

REVISED

USC LUSK/FBE REAL ESTATE SEMINAR
presented by Otto Van Hemert
FRIDAY, February 23, 2007
10:30 am – 12:00 pm, **Room: RGL-215**

Mortgage Timing*

Ralph S.J. Koijen
Tilburg University

Otto Van Hemert
NYU Stern

Stijn Van Nieuwerburgh
NYU Stern and NBER

February 19, 2007

Abstract

Mortgages can be broadly classified into adjustable-rate mortgages (ARMs) and fixed-rate mortgages (FRMs). We document a surprising amount of time variation in the fraction of newly-originated mortgages that are of either type, both in the US and in the UK. A simple utility framework of mortgage choice points to the importance of term structure variables, and in particular bond risk premia. We extract these term structure variables from a flexible VAR-model and relate them to the fraction of newly-originated mortgages that are of the ARM type. In the US, the inflation risk premium accounts for the bulk of the variation in the ARM share. Other term structure variables, such as the yield spread and the long-term interest rate, have a weaker relation to the ARM share than previously thought. The presence of the prepayment option in US FRM contracts does not materially affect these results. Taken together, this suggests that households choose the right mortgage at the right time. We show that a simple rule-of-thumb approximates the bond risk premium well, making mortgage choice a much less daunting task. Finally, we uncover interesting differences between the US and the UK. The real rate risk premium rather than the inflation risk premium is the most strongly related to mortgage choice in the UK.

JEL classification: D14, E43, G11, G12, G21

Keywords: mortgage choice, household finance, bond risk premia

*First draft: November 15, 2006. Koijen: Department of Finance, CentER, Tilburg University, Tilburg, the Netherlands, 5000 LE; r.s.j.koijen@tilburguniversity.nl; Tel: 31-13-4663238; http://center.uvt.nl/phd_stud/koijen/. Van Hemert: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; ovanheme@stern.nyu.edu; Tel: (212) 998-0353; <http://www.stern.nyu.edu/~ovanheme>. Van Nieuwerburgh: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; svnieuwe@stern.nyu.edu; Tel: (212) 998-0673; <http://www.stern.nyu.edu/~svnieuwe>. The authors would like to thank Yakov Amihud, Andrew Ang, Jules van Binsbergen, John Campbell, João Cocco, John Cochrane, Joost Driessen, Ludovic Phalippou, James Vickery, Stan Zin, and seminar participants at Carnegie Mellon University and the University of Amsterdam for comments.

1 Introduction

One of the most important decisions any household has to make during its lifetime is whether to own a house and, if so, how to finance it. The home ownership rate in the US stands at 68% and US residential mortgage debt exceeds \$9 trillion. There are two broad categories of housing finance: adjustable-rate mortgages (ARMs) and fixed-rate mortgages (FRMs). There is a surprisingly large variation in the composition of newly-originated mortgages. Figure 1 plots the share of newly-originated mortgages that is of the ARM-type in the US economy between January 1985 and June 2006. This ARM share varies between 10% and 70%. In this paper we seek to explain this variation.

[Figure 1 about here.]

We claim that a large fraction of the variation in the ARM share can be attributed to time variation in bond risk premia. Consider a simple homoscedastic economy without inflation in which households have mean-variance preferences over consumption, and consume what is left from income after mortgage payments are made. In such an economy, the choice between an ARM and FRM boils down to comparing expected mortgage payments and their (constant) variability. Ignoring the prepayment option, FRMs are long-term loans whose payments are tied to the long-term nominal interest rate. ARM payments are tied to the short-term nominal interest rate instead. The difference in expected payments on the FRM and ARM equals the (long-term) nominal bond risk premium. The payments on the FRM are known at origination, while the ARM payments depend on future short rates. The mortgage choice then reduces to a trade-off between bond risk premia and short rate volatility. For the more realistic economy in which inflation erodes nominal mortgage payments, we decompose the nominal bond risk premium into the real rate premium and expected inflation premium. The FRM-ARM payment differential (approximately) equals the sum of the real premium and the expected inflation premium. An increase in either component of the bond risk premium makes the FRM less desirable, and is predicted to increase the share of ARM originations. In sum, time variation in bond risk premia leads to time variation in the preferred mortgage type.

Figure 2 plots the ARM share (solid line, measured against the left axis) alongside the five-year expected inflation risk premium (dashed line, measured against the right axis). We obtain the inflation risk premium as the difference between the five-year nominal bond yield and the sum of the five-year real bond yield and the five-year expected inflation. The nominal yield data are from the Federal Reserve Bank of New York and real bond yield data from McCulloch. Real data are available as of January 1997 when the US Treasury introduced treasury inflation-protected securities (TIPS). We use the median long-term inflation forecast of the survey of professional forecasters (SPF) to measure expected inflation. Ang, Bekaert, and Wei (2006) argue that such survey data

provides the best inflation forecasts among a wide array of methods. The contemporaneous correlation between the two series is 80%. This suggests that a large fraction of variation in the ARM share can be understood by time variation in the inflation risk premium. To illustrate, in each of the 1998.10-2000.4 and 2003.5-2005.3 periods, the inflation risk premium increased by more than 150 basis points. This made fixed-rate mortgages relatively more expensive, and US households shifted into ARMs. In both episodes, the ARM share tripled.

[Figure 2 about here.]

In Section 2, we formalize the utility-based mortgage choice argument. The analysis points to four yield curve determinants of mortgage choice: the expected inflation risk premium, the real rate risk premium, the variability of expected inflation, and the variability of the real rate. Because the inflation risk premium and the real rate risk premium for a fixed horizon describe nominal risk premia across different horizons well, these two risk premia capture mortgage choice across a variety of households that differ in their mortgage horizon. They not only describe individual mortgage choice, but also aggregate mortgage choice. We develop a vector auto-regression (VAR) model in Section 3 in order to estimate these four components on US data. The VAR structure readily provides a way to measure expected inflation and expected real rates and is an alternative to the professional forecasters data. Based on these measures of expected inflation and real rates, we construct the inflation risk premium and the real rate risk premium, and document its time variation.¹ The regression analysis of Section 4 uncovers that the four term-structure determinants typically enter with the right sign. The inflation risk premium emerges as the dominant explanatory variable for mortgage choice in the US. It alone explains about 60% of the variation in the ARM share. Adding the other term structure variables does not affect this conclusion. Neither does accounting explicitly for the prepayment option that is embedded in US FRM contracts. To analyze the impact of the prepayment option on the preference for mortgage types, we show how to value this option in a model that features time-varying risk premia.² We show that the prepayment option reduces the exposures to the underlying risk factors. However, it continues to hold that higher bond risk premia favor ARMs.

We compare these results with predictors of the ARM share proposed in the literature. Campbell and Cocco (2003) advocate the spread between the yields on a nominal long-term and short-term bond, and Campbell (2006) and Vickery (2006) use the spread between a FRM rate and

¹Fama and French (1989), Campbell and Shiller (1991), Dai and Singleton (2002), Buraschi and Jiltsov (2005), and Cochrane and Piazzesi (2005) were among the first to document and study time variation in bond risk premia.

²We contribute to the large literature on prepayment models which either assumes optimal prepayment (e.g., Dunn and McConnell (1981) and Pliska (2006)) or empirical prepayment behavior (e.g., Schwartz and Torous (1989) and Boudoukh, Whitelaw, Richardson, and Stanton (1997)). We consider a rational prepayment model and abstract from refinancing costs. Longstaff (2005) and Stanton (1995) model refinancing costs explicitly.

an ARM rate, as a determinant of the ARM share.³ We find low explanatory power for these variables over the common sample. Our model suggests why. The yield spread not only measures the nominal bond risk premium but also deviations of expected future nominal short rates from the current nominal short rate. Our VAR model shows that these two components are *negatively correlated*. For example, when expected inflation is high, the inflation risk premium is high as well, but expected future short rates are below the current one because inflation is expected to revert back to its long-term mean. As a result, the yield spread is a poor proxy for bond risk premia, and for mortgage choice. We show that bond risk premia are the relevant theoretical explanatory variables for the ARM share and we confirm their importance in our empirical analysis.

Our results suggest that households may have an ability to optimally time their mortgage choice. This finding contributes to the broader debate in household finance on the degree of financial sophistication of households (Campbell (2006)). At first glance, choosing the right mortgage at the right time is no easy task. Our analysis shows that it requires the ability to calculate inflation and real risk premia. However, we show in Section 5 that a simple rule-of-thumb describes mortgage choice extremely well. This rule-of-thumb approximates bond risk premia as the difference between the current long-term nominal interest rate and a backward-looking average of short-term nominal interest rates. This proxy for bond risk premia is much easier to compute; it only requires calculation of an average short rate over the recent past (2-4 years), and it does not require any real interest rate data. Yet, it closely captures the dynamics of the bond risk premia that we extract from the VAR model. We conclude that optimal mortgage choice maybe easier than previously thought.

In Section 6, we verify the robustness of our results to (i) alternative definitions of the ARM share, (ii) using a different VAR model to construct long-term expectations and risk premia, (iii) real interest rate data generated by the term structure model of Ang, Bekaert, and Wei (2007) rather than using TIPS data, (iv) the persistence in the variables included in the regressions. The analysis leads us to conclude that bond risk premia are a robust determinant of aggregate mortgage choice.

We also study mortgage choice in the United Kingdom. If bond risk premia are an important determinant of aggregate mortgage choice, our results should carry over to another country with another interest rate environment. There are some important differences with the US. FRM contracts in the UK have much shorter maturities than in the US. This implies that inflation risk, which manifests itself predominantly at long horizons, may be less important for choosing between

³Campbell and Cocco (2003) have a rich model of portfolio and mortgage choice for a household that faces persistent labor income shocks, stochastic equity returns and house prices. Risk premia are constant. Our model purposely simplifies the model along several dimensions in order to focus on the role of time-varying risk premia in more detail. Vickery (2006) also finds that household-specific characteristics have little explanatory power for mortgage choice. This is an important finding because it suggests that market-wide variables are the relevant variables to study.

ARMs and FRMs. Also, FRMs do not have a prepayment option. Finally, in the UK we have the benefit of a longer time series of real interest rate data. Our analysis shows that the real rate and expected inflation premium positively predict the ARM share in the UK, just as they did in the US. However, in sharp contrast to the US, we find that it is the real rate premium instead of the inflation risk premium that is the dominant predictor of mortgage choice in the UK. The variation in the ARM share explained by these bond risk premia equals 72% for the 2002-2006 sample for which we have monthly ARM share data available, and 23% for the 1993-2006 sample for which we interpolated quarterly ARM share data. Interestingly, the relative importance of the real rate premium in the UK and the inflation risk premium in the US seems to be captured by the rule-of-thumb proxy for bond risk premia. That proxy is much more strongly correlated with the real rate premium in the UK and with the inflation risk premium in the US.

Our findings resonate with the portfolio literature. Brandt and Santa-Clara (2006), Campbell, Chan, and Viceira (2003), Sangvinatsos and Wachter (2005), and Koijen, Nijman, and Werker (2007) have emphasized that forming portfolios that take into account time-varying risk premia can substantially improve performance for long-term investors.⁴ 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. The data suggest that time variation in these risk premia is also important from a positive point of view, to explain observed variation in mortgage choice.

