued between 1994 and 2007.⁵ Table 1 reports loan-level results of probit regressions with an ARM dummy as left-hand side variable. All right-hand side variables have been scaled by their standard deviation. We report the coefficient estimate,

⁴We will later refer to this as adaptive expectation formation. There is a related literature on extrapolative expectations in behavioral finance, e.g. Andreassen and Krauss (1990), DeBondt (1993), and Choi (2006).

⁵Appendix A provides more detail. We thank Nancy Wallace for graciously making these data available to us.

a robust t-statistic, and the fraction of loans that is correctly classified by the probit model.⁶ We keep the 654,368 loans for which we have all variables of interest available. The first row shows that the household decision rule is a strong predictor of loan-level mortgage choice. It has the right sign, a t-statistic of 253, and it -alone- classifies 69.4% of loans correctly. Its coefficient indicates that a one standard deviation increase in the bond risk premium increases the probability of an ARM choice from 39% to 56%, an increase of more than one-third.

It is interesting to contrast this result with a similar probit regression that has three well-documented indicators of financial constraints on the right-hand side: the loan balance at origination (BAL), the credit score of the borrower (FICO), and the loan-to-value ratio (LTV). The second row, which also includes four regional dummies for the biggest mortgage markets (California, Florida, New York, and Texas), confirms that a lower balance, a lower FICO score, and especially a higher LTV ratio increase the probability of choosing an ARM. However, the (scaled) coefficients on the loan characteristics are smaller than the coefficient on the household decision rule κ , suggesting a smaller economic effect. Furthermore, the three financial constraint variables classify only 59.0% of loans correctly; adding four state dummies increases correct classifications to 61.7%. Adding the three financial constraint proxies and the four regional dummies to the household decision rule does not increase the probability of classified loans (Row 3). The number of classified loans is 68.8%, no bigger than what is explained by κ alone.⁷ Moreover, the household decision rule variable remains the largest and by far the most significant regressor. Its marginal effect on the probability of choosing an ARM is unaffected.

[Table 1 about here.]

The rest of the paper is devoted to understanding why the simple decision rule works. We argue that it is a good proxy for the bond risk premium. The next section develops a rational model of mortgage choice that links time variation in the bond risk premium to time variation in the ARM share. While households might not have the required financial sophistication to solve complex investment problems (Campbell (2006)), the near-optimality of the simple decision rule suggests that close-to-rational mortgage decision making may well be within reach.⁸

The bond risk premium is not to be confused with the yield spread, which is the difference between the current long yield and the current short yield. To illustrate this distinction, the household decision rule in Figure 1 has a correlation of -25% with the 5-1 year yield spread. While κ had a correlation with the ARM share of 81%, the correlation between the yield spread and the

⁶By pure chance, one would classify 50% of the contracts correctly.

⁷Note that the maximum likelihood estimation does not maximize correct classifications, so that adding regressors does not necessarily increase correct classifications.

⁸One branch of the real estate finance literature documents slow prepayment behavior (e.g., Schwartz and Torous (1989)). Brunnermeier and Julliard (2006) study the effect of money illusion on house prices, and Gabaix, Krishnamurthy, and Vigneron (2006) study limits to arbitrage in mortgage-backed securities markets.

aggregate ARM share is -6% over the same sample. This correlation is indicated by the solid line in the left panel of Figure 2. The correlation with the mortgage rate spread, indicated by the solid line in the right panel, is somewhat higher at 33%. However, it remains substantially below the 81% of the simple rule with mortgage rates. The long yield also has a much lower correlation with the ARM share than the household decision rule (dashed lines). The second role of the model is to help clarify the distinction between the bond risk premium and the yield spread or long yield.

2 Model with Time-Varying Bond Risk Premia

Various term structure variables have been suggested in the literature to predict aggregate mortgage choice, such as the yield spread and yields of various maturities.⁹ The question of which term structure variable is the best predictor of individual and aggregate mortgage choice motivates us to set up a model that explores this link. Rather than developing a full-fledged life-cycle model, we study a tractable two-period model that allows us to focus solely on the role of time variation in bond risk premia. This extension of Campbell (2006) is motivated by the empirical evidence pointing to the failure of the expectations hypothesis in US post-war data.¹⁰ We first explore an individual household's choice between a fixed-rate mortgage and an adjustable-rate mortgage (Sections 2.1-2.4). Subsequently, we aggregate mortgage choices across households to link the term structure dynamics to the ARM share (Section 2.6). The model sheds light on the difference between the bond risk premium, the yield spread, and the long yield in Section 2.5. Finally, Section 2.7 discusses extensions of the model and the relationship with the literature.

2.1 Setup

We consider a continuum of households on the unit interval, indexed by j . Households are identical, except in their attitudes toward risk parameterized by γ_j . The cumulative distribution function of risk aversion coefficients is denoted by $F(\gamma)$.

At time 0, households purchase a house and use a mortgage to finance it. The house has a nominal value $H_t^\$$ at time t . For simplicity, the loan is non-amortizing. We assume a loan-to-value ratio equal to 100%, so that the mortgage balance is given by $B = H_0^\$$. The investment horizon and the maturity of the mortgage contract equal 2 periods. Interest payments on the mortgage are made at times 1 and 2. At time $t = 2$, the household sells the house at a price $H_2^\$$ and pays down the mortgage. The household chooses to finance the house using either an ARM or an FRM, with

⁹For instance, Berkovec, Kogut, and Nothhaft (2001), Campbell and Cocco (2003), and Vickery (2007b).

¹⁰Fama and French (1989), Campbell and Shiller (1991), Dai and Singleton (2002), Buraschi and Jiltsov (2005), Ang and Piazzesi (2003), Cochrane and Piazzesi (2005), and Ang, Bekaert, and Wei (2006), among others, document and study time variation in bond risk premia.

associated nominal interest rates q^i , $i \in \{ARM, FRM\}$. In each period, the household receives nominal income $L_t^\$$.

We postulate that the household is borrowing constrained: In each period, she consumes what is left over from the income she receives after making the mortgage payment (equation (2)). Because the constrained household cannot invest in the bond market, she cannot undo the position taken in the mortgage market. Terminal consumption equals income after the mortgage payment plus the difference between the value of the house and the mortgage balance (equation (3)).

Each household maximizes lifetime utility over real consumption streams $\{C/\Pi\}$, where Π is the price index and $\Pi_0 = 1$. Preferences in (1) are of the CARA type with risk aversion parameter γ_j , except for a log transformation. The subjective time discount factor is $\exp(-\beta)$.¹¹ The maximization problem of household j reads:

$$\max_{i \in \{ARM, FRM\}} - \log \left(\mathbb{E}_0 \left[e^{-\beta - \gamma_j \frac{C_1}{\Pi_1}} \right] \right) - \log \left(\mathbb{E}_0 \left[e^{-2\beta - \gamma_j \frac{C_2}{\Pi_2}} \right] \right) \quad (1)$$

$$\text{s.t. } C_1 = L_1^\$ - q_1^i B, \quad (2)$$

$$C_2 = L_2^\$ - q_2^i B + H_2^\$ - B. \quad (3)$$

We assume that real labor income, $L_t = L_t^\$/\Pi_t$, is stochastic and persistent:

$$L_{t+1} = \mu_L + \rho_L (L_t - \mu_L) + \sigma_L \varepsilon_{t+1}^L, \varepsilon_{t+1}^L \sim \mathcal{N}(0, 1).$$

In addition, we assume that the real house value is constant and let $H_t = H_t^\$/\Pi_t$.

2.2 Bond Pricing

The one-period nominal short rate at time t , $y_t^\$(1)$, is the sum of the real rate, $y_t(1)$, and expected inflation, x_t :

$$y_t^\$(1) = y_t(1) + x_t. \quad (4)$$

Denote the corresponding price of the one-period nominal bond by $P_t^\$(1)$. Following Campbell and Cocco (2003), we assume that realized inflation and expected inflation coincide:

$$\pi_{t+1} = \log \Pi_{t+1} - \log \Pi_t = x_t, \quad (5)$$

¹¹This log transformation is reminiscent of an Epstein and Zin (1989) aggregator which introduces a small preference for early resolution of uncertainty (see also Van Nieuwerburgh and Veldkamp (2007)). While this modification is solely made for analytical convenience, it implies that β does not affect mortgage choice. In Section 5.3, we investigate the role of the subjective discount rate in a calibrated, multi-period model with CRRA preferences. We show that the risk-return tradeoff which governs mortgage choice is unaffected for conventional values of β . The same conclusion holds when we introduce a realistic moving rate.

so that there is no unexpected inflation risk.¹² To accommodate the persistence in the real rate and expected inflation, we model both processes to be first-order autoregressive:

$$\begin{aligned} y_{t+1}(1) &= \mu_y + \rho_y (y_t(1) - \mu_y) + \sigma_y \varepsilon_{t+1}^y, \\ x_{t+1} &= \mu_x + \rho_x (x_t - \mu_x) + \sigma_x \varepsilon_{t+1}^x. \end{aligned}$$

Their innovations are jointly Gaussian with correlation matrix R :

$$\begin{pmatrix} \varepsilon_{t+1}^y \\ \varepsilon_{t+1}^x \end{pmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho_{xy} \\ \rho_{xy} & 1 \end{bmatrix} \right) = \mathcal{N}(0_{2 \times 1}, R).$$

We assume that labor income risk is uncorrelated with real rate and expected inflation innovations.

This structure delivers a familiar conditionally Gaussian term structure model. The important innovation in this model relative to the literature on mortgage choice is that the market prices of risk λ_t are *time-varying*. The nominal pricing kernel $M^{\$}$ takes the form:

$$\log M_{t+1}^{\$} = -y_t^{\$(1)} - \frac{1}{2} \lambda_t' R \lambda_t - \lambda_t' \varepsilon_{t+1},$$

with $\varepsilon_{t+1} = [\varepsilon_{t+1}^y, \varepsilon_{t+1}^x]'$ and $\lambda_t = [\lambda_t^y, \lambda_t^x]'$. If we were to restrict the prices of risk to be affine, our model would fall in the class of affine term structure models (see Dai and Singleton (2000)), but no such restriction is necessary.

The no-arbitrage price of a two-period zero-coupon bond is:

$$e^{-2y_0^{\$(2)}} = \mathbb{E}_0 [M_{t+1}^{\$} M_{t+2}^{\$}] = e^{-y_0^{\$(1)} - \mathbb{E}_0(y_1^{\$(1)}) + \lambda_0' R \sigma + \frac{1}{2} \sigma' R \sigma},$$

with $\sigma = [\sigma_y, \sigma_x]'$. This equation implies that the long rate equals the average expected future short rate plus a time-varying nominal bond risk premium $\phi^{\$}$:

$$y_0^{\$(2)} = \frac{y_0^{\$(1)} + \mathbb{E}_0(y_1^{\$(1)})}{2} - \frac{\lambda_0' R \sigma}{2} - \frac{1}{4} \sigma' R \sigma = \frac{y_0^{\$(1)} + \mathbb{E}_0(y_1^{\$(1)})}{2} + \phi_0^{\$(2)}. \quad (6)$$

The long-term nominal bond risk premium $\phi_0^{\$(2)}$ contains the market price of risk λ_0 and absorbs the Jensen correction term.

¹²Brennan and Xia (2002) show that the utility costs induced by incompleteness of the financial market due to unexpected inflation are small. In a previous version of this paper, we have done a numerical, multi-period mortgage choice analysis. We found that unexpected inflation risk did not affect the household's risk-return tradeoff in any meaningful way.

2.3 Mortgage Pricing

A competitive fringe of mortgage lenders prices ARM and FRM contracts to maximize profit, taking as given the term structure of Treasury interest rates generated by $M^\$$.

Denote the ARM rate at time t by q_t^{ARM} . This is the rate applied to the mortgage payment due in period $t + 1$. In each period, the zero-profit condition for the ARM rate satisfies:

$$B = \mathbb{E}_t [M_{t+1}^\$ (q_t^{ARM} + 1) B] = (q_t^{ARM} + 1) B P_t^\$(1).$$

This implies that the ARM rate is equal to the one-period nominal short rate, up to an approximation:

$$q_t^{ARM} = P_t^\$(1)^{-1} - 1 \simeq y_t^\$(1).$$

Similarly, the zero-profit condition for the FRM contract stipulates that the present discounted value of the FRM payments must equal the initial loan balance:

$$B = \mathbb{E}_0 [M_1^\$ q_0^{FRM} B + M_1^\$ M_2^\$ q_0^{FRM} B + M_1^\$ M_2^\$ B] = q_0^{FRM} P_0^\$(1) B + [q_0^{FRM} + 1] P_0^\$(2) B.$$

Per definition, the nominal interest rate on the FRM is fixed for the duration of the contract. We abstract from the prepayment option for now, but examine its role in Section 5.1. The FRM rate, which is a two-period coupon-bearing bond yield, is then equal to:

$$q_0^{FRM} = \frac{1 - P_0^\$(2)}{P_0^\$(1) + P_0^\$(2)} \simeq \frac{2y_0^\$(2)}{2 - y_0^\$(1) - 2y_0^\$(2)} \simeq y_0^\$(2).$$

The FRM rate is approximately equal to the two-period nominal bond rate.

Our setup embeds two assumptions that merit discussion. The first assumption is that the stochastic discount factor $M^\$$ that prices the term structure of interest rates is different from the inter-temporal marginal rate of substitution of the households in section 2.1. Without this assumption, mortgage choice would be indeterminate.¹³ The second assumption is that we price mortgages as derivatives contracts on the Treasury yield curve. Hence, the same sources that drive time variation in the Treasury yield curve will govern time variation in mortgage rates.

¹³Any equilibrium model of the mortgage market requires a second group of unconstrained investors. Time variation in risk premia could then arise from time-varying risk-sharing opportunities between the constrained and the unconstrained agents, as in Lustig and Van Nieuwerburgh (2006). In their model, the unconstrained agents price the assets at each date and state. Such an environment justifies taking bond prices as given when studying the problem of the constrained investors. Lustig and Van Nieuwerburgh (2006) consider agents with (identical) CRRA preferences. In numerical work, presented in Appendix D, we verify that the same risk-return tradeoff that the constrained households face also hold for CRRA preferences. A full-fledged equilibrium analysis of the mortgage market is beyond the scope of the current paper.

2.4 A Household's Mortgage Choice

We now derive the optimal mortgage choice for the household of Section 2.1. The crucial difference between an FRM investor and an ARM investor is that the former knows the value of all nominal mortgage payments at time 0, while the latter knows the value of the nominal payments only one period in advance. The risk-averse investor trades off lower expected payments on the ARM against higher variability of the payments. Appendix B computes the life-time utility under the ARM and the FRM contract. It shows that household j prefers the ARM contract over the FRM contract if and only if

$$\begin{aligned} & q_0^{FRM} - q_0^{ARM} + (q_0^{FRM} - \mathbb{E}_0 [q_1^{ARM}]) e^{-\mathbb{E}_0[x_1]} > \\ & \frac{\gamma_j}{2} B e^{-x_0 - 2\mathbb{E}_0[x_1]} \left[\sigma' R \sigma + (\mathbb{E}_0 [q_1^{ARM}] + 1)^2 \sigma_x^2 - 2 (\mathbb{E}_0 [q_1^{ARM}] + 1) (\sigma_x e'_2 R \sigma) \right] \\ & - \frac{\gamma_j}{2} B e^{-x_0 - 2\mathbb{E}_0[x_1]} (q_0^{FRM} + 1)^2 \sigma_x^2. \end{aligned} \quad (7)$$

The left-hand side measures the difference in expected payments on the FRM and the ARM. All else equal, a household prefers an ARM when the expected payments on the FRM are higher than those on the ARM. Appendix B shows that the difference between the expected mortgage payments on the FRM and ARM contracts approximately equals the two-period bond risk premium $\phi_0^{\$}(2)$. This leads to the main empirical prediction of the model: the ARM contract becomes more attractive in periods in which the bond risk premium is high.

The right-hand side of (7) measures the risk in the payments, where we recall that γ_j controls risk aversion. The first line arises from the variability of the ARM payments, the second line represents the variability of the FRM payments. All else equal, a risk-averse household prefers the ARM when the payments on the ARM are less variable than those on the FRM. The risk in the FRM contract is inflation risk (σ_x^2). The balance and the interest payments erode with inflation. The risk in the ARM contract consists of three terms. ARMs are risky because the nominal contract rate adjusts to the nominal short rate each period. The variance of the nominal short rate is $\sigma' R \sigma$. The second term is expected inflation risk, which enters in the same form as in the FRM contract. However, inflation risk is offset by the third term which arises from the positive covariance between expected inflation and the nominal short rate ($\sigma_x e'_2 R \sigma$). In low inflation states the mortgage balance erodes only slowly, but the low nominal short rates and ARM payments provide a hedge. The appendix shows that the risk in the ARM is approximately equal to the variability of the real rate (σ_y^2). In sum, the risk-return tradeoff of household j in (7), for some generic period t , can be written concisely as:

$$\phi_t^{\$}(2) - \frac{\gamma_j}{2} B \sigma_y^2 + \frac{\gamma_j}{2} B \sigma_x^2 > 0. \quad (8)$$

2.5 Yield Spread and Long Yield are Poor Proxies

We are the first to suggest the long-term bond risk premium as the determinant of household's mortgage choice. It is the risk premium that is earned on investing in a nominal long-term bond and financing this investment by rolling over a short position in a nominal short-term bond.¹⁴ It is important to emphasize that the long-term bond risk premium is markedly different from both the yield spread and the long-term yield, both of which have been used in the literature to predict mortgage choice.

Using equation (6), the difference between the long yield (on the two-period bond) and the short yield (on the one-period bond) can be written as

$$y_0^{\$}(2) - y_0^{\$}(1) = \phi_0^{\$}(2) + \frac{\mathbb{E}_0(y_1^{\$}(1)) - y_0^{\$}(1)}{2}. \quad (9)$$

The multi-period equivalent for some generic date t and generic maturity τ is

$$y_t^{\$}(\tau) - y_t^{\$}(1) = \phi_t^{\$}(\tau) + \left(\frac{1}{\tau} \sum_{j=1}^{\tau} \mathbb{E}_t [y_{t+j-1}^{\$}(1)] - y_t^{\$}(1) \right). \quad (10)$$

In both expressions, the second term on the right introduces an errors-in-variables problem when the yield spread is used as a proxy for the long-term bond risk premium $\phi_0^{\$}(2)$. This errors-in-variables problem turns out to be so severe that the yield spread has no predictive power for mortgage choice. To understand this further, consider two stark cases. First, in a homoscedastic world with zero risk premia ($\phi_t^{\$}(\tau) = 0$), the yield spread equals the difference between the average expected future short rates and the current short rate. Since long-term bond rates are the average of current and expected future short rates, both the FRM and the ARM investor face the same expected payment stream. The yield spread is completely uninformative about mortgage choice. Second, in a world with constant risk premia, *variations* in the yield spread capture *variations* in deviations between expected future short rates and the current short rate. But again, these variations are priced into both the ARM and the FRM contract. It is only the bond risk premium which affects the mortgage choice for a risk-averse investor. The problem with the yield spread as a measure of the relative desirability of the ARM contract is intuitive: The current short yield is not a good measure for the expected payments on an ARM contract because the short rate exhibits mean reversion which changes expected future payments.

¹⁴The strategy holds a τ -period bond until maturity and finances it by rolling over the 1-year bond for τ periods. This definition is different from the one-period bond risk premium in which the long-term bond is held for one period only. Cochrane and Piazzesi (2006) study various definitions of bond risk premia, including ours.

The long yield suffers from a similar errors-in-variables problem:

$$y_0^{\$}(2) = \phi_0^{\$}(2) + \frac{y_0^{\$}(1) + \mathbb{E}_0(y_1^{\$}(1))}{2}, \quad (11)$$

where the second term on the right again introduces noise in the predictor of mortgage choice. The problem with the long yield as a measure of the relative desirability of the ARM contract is intuitive: It contains no information on the difference in expected payments between the two contracts. In conclusion, our simple rational mortgage model suggests that both the yield spread and the long-term yield are imperfect predictors of mortgage choice.

2.6 Aggregate Mortgage Choice

We aggregate the individual households' mortgage choices to arrive at the ARM share. Define the cutoff risk aversion coefficient that makes a household indifferent between the ARM and FRM contract by:

$$\gamma_t^* \equiv \frac{2\phi_t^{\$}(2)}{B(\sigma_y^2 - \sigma_x^2)}.$$

Empirically, we find that $(\sigma_y^2 - \sigma_x^2) > 0$, which guarantees a positive value for the cutoff γ_t^* . Households that are relatively risk tolerant, with $\gamma_j < \gamma_t^*$, prefer the ARM contract. Because F is the cumulative density function of the risk-aversion distribution, the ARM share is given by:

$$ARM_t \equiv F(\gamma_t^*).$$

The complementary fraction of (more risk-averse) households chooses the FRM. The location parameter of the distribution of risk aversion determines the unconditional level of the ARM share. The scale parameter of this distribution governs the sensitivity of aggregate mortgage choice to changes in the bond risk premium. If risk preferences are highly dispersed, the ARM share will be insensitive to changes in the bond risk premium. Conversely, if heterogeneity across households is limited, small changes in the bond risk premium induce large shifts in the ARM share. Hence, the model provides a mapping between the (reduced-form) coefficients of a regression of the ARM share on a constant and the nominal bond risk premium and the two structural parameters that govern the cross-sectional distribution of risk aversion (see also Appendix D).

2.7 Alternative Determinants of Mortgage Choice

Our stylized model of mortgage choice abstracts from several real-life features that are potentially important. Several such features would be straightforward to add to our model, for example

stochastic real house prices, a temporary and a permanent component in labor income, and a more general correlation structure between real rate and expected inflation innovations on the one hand and labor income and house prices on the other hand. We could also extend the model to allow for saving in one-period bonds. For realism, we would then impose borrowing constraints along the lines of the life-cycle literature (Cocco, Gomes, and Maenhout (2005)). The models of Campbell and Cocco (2003) and Van Hemert (2007) allow for such features -and more- in the context of a life-cycle model. Campbell and Cocco (2003) show that households with a large mortgage, risky labor income, high risk aversion, a high cost of default, and a low probability of moving are more likely to prefer an FRM contract. In both studies, bond risk premia are assumed to be constant. Our model's sole purpose is understand the link between the term structure of interest rates and both individual and aggregate mortgage choice. We find that the long-term bond risk premium, and not the yield spread or the long yield, is the key determinant of mortgage choice. This is the hypothesis we test empirically in Section 3.

3 Empirical Results

The main task to render the theory testable is to measure the nominal bond risk premium. The latter is the difference between the current nominal long interest rate and the average expected future nominal short rate (see (6)):

$$\phi_t^{\$}(\tau) = y_t^{\$}(\tau) - \frac{1}{\tau} \sum_{j=1}^{\tau} \mathbb{E}_t [y_{t+j-1}^{\$}(1)]. \quad (12)$$

The difficulty resides in measuring the second term on the right, average expected future short rates.

3.1 Household Decision Rule

If we assume that households measure expected future short rates by forming simple averages of past short rates, we arrive at the household decision rule $\kappa_t(\rho; \tau)$ of Section 1:

$$\phi_t^{\$}(\tau) \simeq y_t^{\$}(\tau) - \frac{1}{\tau} \sum_{j=1}^{\tau} \left\{ \frac{1}{\rho} \sum_{u=0}^{\rho-1} y_{t-u}^{\$}(1) \right\} = y_t^{\$}(\tau) - \frac{1}{\rho} \sum_{u=0}^{\rho-1} y_{t-u}^{\$}(1) \equiv \kappa_t(\rho; \tau). \quad (13)$$

Equation (13) is a model of *adaptive expectations* that only requires knowledge of the current long bond rate, a history of recent short rates (ρ months), and the ability to calculate a simple average. The adaptive expectations scheme delivers a simple proxy $\kappa_t(\rho; \tau)$ for the theoretical bond risk premium $\phi_t^{\$}(\tau)$. Panel A of Figure 3 shows the $\tau = 5$ - and $\tau = 10$ -year time series with a three

year look-back, and computed off Treasury interest rates. The two series have a correlation of 92%.¹⁵

[Figure 3 about here.]

In Section 1 we showed that the household decision rule has a high correlation with the ARM share. We now extend that analysis by regressing the ARM share on the nominal bond risk premium. We lag the predictor variable for one month in order to study what changes in this month's risk premium imply for next month's mortgage choice. In addition, the use of lagged regressors mitigates potential endogeneity problems that would arise if mortgage choice affected the term structure of interest rates.¹⁶ The first two rows of Table 2 shows the slope coefficient, its Newey-West t-statistic using 12 lags, and the regression R^2 for these regressions. Throughout the table, the regressors are normalized by their standard deviation for ease of interpretation. They reinforce the point made in Section 1 that the household decision rule is a highly significant predictor of the ARM share. The 5-year (10-year) bond risk premium proxy has a t-statistic of 7.1 (7.5) and explains 71% (68%) of the variation in the ARM share. A one-standard deviation increase in the risk premium increases the ARM share by 7-8 percentage points. This is a large effect since the average ARM share is 28.7%. Intuitively, an FRM holder has to pay the bond risk premium. An increase in the risk premium increases the expected payments on the FRM relative to the ARM, and makes the ARM more attractive.

[Table 2 about here.]

3.2 Forward-Looking Measures

The household decision rule is a proxy for the theoretical bond risk premium when an adaptive expectations scheme is used to form the conditional expectation in equation (12). From an academic point of view, there are more conventional ways of measuring average expected future short rates. We study two below: one based on forecasters' expectations and one based on a VAR model.

¹⁵Since we consider look-back periods of up to 5 years, we lose the first 5 years of observations, and the series start in 1989.12. This is the same sample as used in Figures 1 and 2. We do not extend the sample before 1985.1 for two reasons. First, the interest rates in the early 1980s were dramatically different from those in the period we analyze. As such, we do not consider it to be plausible that households use adaptive expectations and data from the "Volcker regime" to form κ in the first years of our sample. A second and related reason is that Butler, Grullon, and Weston (2006) argue that there is a structural break in bond risk premia in the early 1980s. To avoid any spurious results due to structural breaks, we restrict attention to the period 1985.1-2006.6.

¹⁶As a robustness check, we have tested for Granger causality. First, we regress the ARM share on its own lag and the lagged bond risk premium; the lagged bond risk premium is statistically significant. Second, we regress the bond risk premium on its lag and the lagged ARM share; the lagged ARM share is statistically insignificant. Therefore, the bond risk premium Granger causes the ARM share, but the reverse is not true.

3.2.1 Forecaster Data

Our forecaster data come from Blue Chip Economic Indicators. Twice per year (March and October), a panel of around 40 forecasters predict the average three-month T-bill rate for the next calendar year, and each of the following four calendar years. They also forecast the average T-bill rate over the ensuing five years. We average the consensus forecast data over the first five, or all ten, years to construct the expected future nominal short rate in (12). This delivers a semi-annual time series from 1985 until 2006 for $\tau = 5$ and one for $\tau = 10$. We use linear interpolation of the forecasts to construct monthly series.¹⁷ Combining the 5-year (10-year) T-bond yield with the 5-year (10-year) expected future short rate from Blue Chip delivers the 5-year (10-year) nominal bond risk premium. Panel B of Figure 3 shows the 5-year (solid line) and 10-year time series (dashed line); they have a correlation of 94%. We then regress the ARM share on the nominal bond risk premium. The 5-year bond risk premium is a highly significant predictor of the ARM share (Row 3). It has a t-statistic of 3.9, and explains 40% of the variation in the ARM share. A one-standard deviation, or one percentage point, increase in the nominal bond risk premium increases the ARM share by 8.6 percentage points. The results with the 10-year risk premium (Row 4) are comparable. The coefficient has a similar magnitude, a t-statistic of 4.2, and an R^2 of 43%.

3.2.2 VAR Model

A second way to form the forward-looking conditional expectation in equation (12) is to use a vector auto-regressive (VAR) term structure model, as in Ang and Piazzesi (2003). The state vector Y contains the 1-year ($y_t^{\$}(1)$), the 5-year ($y_t^{\$}(5)$), and the 10-year nominal yields ($y_t^{\$}(10)$), as well as realized 1-year log inflation ($\pi_t = \log \Pi_t - \log \Pi_{t-1}$). We start the estimation in 1985, near the end of the Volcker period. Our stationary, one-regime model would be unfit to estimate the entire post-war history (see Ang, Bekaert, and Wei (2007) and Fama (2006)). Estimating the model at monthly frequency gives us a sufficiently many observations (258 months). The VAR(1) structure with the 12-month lag on the right-hand side is parsimonious and delivers plausible long-term expectations.¹⁸ We use the letter u to denote time in months, while t continues to denote time in years. The law of motion for the state is

$$Y_{u+12} = \mu + \Gamma Y_u + \eta_{u+12}, \quad \text{with } \eta_{u+12} \mid \mathcal{I}_u \sim D(0, \Sigma_t), \quad (14)$$

¹⁷The correlations with the ARM share are similar using either semi-annual or monthly data.

¹⁸As a robustness check, we considered a VAR(2) model and estimated the model on the basis of quarterly instead of monthly data. The results become even somewhat stronger for a second-order VAR model and we found similar results for quarterly data as for monthly data.

with \mathcal{I}_u representing the information at time u . The VAR structure immediately delivers average expected future nominal short rates:

$$\frac{1}{\tau} \mathbb{E}_u \left[\sum_{j=1}^{\tau} y_{u+(12 \times (j-1))}(1) \right] = \frac{1}{\tau} e_1' \sum_{j=1}^{\tau} \left\{ \left(\sum_{i=1}^{j-1} \Gamma^{i-1} \right) \mu + \Gamma^{j-1} Y_u \right\}. \quad (15)$$

Together with the nominal long yield, this delivers our VAR-based measure of the nominal bond risk premium. Panel C of Figure 3 shows the 5-year and 10-year time series; they have a correlation of 96%.

Rows 5 and 6 of Table 2 show the ARM regression results using the VAR-based 5-year and 10-year bond risk premium. Again, both bond risk premia are highly significant predictors of the ARM share. The t-statistics are 4.2 and 3.9. They explain 32% and 35% of the variation in the ARM share, respectively.¹⁹ The economic magnitude of the slope coefficient is again very close to the one obtained from forecasters and to the one estimated from the household decision rule: In all three cases, a one-standard deviation increase in the risk premium increases the ARM share by about 8 percentage points.

The analysis in Section 2.6 allows us to interpret the 28% average ARM share and the 8% sensitivity of the ARM share to the bond risk premium in terms of the structural parameters of the model. More precisely, the mean and slope of the ARM share regression are a function of the location and scale parameters of the cross-sectional risk aversion distribution, respectively. We assume a log-normal distribution for γ_j and estimate a mean of 5.0 and a standard deviation of 2.9. The implied median level of risk aversion is 155. Appendix D.4 describes the inference procedure in detail.

In sum, the forward-looking measures and the household rule of thumb deliver quantitatively similar sensitivities of the ARM share to the bond risk premium. This suggests that choosing the “right mortgage at the right time” may require less “financial sophistication” of households than previously thought. As evidenced by the higher R^2 in Rows 1 and 2 compared to Rows 3 to 6, the household decision rule turns out to be the strongest predictor. If the adaptive expectations scheme accurately describes households’ behavior, we would expect it to explain more of the variation in households’ mortgage choice. We discuss the differential performance of the backward- and forward-looking measures further in Section 4.

¹⁹We have also considered and estimated a VAR model with heteroscedastic innovations. In such a model, time variation in the volatility of expected inflation and expected real rates delivers two additional channels for variation in mortgage choice. While both conditional volatilities entered with the predicted sign in the regression, neither was statistically significant. Together, these terms added little explanatory power above the nominal bond risk premium.

3.3 Alternative Interest Rate Measures

The household decision rule has the appealing feature that it nests two commonly-used predictors of mortgage choice as special cases. First, when $\rho = 1$, we recover the yield spread:

$$\kappa_t(1; \tau) = y_t^{\$}(\tau) - y_t^{\$(1)}.$$

The yield spread is the optimal predictor of mortgage choice in our model only if the conditional expectation of future short rates equals the current short rate. This is the case only when short rates follow a random walk. Second, when $\rho \rightarrow \infty$, then $\kappa_t(\rho; T)$ converges to the long-term yield in excess of the unconditional expectation of the short rate:

$$\lim_{\rho \rightarrow \infty} \kappa_t(\rho; T) = y_t^{\$(T)} - \mathbb{E} [y_t^{\$(12)}], \quad (16)$$

by the law of large numbers.²⁰ Because the second term is constant, all variation in financial incentives to choose a particular mortgage originates from variation in the long-term yield. This rule is optimal when short rates are constant.

For all cases in between the two extremes, the simple model of adaptive expectations puts some positive and finite weight on average recent short-term yields to form conditional expectations. As Section 2.5 argued, this is why both the yield spread and the long yield suffer from an errors-in-variables problem in the ARM share regressions. To understand this problem, consider the VAR model estimates. They show that the two terms on the right-hand side of (10) are *negatively* correlated (-.57 for 5-year and -.54 for 10-year yield). One reason why the correlation between the nominal bond risk premium and the difference between expected future short rates and the current short rate is negative is the following. When expected inflation is high, the inflation risk premium -and hence the nominal bond risk premium- tends to be high. But at the same time, expected future short rates are below the current short rate because inflation is expected to revert back to its long-term mean. This negative correlation makes the yield spread a very noisy proxy for the nominal bond risk premium, and is responsible for the low R^2 in the regression of the ARM share on the yield spread. Indeed, Rows 7 and 8 of Table 2 confirm that the lagged yield spread explains less than 1% of the variation in the ARM share in the full sample (1985.1-2006.6). The weak case for the yield spread is also evident in the loan-level data. The second panel of Table 1 shows that the yield spread carries a much smaller (normalized) coefficient than the bond risk premium in the top panel, has a much lower t-statistic, and helps classify a lot fewer individual loans correctly.

The long yield suffers from a similar errors-in-variables problem. However, the two terms on the right-hand side of equation (11) are positively correlated (.58 for 5-year and .66 for 10-year

²⁰This requires a stationarity assumption on the short rates.

yield, based on VAR estimation), making the problem less severe. Rows 9 and 10 of Table 2 show that the long yield explains 37-39% of the ARM share, with a sensitivity coefficient of around 8.5%. The loan-level analysis in the third panel of Table 1 shows that the long yield enters the probit regressions with the wrong sign, substantially reducing the appeal of the long yield as a mortgage choice predictor.

An alternative source of interest rate data comes from the mortgage market. We use the 1-year ARM rate as our measure of the short rate and the 30-year FRM rate as our measure of the long rate (see Appendix A). The household decision rule based on mortgage rate data works well. The regression results in Row 11 are for a two-year look-back period, the horizon that maximizes the correlation with the ARM share in the right panel of Figure 2, and deliver an R^2 of 60%. The point estimate of 7.3 is similar to the one from the decision rule based on Treasury rates in Row 1. Row 12 shows similar results for a three-year look-back period. As we did for Treasury yields, we also regress the ARM share on the slope of the yield curve (30-year FRM rate minus 1-year ARM rate) and the long yield (30-year FRM rate). Row 13 shows that the FRM-ARM spread has lower explanatory power than the household decision rule, but much higher explanatory power than the Treasury yield spread. This improvement occurs only because the FRM-ARM spread contains additional information that is not in the Treasury yield spread.²¹ The explanatory power of the FRM rate is similar to that of the long Treasury yield (Row 15 and right panel of Figure 2).

The household decision rule that we introduce in Section 1 is motivated by the theoretical model in Section 2. It computes the expectations of future short rates in (12) as an equally-weighted average of past short rates. An alternative adaptive expectations proxy for the bond risk premium is to subtract from the long rate a geometrically-declining weighted average of past short rates. Earlier work by Shiller (1979) and others used the spread between the long rate and a moving average of the short rate to measure bond risk premia.²² In a different context, a similar weighting scheme was recently used by Malmendier and Nagel (2008). This delivers the alternative adaptive expectations measure of the bond risk premium:

$$\hat{\kappa}_t(\gamma, n) \equiv y_t^{\$}(n) - \mu_t(\gamma), \quad \text{where } \mu_t(\gamma) = (1 - \gamma) \sum_{i=0}^{\infty} \gamma^i y_{t-i}^{\$}(1) \quad \text{and } 0 \leq \gamma < 1. \quad (17)$$

As for the adaptive expectations with equal weights, the rule with geometrically declining weights

²¹The correlation between the FRM-ARM spread and the 10-1-year government bond yield spread is only 32%. This spread also captures the value of the prepayment option, as well as the lenders' profit margin differential on the FRM and ARM contracts. To get at this additional information, we orthogonalize the FRM-ARM spread to the 10-1 yield spread, and regress the ARM share on the orthogonal component (Row 14). For the full sample, we find a strongly significant effect on the ARM share. Partially this is due to the fact that this orthogonal spread component has a correlation of 60% with the fee differential between an FRM and an ARM contract. It only has a correlation of 16% with the rule-of-thumb risk premium.

²²See Sargent (1999) and Evans and Ramey (2006), and the references therein, for a discussion of the (near) rationality of adaptive expectations for forecasting.

nesses the yield spread ($\gamma = 0$) and the long yield ($\gamma \rightarrow 1$). This measure has similarly strong predictive ability for the ARM share as the measure with equal weights for (monthly) values of γ between .95 and .99. Rows 16 and 17 report the results for $\hat{\kappa}_t(.97, 5)$ and $\hat{\kappa}_t(.97, 10)$, initiated in 1985.1. In sum, both equally-weighted and geometrically declining weighting schemes for past short rates deliver measures of the bond risk premium which strongly predict future ARM shares.

Finally, we investigate two additional interest rate-based variables which implement alternative, more ad-hoc, rules-of-thumb. The first rule takes the current FRM rate minus the three-year moving average of FRM rate (row 18 of Table 2). The second rule does the same, but for the ARM rate (row 19). The first rule captures the idea behind the popular investment advice of “locking in a low long-term rate while you can”. The slope coefficients in the FRM and ARM rule are smaller than what we find for the bond risk premium (6.0 and 3.1) and less precisely measured (t-statistics of 3.7 and 2.4). The R^2 in the two regressions are 22% and 6%, respectively. Both alternative rules perform much worse than the household decision rule of Section 3.1, which is guided by theory.

4 The Recent Episode and the Inflation Risk Premium

The previous sections showed that all three estimates of the theoretical bond risk premium are positively and significantly related to the FRM-ARM choice. In this section, we investigate the difference between the household decision rule, which shows the strongest relationship and is based on adaptive expectations, and the forecasters- and VAR-based measures, which show a somewhat weaker relationship and are based on forward-looking expectations. Figure 4 shows that this difference in performance is especially pronounced after 2004. The figure displays the 10-year rolling-window correlation for each of the three measures with the ARM share. While the rule-of-thumb measure has a stable correlation across sub-samples, the performance of the forecasters-based measure as well as the VAR-based measure drop off steeply around 2004.

[Figure 4 about here.]

The reason for this failure is that the ARM share increased substantially between June 2003 and December 2004 with no commensurate increase in the Blue Chip or VAR risk premia measures. A similarly steep drop-off in correlation occurs for the long yield and for the FRM-ARM rate differential, both of which also performed well in the full sample. We explore two possible explanations for why the ARM share was high in 2004 when the forward-looking bond risk premia were low.

4.1 Product Innovation in the ARM Segment

A first potential explanation for the increase in the ARM share between June 2003 and December 2004 is product innovation in the ARM segment of the mortgage markets. An important develop-

ment was the increased popularity of hybrid mortgages: adjustable-rate mortgages with an initial fixed-rate period.²³ Figure 5 shows our benchmark measure of the ARM share (solid line) alongside a measure of the ARM share that excludes all hybrid contracts with initial fixed-rate period longer than three years. We label this measure \widetilde{ARM} . A substantial fraction of the increase in the ARM share in 2003-05 was due to the rise of hybrids. Under this hypothesis the ARM share went up despite the low bond risk premium because new types of ARM mortgage contracts became available that unlocked the dream of home ownership.²⁴

[Figure 5 about here.]

To test this hypothesis, we recompute the rolling correlations for \widetilde{ARM} , which excludes the hybrids. The correlation with the forecasters-based measure over the last 10-year window improves from 23% to 48%. The correlation over the longest available sample (since 1992) improves from 44% to 67%. In sum, the recently increased prevalence of the hybrids is part of the explanation. However, it cannot account for the entire story.

4.2 Forecast Errors

A second potential explanation is that the forecasters made substantial errors in their predictions of future short rates in recent times. We recall that nominal short rates came down substantially from 6% in 2000 to 1% in June 2003. Our Blue Chip data show that forecasters expected short rates to increase substantially from their 1% level in June 2003. Instead, nominal short rates increased only moderately to 2.2% by December 2004. Forecasters substantially over-estimated future short rates starting in the 2003.6-2004.12 period. As a result, the Blue Chip measure of bond risk premia is too low in that episode, and underestimates the desirability of ARMs.

Forecast errors in nominal rates translate in forecast errors for real rates. This is in particular the case when inflation is relatively stable and therefore easier to forecast. Figure 6 shows that the Blue Chip consensus forecast for the average real short rate over the next two years shows large disparities with its realized counterpart. We calculate the average expected future real short rate as the difference between the Blue Chip consensus average expected future nominal short rate and the Blue Chip consensus average expected future inflation rate. We calculate the realized real rate as the difference between the realized nominal rate and expected inflation, which we measure as the one-quarter ahead inflation forecast. The realized average future real short rates are calculated from the realized real rates. Finally, the forecast errors are scaled by the nominal short rate to obtain relative forecasting errors. The figure shows huge forecast errors in the 2000-2003 period,

²³Starting in 1992, we know the decomposition of the ARM by initial fixed-rate period. We are grateful to James Vickery for making these detailed data available to us.

²⁴In addition to the hybrid segment, the sub-prime market segment, which predominantly offers ARM contracts, also grew strongly over that period. However, our ARM sample does not contain this market segment.

relative to the earlier period. The forecast errors are on the order of 1.25 percentage point per year, about 50-75% of the value of the nominal short rate. These large forecast errors motivate the use of the inflation risk premium, as explained below.

[Figure 6 about here.]

A similar problem arises with the VAR-based bond risk premium. The VAR system also fails to pick up the declining short rates in the 2000-2004 period. It therefore also over-predicts the short rate and underestimates the desirability of ARM contracts.

Filtering Out Forecast Errors Forecast errors in the real rate not only help us identify the problem, they also offer the key to the solution. The nominal bond risk premium in the model of Sections 2.1 and 2.2 contains compensation for both real rate risk and expected inflation risk:

$$\phi_t^{\$}(\tau) = \phi_t^y(\tau) + \phi_t^x(\tau). \quad (18)$$

Similar to the nominal risk premium in (12), the real rate risk premium, ϕ_t^y , is the difference between the observed real long rate and the average expected future real short rate:

$$\phi_t^y(\tau) \equiv y_t(\tau) - \frac{1}{\tau} \sum_{j=1}^{\tau} \mathbb{E}_t [y_{t+j-1}(1)], \quad (19)$$

where $y_t(\tau)$ is the real yield of a τ -month real bond at time t . Following Ang, Bekaert, and Wei (2007), we define the inflation premium at time t , ϕ_t^x , as the difference between long-term nominal yields, long-term real yields, and long-term expected inflation:

$$\phi_t^x(\tau) \equiv y_t^{\$}(\tau) - y_t(\tau) - x_t(\tau), \quad (20)$$

where long-term expected inflation is given by:

$$x_t(\tau) \equiv \frac{1}{\tau} \mathbb{E}_t [\log \Pi_{t+\tau} - \log \Pi_t].$$

A key insight is that both the nominal long yield $y_t^{\$}(\tau)$ and the real long yield $y_t(\tau)$ contain expected future real short rates. Thus, their difference does not. Therefore, their difference zeroes out any forecast errors in expected future real short rates. Equation (20) shows that the inflation-risk premium, $\phi_t^x(\tau)$, contains the difference between $y_t^{\$}(\tau) - y_t(\tau)$, and therefore does not suffer from the forecast error problem.²⁵ In short, one way to correct the nominal bond risk premium for the forecast error is to only use the inflation risk premium component.