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.⁵ Our findings suggest that households also have the ability to incorporate information on bond risk premia in their long-term financing decision.

2 Determinants of Mortgage Choice

This section explores the choice between a fixed-rate (FRM) and an adjustable-rate mortgage (ARM). The model in Section 2.2 is kept deliberately simple and serves to motivate the use of term structure variables as determinants of mortgage choice in Section 4. Section 2.3 discusses the use of the yield spread as a determinant of mortgage choice. But first, we introduce some bond pricing notation.

⁴Campbell and Viceira (2001), Brennan and Xia (2002), and van Hemert (2006) derive the optimal portfolio strategy for long-term investors in the presence of stochastic real interest rates and inflation, but these papers assume risk premia to be constant.

⁵See also Butler, Grullon, and Weston (2006) and Baker, Taliaferro, and Wurgler (2006) for a recent discussion of this result.

2.1 Bond Pricing Preliminaries

We denote the nominal price at time t of a nominal τ -month zero-coupon bond by $P_t(\tau)$. Time (t) is expressed in months. The yield $y_t^{\$}(\tau)$, and the 1-year forward rate $f_t^{\$}(\tau)$ are given by⁶

$$y_t^{\$}(\tau) \equiv -\frac{1}{\tau} \log(P_t(\tau)) \quad \text{and} \quad f_t^{\$}(\tau) \equiv -\log\left(\frac{P_t(\tau+12)}{P_t(\tau)}\right). \quad (1)$$

We do not impose the Expectations Hypothesis: $f_t^{\$}(\tau) \neq \mathbb{E}_t[y_{t+\tau}^{\$}(12)]$. Equation (2) defines the nominal risk premium on a $(\tau/12)$ - year bond:

$$\phi_0^{\$}(\tau) \equiv y_0^{\$}(\tau) - \frac{1}{\tau/12} \sum_{t=1}^{\tau/12} \mathbb{E}_0 \left[y_{12 \times (t-1)}^{\$}(12) \right] = \frac{1}{\tau/12} \sum_{t=1}^{\tau/12} f_0^{\$}(12 \times (t-1)) - \frac{1}{\tau/12} \sum_{t=1}^{\tau/12} \mathbb{E}_0 \left[y_{12 \times (t-1)}^{\$}(12) \right] \quad (2)$$

where the second equality uses the fact that the yield on a τ -month zero coupon bond equals the average forward rate. For future use, we rewrite the nominal bond risk premium as the sum of the inflation risk premium and the real rate risk premium

$$\phi_0^{\$}(\tau) = \phi_0^x(\tau) + \phi_0^y(\tau). \quad (3)$$

Analogous to the nominal risk premium $\phi_0^{\$}$ in Equation (2), we define the real rate risk premium at time 0, ϕ_0^y , as the difference between the observed long-term real rate and the expected long-term real rate. The latter is the average of the expected future short real rates

$$\phi_0^y(\tau) \equiv y_0(\tau) - \frac{1}{\tau/12} \sum_{t=1}^{\tau/12} \mathbb{E}_0 \left[y_{12 \times (t-1)}(12) \right], \quad (4)$$

where $y_t(\tau)$ is the real yield of a τ -month real bond at time t . We impose that the yield at time t of an 1-year real bond, $y_t(12)$, is the difference between the 1-year nominal yield, $y_t^{\$}(12)$, and 1-year expected inflation, $x_t = x_t(12)$

$$y_t(12) = y_t^{\$}(12) - x_t(12). \quad (5)$$

Following Ang, Bekaert, and Wei (2007), we define the expected inflation premium at time 0, ϕ_0^x , as the difference between long-term nominal yields, long-term real yields, and long-term expected inflation

$$\phi_0^x(\tau) \equiv y_0^{\$}(\tau) - y_0(\tau) - x_0(\tau). \quad (6)$$

⁶We generally consider τ to be a multiple of 12, which implies that $\tau/12$ is integer-valued.

This uses the decomposition of realized inflation at time t into expected inflation conditional on the time $t - 12$ information, x_{t-12} , and unexpected inflation, ϵ_t

$$\pi_t = x_{t-12} + \epsilon_t, \quad (7)$$

and uses the definition of the long-term expected inflation

$$x_t(\tau) = \frac{1}{\tau/12} \mathbb{E}_t [\log \Pi_{t+\tau} - \log \Pi_t],$$

with Π_t the price index at time t , and $\pi_t = \log \Pi_t - \log \Pi_{t-12}$.

2.2 Optimal Mortgage Choice

We consider a discrete-time setting for an investor with constant relative risk aversion preferences over a real consumption stream $\{C_t\}$. The preference parameter γ summarizes the investor's risk preferences. The subjective time discount factor is 1. The investor receives a stochastic real income stream $\{L_t\}$, which is idiosyncratic (uncorrelated with aggregate variables) and has an unconditional mean ℓ and variance σ_ℓ^2 .

At time 0, the investor buys a house whose real value is normalized to \$1. We assume that the house price has a constant real value. To finance the house, the investor chooses a mortgage of the ARM or FRM type. The face value of the mortgage equals \$1 as well; we assume a 100% loan-to-value ratio. The investment horizon and the maturity of the mortgage contract equal T years. At times 12 through $12 \times T$ the investor pays interest on the mortgage, but no payments on the principal are due.

We think of the nominal interest rate on an FRM contract as the time-zero forward rate in each period on forward contracts with annual delivery dates $t = 12, 24, \dots, 12 \times T$. This assumption captures the essence of a nominal FRM: future mortgage payments are fixed in nominal terms at the origination time 0.⁷ The nominal interest rate on an ARM contract is the short rate in each period. 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.⁸ Denote the stream of real mortgage payments by $\{q_t\}$:

$$q_t^{FRM} = \frac{f_0^s(t-12)}{\Pi_t}, \quad q_t^{ARM} = \frac{y_{t-12}^s(12)}{\Pi_t}. \quad (8)$$

⁷For ease of exposition we do not impose that the FRM interest payments are equal over time, only that they are known at time 0. Constant mortgage payments would be the harmonic mean of all forward rates of maturities $12, \dots, 12 \times T$.

⁸In the background, a competitive fringe of mortgage lenders prices mortgages to maximize profit taking as given the term structure of treasury interest rates.

To keep the problem as simple as possible, we make three further assumptions. First, we postulate that the investor is liquidity constrained: In each period, she consumes what is left over from income after making the mortgage payment.⁹ The mortgage choice problem at time 0 is

$$\max_{h \in \{ARM, FRM\}} \mathbb{E}_0 \left[\sum_{t=1}^T \frac{(C_{12 \times t}^h)^{1-\gamma}}{1-\gamma} \right], \quad (9)$$

$$\text{s.t. } C_{12 \times t}^h = L_{12 \times t} - q_{12 \times t}^h, \quad t = 1, \dots, T-1, \quad (10)$$

$$\text{and } C_{12 \times T}^h = L_{12 \times T} - q_{12 \times T}^h + 1 - 1/\Pi_{12 \times T}. \quad (11)$$

Terminal consumption equals income after the mortgage payment plus the difference between the real value of the house, which is 1, and the real mortgage balance, which is $1/\Pi_{12 \times T}$.

Second, we focus on a second-order Taylor expansion of the CRRA preferences. Third, we approximate around zero inflation. These last two assumptions are for expositional reasons only. Appendix A contains the derivation and shows the numerical solution in a model that does not make these assumptions. Taken together, an investor prefers the T -year ARM contract over the T -year FRM contract at time zero if and only if

$$\sum_{t=1}^T \mathbb{E}_0 [q_{12 \times t}^{FRM}] + \frac{\gamma}{2\ell} \mathbb{E}_0 [(q_{12 \times t}^{FRM})^2] > \sum_{t=1}^T \mathbb{E}_0 [q_{12 \times t}^{ARM}] + \frac{\gamma}{2\ell} \mathbb{E}_0 [(q_{12 \times t}^{ARM})^2], \quad (12)$$

and recall that ℓ is the unconditional average labor income.

The difference between the expected mortgage payments for the FRM and ARM investors equals the bond risk premium

$$\mathbb{E}_0 \left[\sum_{t=1}^T q_{12 \times t}^{FRM} \right] - \mathbb{E}_0 \left[\sum_{t=1}^T q_{12 \times t}^{ARM} \right] = \sum_{t=1}^T f_0^{\$}(12 \times (t-1)) - \sum_{t=1}^T \mathbb{E}_0 \left[y_{12 \times (t-1)}^{\$}(12) \right] = T \phi_0^{\$}(12 \times T), \quad (13)$$

where the first equality uses the same approximations as described in Appendix A, and the second equality uses the definition of the risk premium on a T -year nominal bond in (2). The FRM investor faces no uncertainty over the nominal mortgage payments, whereas the ARM investor faces nominal interest rate risk. The variability of nominal ARM payments is $\sum_{t=1}^T \mathbb{E}_0 \left[(y_{12 \times (t-1)}^{\$}(12))^2 \right]$. Under the approximations made before, the same holds true for the real payment variability. Combining the difference in expected payments and the difference in the variability of the payments, we arrive at (14), which states that the investor prefers an ARM if the nominal bond risk premium exceeds

⁹This seems a plausible assumption because most households are young and not very wealthy at the time of mortgage origination. The implication is that the household does not invest savings in the bond market, and cannot undo the position taken in the mortgage market.

the variability of the nominal interest rate multiplied by the risk aversion coefficient

$$\phi_0^{\$}(12 \times T) > \frac{\gamma}{T} \sum_{t=1}^T \mathbb{E}_0 \left[(y_{12 \times (t-1)}^{\$}(12))^2 \right], \quad (14)$$

$$\phi_0^x(12 \times T) + \phi_0^y(12 \times T) > \frac{\gamma}{T} \sum_{t=1}^T \mathbb{E}_0 \left[(y_{12 \times (t-1)}(12) + x_{12 \times (t-1)}(12))^2 \right]. \quad (15)$$

If the protection that an FRM offers against nominal interest rate volatility to the nominal investor is too expensive, an ARM becomes more attractive. The second inequality exploits the definition of the nominal bond risk premium and that of the nominal short-term interest rate in (5). While the formulations in (14) and (15) are equivalent, Section 3.4 shows that there are several reasons to consider the two components of the nominal risk premium and the two components of the variability separately. Thus, equation (15) points to four term-structure determinants of mortgage choice: the real rate premium, the expected inflation premium, the real rate variance, and the expected inflation variance. In our main exercise, we will assume that the squared expected inflation and squared real rates are constant in expectation. I.e., the right hand side of equation (15) is constant. Then, the main prediction of the model is that an increase in either bond risk premium increases the expected payments on the FRM, makes the ARM more desirable, and should increase the ARM share.

Appendix A describes numerical results that link the utility over consumption streams resulting from FRM and ARM contracts to the real interest and expected inflation premium. In that section we take into account fully the effects of inflation. The results confirm the intuition of this section in the sense that the utility difference between the two contracts is largely explained by the two premia. In the robustness section 6 at the end of the paper, we study the case of time-varying second moments. We model and estimate the variability of expected inflation and the real rate and include them in the ARM share regressions.