²⁵The same Blue Chip forecast data, as well as data from the Survey of Professional Forecasters, indeed show that inflation forecasts do not suffer from the same problem as nominal interest rate forecasts. This is consistent

Measuring the Inflation Risk Premium To implement equation (20), we need a measure of long real yields and a measure of expected future inflation rates. Real yield data are available as of January 1997 when the US Treasury introduced Treasury Inflation-Protected Securities (TIPS). We omit the first six months when liquidity was low, and only a 5-year bond was trading. In what follows, we consider two empirical measures for expected inflation. Our first measure for expected inflation is computed from the same semi-annual Blue Chip long-range consensus forecast data we used for the nominal short rate, using the same method, but using the series for the CPI forecast instead of the nominal short rate.²⁶ The inflation-risk premium is then obtained by subtracting the real long yield and long-term expected inflation from the nominal long yield, as in (20).

Alternatively, we can use the VAR to form expected future inflation rates and thereby the inflation risk premium. We start by constructing the 1-year expected inflation series as a function of the state vector

$$x_t(1) = \mathbb{E}_t[\pi_{t+1}] = e'_4\mu + e'_4\Gamma Y_t, \quad (21)$$

where e_4 denotes the fourth unit vector. Next, we use the VAR structure to determine the τ -year expectations of the average inflation rate in terms of the state variables:

$$\frac{1}{\tau}\mathbb{E}_t\left[\sum_{j=1}^{\tau} e'_4 Y_{t+j-1}\right] = \frac{1}{\tau}e'_4 \sum_{j=1}^{\tau} \left\{ \left(\sum_{i=0}^{j-1} \Gamma^{i-1} \right) \mu + \Gamma^{j-1} Y_t \right\}. \quad (22)$$

With the long-term expected inflation from (22) in hand, we form the inflation risk premium as the difference between the observed nominal yield, the observed real yield, and expected inflation.

Results Figure 7 shows the inflation risk premium (dashed line) alongside the ARM share (solid line). The inflation risk premium is based on Blue Chip forecast data. Between March 2003 and March 2005 (closest survey dates), the inflation risk premium increased by 1.2 percentage points, or two standard deviations. The nominal bond risk premium, in contrast, only increased only by one standard deviation.

[Figure 7 about here.]

Over the period 1997.7-2006.6, the raw correlation between the ARM share and the 5-year (10-year) inflation risk premium is 84% (82%) for the Blue Chip measure and 80% (78%) for the

with Ang, Bekaert, and Wei (2007), who argue that inflation forecasts provide the best predictors of future inflation among a wide set of alternatives.

²⁶We have compared the inflation forecasts from Blue Chip with those from the Survey of Professional Forecasters, the Livingston Survey, and the Michigan Survey, and found them to be very close. Ang, Bekaert, and Wei (2006) argue that such survey data provides the best inflation forecasts among a wide array of methods.

VAR measure. Finally, we regress the ARM share on the 5-year and 10-year inflation risk premium for the period 1997.7-2006.6. For the Blue Chip measure, we find a point estimate of 6.95 (6.97) for the 5-year (10-year) inflation risk premium. The economic effect is therefore comparable to what we find for the nominal bond risk premium (Section 3.2.1). The coefficient is measured precisely; the t-statistic is 8.0 (7.9). The 5-year (10-year) inflation-risk premium alone explains 66% of the variation (67%) in the ARM share. Likewise, for the VAR-based measure, we find a point estimate of 6.80 (6.40) for the 5-year (10-year) inflation risk premium. The coefficient is measured precisely; the t-statistic equals 8.5 (6.8). The inflation risk premium alone explains 64% of the variation (56%) in the ARM share. We conclude that the inflation risk premium has been a very strong determinant of the ARM share in the last ten years.

In conclusion, at the end of the sample, the forward-looking expectations measures of the bond risk premium suffered from large differences between realized average short rates, and what forecasters or a VAR predicted for these same average short rates. The adaptive expectations scheme of the household decision rule did not suffer from the same problem. This explains why it performed much better in predicting the ARM share in the last part of the sample. The inflation risk premium component of the bond risk premium successfully purges that forecast error from the forward-looking bond risk premium measures. We showed that it is a strong predictor of the ARM share in the 1997.7-2006.6 sample.

5 Extensions

In this section, we extend our results along several dimensions. First, we analyze the role of the prepayment option. Second, we study ARM share dynamics for different U.S. states as well as income groups. Third, we revisit the role of financial constraints for mortgage choice. Fourth, we study potential supply-side determinants of the ARM share. Fifth, we study bond return volatility and predictability. Sixth, we show that the household decision rule also predicts mortgage refinancing and portfolio choice decisions. Seventh, we analyze the robustness of the statistical inference. Finally, we study the role of liquidity in the TIPS market for our results.

5.1 Prepayment Option

Sofar the analysis has ignored the prepayment option. In the US, an FRM contract typically has an embedded option which allows the mortgage borrower to pay off the loan at will. We show how the presence of the prepayment option affects mortgage choice within the utility framework of Section 2.²⁷

²⁷We contribute to the large literature on rational prepayment models, e.g., Dunn and McConnell (1981), Stanton and Wallace (1998), Longstaff (2005), and Pliska (2006), by adding time variation in risk premia. Other studies

FRM Rate With Prepayment A household prefers to prepay at time 1 if the utility derived from the ARM contract exceeds that of the FRM contract. Prepayment entails no costs, but this assumption is easy to relax in our framework. It then immediately follows from comparing the time-1 value function that prepayment is optimal if and only if:

$$q_0^{FRMP} > q_1^{ARM},$$

where the superscript P in q_0^{FRMP} indicates the FRM contract with prepayment. The FRM rate with prepayment satisfies the following zero-profit condition. It stipulates that the present value of mortgage payments the lender receives must equate the initial mortgage balance B :

$$\begin{aligned} B &= \mathbb{E}_0 \left[M_1^\$ q_0^{FRMP} B + I_{(q_0^{FRMP} > q_1^{ARM})} M_1^\$ M_2^\$ q_1^{ARM} B + I_{(q_0^{FRMP} \leq q_1^{ARM})} M_1^\$ M_2^\$ q_0^{FRMP} B + M_1^\$ M_2^\$ B \right] \\ &= q_0^{FRMP} P_0^\$(1) B + [q_0^{FRMP} + 1] P_0^\$(2) B - B \mathbb{E}_0 [M_1^\$ M_2^\$ \max \{q_0^{FRMP} - q_1^{ARM}, 0\}], \end{aligned}$$

where the last term represents the value of the embedded prepayment option held by the household. $I_{(x < y)}$ denotes an indicator function that takes a value of one when $x < y$. This option value satisfies:

$$B \mathbb{E}_0 [M_1^\$ M_2^\$ \max \{q_0^{FRMP} - q_1^{ARM}, 0\}] = B (1 + q_0^{FRMP}) \left[P_0^\$(2) \Phi(d_1) - \frac{1}{1 + q_0^{FRMP}} P_0^\$(1) \Phi(d_2) \right],$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution, and the expressions for d_1 and d_2 are provided in Appendix C. The household has $B(1 + q_0^{FRMP})$ European call options on a two-period bond with expiration date $t = 1$ (when it becomes a one-year bond with price $P_1^\$(1) = 1/(1 + q_1^{ARM})$), and with an exercise price of $1/(1 + q_0^{FRMP})$. Substituting the option value into the zero-profit condition we get:

$$B = (q_0^{FRMP} + \Phi(d_2)) P_0^\$(1) B + [q_0^{FRMP} + 1] P_0^\$(2) B (1 - \Phi(d_1)).$$

The mortgage balance equals the sum of (i) the (discounted) payments at time $t = 1$, a certain interest payment and a principal payment with risk-adjusted probability $\Phi(d_2)$, and (ii) the (discounted) payments at time $t = 2$, when both interest and principal payments are received with risk-adjusted probability $1 - \Phi(d_1)$. The no-arbitrage rate q_0^{FRMP} on an FRM with prepayment solves the fixed-point problem:

$$q_0^{FRMP} = \frac{1 - (1 - \Phi(d_1)) P_0^\$(2) - \Phi(d_2) P_0^\$(1)}{P_0^\$(1) + (1 - \Phi(d_1)) P_0^\$(2)},$$

consider reduced-form models that can accommodate slow prepayment (e.g., Schwartz and Torous (1989), Stanton (1995), Boudoukh, Whitelaw, Richardson, and Stanton (1997), and Schwartz (2007)).

which cannot be solved for analytically as q_0^{FRMP} appears in d_1 and d_2 on the right-hand side. For $\Phi(d_1) = \Phi(d_2) = 1$, prepayment is certain, and we retrieve the expression for the year-one ARM rate, q_0^{ARM} . For $\Phi(d_1) = \Phi(d_2) = 0$, prepayment occurs with zero probability, and we obtain the expression for the FRM without prepayment, q_0^{FRM} .

This framework clarifies the relationship between time-varying bond risk premia and the price of the prepayment option. The bond risk premium goes up when the price of interest rate risk goes down. But a decrease in the price of interest rate risk makes prepayment less likely under the risk-neutral distribution. This is because the risk-neutral distribution shifts to the right and makes low interest rate states, where prepayment occurs, less likely. Therefore, the price of the prepayment option is decreasing in the bond risk premium.

Reduced Sensitivity A fixed-rate mortgage *without* prepayment option is a coupon-bearing nominal bond, issued by the borrower and held by the lender.²⁸ An FRM *with* prepayment option resembles a callable bond: the borrower has the right to prepay the outstanding mortgage debt at any point in time. The price sensitivity of a callable bond to interest rate shocks differs from that of a regular bond. This is illustrated in Figure 8. We use the bond pricing setup of Section 2.2 and set $\mu_y = \mu_x = 2\%$, $\rho_y = \rho_x = 0.5$, $\rho_{xy} = 0$, $\sigma_y = \sigma_x = 2\%$, and $\lambda_0 = [-0.4, -0.4]'$. These values imply a two-period nominal bond risk premium of $\phi_0^{\$}(2) = 0.78\%$. We vary the short rate at time zero, $y_0^{\$}(1) = y_0(1) + x_0$, assuming $y_0(1) = x_0$. The callable bond can be called at time one with exercise price of 0.96 (per dollar face value). The non-callable bond price is decreasing and convex in the nominal interest rate. The callable bond price is also decreasing in the nominal interest rate, but, the relationship becomes concave when the call option is in the money (“negative convexity”). This means that the callable bond has positive, but diminished exposure to nominal interest-rate risk.

[Figure 8 about here.]

Utility Implications of the Prepayment Option Next, we study how the prepayment option affects the relationship between the bond risk premium and the ARM-FRM utility differential. We use the same term-structure variables as in Figure 8, but vary the market prices of risk λ_0 . We maintain the assumption of equal prices of inflation risk and real interest rate risk, and fix the initial real interest and inflation rate at their unconditional means, i.e. $y_0(1) = \mu_y$ and $x_0 = \mu_x$. We assume the investor has a mortgage balance and house size normalized to 1, constant real labor income of 0.41, and a risk aversion coefficient $\gamma = 10$. Figure 9 plots the difference between

²⁸This analogy is exact for an interest-only mortgage. When the mortgage balance is paid off during the contractual period (amortizing), the loan can be thought of as a portfolio of bonds with maturities equal to the dates on which the down-payments occur. Acharya and Carpenter (2002) discuss the valuation of callable, defaultable bonds.

the lifetime utility from the ARM contract and the lifetime utility from the FRM contract. The solid line depicts the case without prepayment option; the dashed line plots the utility difference when the FRM has the prepayment option. No approximations are used for this exercise. The utility difference is increasing in the bond risk premium, both with and without prepayment option. However, the sensitivity of the utility difference to changes in the bond risk premium is somewhat reduced in presence of a prepayment option. This is consistent with the fact that a callable bond has diminished interest rate exposure and therefore contains a lower bond risk premium than a non-callable bond. This shows that our main result, a positive relationship between the utility difference of an ARM and an FRM contract and the nominal bond risk premium, goes through.

[Figure 9 about here.]

5.2 Composition Effects

One alternative hypothesis is that variation in the aggregate ARM share is driven by changes in the composition of the types of households originating a mortgage. The types we consider differ by location (U.S. states), and income or wealth. For example, it could be that the ARM shares of different U.S. states are at different levels and constant over time, so that all variation in the aggregate ARM share is due to shifts in the composition of loans across the different states. Similar composition effects could arise with respect to income or wealth groups. We show that this is not the case. Instead, ARM shares fluctuate substantially over time in each group (U.S. state or income/wealth group) and the groups' ARM shares move up and down together.

We first study the geographical composition effect based on a breakdown by U.S. state of our aggregate ARM share data from the Monthly Interest Rate Survey.²⁹ We also obtained the fraction of loan volume of each state, at each date. This allows us to do the state-level analysis for a long time period (1987.1-2006.6) and for the same source of data as our aggregate ARM share. These data have the advantage that when we average (dollar-weighted) ARM shares across all states, we recover the national ARM share. We conduct two different exercises to study the composition effect. First, we run a panel regression with state fixed-effects of the ARM share on the bond risk premium (the household decision rule variable $\kappa_t(3; 5)$). The total number of state-month observations is 11,038. The regression measures whether the bond risk premium significantly affects the ARM share, after controlling for differences in the average ARM share across states. The slope coefficient is 0.083 with a t-stat of 55.7; the overall regression R^2 is 19%. This magnitude coefficient is essentially the same as the one we found in Section 3.1 using the aggregate ARM share. We conclude that the bond risk premium variable is a highly significant determinant of the common variation in the ARM shares of the U.S. states. Second, we illustrate the composition

²⁹We are very grateful to James Vickery for providing us with these MIRS data.

effect graphically in Figure 10. It plots the aggregate ARM share (solid line) alongside the ARM share that arises if we control for composition effects (dotted line) and the ARM share that only varies because of composition effects (dashed line).³⁰ The figure shows that all the action in the aggregate ARM share comes from common variation in the ARM share of each U.S. state, not from compositional shifts in loan volumes between U.S. states.

[Figure 10 about here.]

Next, we study U.S. states based on the loan-level database of Section 1 and Appendix A. We classify each loan by the state in which it was originated, and for each month in the sample, we calculate the average ARM share of each state. The total number of state-month observations is 3,687. We then run a panel regression of the ARM share on the bond risk premium (the household decision rule variable $\kappa_t(3; 5)$) with state fixed-effects. The first row of Table 3 shows that the risk premium coefficient is highly significant, with a t-stat of 46. A one standard deviation increase in the household decision rule increase the ARM share by 16 percentage points from an average of 28 percentage points. The bond risk premium alone explains 36% of the variation in state-level ARM shares. This shows that the bond risk premium significantly affects the ARM share, after controlling for differences in the average ARM share across states. Row 2 adds the average mortgage balance, the average FICO score, and the average loan-to-value ratio, where the average is computed by averaging over all loans in a given state in a given month. The bond risk premium variable remains essentially unaffected by including these additional controls. The R^2 only increases marginally, from 36 to 40%. Since all regressors are scaled by their standard deviation, the magnitudes of the coefficient on the characteristics show that their economic significance is smaller than that of the risk premium variable. To show that the bond risk premium is not simply picking up variation in these characteristics, we omit the risk premium as a regressor in Row 3. The R^2 drops to a mere 11%. We also computed the ARM share \widehat{ARM} , which only fluctuates because of ARM share movements in each state but contains no composition effect, and \widetilde{ARM} which fluctuates only because of composition effects. The figure based on loan-level data essentially looks the same as Figure 10, with a correlation of 99% between the aggregate ARM share and \widehat{ARM} . We conclude that, while U.S. states have different average ARM shares, all time series variation in the aggregate ARM share is driven by common variation in the ARM shares of each state and not by compositional changes between them.

³⁰The aggregate ARM share is the weighted average of the ARM share in each state i : $ARM_t = \sum_{i=1}^{50} w_{it} ARM_{it}$. Weights w_{it} are defined as the dollar volume of mortgages originated in state i at time t as a fraction of the combined dollar origination in all states at time t . Define the time-series average loan-volume share of state i as $\bar{w}_i = \frac{1}{T} \sum_{t=1}^T w_{it}$. Likewise, let the time-series average ARM share in state i is $\overline{ARM}_i = \frac{1}{T} \sum_{t=1}^T ARM_{it}$. We define a second ARM share that shuts down any changes in composition by holding the weights constant at their time-series average: $\widetilde{ARM}_t = \sum_{i=1}^{50} \bar{w}_i ARM_{it}$. This is the dotted line in the figure. Finally, we define a third ARM share that only captures compositional changes across states and holds the ARM share in each state fixed at its average: $\widehat{ARM}_t = \sum_{i=1}^{50} w_{it} \overline{ARM}_i$. This is the dashed line in the figure.

[Table 3 about here.]

Next, we return to our loan-level data to investigate additional composition effects with respect to income or wealth groups. Because we have no direct information on the wealth or income of the borrower, we use four different proxies: the real loan balance, the real house price appraisal value, the FICO credit score, and 1999 ZIP code-level income.³¹ Presumably, wealthier households or households with higher income choose mortgages with a higher loan balance, buy more expensive houses, have higher FICO scores, and cluster in ZIP codes with higher income. While none of these four characteristics is perfect, they all are reasonable proxies to address the wealth/income group composition effect. For each of these four proxies, we present fixed-effect panel regressions in rows 4-7 of Table 3. In all four regressions, the bond risk premium is a highly significant predictor of common variation in ARM shares across groups. Its economic significance exceeds that of group-averaged loan characteristics such as BAL , $FICO$, and LTV (not reported). For each group definition, we also constructed ARM , \widehat{ARM} , and \widetilde{ARM} series, and confirmed that variation in the ARM share is not the result of a composition effect along the dimension of income or wealth.

5.3 Financial Constraints

Another alternative hypothesis is that there is a group of financially-constrained households which postpones the purchase of a house until the ARM rate is sufficiently low to qualify for a mortgage loan. Under this alternative hypothesis, the time series variation in the dollar volume of ARMs would drive the variation in the ARM share, and the dollar volume in FRMs would be relatively constant. Figure 11 plots the dollar volume of ARM and FRM mortgage originations for the entire U.S. market, scaled by the overall size of the mortgage market. The data are compiled by OFHEO. It shows that there are large year-on-year fluctuations in both the ARM and the FRM market segment. This dispels the hypothesis that the variation in the ARM share over the last 20 years is driven by fluctuations in participation in the ARM segment.

[Figure 11 about here.]

Loan-level data provide arguably the best laboratory to test the importance of financial constraints. As we showed in Section 1 and Table 1, loan balance, FICO score, and LTV ratio were all significant predictors of the probability of choosing an ARM. However, they did not drive out the bond risk premium. Rather, the bond risk premium is economically the stronger determinant of mortgage choice based on the size of its coefficient, its t-statistic, and the number of correctly classified loans. The financial constraint variables did not add any explanatory power. This is a

³¹The loan balance, house price appraisal, FICO score, and ZIP code of each loan are part of our loan-level database. For the ZIP code exercise, we use 1999 median household income data by ZIP code from the 2000 decennial Census (31,888 ZIP codes) and assign to each loan a 1999 household income level by matching ZIP codes.

powerful result given that the balance, FICO score, and LTV ratio are cross-sectional variables, while the bond risk premium is a time series variable.

As discussed in the previous section, fluctuations in the aggregate ARM share are not driven by composition effects across U.S. states. Thus while the differences in the average ARM share across states may reflect financial constraint-type variables such as the median house price, the dynamics of the ARM share are not driven by the dynamics of these variables.³²

The importance of the yield spread as a predictor of mortgage choice also relates to the role of financial constraints. Section 3.3 showed that the yield spread did not display a strong comovement with the ARM share. We argued that the yield spread not only captures the bond risk premium, but also deviations of expected future short rates from current short rates, causing the problem. However, when a household is perfectly impatient and only cares about consumption in the current period ($\beta = 0$), only the current period's differential between the long-term and the short-term interest rate matters. The same is true if a household plans to move in the current year.³³ The multi-period model of Appendix D allows us to investigate the quantitative role of the time discount factor and the moving rate for mortgage choice. Our conclusions are that for conventional values of the time discount factor or the moving rate, it is the bond risk premium which matters. Finally, we have investigated the extent to which the yield spread affects mortgage choice in the data, over and above the risk premium. In a multiple time series regression of the ARM share on the risk premium and the yield spread, the latter was typically not significant. Its sign flips across specifications, its t-statistic is low, and it does not contribute to the R^2 of the regression, beyond the effect of the risk premium. In the loan-level data, adding the yield spread to the probit regression with the bond risk premium, the loan balance, FICO score, LTV ratio, and the regional dummies only strengthens the effect of the bond risk premium (Row 6 of Table 1). While the yield spread is highly significant in this regression, it does not seem to be the case that its explanatory power captures the effect of financial constraints. The coefficients and significance level of the loan balance, FICO score, and LTV ratio are not diminished. Put differently, adding the yield spread to the bond risk premium has stronger effects than adding these three loan characteristics.

We conclude that, first, the bond risk premium is a powerful predictor of mortgage choice in these loan-level data. Second, while measures of financial constraints certainly enter significantly in these regressions, both their economic and statistical effect on mortgage choice is smaller.

³²We investigated the effect of the aggregate loan-to-value ratio, aggregate house price-income, and house price-rent ratios on the ARM share, but found no relationship.

³³Mobility in and of itself is an unlikely candidate to explain variation in the ARM share. Current Population Survey data for 1948-2004 from the US Census show that the average annual (monthly) moving rate is 18.1% (1.27%), and the out-of-county moving rate is 6.2% (1.16%). Moreover, these moving rates show no systematic variation over time.

5.4 Supply-Side

A last alternative hypothesis we consider is that mortgage lenders might push borrowers actively into taking one type of mortgage, and that their incentive to do so may be correlated with the ARM share. The most direct way lenders can affect a households' mortgage choice is through prices. Vickery (2007a) presents some evidence that lenders with different ex-ante risk exposures price similar loans differently, i.e., they offer incentives through rates. Next, we discuss alternative channels through which lenders could offer incentives to borrowers. Taking all the evidence together, we conclude that the evidence for the supply-side channel is weak, much weaker than the demand-based evidence we presented before.

The difference between the FRM and the ARM rate contains four components: the spread between the long- and short-term Treasury interest rates, the prepayment risk premium on the FRM, the default risk premium differential between the FRM and the ARM, and a residual. In symbols:

$$y_t^{FRM} - y_t^{ARM} = (y_t^{\$(10)} - y_t^{\$(1)}) + prepay_t^{FRM} + (def_t^{FRM} - def_t^{ARM}) + \varepsilon_t. \quad (23)$$

The prepayment and default risk premium capture compensation for interest rate and credit risk, while the first component captures the funding cost differential in the absence of prepayment and default. The residual ε_t is what summarizes the lenders' incentive to issue one type of contract over the other. One practical challenge in measuring the relative incentive ε is to accurately measure the prepayment and default risk premium terms. The incentive to offer FRMs is captured by the option-adjusted spread (OAS). The OAS measures the difference between the mortgage-backed securities (MBS) spread and the value of the prepayment option. The mortgage-backed securities spread itself is the difference between the current coupon 30-year Fannie Mae MBS rate and the 10-year CMT rate. Since it is an MBS rate, it does not contain default risk.³⁴ In symbols:

$$OAS_t \equiv \underbrace{y_t^{FRM} - y_t^{\$(10)} - def_t^{FRM}}_{MBS \text{ spread}} - prepay_t^{FRM}.$$

In order to measure the relative incentive, we subtract from the incentive to offer FRMs, measured by the OAS, the incentive to offer ARMs, $y_t^{ARM} - y_t^{\$(1)} - def_t^{ARM}$.³⁵ We measure def_t^{ARM} in three different ways: (1) we assume that it is constant, (2) we assume that it is equal to the corporate bond default risk premium, measured as the Moody's Baa-Aaa yield spread, plus or minus a

³⁴Banks can securitize conforming mortgages with Freddie Mac and Fannie Mae, who then assume the default risk in return for a guarantee fee. Data from the annual report of Office of Federal Housing Enterprise Oversight (OFHEO) to Congress show that the guarantee fee is small (around 22 basis points) and not volatile.

³⁵We use the 1-year ARM rate from Freddie Mac which has the advantage that it is indexed to the 1-year CMT rate.

constant, and (3) we posit and estimate a model, detailed in Appendix E, where the default rate depends on lagged interest rates and the house price-to-rent ratio. For each of the three measures of def_t^{ARM} , we form the lenders' relative incentive ε_t . A regression of the ARM share on the lagged incentive measure ε results in insignificant coefficient estimates and R^2 below 1.5% for all three cases. We conclude that, to the extent that the residual ε manages to isolate the supply-side incentives that lenders have to offer the FRM instead of the ARM, the weak predictability of the residual for the ARM share indicates little support for the supply-side story.

A second approach to study whether banks have incentives to push customers into the ARM versus the FRM is to focus on specific channels through which this may occur. First, we used data on market shares of mortgage lenders from Inside Mortgage Finance to compute a Herfindahl-Hirschmann Index of market concentration. We found that the mortgage lending industry is competitive,³⁶ and that differences in the HHI for the ARM and FRM segments do not correlate with the ARM share. To the extent that the degree of competition reflects profit opportunities, this suggests no important differences between profits on ARM and FRMs over time. Second, we studied how fees and charges at mortgage origination differ between ARM and FRM contracts. Both fees have come down substantially over time, consistent with increased competition. However the fee differential is only 18 basis points on average and has limited forecasting power for the ARM share; the regression R^2 is 6%. Finally, lenders may have an incentive to steer borrowers in a particular type of mortgage contract when there is a maturity mismatch between their assets and liabilities. We argue that this is unlikely to drive our results. First, ARM shares of loans originated by mortgage banks and depository institutions (commercial and savings banks) display almost identical time-series variation (Vickery (2007a)). Yet, mortgage banks do not face a maturity mismatch issue because they immediately securitize all loans they originate and end up with small balance sheets; only depository institutions do. Second, if the maturity mismatch were an important consideration one would expect depository institutions to not only steer new customers to ARMs, but also to use interest rate derivatives for hedging. I.e., we expect these two instruments to be complementary. Using Federal Deposit Insurance call report data, we find that large bank holding companies indeed use interest rate derivatives for hedging.³⁷ However, the asset-weighted (equally-weighted) average ratio of derivatives to assets only mildly predict the ARM share, with an R^2 of 10% (1%).

³⁶The HHI is 439, well below the threshold of 1,000 for an unconcentrated industry from the U.S. Department of Justice.

³⁷The data construction follows Purnanandam (2007), who performed a similar analysis for commercial banks instead of the (consolidated) bank holding companies. The derivatives data are the sum of series BHCK8725 and BHCK8729 before 2001 and equal the series BHCK8725 afterwards.

5.5 Bond Return Volatility and Predictability

One possibility is that the household decision rule is more correlated with changes in bond volatility than with changes in expected bond returns, and that households use it because they (mistakenly) think that bond volatility is an appropriate measure of bond risk. Viceira (2007) shows that the yield spread is associated with future excess bond returns, but not with bond volatility, and hence does not suffer from this problem. We find that our household decision rule shares that feature with the yield spread: it does not predict future interest rate volatility. The rule predicts one-period ahead realized bond return volatility with the wrong negative sign and is insignificant; the R^2 is a mere 3.3%. Therefore, we can rule out that the adaptive expectations rule is more correlated with changes in bond volatility (or bond variance, which is what enters in equation (8) and an order of magnitude smaller than bond volatility) than with changes in bond risk premia.

We also study long-horizon bond return predictability. We consider not only the household decision rule, which is an adaptive expectations measure whose second term places equal weights on past short rates, but also adaptive expectations measures which place geometrically-declining weights on past short rates (as introduced in equation 17). In addition, we study predictability by the VAR- and Blue Chip forecaster-based measures of the bond risk premium, the yield spread ($\gamma = 0$), the long yield ($\gamma = .999$), and the Cochrane and Piazzesi (2005) factor. Table 4 shows that the adaptive expectations measure of the bond risk premium with geometrically-declining weights has the strongest predictive ability for future five-year holding period returns (last column) on ten-year (Panel A) and five-year bonds (Panel B) among all measures. R^2 statistics are as high as 48% for ten-year and 40% for five-year bond returns for $\gamma = .99$ and 26% and 20%, respectively, for $\gamma = .97$. Despite its similar ability to predict the ARM share, the household decision rule $\kappa(3, 10)$ is a weaker predictor of bond returns than $\hat{\kappa}(.97, 10)$ with an R^2 at the five-year horizon of 4.5%. The yield spread and the household decision rule have similar predictive ability for long-horizon bond returns: the errors from their respective predictive regressions have a correlation of 98%. One important caveat to these long-horizon predictability results is that inference is substantially hampered by the short sample. R-squared values are not accurately estimated over a short sample period and even values beyond 40% may not be that informative. Yet, the short sample seems inevitable given the well-documented structural break in the bond return time series.³⁸

[Table 4 about here.]

³⁸Ignoring the break would cause coefficient biases in predictability regressions, e.g., Viceira (1996), Baker, Taliferro, and Wurgler (2006), Duffee (2007), and Lettau and Van Nieuwerburgh (2007).

5.6 Other Household Decisions

An interesting question is whether the household decision rule predicts other household decisions besides the choice between an ARM and FRM for new mortgage purchase decisions. We investigate two: mortgage refinancing decisions of existing mortgages and portfolio allocations.

First, we find that our proxy for the bond risk premium also predicts future refinancing activity, using both loan-level and aggregate data. Our previously-introduced data set of individual mortgage loans contains information on whether the mortgage origination was a regular refinancing, a cash-out refinancing, or a purchase-money mortgage. Grouping regular and cash-out refinancing loans (42% of all observations), we find a slope coefficient on the household decision rule of 0.37 with a t-stat of 152. The probit regression with the household decision rule as the only predictor classifies 66.7% of refinancing decisions correctly. For comparison, the slope coefficient for purchase-money loans is 0.48 with a t-stat of 195 and 70.7% of correctly-classified loans.³⁹ Turning to aggregate data, we construct a *refinancing activity index* as the ratio of the number of refinancing applications to the number of purchase-money mortgage applications. Both time series come from the Mortgage Bankers Association (MBA) and are available from January 1990 to June 2006. The incentive to refinance from an FRM to an ARM loan ought to be predicted positively and the incentive to refinance from an ARM to an FRM loan should be predicted with a negative sign by the household decision rule. To capture both directions simultaneously, we use the absolute value of the household decision rule as independent variable. A well known determinant of FRM-to-FRM refinancing activity is the ratio of the weighted-average coupon on all outstanding FRM contracts to the current FRM rate (WAC/MR).⁴⁰ After controlling for the WAC/MR ratio, which enters significantly and with the expected sign, we find that the absolute value of the household decision rule has additional explanatory power for the refinancing activity index. Its slope coefficient has the expected positive sign and a t-stat of 5.2. The regression R^2 is 70.2%, an improvement over the 66.0% for the regression with only the WAC/MR regressor. We conclude that the household decision rule, our adaptive expectations bond risk premium measure, plays an important role explaining aggregate mortgage refinancing activity.

Second, we find evidence that the household decision rule predicts the portfolio share households invest in bonds. We start by analyzing the bond share excluding stocks, defined as the ratio of US households' bond holdings to bond plus cash holdings. Quarterly data come from Table B.100 for U.S. households from the Flow of Funds. The sample period is 1989.IV-2006.II; the same sample as

³⁹This reduced sensitivity is consistent with the theory. It arises because refinancing may occur because of expectations of lower interest rates, not only because of changes in the risk premium. For newly-originated mortgages, the bond risk premium is all that matters.

⁴⁰Following Duarte (2008), we collect monthly data from Bloomberg on the outstanding balance and weighted-average coupon of the 105 most traded generic 30-year agency FRM pools in order to construct a time-series for the WAC. These 30-year FRM contracts are issued by Fannie Mae, Freddie Mac Gold, Freddie Mac Non-Gold, Ginnie I, and Ginnie II. For the current mortgage rate MR, we use the Fannie May 30-year mortgage rate.

for our ARM share regressions. In a regression of the bond share excluding stocks on the lagged and standardized household decision rule $\kappa(3; 5)$ and a constant, the former comes in with the expected positive sign. A one-standard deviation increase in the bond risk premium increases the bond share by 0.88 percentage points; the slope coefficient is statistically significant at the 5% confidence level with a the t-stat of 2.53. The regression R^2 is 9%. By contrast, the yield spread comes in with the wrong sign and is significant. Next, we form a bond share including stocks, defined by the the ratio of U.S. household bond holdings to stock plus bond plus cash holdings. Stock holdings include mutual funds, but the results are not sensitive to this choice. After controlling for time variation in the equity risk premium, the household decision rule predicts the bond share including stocks, positively and significantly.⁴¹ A one-standard deviation increase in the household decision rule increases the bond share by 0.62 percentage points. The yield spread comes in (significantly), but with the wrong sign. In conclusion, we find evidence from U.S. households' portfolio decisions that the allocation to bonds goes up in response to an increase in the adaptive expectations measure of the bond risk premium.

5.7 Persistence of Regressor

In contrast to the bond risk premium, most term structure variables do not explain much of the variation in the ARM share (Table 2). This is especially true in the last ten years of our sample, when the inflation risk premium has strong explanatory power (Section 4), but the real yield or the FRM-ARM rate differential do not. This suggests that our results for the risk premium are not simply an artifact of regressing a persistent regressand on a persistent regressor, because many of the other term structure variables are at least as persistent.⁴² To further investigate this issue, we conduct a block-bootstrap exercise, drawing 10,000 times with replacement 12-month blocks of innovations from an augmented VAR. The latter consists of the four equations of the VAR of Section 3.2.2, and is augmented with an equation for the ARM share. The ARM share equation is allowed to depend on the four lagged VAR elements, as well as on its own lag. The lagged ARM share itself does not affect the VAR elements. The bootstrap estimate recovers the point estimate (no bias), and it leads to a confidence interval that is narrower (6.40) than the Newey-West confidence interval we use in the main text (8.24), but wider than an OLS confidence interval (3.73). We conclude that the Newey-West standard errors we report are conservative.

One further robustness check we performed is to regress quarterly changes in the ARM share (between periods t and $t+3$) on changes in the term structure variables of the benchmark regression specification (between periods $t-1$ and t). We continue to find a positive and strongly significant

⁴¹The equity risk premium is the projection of the CRSP value-weighted stock return on the corresponding dividend-yield and the short rate.

⁴²The ARM share itself is not that persistent. Its annual autocorrelation is 30%, compared to 76% for the one-year nominal interest rate. An AR(1) at an annual frequency only explains 8.8% of the variation in the ARM share.

effect of the risk premium on the ARM share (t-statistic around 5). The effect of a change in the bond risk premium is similar to the one estimated from the level regressions: a one percentage point increase in the bond risk premium leads to a 10 percentage point increase in the ARM share over the next quarter. The R^2 of the regression in changes is obviously lower, but still substantial. For the 5-year (10-year) risk premium based on the VAR, it is 12% (18%), for the forecaster measure it is 25% (30%), and for the rule-of-thumb it is 26% (27%).

5.8 Liquidity and the TIPS Market

The results in Section 4, which use the inflation risk premium, are based on TIPS data. The TIPS markets suffered from liquidity problems during the first years of operation, which may have introduced a liquidity premium in TIPS yields (see Shen and Corning (2001) and Jarrow and Yildirim (2003)). A liquidity premium is likely to induce a downward bias in the inflation risk premium. As long as this bias does not systematically covary with the ARM share, it operates as an innocuous level effect and adds measurement error.

To rule out the possibility that our inflation risk premium results are driven by liquidity premia, we use real yield data backed out from the term structure model of Ang, Bekaert, and Wei (2007) instead of the TIPS yields. We treat the real yields as observed, and use them to construct the inflation risk premium.⁴³ Since the Ang-Bekaert-Wei data are quarterly (1985.IV-2004.IV), we construct the quarterly ARM share as the simple average of the three monthly ARM share observations in that quarter. We then regress the quarterly ARM share on the one-quarter lagged inflation and real rate risk premium. We find that both components of the nominal bond risk premium, the inflation-risk premium, and the real rate risk premium, enter with a positive sign. This is consistent with the theoretical model developed in Section 2. Both coefficients are statistically significant: The Newey-West t-statistic on the inflation risk premium is 3.90 and the t-statistic on the real rate risk premium is 2.12. The regression R-squared is 53%.

As a final robustness check, we repeated our regressions using only TIPS data after 1999.1, after the initial period of illiquidity. We found very similar results to those based on data starting in 1997.7. This suggests that liquidity problems in TIPS markets may have affected the inflation-risk premium, but this does not significantly affect our results. We conclude that our results are robust to using alternative real yield data.

⁴³We thank Andrew Ang for making these data available to us.

6 Conclusion

We have studied the link from the term structure of interest rates to the mortgage choice. A simple theoretical model shows that it is the long-term nominal bond risk premium that arises as the crucial determinant of the relative attractiveness of the adjustable-rate versus the fixed-rate mortgage. Thinking of fixed-rate mortgages as a short position in long-term bonds and adjustable-rate mortgages as rolling over a short position in short-term bonds implies that fixed-rate mortgage holders are paying a nominal bond risk premium. The higher the bond risk premium, the more expensive the FRM, and the higher the ARM share. In the model, time variation in the bond risk premium generates time variation in the share of households that opts for the ARM. The model also shows that both the yield spread and the long yield are poor proxies for the bond risk premium. Our empirical work confirms that these are weaker determinants of mortgage choice.

In aggregate time-series data, the same long-term bond risk premium is strongly related to the ARM share. Our results are consistent across three different methods of computing bond risk premia. We used forecasters' expectations, a VAR-model, and a simple adaptive expectation scheme, or "household decision rule." For all three measures of the bond risk premium, a one standard deviation increase leads to an eight percentage point increase in the ARM share. The simple household decision rule not only explains 70% of the variation in the aggregate ARM share, it also works well in explaining individual household mortgage choices. We used a new loan-level data set to show that the household decision rule is a highly significant predictor of the probability of choosing an ARM; it alone classifies 70% of loans correctly. Its explanatory power exceeds that of loan-level characteristics such as the mortgage balance, credit score or loan-to-value ratio, and that power is unaffected by the inclusion of these proxies for financial constraints.

In a previous version of the paper, we also studied mortgage choice in the United Kingdom. Long-term fixed-rate mortgages are a lot less prevalent in the UK than in the US and only a recent addition to the market. While this makes the maturity choice somewhat less relevant, we still found a significant positive covariation between the ARM share and the bond risk premium.

Taken together, our findings suggest that households may be making close-to-optimal mortgage choice decisions. Capturing the relevant time variation in bond risk premia is feasible by using a simple rule. This paper contributes to the growing household finance literature (Campbell (2006)), which debates the extent to which households make rational investment decisions. Given the importance of the house in the median household's portfolio and the prevalence of mortgages to finance the house, the problem of mortgage origination deserves a prominent place in this debate.

References

- ACHARYA, V. V., AND J. N. CARPENTER (2002): “Corporate Bond Valuation and Hedging with Stochastic Interest Rates and Endogenous Bankruptcy,” *Review of Financial Studies*, 15, 1355–1383.
- ANDREASSEN, P. B., AND S. J. KRAUSS (1990): “Judgemental Extrapolation and the Salience of Change,” *Journal of Forecasting*, 9(4), 771–792.
- ANG, A., G. BEKAERT, AND M. WEI (2006): “Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?,” *Forthcoming Journal of Monetary Economics*.
- (2007): “The Term Structure of Real Rates and Expected Inflation,” *Forthcoming Journal of Finance*.
- ANG, A., AND M. PIAZZESI (2003): “A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables,” *Journal of Monetary Economics*, 50, 745–787.
- BAKER, M., R. GREENWOOD, AND J. WURGLER (2003): “The Maturity of Debt Issues and Predictable Variation in Bond Returns,” *Journal of Financial Economics*, 70, 261–291.
- BAKER, M., R. TALIAFERRO, AND J. WURGLER (2006): “Predicting Returns with Managerial Decision Variables: Is There a Small-sample Bias?,” *Journal of Finance*, 61, 1711–1729.
- BERKOVEC, J. A., D. J. KOGUT, AND F. E. NOTHAFT (2001): “Determinants of the ARM Share of FHA and Conventional Lending,” *The Journal of Real Estate Finance and Economics*, 22, 23–41.
- BLACK, F., AND M. SCHOLES (1973): “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy*, 81, 637–654.
- BOUDOUGH, J., R. F. WHITELAW, M. RICHARDSON, AND R. STANTON (1997): “Pricing Mortgage-Backed Securities in a Multi-Factor Interest-Rate Environment: A Multivariate Density Approach,” *Review of Financial Studies*, 10, 405–446.
- BRANDT, M. W., AND P. SANTA-CLARA (2006): “Dynamic Portfolio Selection by Augmenting the Asset Space,” *The Journal of Finance*, 61, 2187–2218.
- BRENNAN, M. J., AND Y. XIA (2002): “Dynamic Asset Allocation under Inflation,” *Journal of Finance*, 57, 1201–1238.
- BRUNNERMEIER, M., AND C. JULLIARD (2006): “Money Illusion and Housing Frenzies,” *Review of Financial Studies*, Forthcoming.
- BURASCHI, A., AND A. JILTSOV (2005): “Time-varying inflation risk premia and the expectations hypothesis: A monetary model of the treasury yield curve,” *Journal of Financial Economics*, 75, 429–490.
- BUTLER, A. W., G. GRULLON, AND J. P. WESTON (2006): “Can Managers Successfully Time the Maturity Structure of Their Debt Issues?,” *Journal of Finance*, 61, 1731–1758.

- CAMPBELL, J. Y. (2006): “Household Finance,” *Journal of Finance*, 61, 1553–1604.
- CAMPBELL, J. Y., Y. L. CHAN, AND L. VICEIRA (2003): “A Multivariate Model of Strategic Asset Allocation,” *Journal of Financial Economics*, 67, 41–80.
- CAMPBELL, J. Y., AND J. COCCO (2003): “Household Risk Management and Optimal Mortgage Choice,” *Quarterly Journal of Economics*, 118, 1449–1494.
- CAMPBELL, J. Y., AND R. J. SHILLER (1991): “Yield Spreads and Interest Rate Movements: A Bird’s Eye View,” *Review of Economic Studies*, 58, 495–514.
- CAMPBELL, J. Y., AND L. VICEIRA (2001): “Who Should Buy Long-Term Bonds?,” *American Economic Review*, 91, 99–127.
- CHOI, J. (2006): “Extrapolative Expectations and the Equity Premium,” Yale University Mimeo.
- COCCO, J. (2005): “Portfolio Choice in the Presence Housing,” *The Review of Financial Studies*, 18, 535–567.
- COCCO, J. F., F. J. GOMES, AND P. J. MAENHOUT (2005): “Consumption and Portfolio Choice over the Life-Cycle,” *The Review of Financial Studies*, 18, 491–533.
- COCHRANE, J. H., AND M. PIAZZESI (2005): “Bond Risk Premia,” *American Economic Review*, 95, 138–160.
- (2006): “Decomposing the Yield Curve,” Working Paper, University of Chicago.
- DAI, Q., AND K. J. SINGLETON (2000): “Specification Analysis of Affine Term Structure Models,” *Journal of Finance*, 55, 1943–1978.
- (2002): “Expectation Puzzles, Time-varying Risk Premia, and Affine Models of the Term Structure,” *Journal of Financial Economics*, 63, 415–441.
- DEBONDT, W. (1993): “Betting on Trends: Intuitive Forecasts of Financial Risk and Return.,” *International Journal of Forecasting*, 9(3), 355–371.
- DUARTE, J. (2008): “The Causal Effect of Mortgage Refinancing on Interest-Rate Volatility: Empirical Evidence and Theoretical Implications,” *Review of Financial Studies*, Forthcoming.
- DUFFEE, G. R. (2007): “Are Variations in Term Premia Related to the Macroeconomy?,” Working Paper Haas School of Business.
- DUNN, K. B., AND J. J. MCCONNELL (1981): “Valuation of Mortgage-Backed Securities,” *Journal of Finance*, 36, 599–617.
- EPSTEIN, L. G., AND S. ZIN (1989): “Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–969.