2.3 The Yield Spread as a Predictor of the ARM Share

The above framework is well suited to understand the role of the yield spread as a predictor of mortgage choice. Campbell and Cocco (2003) and Campbell (2006) have argued that the slope of the yield curve is a key determinant of mortgage choice. They argue that when nominal long-term interest rates are high compared to nominal short-term rates, ARMs seem attractive relative to FRMs.

Condition (16) shows why the yield spread may be an imperfect measure of the relative attractiveness of both mortgage types. Consider the following decomposition of the nominal yield spread into the nominal bond risk premium and the deviations of average expected future short rates and

the current nominal short rate,

$$y_0^{\$}(\tau) - y_0^{\$(12)} = \phi_0^{\$(\tau)} + \left(\frac{1}{\tau/12} \sum_{t=1}^{\tau/12} \mathbb{E}_0 \left[y_{12 \times (t-1)}^{\$(12)} \right] - y_0^{\$(12)} \right). \quad (16)$$

In a homoscedastic world with zero risk premia ($\phi_0^{\$(\tau)} = 0$), 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 will face the same expected payment stream in this world. The yield spread is completely uninformative about mortgage choice. Likewise, 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. A second way of seeing what goes wrong is to think of the current FRM-ARM rate spread as the determinant of mortgage choice. This measure deducts from the current FRM rate (long-term bond rate) the current ARM rate (one-period interest rate). Equation (16) shows that the correct proxy for the bond risk premium, and hence for mortgage choice, subtracts from the FRM rate the *average* of the current and future ARM rates (expected future one-period interest rate). Indeed, the latter is the actual rate that the ARM investor will have to pay over the life of the mortgage. It is the rate on rolling over a short-term bond position.

In our model with time-varying risk premia, estimated below, it turns out that the two terms on the right-hand side of (16) are *negatively* correlated. This makes the yield spread a contaminated proxy for the nominal bond risk premium, and as we show empirically below, a weak determinant of mortgage choice.

3 VAR Model

We now set up a VAR model to construct long-term inflation and real interest rate expectations that are needed to estimate real interest rate and expected inflation risk premia. The VAR offers an alternative way to form inflation expectations to the professional analyst survey data, used in the introduction. In addition, it allows us to form real rate risk premia. In a first step, we work with a homoscedastic term structure model. The structure that the VAR imposes will turn out to be valuable to understand how exactly the two risk premia affect mortgage choice, analyzed later in Section 4. At the end of the paper, we study an extension with heteroscedastic innovations.

3.1 VAR Setup

Our state vector Y contains the one-year ($y_t^{\$}(12)$), the five-year ($y_t^{\$}(60)$), and the ten-year nominal yields ($y_t^{\$}(120)$), as well as realized, one-year log inflation ($\pi_t = \log \Pi_t - \log \Pi_{t-12}$). On the right-hand side of the VAR(1) is the 12-month lag of the state variables. Time (t) is expressed in months and we use overlapping monthly observations.¹⁰ The law of motion for the state is

$$Y_{t+12} = \mu + \Gamma Y_t + \eta_{t+12}, \quad \text{with } \eta_{t+12} \mid \mathcal{I}_t \sim D(0, \Sigma), \quad (17)$$

with \mathcal{I}_t representing the information at time t . For now, we assume that the innovation covariance matrix is constant. Section 6.1.1 specifies a VAR model with heteroscedastic innovations.

We start by constructing the 1-year expected inflation series as a function of the state vector

$$x_t(12) = \mathbb{E}_t[\pi_{t+12}] = e'_4 \mu + e'_4 \Gamma Y_t, \quad (18)$$

where e_4 is the fourth unit vector. We construct the 1-year real short rate by subtracting expected inflation from the 1-year nominal rate (see (5))

$$y_t(12) = y_t^{\$}(12) - x_t(12) = -e'_4 \mu + (e'_1 - e'_4 \Gamma) Y_t. \quad (19)$$

Next, we use the VAR structure to determine the n -year expectations of the average inflation and the average real rate in terms of the state variables. For expected average inflation this becomes

$$x_t(12 \times n) \equiv \frac{1}{n} \mathbb{E}_t \left[\sum_{i=1}^n e'_4 Y_{t+(12 \times i)} \right] = \left(\frac{1}{n} \right) e'_4 \left\{ \sum_{i=1}^n \left(\sum_{j=0}^{i-1} \Gamma^j \mu \right) + \sum_{i=1}^n \Gamma^i Y_t \right\}. \quad (20)$$

The long-run expected average real rate is also a function of the current state

$$\begin{aligned} y_t(12 \times n) &\equiv \frac{1}{n} \mathbb{E}_t \left[\sum_{i=0}^{n-1} y_{t+(12 \times i)}(12) \right] \\ &= \left(\frac{1}{n} \right) e'_1 \left\{ \sum_{i=1}^{n-1} \left(\sum_{j=0}^{i-1} \Gamma^j \mu \right) + \sum_{i=1}^{n-1} \Gamma^i Y_t \right\} + \frac{e'_1 Y_t}{n} - x_t(12 \times n). \end{aligned} \quad (21)$$

With the long-term expected real rate from (21) in hand, we can form the real risk premium by subtracting this expectation from the observed real rate (as in (4)). Similarly, with the long-term expected inflation from (20) in hand, we form the inflation risk premium as the difference between the observed nominal yield, the observed real yield, and expected inflation (as in (6)).

¹⁰We have also estimated the model on quarterly data and found very similar results.

3.2 VAR Estimation Results

We estimate a VAR-model with monthly observations for the period 1985.1-2006.6. Monthly nominal yield data are from the Federal Reserve Bank of New York.¹¹ The inflation rate is based on monthly CPI-U available from the Bureau of Labor Statistics.¹² We start the model in 1985, near the end of the Volcker deflation. 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.¹³

Figure 3 shows the estimation results. The top left panel shows the 1-year expected inflation x_t as well as the 1-year real rate y_t , computed from (18) and (19). The bottom two panels show the long-term expectations of the same variables at the five- and ten-year horizons, computed from (20) and (21) respectively. Expected inflation is relatively smooth at all horizons; its values are nearly identical at the five-year and ten-year horizons. It is 2.9% per year on average; higher at the beginning of the sample (3.48% in 1985.2) and lower near the end of the sample (2.46% in 2004.3). Interestingly, the survey data on long-term expected inflation, which we used in the introduction, show a similar pattern. They are also nearly constant, albeit at a slightly lower level of 2.5%. Real rate expectations display more variation over time. At the one-year horizon, real yields hover between -2% (2004) and 6% per year (1984). At the ten-year horizon, these expectations are smoother. They hover between 0.5% and 3.5%, but show the same pattern of fluctuations.

[Figure 3 about here.]

Combining data on nominal and real five-year and ten-year yields, we form the real rate and expected inflation risk premia. The real yield data are from McCulloch.¹⁴ The left panel of Figure 4 plots the risk premia at a five-year horizon, while the right panel plots the ten-year horizon premia. The figure starts in July of 1997, the first period for which five-year and ten-year real yield data are available in the US.¹⁵ Expected inflation risk premia in both panels are negative until 2004. This negative risk premium is not surprising given that the observed spread between

¹¹The nominal yield data are available at <http://www.federalreserve.gov/pubs/feds/2006>.

¹²The inflation data are available at <http://www.bls.gov>.

¹³As a robustness check, we also considered a VAR(2)-model. In section 4, we redo the ARM share regressions for the term structure variables arising from that model.

¹⁴The real yield data are available at <http://www.econ.ohio-state.edu/jhm/ts/ts.html>. At the end of the paper, we also perform our analysis with real yield data generated by the term structure model of Ang, Bekaert, and Wei (2007). We show that our main conclusions are unaffected.

¹⁵We do not use the first six months of 1997, in which only a five-year TIPS was available. The TIPS market may have suffered from liquidity problems in its first two of years of operation (see Shen and Corning (2001), Jarrow and Yildirim (2003), and Ang, Bekaert, and Wei (2007)). As a robustness check, we repeated all regressions of Section 4 using only TIPS data after 1999.1. We found very similar results.

nominal and real yields is often below 2% and inflation expectations are always above 2%. Most of the action in the nominal-real spread is inherited by the inflation risk premium because expected inflation varies much less. The ten-year risk premium varies between -1.65% in 1998.8 and +0.35% in 2004.4. The real rate premium on the other hand is estimated to be positive, and varies between 0.8% per year in 2005.5 and 2.9% in 2002.1 at the ten-year horizon.

[Figure 4 about here.]

The two risk premia have a negative correlation of -0.64 and -0.59 at the five- and ten-year horizons, respectively. Because of this negative correlation, their sum, the nominal risk premium, cancels out a lot of interesting variation. Unsurprisingly, this sum will turn out to be less informative for mortgage choice than its components.

3.3 Extending the Sample of Bond Risk Premia

The data on nominal yields and realized inflation, but also on the nominal bond risk premium (obtained from the VAR) go back to 1985. However, the unavailability of real yield data before 1997.7 prevents us from *decomposing* the nominal bond risk premium into its two components: the inflation risk premium and the real rate risk premium. In order to study mortgage choice in the US using the two separate risk premia, we develop a projection method that allows us to extend the sample back to 1985.

We construct a long time series for the real rate risk premium by first regressing the real rate risk premium on a set of state variables z_t that are observable over the complete sample period. Specifically, we estimate the regression

$$\phi_t^y = \alpha + \beta' z_t + \epsilon_t, \quad (22)$$

over the period 1997:7-2006:6, and construct the real rate risk premium for the full sample period using the estimated coefficients $\hat{\phi}_t^y = \hat{\alpha} + \hat{\beta}' z_t$. We back out the inflation rate risk premium as the difference between the nominal risk premium and the projected real rate risk premium. This method gives reliable answers as long as (i) the relationship between risk premia and the state variables z_t does not change dramatically after 1997:7 and (ii) the state variables capture most of the variation in the risk premia. With these considerations in mind, we select $z_t = (Y_t', Y_{t-12}')'$, where Y_t contains the VAR variables. A regression of the ten-year (five-year) real rate premium on z gives an in-sample R^2 of 90% (86%).

Figure 5 shows the observed nominal bond risk premium $\{\phi_t^{\$}\}$ (solid line) together with its projected components (lines with circles) at the ten-year horizon. It also overlays the risk premia shown in the left panel of Figure 4 for the 1997.7-2006.6 period. The projections fit these risk

premia closely beyond 1997.7. Interestingly, the projections indicate that inflation risk premia were higher (and often positive) before 1997. Real rate risk premia came down from 4% in 1985 to 2% in 1997.

[Figure 5 about here.]

3.4 Aggregation

Equation (14) in Section 2.2 shows that *individual* mortgage choice depends on the nominal bond risk premium, which is the sum of the expected inflation risk premium and the real rate risk premium. See equation (3) and also Campbell and Viceira (2001), Brennan and Xia (2002), and Ang, Bekaert, and Wei (2007). We are interested in explaining *aggregate* mortgage choice and argue that it depends on the two component risk premia, rather than on their sum alone. We envision households that are heterogeneous in the *effective* mortgage maturity. This heterogeneity arises either from different contract length or from heterogeneity in when exogenous moving shocks hit (and triggers prepayment of the loan). If τ_j denotes the effective contract length of household j , the nominal bond risk premium at that horizon determines its mortgage choice.