- EVANS, G. W., AND G. RAMEY (2006): “Adaptive expectations, Underparameterization and the Lucas Critique,” *Journal of Monetary Economics*, 53, 249–264.
- FAMA, E. F. (2006): “The Behavior of Interest Rates,” *Review of Financial Studies*, 19, 359–379.
- FAMA, E. F., AND K. R. FRENCH (1989): “Business Conditions and Expected Returns on Stocks and Bonds,” *Journal of Financial Economics*, 25, 23–49.
- GABAIX, X., A. KRISHNAMURTHY, AND O. VIGNERON (2006): “Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market,” *The Journal of Finance*, Forthcoming.
- GREENWOOD, R., AND D. VAYANOS (2007): “Bond Supply and Excess Bond Returns,” Working Paper, London School of Economics.
- JAMSHIDIAN, F. (1989): “An Exact Bond Option Formula,” *Journal of Finance*, 44, 205–209.
- JARROW, R., AND Y. YILDIRIM (2003): “Pricing Treasury Inflation Protected Securities and Related Derivatives Using an HJM Model,” *Journal of Financial and Quantitative Analysis*, 38, 337–358.
- KOIJEN, R. S. J., T. E. NIJMAN, AND B. J. M. WERKER (2007): “When Can Life-cycle Investors Benefit from Time-varying Bond Risk Premia?,” Working Paper, Tilburg University.
- LETTAU, M., AND S. VAN NIEUWERBURGH (2007): “Reconciling the Return Predictability Evidence,” Forthcoming Review of Financial Studies.
- LONGSTAFF, F. A. (2005): “Borrower Credit and the Valuation of Mortgage-Backed Securities,” *Real Estate Economics*, 33, 619–661.
- LUSTIG, H., AND S. VAN NIEUWERBURGH (2006): “Can Housing Collateral Explain Long-Run Swings in Asset Returns?,” Working Paper NYU Stern and UCLA.
- MALMENDIER, U., AND S. NAGEL (2008): “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?,” working Paper Stanford University.
- MERTON, R. C. (1973): “Theory of Rational Option Pricing,” *Bell Journal of Economics and Management Science*, 4, 141–183.
- PLISKA, S. R. (2006): “Optimal Mortgage Refinancing with Endogenous Mortgage Rates: an Intensity Based Equilibrium Approach,” Working Paper University of Illinois at Chicago.
- PURNANANDAM, A. (2007): “Interest Rate Derivatives at Commercial Banks: An Empirical Investigation,” *Journal of Monetary Economics*, 54, 1769–1808.
- QI, M., AND X. YANG (2007): “Loss Given Default of High Loan-to-Value Residential Mortgages,” Office of the Comptroller of the Currency, Economics and Policy Analysis Working Paper 2007-4.

- SANGVINATSOS, A., AND J. A. WACHTER (2005): “Does the Failure of the Expectations Hypothesis Matter for Long-Term Investors?,” *Journal of Finance*, 60, 179–230.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press.
- SCHWARTZ, A. (2007): “Household Refinancing Behavior in Fixed Rate Mortgages,” Working Paper, Harvard University.
- SCHWARTZ, E. S., AND W. N. TOROUS (1989): “Prepayment and the Valuation of Mortgage-Backed Securities,” *Journal of Finance*, 44, 375–392.
- SHEN, P., AND J. CORNING (2001): “Can TIPS Help Identify Long-Term Inflation Expectations,” *Economic Review, Federal Reserve Bank of Kansas City*, Fourth Quarter, 61–87.
- SHILLER, R. (1979): “The Volatility of Long-Term Interest Rates and Expectations models of the Term Structure,” *Journal of Political Economy*, 87, 1190–1219.
- STANTON, R. (1995): “Rational Prepayment and the Valuation of Mortgage-Backed Securities,” *Review of Financial Studies*, 8, 677–708.
- STANTON, R., AND N. WALLACE (1998): “Mortgage Choice: What is the Point?,” *Real Estate Economics*, 26(2), 173–205.
- VAN HEMERT, O. (2007): “Life-Cycle Housing and Portfolio Choice with Bond Markets,” Working Paper, NYU Stern School of Business.
- VAN NIEUWERBURGH, S., AND L. VELDKAMP (2007): “Information Acquisition and Under-Diversification,” Working Paper New York University.
- VICEIRA, L. (1996): “Testing For Structural Change in the Predictability of Asset Returns,” Unpublished manuscript, Harvard University.
- (2007): “Bond Risk, Bond Return Volatility, and the Term Structure of Interest Rates,” manuscript Harvard University.
- VICKERY, J. (2007a): “How Do Financial Frictions Shape the Product Market? Evidence from Mortgage Originations,” Working Paper New York Federal Reserve Bank.
- (2007b): “Interest Rates and Consumer Choice in the Residential Mortgage Market,” Working Paper, Federal Reserve Bank of New York.
- YAO, R., AND H. H. ZHANG (2005): “Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraint,” *Review of Financial Studies*, 18(1), 197–239, *Review of Financial Studies*.

A Data

Aggregate Time Series Data for ARM Share Our baseline data series is from the Federal Housing Financing Board. It is based on the Monthly Interest Rate Survey (MIRS), a survey sent out to mortgage lenders. The monthly data start in 1985.1 and run until 2006.6, and we label this series $\{ARM_t\}$. Major lenders are asked to report the terms and conditions on all conventional, single-family, fully-amortizing, purchase-money loans closed the last five working days of the month. The data thus excludes FHA-insured and VA-guaranteed mortgages, refinancing loans, and balloon loans. The data for our last sample month, June 2006, are based on 21,801 reported loans from 74 lenders, representing savings associations, mortgage companies, commercial banks, and mutual savings banks. The data are weighted to reflect the shares of mortgage lending by lender size and lender type as reported in the latest release of the Federal Reserve Board's Home Mortgage Disclosure Act data. They are available at <http://www.fhfb.gov/Default.aspx>.

These MIRS data include only new house purchases (for both newly-constructed homes and existing homes), not refinancings. Freddie Mac publishes a monthly index of the share of refinancings in mortgage originations. The average refi share over the 1987.1-2007.1 period is 39.3%. So, purchase-money loans accounts for approximately 60% of the mortgage flow. The sample consists predominantly of conforming loans, only a very small fraction is jumbo mortgages. The ARM share for jumbos in the MIRS sample is much higher on average, but has a 70% correlation with the conforming loans in the sample. While the data do not permit precise statements about the representativeness of the MIRS sample, its ARM share has a correlation of 94% with the ARM share in the Inside Mortgage Finance data. The comparison is for annual data between 1990 and 2006, the longest available sample. We thank Nancy Wallace for making the IMF data available to us.

There is an alternative source of monthly ARM share data available from Freddie-Mac, based on the Primary Mortgage Market Survey. This survey goes out to 125 lenders. The share is constructed based on the dollar volume of conventional mortgage originations within the 1-unit Freddie Mac loan limit as reported under the Home Mortgage Disclosure Act (HMDA) for 2004. Given that Freddie Mac also publishes the aforementioned refinancing share of originations based on the same Primary Mortgage Market Survey, it appears that this series includes not only purchase mortgages but also refinancings. This series is available from 1995.1 and has a correlation with our benchmark measure of 90%.

Loan-level Mortgage Data We explore a new data set which contains information on 911,000 loans from a large mortgage trustee for mortgage-backed security special purpose vehicles. It contains data from many of the largest mortgage lenders such as Aames Capital, Bank of America, Citi Mortgage, Countrywide, Indymac, Option One, Ownit, Wells Fargo, Washington Mutual. We use information on the loan type, the loan origination year and month, the balance, the loan-to-value ratio, the FICO score, and the contract rate at origination. We also have geographic information on the region of origination. We merge these data with our bond risk premium and interest rate variables, with matching based on month of origination. While the sample spans 1994-2007, 95% of mortgage contracts are originated between 2000 and 2005.

Treasury and Mortgage Yields and Inflation Monthly nominal yield data are obtained from the Federal Reserve Bank of New York. They are available at <http://www.federalreserve.gov/pubs/feds/2006>. We use the 1-year ARM rate as our measure of the short mortgage rate and the 30-year FRM rate as our measure of the long-term mortgage rate. We use the effective rate data from the Federal Housing Financing Board, Table 23. The effective rate adjusts the contractual rate for the discounted value of initial fees and charges. The FRM-ARM spreads with and without fees have a correlation of .998. The inflation rate is based on the monthly Consumer Price Index for all urban consumers from the Bureau of Labor Statistics. The inflation data

are available at <http://www.bls.gov>. Real yield data are available as of January 1997 when the US Treasury introduced Treasury Inflation-Protected Securities (TIPS). The real yield data are available from McCulloch at <http://www.econ.ohio-state.edu/jhm/ts/ts.html>.

B Risk-Return Tradeoff

This appendix computes the expected utility from time-1 and time-2 consumption for each of the contracts. We first compute the utility without log transformation, and only at the end, when comparing the two mortgage contracts, reintroduce this log transformation.

Utility from time-1 consumption The (exponent of) utility from time-1 consumption on the FRM contract is:

$$\begin{aligned}\mathbb{E}_0 \left(e^{-\beta-\gamma \frac{C_1}{\Pi_1}} \right) &= \mathbb{E}_0 \left(e^{-\beta-\gamma \left[\frac{L_1^S - q_0^{FRM} B}{\Pi_1} \right]} \right) = \mathbb{E}_0 \left(e^{-\beta-\gamma \left[L_1 - \frac{q_0^{FRM} B}{\Pi_1} \right]} \right) \\ &= e^{-\beta-\gamma \left(\mathbb{E}_0(L_1) - \frac{\gamma \sigma_L^2}{2} - \frac{q_0^{FRM} B}{\Pi_1} \right)}.\end{aligned}$$

For the ARM contract it is:

$$\mathbb{E}_0 \left(e^{-\beta-\gamma \frac{C_1}{\Pi_1}} \right) = \mathbb{E}_0 \left(e^{-\beta-\gamma \left[\frac{L_1^S - q_0^{ARM} B}{\Pi_1} \right]} \right) = e^{-\beta-\gamma \left(\mathbb{E}_0(L_1) - \frac{\gamma \sigma_L^2}{2} - \frac{q_0^{ARM} B}{\Pi_1} \right)}.$$

Utility from time-2 consumption Under the FRM, the time-1 value of the time-2 utility equals:

$$\mathbb{E}_1 \left[e^{-2\beta-\gamma \frac{C_2}{\Pi_2}} \right] = e^{-2\beta-\gamma \left(H_2 + \mathbb{E}_1[L_2] - \frac{\gamma \sigma_L^2}{2} - \frac{(q_0^{FRM} + 1)B}{\Pi_2} \right)},$$

using the same argument as in the period-1 utility calculations.

Next, we calculate the time-0 utility of this time-2 utility:

$$\begin{aligned}\mathbb{E}_0 \left[e^{-2\beta-\gamma \frac{C_2}{\Pi_2}} \right] &= \mathbb{E}_0 \left[e^{-2\beta-\gamma \left(H_2 + \mathbb{E}_0[L_2] + \rho_L \sigma_L \varepsilon_1^L - \frac{\gamma \sigma_L^2}{2} - (q_0^{FRM} + 1)B e^{-x_0 - x_1} \right)} \right] \\ &\simeq \mathbb{E}_0 \left[e^{-2\beta-\gamma \left(H_2 + \mathbb{E}_0[L_2] + \rho_L \sigma_L \varepsilon_1^L - \frac{\gamma \sigma_L^2}{2} - (q_0^{FRM} + 1)B e^{-x_0 - \mathbb{E}_0[x_1] (1 - (x_1 - \mathbb{E}_0[x_1]))} \right)} \right] \\ &= \mathbb{E}_0 \left[e^{-2\beta-\gamma \left(H_2 + \mathbb{E}_0[L_2] + \rho_L \sigma_L \varepsilon_1^L - \frac{\gamma \sigma_L^2}{2} - (q_0^{FRM} + 1)B e^{-x_0 - \mathbb{E}_0[x_1] (1 - \sigma_x \varepsilon_1^x)} \right)} \right] \\ &= e^{-2\beta-\gamma \left(H_2 + \mathbb{E}_0[L_2] - (q_0^{FRM} + 1)B e^{-x_0 - \mathbb{E}_0[x_1]} \right) + \frac{\gamma^2}{2} \left((1 + \rho_L^2) \sigma_L^2 + (q_0^{FRM} + 1)^2 B^2 e^{-2x_0 - 2\mathbb{E}_0[x_1]} \sigma_x^2 \right)}.\end{aligned}$$

In these steps, we used:

$$\begin{aligned}\Pi_2 &= \Pi_1 e^{x_1}, \quad \Pi_1 = e^{x_0}, \\ \mathbb{E}_1(L_2) &= \mu_L + \rho_L (L_1 - \mu_L) = \mu_L + \rho_L^2 (L_0 - \mu_L) + \rho_L \sigma_L \varepsilon_1^L = \mathbb{E}_0(L_2) + \rho_L \sigma_L \varepsilon_1^L, \\ e^{-x_1} &\simeq e^{-\mathbb{E}_0(x_1)} - e^{-\mathbb{E}_0(x_1)} [x_1 - \mathbb{E}_0(x_1)].\end{aligned}$$

For the ARM contract, the time-1 value of the time-2 utility equals:

$$\mathbb{E}_1 \left[e^{-2\beta - \gamma \frac{C_2}{\Pi_2}} \right] = e^{-2\beta - \gamma \left(H_2 + \mathbb{E}_1(L_2) - \frac{\gamma \sigma_L^2}{2} - \frac{(1 + q_1^{ARM})B}{\Pi_2} \right)}.$$

Then for the time-0 value function, it holds:

$$\begin{aligned} \mathbb{E}_0 \left[e^{-2\beta - \gamma \frac{C_2}{\Pi_2}} \right] &\simeq \mathbb{E}_0 \left[e^{-2\beta - \gamma \left(H_2 + \mathbb{E}_0[L_2] + \rho_L \sigma_L \varepsilon_1^L - \frac{\gamma \sigma_L^2}{2} - (1 + q_1^{ARM})B e^{-x_0 - \mathbb{E}_0[x_1]} (1 - \sigma_x \varepsilon_1^x) \right)} \right] \\ &= \mathbb{E}_0 \left[e^{-2\beta - \gamma \left(H_2 + \mathbb{E}_0[L_2] + \rho_L \sigma_L \varepsilon_1^L - \frac{\gamma \sigma_L^2}{2} - B(\mathbb{E}_0[q_1^{ARM}] + 1 + q_1^{ARM} - \mathbb{E}_0[q_1^{ARM}]) e^{-x_0 - \mathbb{E}_0[x_1]} (1 - \sigma_x \varepsilon_1^x) \right)} \right] \\ &\simeq \mathbb{E}_0 \left[e^{-2\beta - \gamma \left(H_2 + \mathbb{E}_0[L_2] + \rho_L \sigma_L \varepsilon_1^L - \frac{\gamma \sigma_L^2}{2} - B(\mathbb{E}_0[q_1^{ARM}] + 1 + \sigma' \varepsilon_1) e^{-x_0 - \mathbb{E}_0[x_1]} (1 - \sigma_x \varepsilon_1^x) \right)} \right] \\ &= e^{-2\beta - \gamma (H_2 + \mathbb{E}_0[L_2] - B(\mathbb{E}_0[q_1^{ARM}] + 1) e^{-x_0 - \mathbb{E}_0[x_1]}) + \frac{\gamma^2}{2} [(1 + \rho_L^2) \sigma_L^2]} \times \\ &\quad \mathbb{E}_0 \left[e^{-\gamma B(\mathbb{E}_0[q_1^{ARM}] + 1) e^{-x_0 - \mathbb{E}_0[x_1]} \sigma_x \varepsilon_1^x + \gamma B e^{-x_0 - \mathbb{E}_0[x_1]} \sigma' \varepsilon_1 - \gamma B e^{-x_0 - \mathbb{E}_0[x_1]} (\sigma' \varepsilon_1) \sigma_x \varepsilon_1^x} \right] \\ &\simeq e^{-2\beta - \gamma (H_2 + \mathbb{E}_0[L_2] - B(\mathbb{E}_0[q_1^{ARM}] + 1) e^{-x_0 - \mathbb{E}_0[x_1]})} \times \\ &\quad e^{\frac{\gamma^2}{2} [(1 + \rho_L^2) \sigma_L^2 + B^2 (\mathbb{E}_0[q_1^{ARM}] + 1)^2 e^{-2x_0 - 2\mathbb{E}_0[x_1]} \sigma_x^2 + B^2 e^{-2x_0 - 2\mathbb{E}_0[x_1]} \sigma' R \sigma - 2B^2 (\mathbb{E}_0[q_1^{ARM}] + 1) e^{-2x_0 - 2\mathbb{E}_0[x_1]} (\sigma_x e_2' R \sigma)]}. \end{aligned}$$

The last approximation assumes that $\gamma e^{-x_0 - \mathbb{E}_0[x_1]} (\sigma' \varepsilon_1) \sigma_x \varepsilon_1^x$ is zero (a shock times a shock). $(\sigma_x e_2' R \sigma)$ is the covariance of x and y^S , where we defined $e_2 = [0, 1]'$. In the third line of the approximation, we use $q_1^{ARM} \simeq y_1^S(1)$.

Now we reintroduce the log transformation to the exponential preferences. Households prefer the ARM if and only if the life-time utility of the ARM contract exceeds that of the FRM contract:

$$\begin{aligned} &\beta + \gamma \left(\mathbb{E}_0(L_1) - \frac{\gamma \sigma_L^2}{2} - \frac{q_0^{ARM} B}{\Pi_1} \right) \\ &+ 2\beta + \gamma \left(H_2 + \mathbb{E}_0[L_2] - B(\mathbb{E}_0[q_1^{ARM}] + 1) e^{-x_0 - \mathbb{E}_0[x_1]} \right) \\ &- \frac{\gamma^2}{2} \left[(1 + \rho_L^2) \sigma_L^2 + B^2 (\mathbb{E}_0[q_1^{ARM}] + 1)^2 e^{-2x_0 - 2\mathbb{E}_0[x_1]} \sigma_x^2 \right. \\ &\quad \left. + B^2 e^{-2x_0 - 2\mathbb{E}_0[x_1]} \sigma' R \sigma - 2B^2 (\mathbb{E}_0[q_1^{ARM}] + 1) e^{-2x_0 - 2\mathbb{E}_0[x_1]} (\sigma_x e_2' R \sigma) \right] \\ &> \\ &\beta + \gamma \left(\mathbb{E}_0(L_1) - \frac{\gamma \sigma_L^2}{2} - \frac{q_0^{FRM} B}{\Pi_1} \right) \\ &+ 2\beta + \gamma \left(H_2 + \mathbb{E}_0[L_2] - (q_0^{FRM} + 1) B e^{-x_0 - \mathbb{E}_0[x_1]} \right) - \frac{\gamma^2}{2} \left((1 + \rho_L^2) \sigma_L^2 + (q_0^{FRM} + 1)^2 B^2 e^{-2x_0 - 2\mathbb{E}_0[x_1]} \sigma_x^2 \right). \end{aligned}$$

This simplifies to:

$$\begin{aligned} &q_0^{FRM} - q_0^{ARM} + (q_0^{FRM} - \mathbb{E}_0[q_1^{ARM}]) e^{-\mathbb{E}_0[x_1]} \\ &> \\ &\frac{\gamma}{2} B e^{-x_0 - 2\mathbb{E}_0[x_1]} \left[\sigma' R \sigma + (\mathbb{E}_0[q_1^{ARM}] + 1)^2 \sigma_x^2 - 2(\mathbb{E}_0[q_1^{ARM}] + 1) (\sigma_x e_2' R \sigma) \right] \\ &- \frac{\gamma}{2} B e^{-x_0 - 2\mathbb{E}_0[x_1]} (q_0^{FRM} + 1)^2 \sigma_x^2. \end{aligned}$$

Simplifying Expressions The first term on the right-hand side of the inequality, i.e., the risk induced by the ARM contract, can be rewritten as:

$$\begin{aligned}
& \frac{\gamma}{2} B e^{-x_0 - 2\mathbb{E}_0[x_1]} \left[\sigma' R \sigma + (\mathbb{E}_0 [q_1^{ARM}] + 1)^2 \sigma_x^2 - 2 (\mathbb{E}_0 [q_1^{ARM}] + 1) (\sigma_x e'_2 R \sigma) \right] \\
= & \frac{\gamma}{2} B e^{-x_0 - 2\mathbb{E}_0[x_1]} \left[\sigma_y^2 - 2\sigma_x \sigma_y \rho_{xy} \mathbb{E}_0 [q_1^{ARM}] + \mathbb{E}_0 [q_1^{ARM}]^2 \sigma_x^2 \right] \\
\simeq & \frac{\gamma}{2} B e^{-x_0 - 2\mathbb{E}_0[x_1]} \sigma_y^2,
\end{aligned}$$

in which we use that $2\sigma_x \sigma_y \rho_{xy} \mathbb{E}_0 [q_1^{ARM}]$ and $\mathbb{E}_0 [q_1^{ARM}]^2 \sigma_x^2$ are an order of magnitude smaller than σ_y^2 , which motivates the approximation in the third line. This in turn implies that the ARM contract primarily carries real rate risk, while, in contrast, the FRM contract carries only inflation risk. This is the risk-return trade-off discussed in the main text.

Ignoring the $e^{-\mathbb{E}_0[x_1]}$ inflation term, the left-hand side of above inequality is the difference in expected nominal payments per dollar mortgage balance. We have:

$$2q_0^{FRM} - q_0^{ARM} - \mathbb{E}_0 (q_1^{ARM}) \simeq 2y_0^\$(2) - y_0^\$(1) - \mathbb{E}_0 [y_1^\$(1)] = 2\phi_0^\$(2)$$

where we use the approximations of Section 2.3.

C Derivation of the Prepayment Option Formula

The value of the prepayment option is given by:

$$\begin{aligned}
B\mathbb{E}_0 \left[M_1^\$ M_2^\$ \max \{ (q_0^{FRMP} - q_1^{ARM}), 0 \} \right] &= B\mathbb{E}_0 \left[\mathbb{E}_1 \left[M_1^\$ M_2^\$ \max \{ (q_0^{FRMP} - q_1^{ARM}), 0 \} \right] \right] \\
&= B\mathbb{E}_0 \left[M_1^\$ \max \{ (q_0^{FRMP} - q_1^{ARM}) P_1^\$(1), 0 \} \right] \\
&= B\mathbb{E}_0 \left[M_1^\$ \max \left\{ \left(1 + q_0^{FRMP} - P_1^\$(1)^{-1} \right) P_1^\$(1), 0 \right\} \right] \\
&= B\mathbb{E}_0 \left[M_1^\$ \max \left\{ \left((1 + q_0^{FRMP}) P_1^\$(1) - 1 \right), 0 \right\} \right] \\
&= B (1 + q_0^{FRMP}) \mathbb{E}_0 \left[M_1^\$ \max \left\{ \left(P_1^\$(1) - \frac{1}{1 + q_0^{FRMP}} \right), 0 \right\} \right]
\end{aligned}$$

where we use that $q_1^{ARM} = P_1^\$(1)^{-1} - 1$. The pricing kernel and the one-year bond price at time $t = 1$ are given by:

$$\begin{aligned}
M_1^\$ &= e^{-y_0^\$(1) - \frac{1}{2} \lambda'_0 R \lambda_0 - \lambda'_0 \varepsilon_1} \\
P_1^\$(1) &= e^{-y_1^\$(1)} = e^{-\mathbb{E}_0[y_1^\$(1)] - \sigma' \varepsilon_1}
\end{aligned}$$

We project the innovation to the pricing kernel on the innovation to the nominal short rate:

$$\begin{aligned}
\eta_1 &\equiv \sigma' \varepsilon_1 \\
\eta_2 &\equiv \lambda'_0 \varepsilon_1 - \frac{\text{Cov}(\eta_1, \lambda'_0 \varepsilon_1)}{\text{Var}(\eta_1)} \eta_1 = \lambda'_0 \varepsilon_1 - \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \eta_1
\end{aligned}$$

with η_1 and η_2 orthogonal and variances given by:

$$\text{Var}[\eta_1] = \sigma' R \sigma, \quad \text{Var}[\eta_2] = \lambda'_0 R \lambda_0 - \frac{(\sigma' R \lambda_0)^2}{\sigma' R \sigma}$$

We first solve for the value of one call option for a general exercise price K , denoted by $C_0(K)$ (see Black and Scholes (1973), Merton (1973), and Jamshidian (1989)):

$$\begin{aligned}
C_0(K) &= \mathbb{E}_0 \left[M_1^{\mathbb{S}} \max \left\{ \left(P_1^{\mathbb{S}}(1) - K \right), 0 \right\} \right] \\
&= \mathbb{E}_0 \left[e^{-y_0^{\mathbb{S}}(1) - \frac{1}{2} \lambda_0' R \lambda_0 - \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \eta_1 - \eta_2} \max \left\{ \left(e^{-\mathbb{E}_0[y_1^{\mathbb{S}}(1)] - \eta_1} - K \right), 0 \right\} \right] \\
&= \mathbb{E}_0 \left[e^{-y_0^{\mathbb{S}}(1) - \frac{1}{2} \lambda_0' R \lambda_0 - \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \eta_1} \max \left\{ \left(e^{-\mathbb{E}_0[y_1^{\mathbb{S}}(1)] - \eta_1} - K \right), 0 \right\} \right] e^{\frac{1}{2} \left(\lambda_0' R \lambda_0 - \frac{(\sigma' R \lambda_0)^2}{\sigma' R \sigma} \right)}
\end{aligned}$$

The option will be exercised if and only if the following holds

$$\eta_1 < -\log(K) - \mathbb{E}_0 \left[y_1^{\mathbb{S}}(1) \right],$$

which occurs with probability

$$\Phi \left(\frac{-\log(K) - \mathbb{E}_0 \left[y_1^{\mathbb{S}}(1) \right]}{\sqrt{\sigma' R \sigma}} \right) \equiv \Phi(x^*).$$

We proceed:

$$\begin{aligned}
C_0(K) &= e^{\frac{1}{2} \left(\lambda_0' R \lambda_0 - \frac{(\sigma' R \lambda_0)^2}{\sigma' R \sigma} \right)} \mathbb{E}_0 \left[e^{-y_0^{\mathbb{S}}(1) - \frac{1}{2} \lambda_0' R \lambda_0 - \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \eta_1} \left(e^{-\mathbb{E}_0[y_1^{\mathbb{S}}(1)] - \eta_1} - K \right) I_{(\eta_1 / \sqrt{\sigma' R \sigma} < x^*)} \right] \\
&= e^{\frac{1}{2} \left(\lambda_0' R \lambda_0 - \frac{(\sigma' R \lambda_0)^2}{\sigma' R \sigma} \right)} \int_{-\infty}^{x^*} e^{-y_0^{\mathbb{S}}(1) - \mathbb{E}_0[y_1^{\mathbb{S}}(1)] - \frac{1}{2} \lambda_0' R \lambda_0 - \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \sqrt{\sigma' R \sigma} x - \sqrt{\sigma' R \sigma} x} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2} x^2} dx \\
&\quad - e^{\frac{1}{2} \left(\lambda_0' R \lambda_0 - \frac{(\sigma' R \lambda_0)^2}{\sigma' R \sigma} \right)} \int_{-\infty}^{x^*} K e^{-y_0^{\mathbb{S}}(1) - \frac{1}{2} \lambda_0' R \lambda_0 - \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \sqrt{\sigma' R \sigma} x} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2} x^2} dx,
\end{aligned}$$

where we use that $\eta_1 / \sqrt{\sigma' R \sigma}$ is standard normally distributed. Rewriting and using that:

$$-2y_0^{\mathbb{S}}(2) = -y_0^{\mathbb{S}}(1) - \mathbb{E}_0 \left[y_1^{\mathbb{S}}(1) \right] + \frac{1}{2} \sigma' R \sigma + \sigma' R \lambda_0,$$

we obtain:

$$\begin{aligned}
C_0(K) &= P_0^{\mathbb{S}}(2) \int_{-\infty}^{x^*} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2} \left(x + \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \sqrt{\sigma' R \sigma} + \sqrt{\sigma' R \sigma} \right)^2} dx - K P_0^{\mathbb{S}}(1) \int_{-\infty}^{x^*} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2} \left(x + \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \sqrt{\sigma' R \sigma} \right)^2} dx \\
&= P_0^{\mathbb{S}}(2) \Phi \left(x^* + \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \sqrt{\sigma' R \sigma} + \sqrt{\sigma' R \sigma} \right) - K P_0^{\mathbb{S}}(1) \Phi \left(x^* + \frac{\sigma' R \lambda_0}{\sigma' R \sigma} \sqrt{\sigma' R \sigma} \right),
\end{aligned}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. Using the definition of x^* , we conclude that the option value is given by:

$$\begin{aligned}
C_0(K) &= P_0^{\mathbb{S}}(2) \Phi(d_1) - K P_0^{\mathbb{S}}(1) \Phi(d_2), \\
d_1 &\equiv \frac{-\log(K) - \mathbb{E}_0 \left[y_1^{\mathbb{S}}(1) \right] + \sigma' R \sigma + \sigma' R \lambda_0}{\sqrt{\sigma' R \sigma}}, \\
&= \frac{\log(P_0^{\mathbb{S}}(2)/K) + y_0^{\mathbb{S}}(1) + \frac{1}{2} \sigma' R \sigma}{\sqrt{\sigma' R \sigma}}, \\
d_2 &\equiv d_1 - \sqrt{\sigma' R \sigma},
\end{aligned}$$

where the second line for d_1 uses the pricing formula of a two-period bond. Now using $K = 1/(1 + q_0^{FRMP})$ and the fact that the investor has $B(1 + q_0^{FRMP})$ of these options, yields the value of the prepayment option:

$$B\mathbb{E}_0 \left[M_1^\$ M_2^\$ \max \{ (q_0^{FRMP} - q_1^{ARM}), 0 \} \right] = B(1 + q_0^{FRMP}) C_0 (1 / (1 + q_0^{FRMP})). \quad (24)$$

D Multi-Period Model

In this appendix, we consider a more realistic, multi-period extension of the simple model in Section 2. It has power utility preferences and features an exogenous moving probability. We use this model (i) to study the role of the time discount factor and the moving rate, and (ii) to solve for the relationship between the cross-sectional distribution over risk aversion parameters and the aggregate ARM share.

D.1 Setup

The household problem Household j chooses the mortgage contract, $i \in \{FRM, ARM\}$, to maximize expected lifetime utility over real consumption:

$$U_j^i = E_0 \sum_{t=1}^T \beta^t (1 - \xi)^{t-1} \frac{(C_t^i)^{1-\gamma_j}}{1 - \gamma_j}, \quad (25)$$

$$C_t^i = L - q_t^i / \Pi_t, \text{ for } t \in \{1, \dots, T-1\} \quad (26)$$

$$C_T^i = 1 + L - (1 + q_T^i) / \Pi_T \quad (27)$$

where β is the (monthly) subjective discount rate, ξ is the (monthly) exogenous moving rate, and γ_j is the coefficient of relative risk aversion. We consider constant real labor income L . We normalize the nominal outstanding balance to one, which makes q_t^i both the nominal mortgage rate and nominal mortgage payment at time t for contract i . This setup incorporates utility up until a move. The certainty-equivalent consumption, \tilde{C}^i , is given by:

$$\tilde{C}_j^i = \left(U_j^i / \sum_{t=1}^T \frac{\beta^t (1 - \xi)^{t-1}}{1 - \gamma_j} \right)^{1/(1-\gamma_j)}. \quad (28)$$

We are interested in the certainty-equivalent consumption differential $\tilde{C}_j^{ARM} - \tilde{C}_j^{FRM}$.

Bond Pricing Following Kojien, Nijman, and Werker (2007), we consider a continuous-time, two-factor essentially affine term structure model. The factors $X_t = [Z_{1t}, Z_{2t}]'$ are identified with the real rate and expected inflation, respectively. The model can be discretized exactly to a VAR(1)-model:

$$Z_t = \mu + \Phi Z_{t-1} + \Sigma \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, I_{3 \times 3}), \quad (29)$$

where the third element of the state is realized inflation, $Z_{3t} = \log \Pi_t - \log \Pi_{t-1}$. The τ -month bond price at time t is exponentially affine in X_t :

$$P_t^\$(\tau) = \exp \{ A_\tau + B_\tau' X_t \}, \quad (30)$$

where $A_\tau = A(\tau/12)$ and $B_\tau = B(\tau/12)$, with $A(\cdot)$ and $B(\cdot)$ derived in Appendix A of Kojien, Nijman, and Werker (2007).

Mortgage Pricing At time t the lender of the FRM receives

$$q^{FRM} (1 - \xi)^{t-1} + (1 - \xi)^{t-1} \xi, \quad (31)$$

where $(1 - \xi)^{t-1}$ is the probability that loan has not been prepaid before time t and $(1 - \xi)^{t-1} \xi$ is the probability it is prepaid at time t . Imposing a zero-profit condition, a mortgage contract of T periods has the following FRM rate:

$$q^{FRM} = \frac{1 - \sum_{t=1}^{T-1} (1 - \xi)^{t-1} \xi P_0^s(t) - (1 - \xi)^{T-1} P_0^s(T)}{\sum_{t=1}^{T-1} (1 - \xi)^{t-1} P_0^s(t) + (1 - \xi)^{T-1} P_0^s(T)}. \quad (32)$$

For the monthly ARM rate we have $q_t^{ARM} = P_t^s(1)^{-1} - 1$.

D.2 Calibration

The term structure parameters are taken from Koijen, Nijman, and Werker (2007). As is the case for the VAR estimates in the main text, the correlation between the yield spread and the bond risk premium is low in the model (-7%). Real labor income, L , is held constant at 0.42. To obtain a theoretically well-defined problem we assume a minimum subsistence consumption level of 0.05/12 per month. The exogenous monthly moving probability is set at 1% per month ($(1 - \xi)^{12} - 1 = 11.36\%$ per year). We consider different values for the coefficient of relative risk aversion, γ , and the monthly subjective discount factor, β .

D.3 Effect of the Subjective Discount Factor and Moving Rates

We generate $N = 1000$ starting values for the state vector at time zero, Z_0 , by simulating forward $M = 60$ months from the unconditional mean for the state vector ($0_{4 \times 1}$) for each of the N paths. Next, we compute the expected utility differential of the ARM and FRM contracts. Expected utilities are computed by averaging realized utilities in $K = 100$ simulated paths (where the same shocks apply to all $N = 1000$ starting values).

Figure 12 plots the R^2 of regressing the model's certainty-equivalent consumption differential between the ARM and FRM contracts on the model's bond risk premium (solid line) or on the model's yield spread (dashed line). Each point corresponds to a different value of the annualized subjective time discount factor β^{12} , between 0.5 and 1. The coefficient of relative risk aversion is set at $\gamma = 5$. For low values of the subjective discount factor ($\beta < .70$), the slope of the yield curve has a stronger relationship to the relative desirability of the ARM. However, for more realistic and more conventional values of the subjective discount factor, say between 0.9 and 1.0, the bond risk premium is the key determinant of mortgage choice. We have also experimented with an upward sloping labor income profile, as in Cocco, Gomes, and Maenhout (2005), and found a similar cut-off rule. A similar result holds when we vary the moving rate instead of the subjective time discount factor: below 10% per month, the risk premium is the more important predictor. For empirically relevant moving rates below 2%, the risk premium is the only relevant predictor.

[Figure 12 about here.]

D.4 Heterogeneous Risk Aversion Level

For each month in our sample period we determine the level of risk aversion that makes an investor indifferent between the ARM and the FRM. Starting values for the vector of state variables, Z , are from Koijen, Nijman, and Werker (2007). The utility differential of an ARM and an FRM is computed as described above. The monthly

subjective discount factor is set at $\beta = 0.96^{1/12} \approx 0.9966$. We assume a log-normal cross-sectional distribution for the risk aversion level:

$$\log(\gamma) \sim \mathcal{N}(\mu_\gamma, \sigma_\gamma^2), \quad (33)$$

which implies that our model predicts the following ARM share:

$$ARM_t^{pred}(\log(\gamma_t^*); \mu_\gamma, \sigma_\gamma) = \Phi\left(\frac{\log(\gamma_t^*) - \mu_\gamma}{\sigma_\gamma}\right) \quad (34)$$

where Φ is the standard normal cumulative density function and where households with a risk aversion smaller than the cutoff γ_t^* choose the ARM. More conservative households choose the FRM.

We determine μ_γ and σ_γ by minimizing the squared prediction error over the sample period (1985:1-2005:12) and estimate a location parameter $\hat{\mu}_\gamma = 5.0$ and a scale parameter $\hat{\sigma}_\gamma = 2.9$. The median level of risk aversion implied by this distribution equals $\exp(\hat{\mu}_\gamma) = 155$. Interestingly, regressing the actual ARM share on the predicted ARM share yields a constant and slope coefficient of 0.03 and 0.90 respectively, which are not significantly different from theoretical implied values of 0 and 1 respectively.

The cutoff log risk aversion level has a sample mean of $\mu_{\gamma^*} = 3.37$ and a sample standard deviation of $\sigma_{\gamma^*} = 0.73$. The predicted increase in the ARM share from a one standard deviation increase in the log indifference risk aversion level around its mean is given by:

$$ARM^{pred}(\mu_{\gamma^*} + 0.5\sigma_{\gamma^*}; \hat{\mu}_\gamma, \hat{\sigma}_\gamma) - ARM^{pred}(\mu_{\gamma^*} - 0.5\sigma_{\gamma^*}; \hat{\mu}_\gamma, \hat{\sigma}_\gamma) = 8.6\% \quad (35)$$

This 8.6% is very close to the slope coefficient we reported in Table 2, Rows 3-6. In conclusion, the model can explain the observed average 28% ARM share and the observed sensitivity of the ARM share to the bond risk premium with a mean log risk aversion of 5 and a standard deviation of log risk aversion of 2.9.

We conjecture that these values would be lower in a model where labor income or house price risk were negatively correlated with the real rate. In that case, the ARM would be more risky because ARM payments would be high when labor income or the house price is low. A lower risk aversion would be needed to choose the FRM. Put differently, the (relatively low) observed ARM share could be justified with a lower mean risk aversion.

E Model for Default Risk Premium on the ARM

We consider a continuum of households with mass unity and associated measure $\nu(i)$. Default of household i is given by:

$$I_{Default,it} = \Phi(\beta_0 + \beta' X_t) + \eta_t + \zeta_{it}, \quad (36)$$

where $E[\eta_t] = E[\zeta_{it}] = E[\eta_t \zeta_{it}] = 0$. According to this model, the default is governed by a common component that captures covariance with the aggregate state X_t ($\Phi(\beta_0 + \beta' X_t)$), a common component that is independent of the state (η_t), and a household-specific component that is independent of all other terms (ζ_{it}). The function $\Phi(\cdot)$ is the cumulative density function of a normal distribution. It is a non-linear function of the covariates X_t , with the nice feature that its range is between zero and one.

We collect foreclosure rate data on ARM contracts from the Mortgage Bankers Association. The data are quarterly from 2002.I-2007.IV. We use spline interpolation to obtain a monthly series over this period. We studied various candidates for the vector X_t . The 1-year lagged 1-year CMT rate and the log house price-to-rent ratio turn out to be the best predictors of the ARM foreclosure rate. A non-linear least-squares estimation results in coefficient

estimates of $\beta_0 = -2.87$ and $\beta_1 = [9.97, 0.06]$. The R^2 , which we define as the variance of $\Phi(\beta_0 + \beta' X_t)$ divided by the variance of the foreclosure rate is 57.6%. The main effect is increasing ARM foreclosure rates 12 months after a rise in the interest rate. Contemporaneous interest rates and the house price-to-rent ratio have essentially no explanatory power for ARM foreclosures.

The yield on an ARM contract is the yield on a defaultable bond which pays \$1 when no default takes place and $\$(1-LGD)$ when default takes place, where LGD stands for loss-given-default:

$$\begin{aligned} \exp(-y_t^{ARM}) &= \int_{\Omega} E_t [M_{t+1} (I_{Default,i,t+1}(1 - LGD) + (1 - I_{Default,i,t+1}))] d\nu(i) \\ &= E_t \left[M_{t+1} \left(\int_{\Omega} I_{Default,i,t+1}(1 - LGD) d\nu(i) + (1 - \int_{\Omega} I_{Default,i,t+1} d\nu(i)) \right) \right] \\ &= E_t [M_{t+1} (\Phi(\beta_0 + \beta' X_{t+1})(1 - LGD) + (1 - \Phi(\beta_0 + \beta' X_{t+1})))] , \end{aligned}$$

where M is the stochastic discount factor. We assume that the LGD is constant over time, and we vary its magnitude to examine the robustness of our results. The key observation is that because the predictors in X_{t+1} are the one-year lagged interest rate and price-rent ratio, they are known at time t , and they can be brought outside of the time- t expectation operator:

$$\begin{aligned} \exp(-y_t^{ARM}) &= E_t [M_{t+1} (\Phi(\beta_0 + \beta' X_{t+1})(1 - LGD) + (1 - \Phi(\beta_0 + \beta' X_{t+1})))] \\ &= (\Phi(\beta_0 + \beta' X_{t+1})(1 - LGD) + (1 - \Phi(\beta_0 + \beta' X_{t+1}))) E_t [M_{t+1}] \\ &= (\Phi(\beta_0 + \beta' X_{t+1})(1 - LGD) + (1 - \Phi(\beta_0 + \beta' X_{t+1}))) \exp(-y_t(1)), \end{aligned}$$

showing that:

$$y_t^{ARM} = y_t(1) + \underbrace{\log(1 - \Phi(\beta_0 + \beta' X_{t+1})LGD)}_{def_t^{ARM}}. \quad (37)$$

To develop intuition for equation (37), we consider a first-order Taylor expansion around $LGD = 0$:

$$y_t^{ARM} = y_t(1) + \Phi(\beta_0 + \beta' X_{t+1})LGD.$$

Given that $\Phi(\cdot)$ is 0.008, averaged over 2002.1-2007.12, and a typical LGD for mortgages is 30%, the approximation implies a default risk premium of 24 basis points. Since we are dealing with prime mortgages, the risk premia are naturally small. For our main exercise, we use the coefficient estimates β_0 and β_1 and data on the predictors X_t for the full sample to form a time-series for the default risk premium from 1985.1 until 2006.6, using (the exact) equation (37).

We consider loss-given-default rates between 20% and 40%, at the upper end of the range of estimates found in the literature. Qi and Yang (2007) estimate loss-given-default (LGD) between 1990 and 2003 for a large loan-level data set. The nationwide mean LGD for loans with a combined LTV ratio of 100% or less is 20.2%, while it is only 2.1% with a CLTV ratio of 80% or less. The LGD is 24.3% instead of 20.2% in 1990-1994, a period with falling real house prices.

Table 1: Probit Regressions of the ARM Share in Loan-Level Data

This table reports slope coefficients, robust t-statistics (in brackets), and R^2 statistics for probit regressions of an ARM dummy on a constant and one or more regressors, reported in the first column. The regressors are $\kappa(3, 5)$, the household decision rule formed with a 5-year Treasury yield and a 3-year average of past 1-year Treasury yield data, the loan balance at origination (BAL), the loan's credit score at origination (FICO), the loan's loan-to-value ratio (LTV), the long-term interest rate (5-year Treasury yield), and the 5-1 year Treasury yield spread. The seventh column indicates when we include four regional dummies for the biggest mortgage markets (California, Florida, New York, and Texas). All independent variables have been normalized by their standard deviation. The sample consists of 654,368 mortgage loans originated between 1994-2006.6.