Figure 6 shows that the real interest rate premium and the inflation risk premium explain most of the variation in total bond risk premia across different maturities. It plots the R^2 of a regression of the nominal risk premium $\phi_t^{\$}(\tau)$, with $\tau = 24, \dots, 120$ on our two risk factors $\phi_t^y(120)$ and $\phi_t^x(120)$ at the ten-year horizon

$$\phi_t^{\$}(\tau) = \alpha_0^{\tau} + \alpha_1^{\tau} \phi_t^y(120) + \alpha_2^{\tau} \phi_t^x(120) + \epsilon_t^{\tau}. \quad (23)$$

The R^2 never goes below 70%. This implies that the inflation risk premium and the real rate risk premium together capture the entire term structure of risk premia.¹⁶ Therefore, they capture the relevant risk premia for a whole range of households who differ in their effective mortgage maturity.

[Figure 6 about here.]

Second, the figure also decomposes the R^2 into a piece that is due to the inflation risk premium, a piece that is due to the real rate risk premium, and a piece due to their covariance. What is important is that the real rate risk premium and the inflation risk premium are not even close to perfectly correlated, and that nominal bond risk premia at different maturities have different loadings on these two risk factors. As a result, aggregation forces us to use both of them separately.¹⁷

¹⁶Cochrane and Piazzesi (2005) argue that a single factor, which is a linear combination of forward rates, can capture most of the variation in single-period expected excess bond returns. Instead, we are interested in the expected excess return of holding the bond for multiple periods. Our nominal bond risk premium is the risk premium on a strategy that holds a τ -period bond until maturity and finances it by rolling over the 1-year bond. Cochrane and Piazzesi (2006) study different definitions of bond risk premia.

¹⁷At horizon 10, the loadings on the real rate and inflation risk premium have to sum to 1 by construction. Because

4 ARM Share Regressions

We are interested in explaining time variation in the fraction of all newly-originated mortgages that is of the adjustable-rate type. In this section, we regress the ARM share on the bond risk premia, motivated in Section 2 and computed from the VAR in Section 3. We lag the predictor variables for one period in order to study what changes in this month’s risk premia and volatilities 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.

4.1 Data on the ARM Share in the U.S.

Our baseline data series is from the Federal Housing Financing Board. It is based on the Monthly Interest Rate Survey, a survey sent out to mortgage lenders.¹⁸ These data include loan originations for both newly constructed homes and existing homes. The monthly data start in 1985.1 and run until 2006.6, and we label this series $\{ARM_t^1\}$. Our baseline measure of the ARM share includes all adjustable mortgages. In particular, it includes hybrid mortgages which have an initial fixed-interest rate payment period. Starting in 1992, we also know the decomposition of the ARM by initial fixed-rate period.¹⁹ This allows us to construct two “stricter” measures of the ARM share. The first alternative measure includes only those ARMs with an initial fixed-rate period of five years or less. It omits the ARMs with an initial fixed-rate period of seven and ten years, so called 7/1 and 10/1 hybrids, as well as miscellaneous loans with initial fixed-rate period greater than 5 years. We label this series $\{ARM_t^2\}$. The second alternative measure, $\{ARM_t^3\}$, contains only ARMs with initial fixed-rate period of 3 years (3/1), one year (1/1), and miscellaneous loans with initial fixed-rate period less than one year. Finally, there is an alternative source of ARM share data available from Freddie-Mac, which constructs a monthly ARM share based on the Primary Mortgage Market Survey.²⁰ This series, which we label $\{ARM_t^4\}$, conceptually measures the same

the real rate risk premium is more volatile than the inflation risk premium, it explains more of the variation in the long-term bond risk premium. This is by construction, and does not imply that it is the more important determinant of long-term risk premia, only that it is the more volatile one.

¹⁸Major 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.

¹⁹We are grateful to James Vickery for making these detailed data available to us.

²⁰This 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.

as $\{ARM_t^1\}$, and is available from 1995.1. Figure 7 plots all four series together, starting in 1992.1. The correlation between measure 2 (measure 3) and our benchmark measure 1 is 98.6% (86.3%). The correlation between measure 4 and our benchmark is 89.9%.

[Figure 7 about here.]

4.2 Main Regression Results

We start by reporting univariate regressions of the benchmark ARM share on the one-period lag of the bond risk premia we identified. The first panel of Table 1 shows the slope coefficient, its Newey-West t-statistic using 12 lags, and the regression R^2 for these regressions. The other panels will be discussed later.

Our main focus is on the 1997.7-2006.6 sample, for which we have real term structure data. The single strongest explanatory variable of variation in the ARM share is the inflation risk premium at the five-year horizon (first row). It has a t-statistic of 8.49, and explains 63.5% of the variation in the ARM share. A 1 percentage point, or two-standard deviation, increase in the expected inflation risk premium increases the ARM share by 12.7 percentage points. The inflation risk premium has to be paid by the FRM holder (the investor). An increase in the inflation risk premium makes the FRM relatively less attractive and increases the ARM share. Figure 2 in the introduction confirms that the two variables co-move strongly. The ten-year inflation risk premium (second row) looks very similar to the five-year risk premium (see Figure 4) and has a similar explanatory power of 56.2%.

The inflation risk premium continues to be strongly related to the ARM share in the full sample 1985.1-2006.6 (left columns), despite the fact that risk premia are constructed from the projection method detailed in Section 3.3.²¹ The larger point estimate suggests an even larger sensitivity of the ARM share to the inflation risk premium over the full sample. The t-statistic of $\phi_t^x(5)$ equals 5.9, and the regression R^2 is still 44%.

The real rate risk premium explains a much smaller fraction of the variation in the ARM share in the US (rows three and four). First, the real rate risk premium has the right sign in the full sample, but its correlation with the ARM share is lower. Only the real rate premium at the ten-year horizon is statistically significantly related to the ARM share; the R^2 is 12%. This correlation has the wrong sign in the 1997.7-2006.6 sample. This may be due to the strong negative correlation between the real rate risk premium and the inflation risk premium.

The nominal bond risk premium, which is the sum of the expected inflation and real rate risk premia, is a much weaker determinant than its components (row five). This is especially true in the 1997-2006 sample. The reason is that its two components are negatively correlated in that

²¹We use the projection for the entire 1985-2006 sample.

sample. In the full sample, the sum performs somewhat better, but only because it is more strongly correlated with the inflation risk premium (70% correlation versus 46% in the shorter sample) and because real rate and expected inflation premia now have a zero correlation. This result underscores the importance of considering both components of the nominal risk premium separately (See also Section 3.4).

[Table 1 about here.]

Next, we include both risk premia on the right-hand side of the ARM share regression. All regressors are demeaned so that the constant reflects the average ARM share. Table 2 shows that the importance of the inflation risk premium as a determinant of the ARM share remains unchanged. Column (5) reports the results for the sample for which we have real yield data, while Column (1) reports the full sample results. Both variables enter with the right sign in the full sample, but only the inflation risk premium is significant. In the later sample, the real rate risk premium enters negatively but is not significant. Compared to the univariate regression, the R^2 improves marginally: from 56.2 to 56.8% for the 1997-2006 sample and from 44.6 to 46.3% for the full sample, respectively. Noteworthy is that the coefficient on the inflation risk premium is stable across both samples; it is always around 15. The results with five-year risk premia instead of ten-year risk premia are very similar (not reported in the table). The R^2 with five-year risk premia is 63.6% in the 97-06 sample and 46% in the full sample. Again, for the 97-06 sample, adding the real rate risk premium barely improves on the fit of the regression with only the expected inflation risk premium. In the US, the inflation risk premium turns out to be the most important determinant of mortgage choice.

[Table 2 about here.]

4.3 Prepayment Option

Sofar we have ignored one other potentially important determinant of mortgage choice: the prepayment option. In the US, an FRM contract typically has an embedded prepayment 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. In the process, we solve for the price of the prepayment option in a model that accommodates time-varying bond risk premia, an innovation in the prepayment literature.

4.3.1 Reduced Sensitivity

A fixed-rate mortgage *without* prepayment option is a coupon-bearing nominal bond, issued by the borrower and held by the lender.²² An FRM *with* prepayment option is a callable bond. The borrower has the right to prepay the outstanding mortgage debt at any point in time; the prepayment option is of the American type. The price sensitivity of a callable bond to interest rate shocks differs from that of a regular bond. Figure 8 plots the price sensitivity of an FRM without prepayment (regular bond) and an FRM with prepayment (callable bond) to changes in the real rate (left panel) and expected inflation (right panel). The model that produces this figure is detailed in Appendix B. The regular bond price is decreasing and convex in both the real interest rate y and expected inflation x . The same is true for the callable bond price, but the relationship becomes concave (also referred to as negative convexity) when the call option is in the money. This happens when the real rate or expected inflation are low. This implies that the price of a callable bond is less sensitive to interest rate changes. This reduced exposure is most pronounced with respect to expected inflation (right panel). In sum, the FRM with prepayment has positive, but diminished exposure to real rate and expected inflation. As such, the expected payments to the FRM with prepayment increase with the real rate and inflation risk premium, but not as much as the FRM without prepayment.

[Figure 8 about here.]

In Appendix B, we detail the calculation of the prepayment option value, which is the difference between the callable and the non-callable bond price. It builds on the term structure model of Appendix A.2. In addition to this option value, we also determine the expected real payments stream that the borrower makes on the FRM contract with and without prepayment, i.e. $\sum_{s=1}^T \mathbb{E}_t [q_{t+12 \times s}^h]$, for $h \in \{p, np\}$. The latter computation uses the zero-profit rates and employs a forward simulation technique for the state variables. We set T equal to 10 years. We then regress the average of expected real payments on a ten-year FRM with and without prepayment on the ten-year real rate risk premium $\phi_t^y(120)$ and expected inflation premium $\phi_t^x(120)$. What is key to our mortgage choice analysis is that the expected payments on the FRM with prepayment continue to increase in both risk premia. Consistent with Figure 8, we find that the sensitivity of the expected payments on an FRM with prepayment option is smaller than the sensitivity for an FRM option without prepayment, but all exposures remain highly statistically significant (all

²²This 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 downpayments occur. Acharya and Carpenter (2002) discuss the valuation of callable, defaultable bonds.

t-statistics are above 3.7).

$$\frac{1}{10} \sum_{s=1}^{10} \mathbb{E}_t [q_{t+12 \times s}^p] = .11 + .60 \times \phi_t^y + .52 \times \phi_t^x + \epsilon_t^p, \quad (24)$$

$$\frac{1}{10} \sum_{s=1}^{10} \mathbb{E}_t [q_{t+12 \times s}^{np}] = .10 + 1.46 \times \phi_t^y + .96 \times \phi_t^x + \epsilon_t^{np}, \quad (25)$$

We find a similar reduction in the sensitivity to both risk premia for the real consumption streams that are associated with the FRM contract with prepayment.

4.3.2 The ARM Share and the Prepayment Option

Finally, we revisit the ARM share regressions and ask whether the prepayment feature of the FRM contract contains any additional explanatory power for the ARM share. We compute the utility difference (measured as certainty-equivalent consumption difference) between the ARM contract and the FRM contract with prepayment; it is a non-linear function of the state variables in our model. We orthogonalize this measure to other term structure variables of Columns (2) and (6), and include the orthogonal component in the regressions. Columns (4) and (8) of Table 2 show that the measure has no additional explanatory power for the ARM share. This is supporting evidence that the prepayment feature plays no significant role in understanding the choice between an ARM and FRM contract in the US.

5 Households' Ability to Estimate Bond Risk Premia

Section 2 developed a model of rational mortgage choice where time variation in mortgage choice was driven by time variation in bond risk premia. Section 3 developed a VAR model to compute the conditional expectations in (26) and therefore bond risk premia. The empirical evidence documented in Section 4 supported the claim that bond risk premia were related to the ARM share. One potential concern with this explanation for mortgage choice is that it requires substantial “financial sophistication” on the part of the households to choose the “right mortgage at the right time”. Campbell (2006) presents examples of financial illiteracy.²³ Even though mortgage choice is one of the most important financial decisions, and even though households may obtain advice from financial professionals or mortgage lenders, having access to nominal and real interest data and estimating a VAR model to form conditional expectations may be beyond reach for the average household. In this section, we show that this concern is unfounded. A simple rule-of-thumb

²³A related literature in real estate documents sub-optimally slow prepayment decisions by households, see for instance Schwartz and Torous (1989), Stanton (1995), and Boudoukh, Whitelaw, Richardson, and Stanton (1997).

captures most of the variation in mortgage choice and is strongly related to our measures of bond risk premia.

In particular, we develop an approximation to the expression for bond risk premia. We assume that households approximate conditional expectations of future short rates by forming simple averages of past short rates, going back ρ months in time:

$$\phi_t^{\$}(T) = y_t^{\$}(T) - \frac{1}{T/12} \sum_{s=1}^{T/12} \mathbb{E}_t \left[y_{t+12 \times (s-1)}^{\$}(12) \right] \quad (26)$$

$$\begin{aligned} &\simeq y_t^{\$}(T) - \frac{1}{T/12} \sum_{s=1}^{T/12} \left\{ \frac{1}{\rho} \sum_{u=0}^{\rho-1} y_{t-u}^{\$}(12) \right\} \\ &= y_t^{\$}(T) - \frac{1}{\rho} \sum_{u=0}^{\rho-1} y_{t-u}^{\$}(12) \equiv \kappa_t(\rho; T). \end{aligned} \quad (27)$$

Equation (27) is a model of adaptive expectations that only requires knowledge of the current long bond rate, a history of recent short rates, and the ability to calculate a simple average. It 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 proposed by Campbell and Cocco (2003) and Campbell (2006):

$$\kappa_t(1; T) = y_t^{\$}(T) - y_t^{\$}(12).$$

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. This complements the discussion of Section 2.3. Second, when $\rho \rightarrow \infty$ and short rates are stationary, then $\kappa_t(\rho; T)$ converges to the long-term yield in excess of the unconditional expectation of the short rate:

$$\lim_{\rho \rightarrow \infty} \kappa_t(\rho; T) = y_t^{\$}(T) - \mathbb{E} \left[y_t^{\$}(12) \right], \quad (28)$$

by the law of large numbers. Because the second term is constant, all variation in financial incentives to choose a particular mortgage originates from variation in the long-term yield. For all cases in between the two extremes, the simple model of adaptive expectations has the household put some positive and finite weight on average recent short-term yields to form conditional expectations.

Figure 9 displays the ARM share (solid) alongside the rule-of-thumb for 10-year bond risk premia that uses three years of past data, $\kappa_t(36; 120)$. The results are portrayed over the period 1987.12-2006.6, which uses the full sample (1985.1-2006.6) to construct κ_t . The figure documents a striking co-movement between the ARM share (solid line, right axis) and the rule-of-thumb for bond risk premia (dashed line, left axis).

[Figure 9 about here.]

Figure 10 shows the correlation of $\kappa_t(\rho, 120)$ for different values of ρ . The bars correspond to $\rho = 12, 24, 36, 48,$ and 60 . The solid line depicts the correlation between the yield spread and the ARM share ($\rho = 1$). The dashed line corresponds to the correlation between the long-term rate and the ARM share ($\rho = \infty$). In the left panel, $\kappa_t(\rho, 120)$ is computed based on treasury yield data, while in the right panel it is computed from the 1-year ARM rate and the 10-year FRM rate.²⁴ The results are shown for the period 1989.12-2006.6, the longest sample for which all measures are available.²⁵

[Figure 10 about here.]

Figure 10 contains three important results. First, both the yield spread (solid line) and the long-term yield (dashed line) have a weak contemporaneous correlation with the ARM share. The second panel of Table 1 confirms that the lagged yield spread explains little variation in the ARM share in both samples (rows 7 and 8); the R^2 is less than 1 percent. This arises because the yield spread not only measures the nominal risk premium, but also the deviation of the expected future short rate from the current short rate (see Section 2.3). Our VAR analysis shows that these two components are negatively related. This makes the yield spread a weak predictor of mortgage choice. Rows 9 and 10 of the same table show that the lagged long-term interest rate is weakly related to the ARM share variation between 1997-2006, but stronger in the full sample. The third panel of Table 1 shows that the FRM-ARM spread explains the ARM share variation better in the full sample, but not in the 1997-2006 sample (row 11). Clearly, the FRM-ARM spread contains additional information unrelated to the treasury yield spread. 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 12). For the full sample, we find a strongly significant effect.²⁶ The explanatory power of the FRM rate is similar to that of the long treasury yield (row 13).²⁷ The same can be seen in the right panel of 10.

²⁴We 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.

²⁵We 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.

²⁶The correlation between the FRM-ARM spread and the ten-one-year government bond yield spread is only 32% over the full sample, and 11% over the 1997-2006 sample. This spread also captures the value of the prepayment option, as well as the lenders’ profit margin differential on the FRM and ARM contracts.

²⁷Furthermore, the explanatory power of long bond yields and FRM rates seems related to that of the expected inflation risk premium. They are positively correlated with the inflation risk premium in the full sample, but negatively correlated in the later sample.

Second, the correlation is hump-shaped in ρ in both panels of Figure 10. The highest correlation with observed mortgage choice is obtained when households use 2-4 years of short rate data in their computation. Third, the correlation peaks at 80%. For comparison, the correlation between the inflation risk premium and the ARM share over the same 17-year period is 58%.²⁸ In sum, this simple way of computing bond risk premia explains most of the variation in the ARM share. This lends further support to our claim that bond risk premia are the key determinant of mortgage choice variation.

Finally, we ask what fraction of the rule-of-thumb bond risk premium, κ_t , is captured by the expected inflation and real rate risk premia that we extract from the VAR. To that end, we regress $\kappa_t(\rho; 120)$ on a constant and both risk premia, $\phi_t^x(120)$ and $\phi_t^y(120)$. For the full sample period for which all measures are available (1989.12-2006.6), we find that the real rate and inflation risk premium explain about 65% of the variation in $\kappa_t(\rho; 120)$. This number increases to almost 85% for the period 1997.7-2006.6. The inflation risk premium always enters positively, whereas the loading on the real rate premium is typically negative, consistent with the results for the ARM share regression in Section 4. The R-squared peaks around 2-4 years, consistent with Figure 10. The slightly weaker correlations over the longer sample may well reflect the fact that the risk premia are based on a projection; we only measure $\phi_t^x(120)$ and $\phi_t^y(120)$ over 1997.7-2006.6, and that is the period for which we find the highest correlation. Panels A and B of Table 3 report results of regressions of the $\kappa_t(\rho; 120)$ measure on the two component risk premia as well as on the component of the yield spread that is orthogonal to the two risk premia. They convey the same message. In conclusion, our VAR-based estimate of bond risk premia explains a large part of the variation in the simple estimate of bond risk premia.

[Table 3 about here.]

Why do households seem to be able to time mortgage markets so well? Because they are able to approximate time-varying bond risk premia by a simple metric which only requires calculating the average of recent short rates.

6 Robustness

In this section we first perform several robustness checks for the US. Second, we discuss mortgage choice in the UK.

²⁸For the optimal value of ρ , we expect $\kappa_t(\rho; T)$ to explain the variation in the ARM share better than using the conditional expectations that follow from the VAR model. After all, we now use a simpler model of expectations that can easily be implemented by households. If this model accurately describes households' behavior, we expect it to explain more of the variation in households' mortgage choice.

6.1 Analysis for the United States

We discuss a rich set of alternative model assumptions and variable definitions for the US. We find that our main finding is robust to these alternative specifications; the bond risk premium, and in particular the inflation risk premium component, remains an important determinant of mortgage choice.

6.1.1 Heteroscedasticity

We now extend the VAR model to allow for heteroscedastic innovations. In particular, we allow for time-varying volatility in the real interest rate y and expected inflation x . Long-term expectations are unaffected by the switch from homoscedastic to heteroscedastic model, so that the term structure dynamics presented before remain identical.

We first estimate the innovations ($\hat{\eta}_t, t = 1, \dots, T$) from the VAR-model and construct the implied innovations to the real rate and expected inflation according to (29) and (30),

$$\eta_{t+12}^x = x_{t+12}(12) - \mathbb{E}_t[x_{t+12}(12)] = e_4' \Gamma \eta_{t+12}, \quad (29)$$

$$\eta_{t+12}^y = y_{t+12}(12) - \mathbb{E}_t[y_{t+12}(12)] = (e_1' - e_4' \Gamma) \eta_{t+12}. \quad (30)$$

Next, we model both conditional variances as an exponentially affine function in their own level

$$V_t^x \equiv \text{Var}_t[x_{t+12}(12)] = \text{Var}_t[\eta_{t+12}^x] = \exp(\alpha_x + \beta_x x_t(12)), \quad (31)$$

$$V_t^y \equiv \text{Var}_t[y_{t+12}(12)] = \text{Var}_t[\eta_{t+12}^y] = \exp(\alpha_y + \beta_y y_t(12)). \quad (32)$$

The coefficients α_i and β_i , $i = x, y$, are estimated consistently via non-linear least squares

$$(\hat{\alpha}_i, \hat{\beta}_i) = \arg \min_{\alpha_i, \beta_i} \frac{1}{T} \sum_{t=1}^T \left([\hat{\eta}_{t+12}^i]^2 - \exp(\alpha_i + \beta_i i_t(12)) \right)^2.$$

The top right panel of Figure 3 plots the conditional volatilities of expected inflation and the real rate (see Equations (31) and (32)). Conditional real rate volatility is 1.06% per year on average, while expected inflation volatility is three times lower at 0.35% per year on average. There is some time variation in these one-year ahead conditional volatilities. The two conditional volatilities co-move strongly negatively; their correlation is -0.71. For example, real rate volatility is high in 2004, when the real rate is low, and low in the 1985, when the real rate is high. In contrast, expected inflation volatility is at its highest level in 1991, when expected inflation is high, and low in 2002, when expected inflation is low.