$\kappa_t(3; 5)$	$y^s(5) - y^s(1)$	$y^s(5)$	BAL	FICO	LTV	Regional dummies	% correctly classified
0.43						No	69.4
[253]							
			-0.05	-0.05	0.17	Yes	61.7
			[21]	[28]	[100]		
0.42			-0.01	-0.08	0.13	Yes	68.8
[244]			[4]	[45]	[72]		
	0.06					No	59.8
	[38]						
	0.09		-0.05	-0.06	0.19	Yes	62.1
	[53]		[23]	[30]	[106]		
0.65	0.43		-0.00	-0.11	0.17	Yes	70.9
[299]	[206]		[2]	[58]	[90]		
		-0.30				No	64.7
		[171]					
		-0.33	-0.05	-0.09	0.20	Yes	66.6
		[179]	[22]	[46]	[110]		
0.54		-0.47	-0.00	-0.15	0.16	Yes	71.6
[290]		[237]	[1]	[71]	[80]		

Table 2: The ARM Share and the Nominal Bond Risk Premium

This table reports slope coefficients, Newey-West t-statistics (12 lags), and R^2 statistics for regressions of the ARM share on a constant and the regressors reported in the first column. The regressors are the τ -year nominal bond risk premium $\phi_t^{\$}(\tau)$, measured in three different ways. We consider $\tau = 5$ and $\tau = 10$ years. The first measure is based on the household decision rule with a 3-year look-back period (rows 1-2). The second measure is based on Blue Chip forecast data (rows 3 and 4) and the third measure is based on the VAR (rows 5-6). Rows 7 and 8 show regressions of the ARM share on the τ -1-year yield spread $y_t^{\$(\tau)} - y_t^{\$(12)}$. Rows 9 and 10 use the τ -year nominal yield, $y_t^{\$(\tau)}$, as predictor. Rows 11 and 12 use the household decision rule computed using the effective 30-year FRM rate and the effective 1-year ARM rate, with a look-back period of 2 years in Row 11 and three years in Row 12. Row 13 uses the difference between the FRM rate $y_t^{\$(FRM)}$ and the ARM rate $y_t^{\$(ARM)}$, while row 15 uses $y_t^{\$(FRM)}$ as independent variable. Row 14 uses the component of the FRM-ARM spread that is orthogonal to the 10-1 Treasury bond spread. Rows 16 and 17 consider two other adaptive expectation measures of the bond risk premium which use geometrically-declining weights on past short rates ($\gamma = .97$). Rows 18 and 19 consider two other rules-of-thumb based on mortgage rates. The FRM rule takes the current FRM rate minus the three-year moving average of the FRM rate (row 18). The ARM rule in Row 19 does the same for the ARM rate. In all rows, the regressor is lagged by one period, relative to the ARM share. All independent variables have been normalized by their standard deviation. The sample is 1985.1-2006.6, except for rows 1 and 2, 11 and 12, and 16 and 17, where we use 1989.12-2006.6, the sample for which the household decision rules are available.

			slope	t-stat	R^2
1.	Househ. Decis. Rule	$\kappa_t(3, 5)$	7.88	7.08	71.23
2.		$\kappa_t(3, 10)$	7.70	7.47	68.03
3.	Blue Chip	$\phi_t^{\$(5)}$	8.67	3.91	40.25
4.		$\phi_t^{\$(10)}$	8.89	4.22	42.62
5.	VAR	$\phi_t^{\$(5)}$	7.73	4.16	32.21
6.		$\phi_t^{\$(10)}$	8.07	3.91	35.13
7.	Slope	$y_t^{\$(5)} - y_t^{\$(1)}$	0.46	0.21	0.11
8.		$y_t^{\$(10)} - y_t^{\$(1)}$	-0.66	-0.32	0.23
9.	Long yield	$y_t^{\$(5)}$	8.37	3.76	37.76
10.		$y_t^{\$(10)}$	8.53	3.85	39.26
11.	Mortgage rates	$\kappa_t(2, FRM)$	7.26	9.37	60.40
12.		$\kappa_t(3, FRM)$	6.28	4.99	45.28
13.		$y_t^{\$(FRM)} - y_t^{\$(ARM)}$	8.09	3.17	35.31
14.		$y_t^{\$(FRM)} - y_t^{\$(ARM)}$ orth.	8.75	3.86	41.28
15.		$y_t^{\$(FRM)}$	7.81	3.71	32.87
16.	Alt. adap. exp. rule	$\hat{\kappa}_t(.97, 5)$	7.60	6.25	66.24
17.		$\hat{\kappa}_t(.97, 10)$	8.14	8.93	75.94
18.	Other Rules-of-Thumb	FRM rule	6.00	3.74	22.54
19.		ARM rule	3.13	2.42	6.12

Table 3: Analysis of Composition Effects in Loan Level Data

This table reports slope coefficients, t-statistics (in brackets), and R^2 statistics from fixed-effect panel regressions. the left-hand side variable is the ARM share in a given month for a given group. The right-hand side variables are the household decision rule $\kappa_t(3; 5)$ and/or group-averaged loan characteristics (loan balance BAL, the FICO score, and the loan-to-value ratio LTV). The last two columns report the number of groups N and the average number of months of data T per group in the sample. The first three rows use U.S. states as groups. Row four is for real loan balance groups, row five is for real house price groups, row six is for FICO score groups, and row seven is for real (ZIP-code level) income groups.

Row	constant	$\kappa_t(3; 5)$	BAL	FICO	LTV	R^2	N	T
1.	.280 [66]	.162 [46]				36.1	50	73.7
2.	-.50 [-7]	.159 [43]	.006 [0.9]	.089 [13]	-.040 [-4.8]	40.1	50	73.7
3.	-1.254 [-14.8]		-0.053 [-6.2]	.140 [17]	0.048 [4.8]	10.8	50	73.7
4.	.292 [40]	.123 [20]				19.5	14	107.6
5.	.282 [32]	.144 [19]				23.1	10	108.3
6.	.248 [31]	.151 [21]				26.9	10	118.0
7.	.260 [34]	.159 [24]				31.8	10	114.3

Table 4: Bond Return Predictability

This table reports R^2 statistics from OLS regressions. The dependent variable in the first column is the ARM share. The dependent variable in Columns 2-6 is the excess return on an n -year bond held for τ years, where $\tau = 1, \dots, 5$. In Panel A, $n = 10$ and in Panel B, $n = 5$. Along the rows in each panel, we vary the independent variable, the adaptive expectations proxy for the bond risk premium $\hat{\kappa}_t(\gamma, n) \equiv y_t^S(n) - \mu_t(\gamma)$, where $\mu_t(\gamma) = (1 - \gamma) \sum_{i=0}^{\infty} \gamma^i y_{t-i}^S(1)$. We do so by varying its parameter γ . In Panel A, the ten-year yield ($n = 10$) is used in the proxy, whereas in Panel B, the five-year yield ($n = 5$) is used. The sample for the ARM share regression starts in 1989.12 for the adaptive expectations measures and in 1985.1 for the other forecast variables, as in the paper. The sample for the excess bond return predictability regressions starts in 1985.1 in all rows, except in 1989.12 for $\kappa(3, 5)$ and $\kappa(3, 10)$.

Panel A: 10-year bond						
dependent var:	ARM share	Excess bond returns				
γ	Column 1	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$	$\tau = 5$
0	0.32%	8.04%	2.96%	0.04%	0.72%	0.97%
0.9	28.46%	11.60%	3.94%	1.02%	2.32%	3.22%
0.95	61.56%	18.65%	10.55%	8.54%	11.74%	12.19%
0.97	75.94%	30.49%	21.97%	21.39%	25.37%	25.92%
0.99	58.10%	49.83%	38.32%	38.46%	43.48%	47.90%
0.999	11.14%	21.36%	14.82%	15.21%	16.21%	23.23%
$\kappa(3, 10)$	68.03%	2.53%	0.06%	0.57%	0.08%	4.48%
Blue Chip	42.62%	16.89%	8.86%	6.67%	6.47%	10.20%
VAR	35.13%	29.57%	22.75%	16.33%	12.78%	15.03%
CP	0.62%	34.86%	22.44%	10.57%	6.31%	5.03%
Panel B: 5-year bond						
dependent var:	ARM share	Excess bond returns				
γ	Column 1	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$	$\tau = 5$
0	0.01%	9.22%	0.28%	1.19%	1.81%	1.81%
0.9	57.8%	3.38%	0.06%	0.29%	0.69%	0.69%
0.95	77.45%	5.4%	3.7%	7.92%	8.4%	8.4%
0.97	66.24%	10.68%	12.09%	19.63%	20.33%	20.33%
0.99	36.26%	23.82%	28.02%	36.78%	39.79%	39.79%
0.999	11.59%	17.03%	17.2%	19.8%	18.72%	18.72%
$\kappa(3, 5)$	71.23%	0.31%	0.40%	0.66%	0.03%	0.03%
Blue Chip	40.25%	8.05%	6.59%	6.97%	4.14%	4.14%
VAR	32.21%	33.48%	28.36%	21.26%	15.88%	15.88%
CP	0.62%	34.84%	18.60%	8.07%	3.58%	3.58%

Figure 1: Household Decision Rule and the ARM Share.

The solid line corresponds to the ARM share in the US, and its values are depicted on the left axis. The dashed line displays the household decision rule $\kappa_t(3, 5)$. It is computed as the difference between the 5-year Treasury yield and the 3-year moving average of the 1-year Treasury yield. The time series is monthly from 1989.12 to 2006.6.

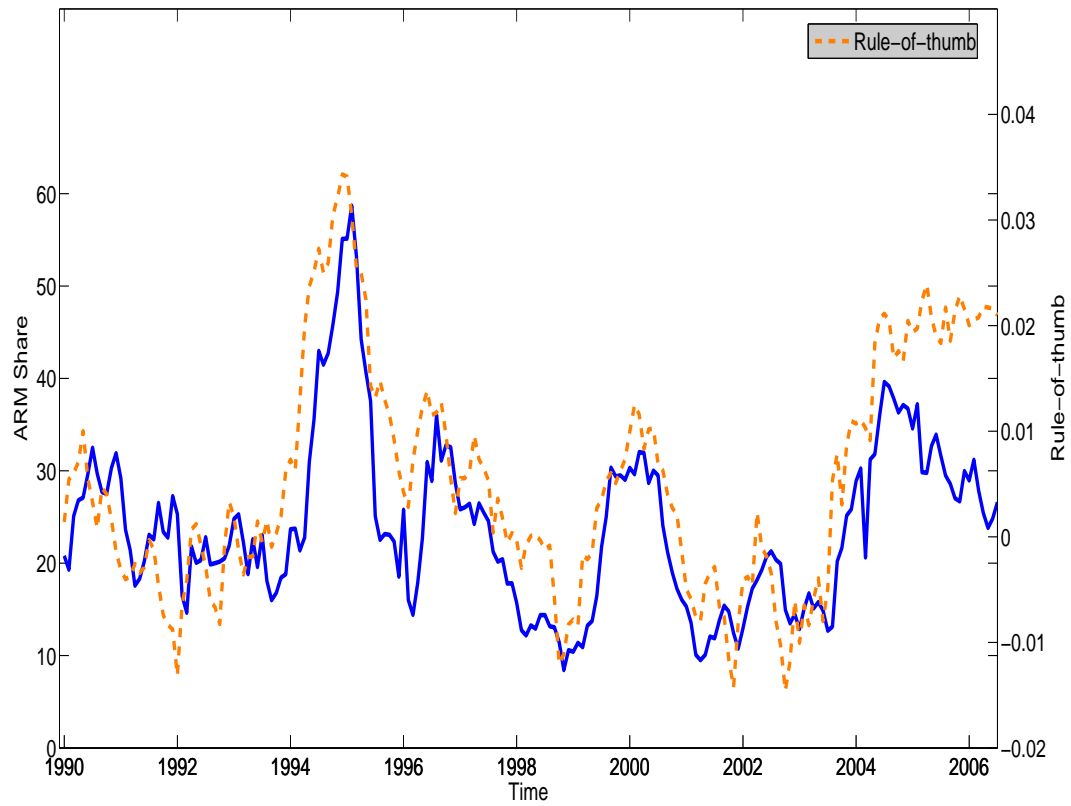


Figure 2: Correlation of the Household Decision Rule and the ARM Share for Different Look-Back Horizons ρ .

The figure plots the correlation of the household decision rule $\kappa_t(\rho; \tau)$ with the ARM share. The blue bars correspond to $\rho = 1, 2, 3, 4,$ and 5 years. The red line corresponds to the correlation between the 5-1 year yield spread (i.e., $\tau = 5$ and $\rho = 1$) and the ARM share. The red dashed line depicts the correlation between the 5-year yield and the ARM share (i.e., $\tau = 5$ and $\rho = \infty$). The left panel uses Treasury yields as yield variable ($\tau = 5$), while the right panel uses the effective 1-year ARM and effective 30-year FRM rates ($\tau = FRM$). The results are shown for the period 1989.12-2006.6, the longest sample for which all measures are available.

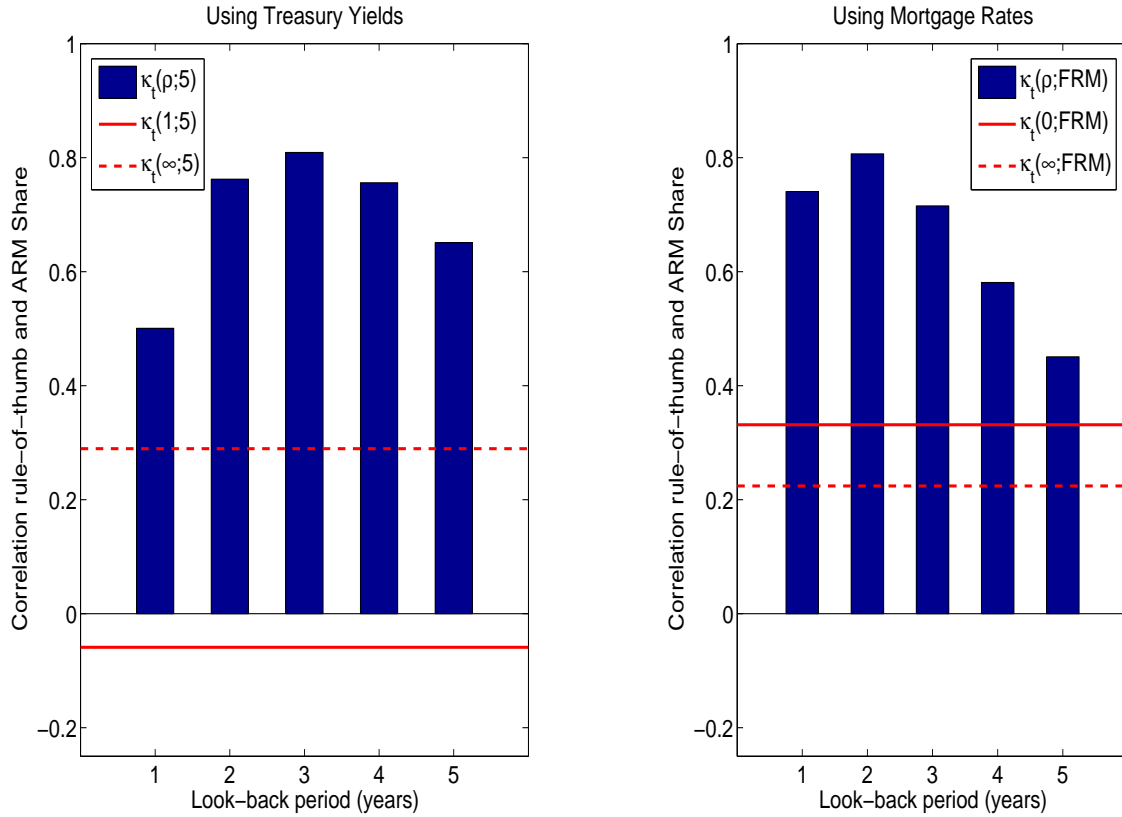
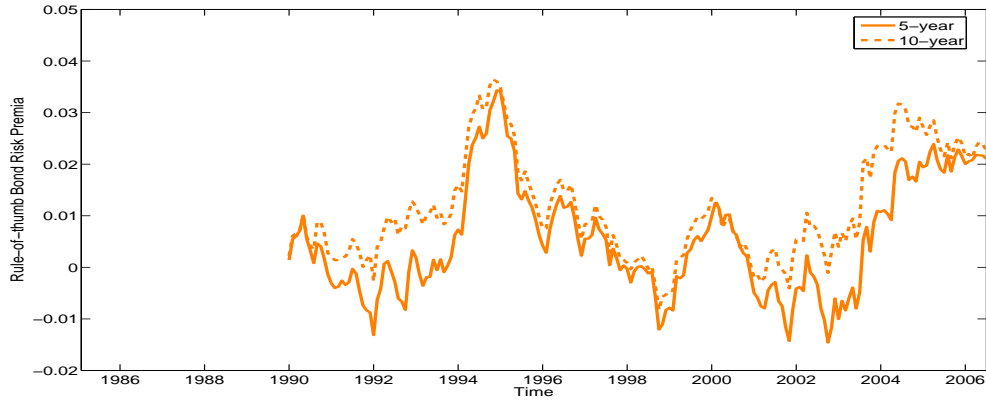


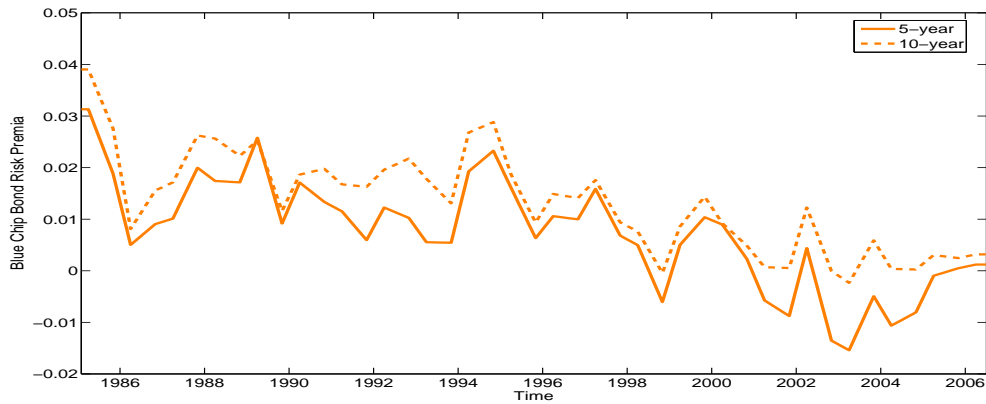
Figure 3: Three Measures of the Nominal Bond Risk Premium

Each panel plots the 5-year and the 10-year nominal bond risk premium. The average expected future nominal short rates that go into this calculation differ in each panel. In the top panel we use adaptive expectations with a three-year look-back period. In the middle panel we use Blue Chip forecasters data. In the bottom panel we use forecasts formed from a VAR model.

Panel A: Household Decision Rule



Panel B: Forward-Looking: Blue Chip Data



Panel C: Forward-Looking: VAR Model

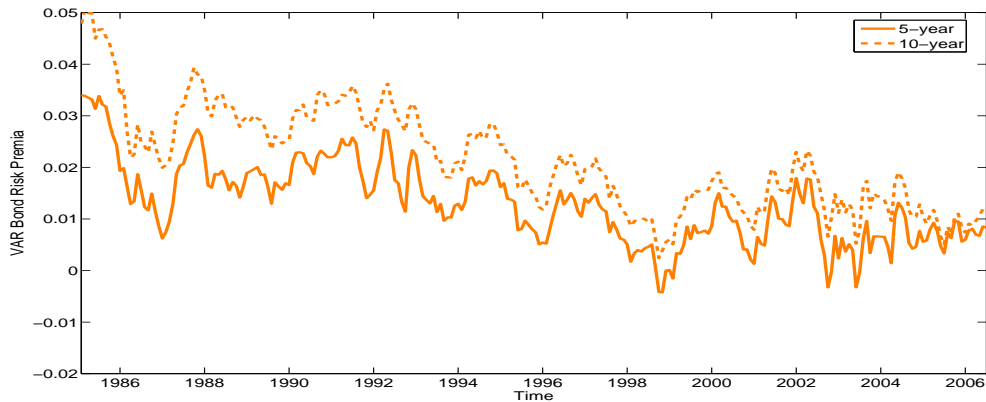


Figure 4: Rolling Window Correlations

The figure plots 10-year rolling window correlations of each of the three bond risk premium measures with the ARM share. The top line is for the household decision rule (dotted), the middle line is for the measure based on Blue Chip forecasters data (solid), and the bottom line is based on the VAR (dashed). The first window is based on the 1985-1995 data sample.

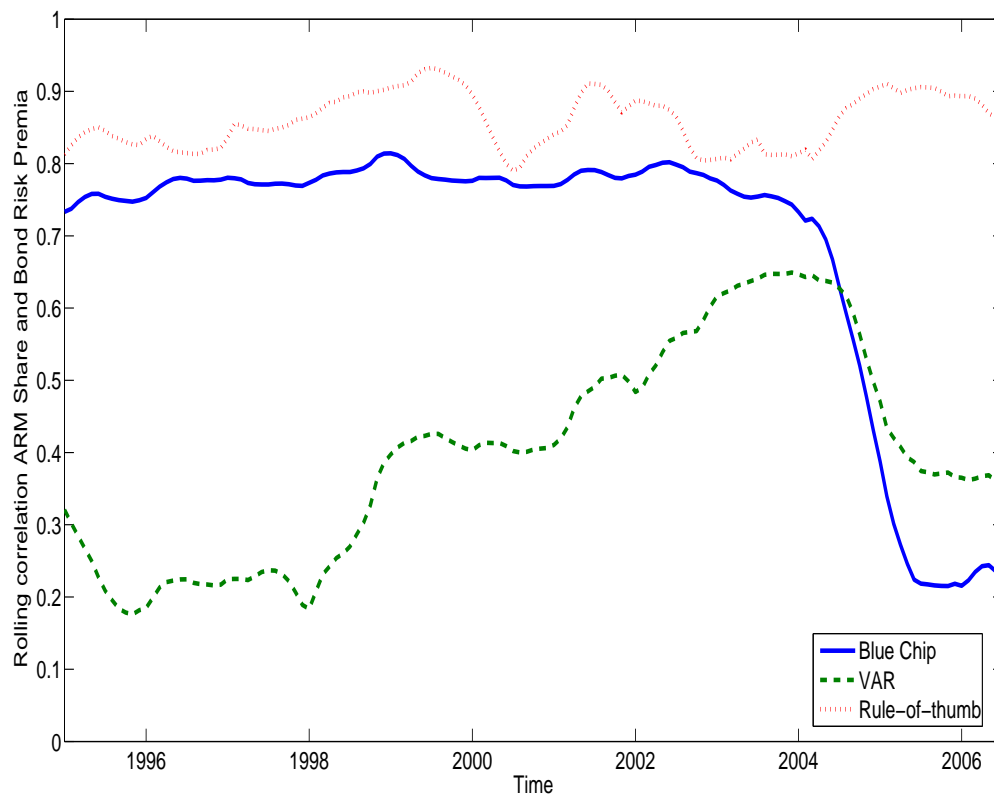


Figure 5: Product Innovation in the Mortgage Market

The solid line plots our benchmark ARM share, which includes all hybrid mortgage contracts, between 1992.1 and 2006.6. The dashed line excludes all hybrids with an initial fixed-rate period of more than three years. The data are from the Monthly Interest Rate Survey compiled by the Federal Housing Financing Board.