The next step is to include the 1-year ahead conditional variances V_t^x and V_t^y in the ARM share

regression.²⁹ In contrast to the risk premia, these conditional variances are available for the entire 1985-2006 sample. Row 14 of Table 1 shows that a higher volatility of the real rate makes the ARM less desirable. Row 15 shows that a higher inflation volatility makes the ARM more desirable. We then add the two volatility terms to the two risk premia as regressors in columns (2) and (6) of Table 2. They have the same signs as in the univariate regressions. The real rate volatility is significant in the full sample. The R^2 improves by 7.5% in the full sample and by 4.1% in the 1997-2006 sample. The negative sign on the volatility of the real rate is predicted by the model. The FRM contract has no real rate risk, so more volatility makes the ARM less desirable relative to the FRM. The positive sign on the conditional volatility of expected inflation is consistent with a rational investor who is constrained with a short to intermediate horizon or an unconstrained investor with any horizon,³⁰ but not with an investor with money illusion (Brunnermeier and Julliard (2006)). High expected inflation volatility (V_t^x) makes the ARM more risky for investors who are unable to disentangle real rates and expected inflation. We note that the volatility of the *nominal* interest rate is also significantly negatively related to the ARM share in the full sample (not reported). Most of the conditional variance of the nominal rate is inherited by the real rate; the two have a correlation of 95%.

Finally, Columns (3) and (7) of Table 2 show that there is no extra information in the yield spread that is useful for predicting the ARM share, and not already present in the term structure variables. They add the orthogonal component of the yield spread as an explanatory variable of the ARM share. The R^2 does not increase and its coefficient is insignificant.

6.1.2 Other ARM Share Measures

As a robustness exercise, we repeat the analysis in Column (6) of Table 2 for the alternative measures of the ARM share discussed above. The second and third column of Table 4 show that the explanatory power of the term structure variables is very similar without the hybrid mortgages. In the second column we count the hybrid ARM contracts with initial fixed-rate period greater

²⁹The theory calls for the average of the 1-period-ahead to T-period-ahead conditional variances. Because long-term average variances are positively correlated with the 1-period-ahead conditional variance, the sign on V_t^x and V_t^y should be the same. The reason is that the term-structure of volatilities is upward sloping. It converges to the unconditional variance, which is higher than the 1-period-ahead conditional variance.

³⁰An increase in inflation uncertainty increases the variance of the intermediate payments on the ARM if the investor is constrained, but not on the FRM (using the approximations in Section 2). In contrast, the terminal payment of the ARM is hedged against expected inflation from period $T - 1$ to T , while the terminal payment for the FRM is not. To understand quantitatively the impact of inflation risk for a borrowing-constrained household, we simulated from the model in Appendix A.2. The simulation results indicate that for short to intermediate horizons (T), the FRM contract carries most inflation risk, but it is the ARM contract that is most exposed to inflation risk for horizons close to 30 years. An unconstrained ARM investor, instead, can shift forward the increase in intermediate mortgage payments, arising from increased expected inflation, to time T . The additional amount borrowed exactly cancels against the erosion of the nominal mortgage balance due to expected inflation. In that case, the FRM carries most inflation risk (see also Campbell (2006)).

than five years as FRMs. The average fraction of ARMs falls from 21.6% to 18.1%, but the results remain virtually unchanged. In the third column, we also eliminate the mortgages with an initial fixed-rate period greater than three years. The average fraction of ARMs drops to 11%, half as much as in Column (1). The R^2 drops by almost 20 percentage points, but the inflation risk premium remains highly significant. The real rate volatility now also becomes significant. We conclude that our results are robust to the classification of the hybrid mortgage contracts. In Column (4), we use the Freddie Mac data instead of the FHFB data. This makes little difference compared to Column 1. The R^2 is the highest for this measure, and equal to 64%. Finally, we also experiment with splitting up the ARMs in Column (1) into mortgage contracts for newly constructed homes and mortgages for existing structures (not reported in the table). The average fraction of ARM shares is 28% for the former group and 20% for the latter. The inflation risk premium is a highly significant determinant of mortgage choice in both groups. The R^2 , however, is twice as high for the existing structures (68%) than for the new-construction group (34%).

[Table 4 about here.]

6.1.3 Alternative Interest Rate Models

We have studied how the ARM share regressions are affected when we change the underlying term structure model. We have estimated a VAR(2)-model for Y instead of a VAR(1). The inflation and real rate risk premia in the VAR(2) model look qualitatively similar to those in Figure 4. The only small difference is that the long-term expected inflation is estimated to be somewhat lower (around 2% per year), so that the inflation risk premium is higher near the end of the sample, and the increase in the inflation risk premium since 2003 is more pronounced. We then rerun the ARM share regressions, corresponding to Columns (2) and (6) in Table 2. The inflation risk premium is as prominent an explanatory variable as in the benchmark analysis. The R^2 on the regression further improves to 70.5% in the second sample (from 61%), and to 57.6% in the full sample (from 53.9%).

We also verified the robustness of our results to alternative volatility models for the expected inflation and the real rate in (31) and (33). This comparison showed that the benchmark volatility model has the lowest sum of squared deviation between the squared innovations and the proposed conditional volatility. Further, because the difference between alternative volatility models is relatively small, the ARM share regressions provide similar results for alternative conditional volatility series.

6.1.4 Alternative Real Yield Measures

Due to the availability of TIPS data, our main results are for the 1997-2006 sample. We have used the projection method as a first robustness exercise to verify whether our results extend to

the longer 1985-2006 sample. As a further robustness check, we use the real yield data backed-out from the term structure model of Ang, Bekaert, and Wei (2007) instead of the TIPS yields. We proceed as in Section 3, forming inflation and real rate expectations in the same way. We treat the real yields as observed, and use them to construct the inflation risk premium and the real rate premium in turn. Since the Ang-Bekaert 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 enter with the correct, positive sign, and both coefficients are statistically significant. The Newey-West t-statistic on the inflation risk premium is 3.9 and the t-statistic on the real rate risk premium is 2.12. The regression R^2 is 53%. Including the quarterly-sampled conditional volatility terms increases the R^2 further to 70%, while increasing the importance of the inflation risk premium. Our results are therefore robust to an alternative, model-implied measurement of real yields.

6.1.5 Robustness: Persistence of Regressor

In contrast to the inflation risk premium, most term structure variables in Table 1 do not explain much of the variation in the ARM share. This is especially true in the 1997-2006 sample. This suggests that our results for the inflation 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. 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 four term structure variables of the benchmark regression specification (between periods $t - 1$ and t). We continue to find a positive and strongly significant effect of the inflation risk premium on the ARM share. The magnitude of the regression coefficients implies that a one percentage point increase in the inflation risk premium increases the ARM share by 15.5 percentage points in the full sample and 11.6 in the 1997-2006 sample, all else equal. These sensitivities are consistent with our findings for the level regressions. The R^2 of the regression in changes is obviously lower, but still substantial: 20.8% in the full sample and 17.0% in the 1997-2006 sample.

6.2 Analysis for the United Kingdom

As a final robustness check, we repeat the full analysis for the UK. This is important for at least three reasons. First, to show that the same yield curve variables also explain a substantial fraction of the ARM share in a different country is an important robustness check. Second, to the extent that bond risk premia look different in the UK than in the US and to the extent that risk premia are still related to variation in the ARM share, the term structure determination theory of mortgage

choice gains further credibility. Third, we have a longer, and potentially better time series of real bond yields available in the UK than in the US.

6.2.1 VAR Results

We repeat the term structure analysis for the UK. That is, we estimate a monthly VAR with 12-month lagged bond yields of one-, five-, and ten-year maturity and realized inflation on the right-hand side. The nominal yields are from the Bank of England, the inflation rate is the 12-month log difference in Retail Price Index (RPI), the counterpart to the American CPI. As for the US, we estimate the VAR on the period 1985.1-2006.6. After we form long-term expected inflation and long-term expected real rates, we construct the inflation risk premium and real rate risk premium using real yield data on five- or ten-year bonds. The UK has much longer time series for inflation-linked bond yields; the Bank of England data start in 1985.1.

[Figure 11 about here.]

There are substantial differences between the evolution of the term structure in the US and in the UK. Figures 11 and 12 are the UK counterparts to Figures 3 and 4. Figure 11 shows that long-term expected inflation and real rate trend downwards over the sample, aside from an increase in the late 1980s. The decline in expected inflation after 1991 tracks the decline in realized inflation; likewise, realized inflation was high in the late 1980s. Inflation expectations stabilize around 2% per year at the end of the sample, lower than in the US. As a result of the bigger nominal-real yield spread and the lower inflation expectations, the inflation risk premium is much larger in the UK. Figure 12 shows that it is mostly positive, and it goes down from 3% to 1% per year over the sample period. Conversely, the real rate premium is mostly negative in the UK; it was positive in the US. After reaching a low around 1990, it stabilizes around 0% after 1995. Just as in the US, the two risk premia are strongly negatively correlated (-78% at the five- and -48% at the ten-year horizon). The bottom two rows zoom in on the 1997.7-2006.6 subsample, the same sample we had in Figure 4 for the US. In the UK, the two risk premia drift away from each other after 2002, whereas in the US, they both seem to converge to zero. Finally, the top right panel of Figure 11 shows that the conditional volatilities of the real rate and expected inflation co-move positively in the UK, whereas they co-move negatively in the US. In short, the term structure of interest rates in the UK looks dramatically different from the term structure in the US.

[Figure 12 about here.]

6.2.2 ARM Share Regression Results

The mortgage market in the UK has some important differences with the US. First, long-term fixed-rate mortgages are a lot less prevalent than in the US. The most prevalent contract is a

standard variable rate contract, for which the interest rate is adjusted several times per year. Most fixed-rate contracts have one-to-three-year fixed interest rate periods. Ten-year fixed-rate contracts are relatively new.³¹ Second, fixed-rate mortgages have no embedded prepayment option. We have data on the mortgage composition in the UK from the Council of Mortgage Lenders starting in 1993. The data are monthly from 2002.1 onwards, and quarterly from 1993.1 onwards. The quarterly data are averages of the three months of the quarter. We use a Kalman filtering procedure, suggested by Hansen and Sargent (2004), to undo the temporal aggregation, which results in a monthly time-series that starts in 1993.1. Appendix C contains the details.

Despite the differences between the UK and the US, Figure 13 shows that there is a lot of variation in mortgage composition in the UK as well. The ARM share varies between 85% and 25%. The overall fraction of adjustable rate contracts is higher than in the US. The ARM share decreases near the end of the sample because of the increased availability and popularity of longer-term fixed-rate contracts.

[Figure 13 about here.]

We repeat the regressions of the ARM share on the one-month lagged term structure variables in Table 5. The left columns report the results for the 1993.1-2006.6 sample, using a monthly ARM series based on the quarterly ARM share data, whereas the right columns report the results for the sample 2002.1-2006.6, for which we have actual monthly ARM shares. In sharp contrast to the analysis for the US, it is the real rate risk premium that is the key driver of mortgage choice in the UK. In both Columns (1) and (5), the ten-year real rate risk premium is significant. The expected inflation risk premium has the right sign, but is only significant in the second sample. The part of the inflation risk premium that is orthogonal to the real rate risk premium does not add to the explanation of the ARM share. I.e., the R^2 increases only marginally compared to the univariate regression of the ARM share on the real rate premium. The effects of the real rate premium are economically large. A 0.35 percentage point, or two-standard deviation, increase in the ten-year real rate risk premium in the 2002-2006 sample increases the ARM share by 16 percentage points, from 52% to 68%. The R^2 is on the same order of magnitude as the one we found for the US in the later sample, but this time the real rate premium is the key variable instead of the inflation risk premium. Just as in the US, adding the variance terms in Column (2) and (6) improves the explanatory power only a little. The R^2 increases to 75% in the later sample. In the specification in column (6), all variables have the right sign, and three are significant. Finally, adding orthogonal information in the yield spread does not improve the R^2 in the full sample, but it helps in the later

³¹The Miles report (2004) contains a detailed overview of the UK mortgage market, and its recommendation is to deepen the long-term fixed-rate mortgage market. Coles and Hardt (2000) discuss differences between US and European mortgage markets.

sample. The R^2 improves by another 1.4% in Column (7). We do not consider the option value effects because there is generally no prepayment option in the UK.

[Table 5 about here.]

6.2.3 Understanding the Difference Between the UK and the US

The rule-of-thumb proxy for bond risk premia provides a hint as to why the inflation risk premium is the dominant explanatory variable for the ARM share in the US, while the real rate risk premium matters most in the UK.

Table 3 reports the results for a regression of $\kappa_t(\rho; 120)$, i.e., the bond risk premia computed as in Section 5, on a constant, both component risk premia, and the yield spread, which we first orthogonalize to the other factors (as in Tables 2 and 5). The ‘★’ indicates the highest correlation between the rule-of-thumb measure and the ARM share. It can be interpreted as the model that best represents the behavior of households. We focus on these columns in our discussion. For the US, we find that $\kappa_t(36; 120)$ loads positively on the inflation risk premium, but negatively on the real rate premium and the yield spread. The yield spread is also insignificant. These results are consistent with the findings in Table 2. For the UK, we find for the period 1993.1-2006.6 that both risk premia load positively, and the yield spread negatively. However, the explanatory power of this regression is relatively low.³² For the more recent period 2002.1-2006.6, in contrast, we find that both risk premia and the yield spread are significantly related to $\kappa_t(12; 120)$. In addition, the signs of all variables are as in Table 5. The R-squared for this regression equals 87%. Hence, the signs on the bond risk premia that can be reliably estimated in the ARM share regressions of Tables 2 and 5 are identical to the signs of a projection of the rule-of-thumb proxies for bond risk premia on our VAR-based measures of risk premia and the yield spread.

If households implement the simple rule-of-thumb to measure risk premia, and that is what the high correlation between the proxy and the ARM share seems to suggest, then we have found an explanation for the relative importance of the real rate risk premium in the UK and the inflation risk premium in the US.

7 Conclusion

We have shown that the time variation in the risk premium on a long-term nominal bond can explain a large fraction of the variation in the share of newly-originated mortgages that are of the adjustable-rate type. Thinking of fixed-rate mortgages as a short position in long-term bonds and

³²The correlation of the ARM share with $\kappa_t(48; 120)$ equals 42% for the period 1993.1-2006.6. The correlation of $\kappa_t(12; 120)$ with the ARM share is 55% instead over the sample 2002.1-2006.6.

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, which consists of an expected inflation premium and a real rate premium. In the US, fixed-rate mortgages tend to have long maturities and are therefore very sensitive to inflation risk. We have shown that the inflation risk premium alone can explain more than sixty percent in the time variation of the mortgage composition. These results are not sensitive to how expected inflation is measured: we studied both a direct measure from the survey of professional forecasters, and an indirect measure from a VAR model. Other, perhaps more straightforward, term structure variables such as the slope of the yield curve, have much lower explanatory power for the ARM share. As an additional check on the validity of the term structure determination of mortgage choice, we have also studied the UK. While the term structure variables of interest have dynamics quite different from those in the US, bond risk premia are still linked to mortgage choice. In the UK, where FRMs are a lot less prevalent and of much shorter maturity, it is the real rate risk premium that drives the variation in mortgage choice instead.

We have also shown that a simple rule-of-thumb approximates the VAR-based risk premia rather well. The proxy is the difference between the long-term interest rate and a backward-looking average of short-term interest rates, where the average is calculated over 2-4 years. This measure is not only strongly related to our bond risk premia, but also to observed mortgage choice. Taken together, our findings suggest that households may be making close-to-optimal mortgage choice decisions, because capturing the relevant time-variation in bond risk premia may be easier than previously thought. 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 should take 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.
- ANG, A., G. BEKAERT, AND M. WEI (2006): “Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?,” *Journal of Monetary Economics*, Forthcoming.
- (2007): “The Term Structure of Real Rates and Expected Inflation,” Working Paper Columbia University.
- 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.
- 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.
- 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.
- COLES, A., AND J. HARDT (2000): "Mortgage Markets: Why the US and EU Markets Are So Different," *Housing Studies*, 15, 775–783.
- DAI, Q., AND K. J. SINGLETON (2002): "Expectation Puzzles, Time-varying Risk Premia, and Affine Models of the Term Structure," *Journal of Financial Economics*, 63, 415–441.
- DE JONG, F. (2000): "Time Series and Cross Section Information in Affine Term Structure Models," *Journal of Business and Economic Statistics*, 18(3), 300–314.
- DUFFEE, G. R. (2002): "Term Premia and Interest Rate Forecasts in Affine Models," *Journal of Finance*, 57, 405–443.
- DUNN, K. B., AND J. J. MCCONNELL (1981): "Valuation of Mortgage-Backed Securities," *Journal of Finance*, 36, 599–617.
- 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.
- HANSEN, L. P., AND T. J. SARGENT (2004): *Recursive Models of Dynamic Linear Economies*. Mimeo University of Chicago and New York University, first edn.
- 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.
- LONGSTAFF, F. A. (2005): "Borrower Credit and the Valuation of Mortgage-Backed Securities," *Real Estate Economics*, 33, 619–661.

- LONGSTAFF, F. A., AND E. S. SCHWARTZ (2001): "Valuing American Option by Simulation: A Simple Least-Squares Approach," *Review of Financial Studies*, 14(1), 113-147.
- PLISKA, S. R. (2006): "Optimal Mortgage Refinancing with Endogenous Mortgage Rates: an Intensity Based Equilibrium Approach," Working Paper University of Illinois at Chicago.
- 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.
- 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.
- STANTON, R. (1995): "Rational Prepayment and the Valuation of Mortgage-Backed Securities," *Review of Financial Studies*, 8, 677-708.
- STENTOFT, L. (2004): "Assessing the Least Squares Monte-Carlo Approach to American Option Valuation," *Review of Derivatives Research*, 7, 129-168.
- VAN HEMERT, O. (2006): "Life-Cycle Housing and Portfolio Choice with Bond Markets," Working Paper, NYU Stern School of Business.
- VICKERY, J. (2006): "Interest Rates and Consumer Choice in the Residential Mortgage Market," Working Paper, Federal Reserve Bank of New York.

A Utility Framework

A.1 Taylor Expansion

In the main text, we focus on a second-order Taylor expansion of the CRRA preferences. We recall that the real labor income process $\{L_{12 \times t}\}_{t=1}^T$ is i.i.d. Denote its first moment by ℓ and its second central moment by σ_ℓ^2 . Each of the intermediate real consumption streams $\{C_{12 \times t}\}_{t=1}^{T-1}$ are approximated around $C_{12 \times t} = \ell$. This second-order approximation delivers:

$$\begin{aligned} \frac{C_{12 \times t}^{1-\gamma}}{1-\gamma} &\approx \frac{\ell^{1-\gamma}}{1-\gamma} + \ell^{-\gamma}(L_{12 \times t} - \ell - q_{12 \times t}) - \frac{\gamma}{2}\ell^{-\gamma-1}(L_{12 \times t} - \ell - q_{12 \times t})^2 \\ \mathbb{E}_0 \left[\frac{C_{12 \times t}^{1-\gamma}}{1-\gamma} \right] &\approx \frac{\ell^{1-\gamma}}{1-\gamma} + \ell^{-\gamma}\mathbb{E}_0[-q_{12 \times t}] - \frac{\gamma}{2}\ell^{-\gamma-1}\mathbb{E}_0[(L_{12 \times t} - \ell)^2 + q_{12 \times t}^2 - 2q_{12 \times t}(L_{12 \times t} - \ell)] \\ -\ell^\gamma \mathbb{E}_0 \left[\frac{C_{12 \times t}^{1-\gamma}}{1-\gamma} \right] &\approx \left\{ \frac{-\ell}{1-\gamma} + \frac{\gamma}{2\ell}\sigma_\ell^2 \right\} + \mathbb{E}_0[q_{12 \times t}] + \frac{\gamma}{2\ell}\mathbb{E}_0[q_{12 \times t}^2] \end{aligned}$$

In the second line, we take the conditional expectation as of time 0. We use that $\mathbb{E}_0[L_{12 \times t} - \ell] = 0$. In the third line, we use that $\mathbb{E}_0[(L_{12 \times t} - \ell)^2] = \sigma_\ell^2$ and that mortgage payments are orthogonal to labor income, because the latter is i.i.d. We also multiply through by $-\ell^\gamma$.

The terminal consumption stream $C_{12 \times T}$ is approximated around $C_{12 \times T} = \ell$, i.e., it approximates around a zero inflation rate: $1/\Pi_{12 \times T} \approx 1$. Going through the same steps as above, we obtain

$$\begin{aligned} -\ell^\gamma \mathbb{E}_0 \left[\frac{C_{12 \times T}^{1-\gamma}}{1-\gamma} \right] &\approx \left\{ \frac{-\ell}{1-\gamma} + \frac{\gamma}{2\ell}\sigma_\ell^2 \right\} + \mathbb{E}_0[q_{12 \times T}] + \frac{\gamma}{2\ell}\mathbb{E}_0[q_{12 \times T}^2] + \\ &\quad \left\{ -\mathbb{E}_0[1 - 1/\Pi_{12 \times T}] + \frac{\gamma}{2\ell}\mathbb{E}_0[(1 - 1/\Pi_{12 \times T})^2] - \frac{\gamma}{\ell}\mathbb{E}_0[q_{12 \times T}(1 - 1/\Pi_{12 \times T})] \right\} \end{aligned}$$

We have used the independence of real labor income from mortgage payments and from inflation. The terms in braces are new; they capture the effect of inflation on the terminal mortgage balance.