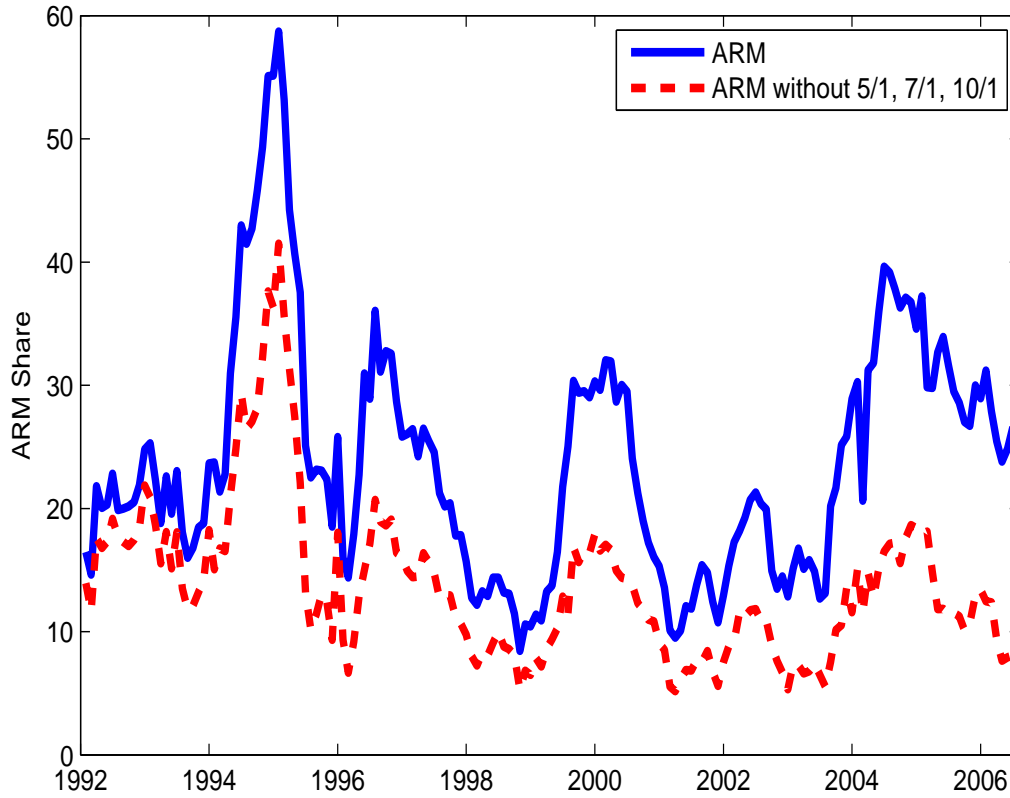


Figure 6: Errors in Predicting Future Real Rates

The figure plots forecast errors in expected future real short rates. The forecast error is computed using Blue Chip forecast data. The average expected future real short rate is calculated as the difference between the Blue Chip consensus average expected future nominal short rate and the Blue Chip consensus average expected future inflation rate. The realized real rate is computed as the difference between the realized nominal rate and the realized expected inflation, which are measured as the one-quarter ahead inflation forecast. The realized average future real short rates are calculated from the realized real rates. The forecast errors are scaled by the nominal short rate to obtain relative forecasting errors. The forecast errors are based on two-year ahead forecasts.

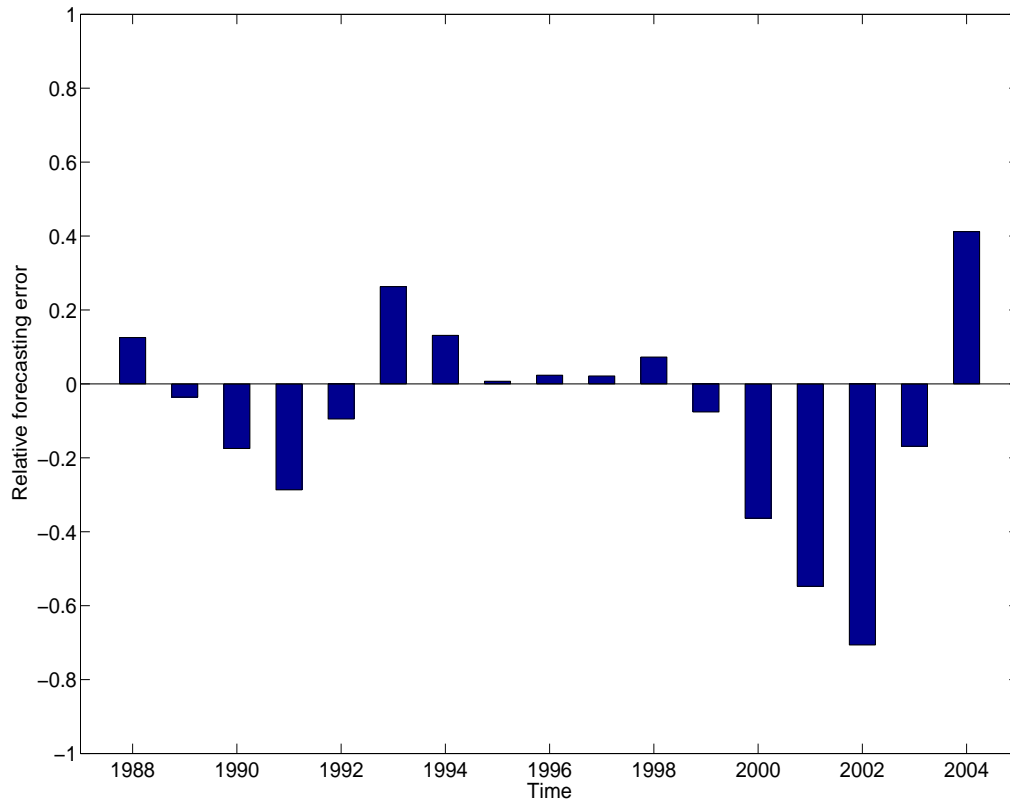


Figure 7: The Inflation Risk Premium and the ARM Share.

The figure plots the fraction of all mortgages that are of the adjustable-rate type against the left axis (solid line), and the inflation risk premium (dashed line) against the right axis. The inflation risk premium is computed as the difference between the 5-year nominal bond yield, the 5-year real bond yield and the expected inflation. The real 5-year bond yield data are from McCulloch and start in January 1997. The inflation expectation is the Blue Chip consensus average future inflation rate over the next 5 years.

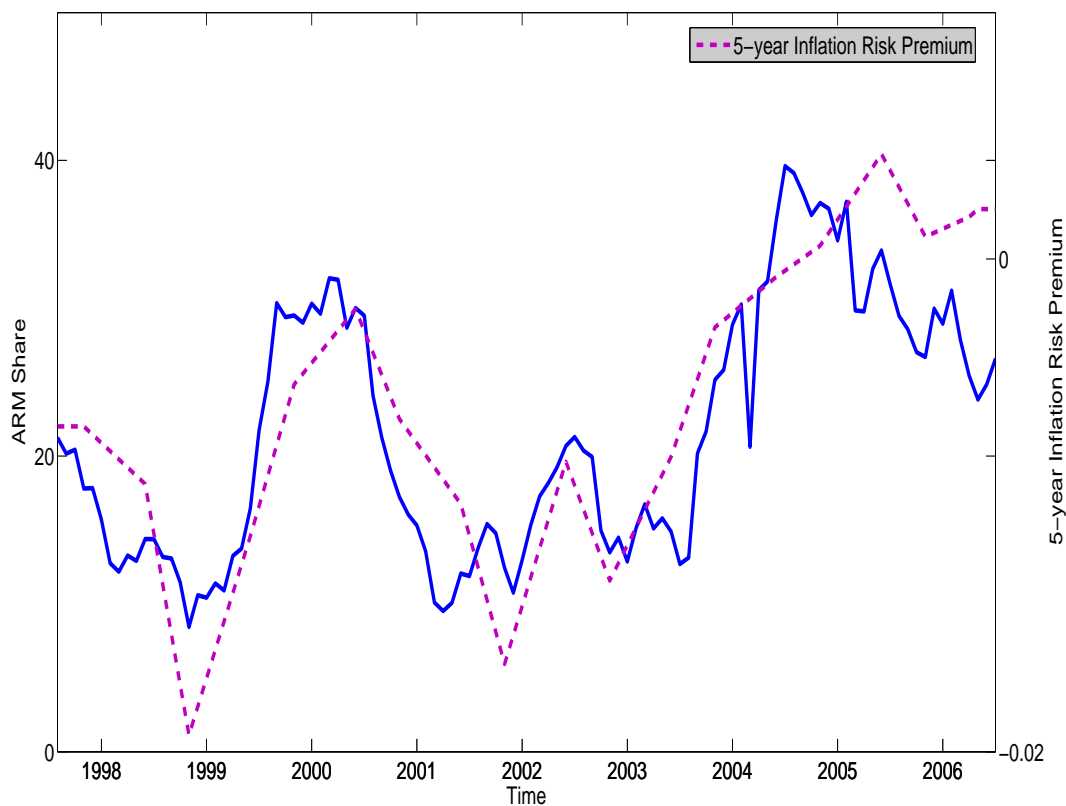


Figure 8: Price Sensitivity to Changes in the Nominal Interest Rate.

The figure plots the price sensitivities of the FRM contract with and without prepayment to the nominal interest rate, $y_0^{\$(1)}$. The mortgage values are determined within the model of Section 5.1. The analogous fixed-income securities are a regular bond (FRM without prepayment) and a callable bond (FRM with prepayment).

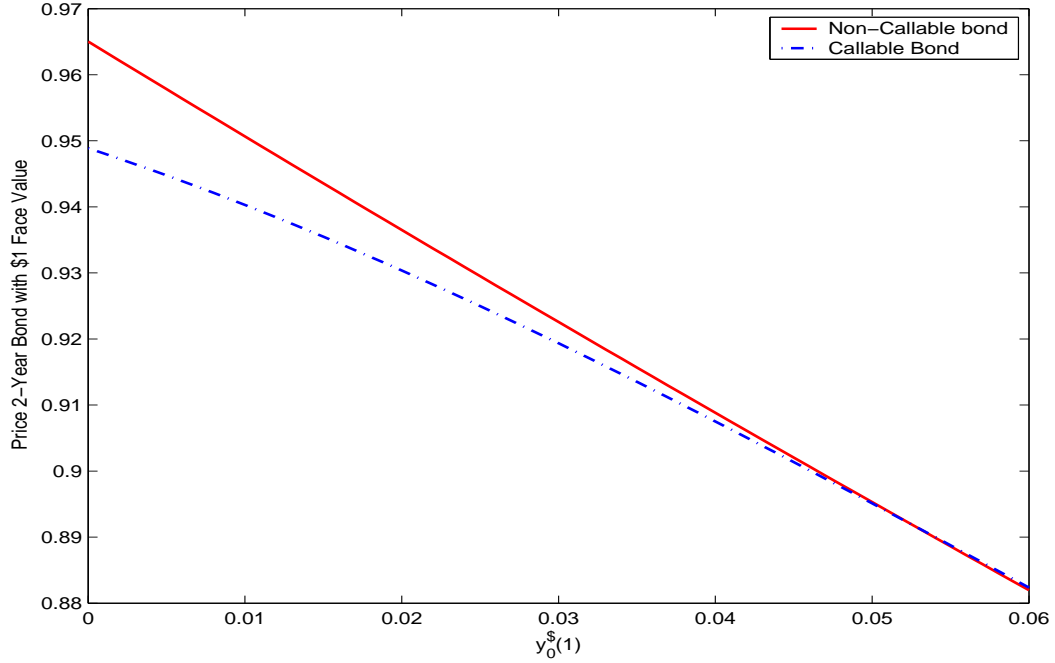


Figure 9: Utility Difference Between ARM and FRM - Prepayment

The figure plots the utility difference between an ARM contract and an FRM contract without prepayment as well as the utility difference between an ARM contract and an FRM contract with prepayment.

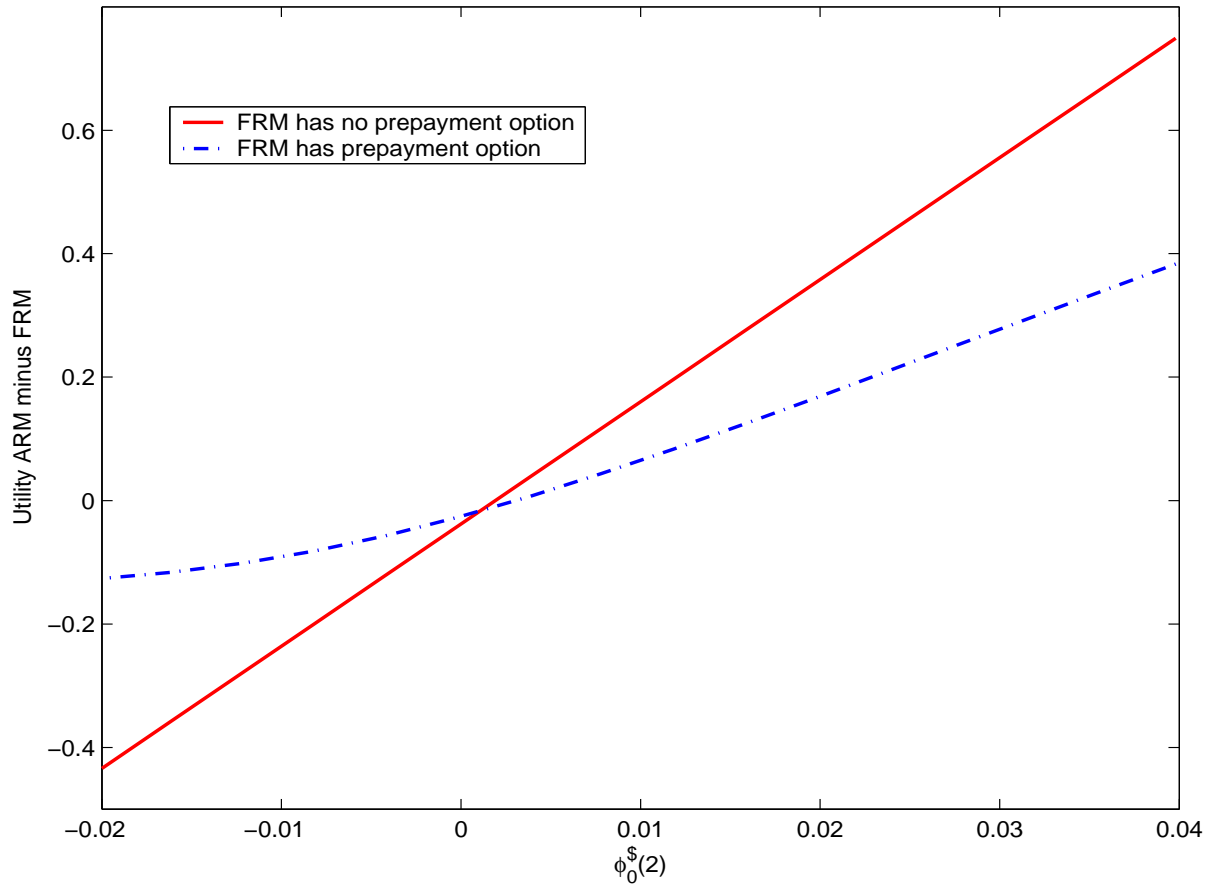


Figure 10: Composition Effect: States - MIRS data

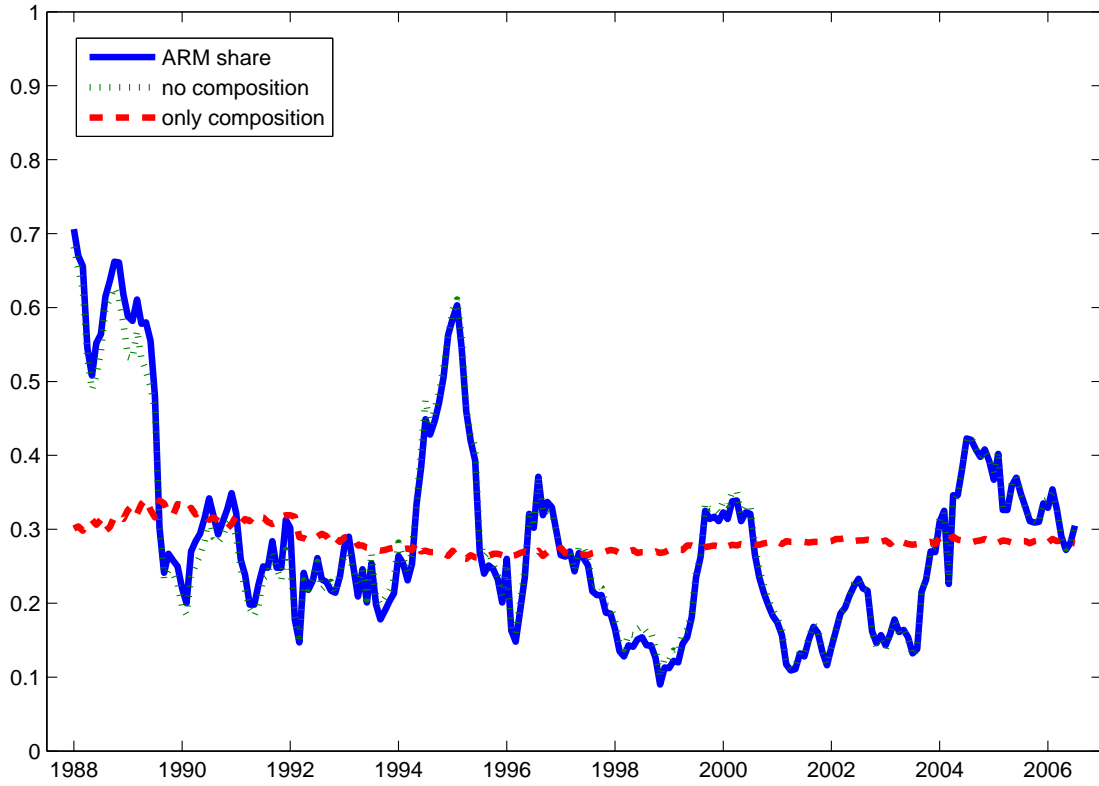


Figure 11: Mortgage Originations in the US.

The figure plots the volume of conventional ARM and FRM mortgage originations in the US between 1990 and 2005, scaled by the overall size of the mortgage market. Data are from the Office of Federal Housing Finance Enterprise Oversight (OFHEO).

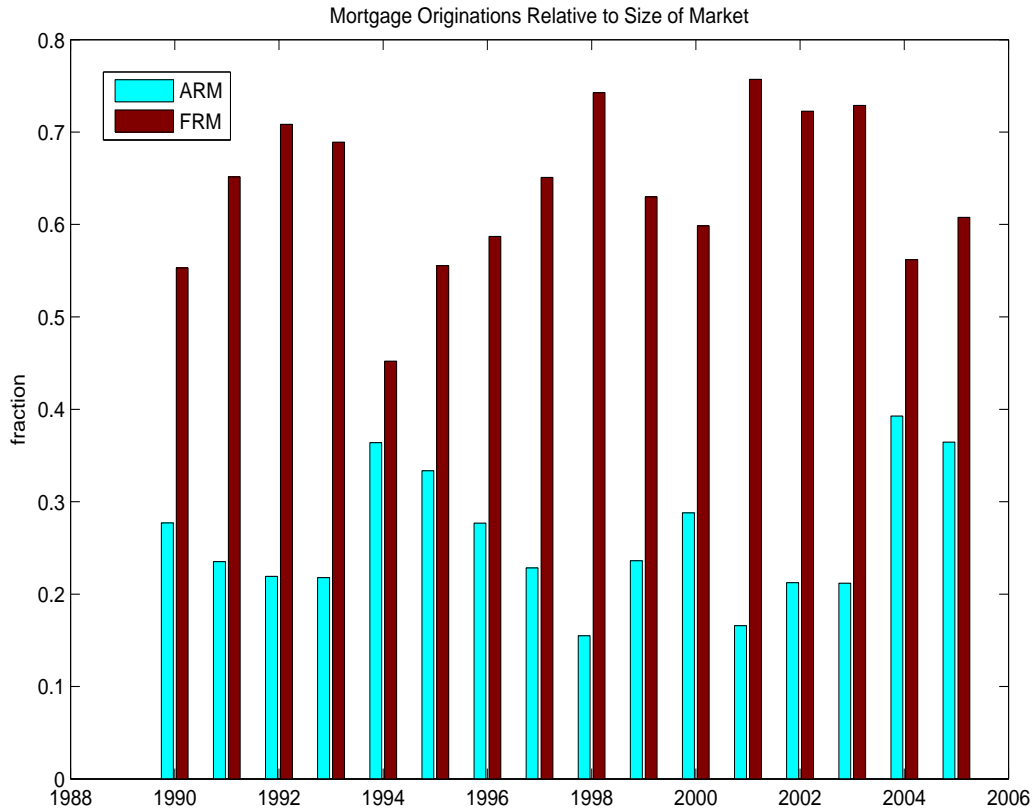


Figure 12: Effect of the Rate of Time Preference

Each point in the figure corresponds to the R^2 of a regression of the certainty equivalent consumption difference between an ARM contract and an FRM contract on either the bond risk premium (solid line) or one the yield spread (dashed line). The annualized subjective discount factor β^{12} , on the horizontal axis, is varied between 0.5 and 1. The time series are generated from a model, which is a multi-period extension of the model in Section 2. The coefficient of relative risk aversion is $\gamma = 5$. The exogenous moving probability is held constant at 1% per month.