Maximizing the utility from real consumption streams is therefore (approximately) equivalent to minimizing these functions of expected mortgage payments, expected squared mortgage payments, and expected inflation. Taken together, an investor prefers the T -year ARM contract over the T -year FRM contract at time zero if and only if

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}_0[q_{12 \times t}^{FRM}] + \frac{\gamma}{2\ell}\mathbb{E}_0[(q_{12 \times t}^{FRM})^2] - \frac{\gamma}{\ell}\mathbb{E}_0[q_{12 \times T}^{FRM}(1 - 1/\Pi_{12 \times T})] &> \\ \sum_{t=1}^T \mathbb{E}_0[q_{12 \times t}^{ARM}] + \frac{\gamma}{2\ell}\mathbb{E}_0[(q_{12 \times t}^{ARM})^2] - \frac{\gamma}{\ell}\mathbb{E}_0[q_{12 \times T}^{ARM}(1 - 1/\Pi_{12 \times T})] & \end{aligned}$$

Finally, we assume that for a generic interest rate r and period t

$$r_t/\Pi_t \approx r_t \left(1 - \sum_{s=1}^t \pi_s \right) \approx r_t$$

The first approximation is a first-order Taylor expansion around $r = 0$. The second approximation says that an interest rate times aggregate inflation is an order of magnitude smaller than the rate itself, if T is not too large. Using the definition of the real payments in (8), this approximation implies that

$$\frac{f_0^\$(12 \times (T - 1))}{\Pi_{12 \times T}} (1 - 1/\Pi_{12 \times T}) \approx 0, \quad \text{and} \quad \frac{y_{12 \times (T-1)}^\$(12)}{\Pi_{12 \times T}} (1 - 1/\Pi_{12 \times T}) \approx 0$$

This delivers the mortgage choice expression (12) in the main text.

Next, we analyze the difference in expected mortgage payments on the FRM and the ARM contracts:

$$\sum_{t=1}^T \mathbb{E}_0 [q_{12 \times t}^{FRM}] - \sum_{t=1}^T \mathbb{E}_0 [q_{12 \times t}^{ARM}] = \sum_{t=1}^T f_0^\$(12 \times (t - 1)) \mathbb{E}_0 \left[\frac{1}{\Pi_{12 \times t}} \right] - \sum_{t=1}^T \mathbb{E}_0 \left[\frac{y_{12 \times (t-1)}^\$(12)}{\Pi_{12 \times t}} \right]$$

Under the same assumption that an interest rate times an inflation rate is approximately equal to the interest rate, we get expression (13) in the main text. Under the same approximation, the expectations of *squared* real and nominal mortgage payments is approximately the same.

A.2 Exact Numerical Solution

Above we made approximations that were meant to cleanly expose the role of bond risk premia as a determinant of mortgage choice. Here, we set up a term structure model that does not make these approximations. We numerically solve for the utility over the real consumption stream arising from an FRM contract and compare it to the utility stream from an ARM contract. We show that the utility difference between the FRM and the ARM is strongly correlated with bond risk premia. This justifies our emphasis on the mortgage payment differential, and hence on time-varying bond risk premia, as the main driver of mortgage choice. In Appendix B, we use that same term structure model to solve for the value of the prepayment option.

A.2.1 Term Structure Model

We use a five-dimensional model and denote the corresponding variables with a ‘ \star ’ superscript:

$$Y_{t+12}^\star = \mu^\star + \Gamma^\star Y_t^\star + \epsilon_{t+12}^\star, \quad (33)$$

$\epsilon_{t+12}^* \sim N(0, \Sigma^*)$. The state vector includes the real rate, expected inflation, realized log-inflation, the 10-year real rate premium, and the 10-year expected inflation premium:

$$Y_t^* = (y_t(12), x_t(12), \pi_t(12), \phi_t^y(120), \phi_t^x(120)).$$

Next, we postulate a nominal log pricing kernel of the form:

$$-m_{t+12} = y_t(12) + x_t(12) + \frac{1}{2}\Lambda_t'\Sigma^*\Lambda_t + \Lambda_t'\epsilon_{t+12}^*. \quad (34)$$

Following the literature on affine term structure models (e.g., Dai and Singleton (2002) and Duffee (2002)), the market prices of risk Λ_t , are assumed to be affine functions of the state vector:

$$\Lambda_t = \lambda_0 + \Lambda_1 Y_t^*. \quad (35)$$

To avoid that the model is over-parameterized, we allow only:

$$\lambda_{0(1:2)}, \Lambda_{1(1:2,1:2)}, \text{ and } \Lambda_{1(4:5,4:5)},$$

to be non zero.

This model implies an affine model for the term structure of interest rates, in which yields are given by:

$$y_t^s(n) = -\frac{A_n}{n} - \frac{B_n'}{n} Y_t^*,$$

with A_n and B_n solutions to the following set of differential equations:³³

$$A_n = A_{n-12} + B_{n-12}'\mu - \Lambda_0'\Sigma^*B_{n-12} + \frac{1}{2}B_{n-12}'\Sigma^*B_{n-12}, \quad (36)$$

$$B_n = \Gamma' B_{n-12} - \left(\begin{array}{ccccc} 1 & 1 & 0 & 0 & 0 \end{array} \right)' - \Lambda_1'\Sigma^*B_{n-12}, \quad (37)$$

and starting values:

$$A_0 = 0 \text{ and } B_0 = 0_{5 \times 1}. \quad (38)$$

³³These equations can be derived using the results in Ang and Piazzesi (2003).

A.2.2 Estimating the Model

Γ^* is constrained to ensure that the conditional expectation of expected inflation equals the second element of Y_t , i.e.:

$$\mu_{(3)}^* = 0 \text{ and } \Gamma_{(3,:)}^* = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \end{pmatrix}.$$

The remaining parameters are estimated by OLS per equation. The time series we use for x_t and y_t come from estimating the VAR model of Section 3. The two bond risk premia time-series are those constructed in Section 3.3. This results in the following estimates:

$$\begin{aligned} \mu^* &= [-0.0257, 0.0153, 0.0000, -0.0002, -0.0212]' \\ \Gamma^* &= \begin{bmatrix} 0.4909 & 2.1981 & -0.7750 & -0.2580 & 0.5634 \\ 0.0739 & 0.6981 & -0.2242 & -0.0541 & 0.2747 \\ 0.0000 & 1.0000 & 0.0000 & 0.0000 & 0.0000 \\ 0.1425 & 0.5938 & -0.1915 & 0.2867 & -0.0279 \\ -0.2385 & 0.6234 & 0.0306 & 0.2668 & 0.3122 \end{bmatrix} \\ \Sigma^*(\times 10^6) &= \begin{bmatrix} 106.0 & 15.8 & 27.3 & -8.0 & 22.8 \\ 15.8 & 11.3 & 25.3 & -1.2 & 6.1 \\ 27.3 & 25.3 & 93.6 & 5.1 & 22.3 \\ -8.0 & -1.2 & 5.1 & 8.4 & 6.2 \\ 22.8 & 6.1 & 22.3 & 6.2 & 14.2 \end{bmatrix} \end{aligned}$$

The unconditional mean of the state vector is given by:

$$(I_5 - \Gamma^*)^{-1} \mu^* = [0.0160, 0.0281, 0.0281, 0.0189, -0.0023]',$$

whose second and third element are identical.

The only parameters that remain to be estimated are the market prices of risk λ_0 and Λ_1 in (35). They are not pinned down by the VAR. Using a large cross-section of bonds allows a more accurate estimation of the market prices of risk (De Jong (2000)). We therefore use yield data of bonds with 2-, 4-, 6-, 8-, and 10-year maturities. Since we cannot fit all yields exactly, we adopt the standard approach of assuming that the yields are observed with measurement error $\eta_{it} \sim N(0, \sigma_i^2)$:

$$y_t^{\$}(n_i) = -\frac{A_{n_i}}{n_i} - \frac{B'_{n_i}}{n_i} Y_t^* + \eta_{it},$$

for all 5 yields ($i = 1, \dots, 5$). The measurement error is assumed to be independent of other innovations, both cross-sectionally and sequentially. The estimation employs a maximum likelihood

procedure. We maximize

$$\prod_{t=1}^T \ell(y_t^{\$}(n_1) | Y_t^*; \lambda_0, \Lambda_1) \cdots \ell(y_t^{\$}(n_5) | Y_t^*; \lambda_0, \Lambda_1),$$

conditional on the estimated VAR parameters. We estimate the following market prices of risk:

$$\lambda_0 = [-226.31, 216.39, 0, 0, 0]',$$

$$\Lambda_1 = \begin{bmatrix} -1300 & 11290 & 0 & 0 & 0 \\ 127 & 3325 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 25678 & -3355 \\ 0 & 0 & 0 & -33018 & 17709 \end{bmatrix}.$$

The (root mean squared) pricing errors on the five bonds range between 16.2 and 32.6 basis points.

A.2.3 Comparing Utility Differences and Bond Risk Premia

With all parameters in place, we use a forward simulation technique to simulate the expected average real payments streams that the borrower makes on an ARM contract and on an FRM contract, $1/T \sum_{t=1}^T \mathbb{E}_0 [q_{12 \times t}^h]$ for $h \in \{ARM, FRM\}$. For our computations, we set T equal to ten years. We also determine the expected utility over the real consumption stream that a household receives from the ARM and FRM contracts. For contract h and risk aversion parameter γ , we denote this by $U_0^{h,\gamma} = \mathbb{E}_0 \left[\sum_{t=1}^T (C_{12 \times t}^h)^{1-\gamma} / (1-\gamma) \right]$, where $C_{12 \times t}^h$ is defined as before (10-11). To ensure the problem is well defined, we assume a lower bound on consumption of 0.1, to be interpreted as a subsistence level guaranteed by the government. The lower bound is never attained in the simulation exercise. For simplicity we assume labor income is constant and equal to $L = 0.41$. We define the certainty-equivalent consumption by $CEQ^{h,\gamma} = \left\{ U_0^{h,\gamma} (1-\gamma) / T \right\}^{1/(1-\gamma)}$. Notice that for a risk-neutral $\gamma = 0$ investor the certainty-equivalent consumption is $CEQ^{h,0} = L + 1/T \mathbb{E}_0 [1 - 1/\Pi_{12 \times T}] - 1/T \sum_{t=1}^T \mathbb{E}_0 [q_{12 \times t}^h]$. That is, it is inversely related to the expected average real payments.

We determined the certainty-equivalent consumption from 1985:1 to 2006:6. Below we regress the difference in the certainty-equivalent consumption between the two contracts on the real interest and expected inflation premium. The first regression is for a risk-neutral $\gamma = 0$ investor, who only cares about expected consumption (and hence expected payments). The second regression is for a risk-averse $\gamma = 5$ investor, who also cares about the variability of his consumption stream.

$$CEQ_t^{ARM,0} - CEQ_t^{FRM,0} = -0.01 + 1.08\phi_t^y + 0.57\phi_t^x + \epsilon_t \quad (R^2 = 0.93),$$

$$CEQ_t^{ARM,5} - CEQ_t^{FRM,5} = -0.01 + 1.07\phi_t^y + 0.52\phi_t^x + \epsilon_t \quad (R^2 = 0.95).$$

For both the $\gamma = 0$ and $\gamma = 5$ case the difference in certainty-equivalent consumption is rising in the premia. The high R^2 -statistic indicates that this effect explains most of the variation. Notice also that the size of the slope coefficients are very similar for the two cases. All this points to the risk premia being the main determinants of mortgage choice, also for risk-averse investors.

B Valuing the Prepayment Option

We now turn to the valuation the prepayment option. Valuation of this option is based on a numerical dynamic programming algorithm that determines optimal refinancing decisions (see also Pliska (2006)).

B.1 Prepayment Model

The price of the prepayment option is the difference between the rate on a fixed-rate mortgage with prepayment and a (hypothetical) rate on an FRM without prepayment. Let the nominal value to the lender of the former contract be $V_t^p(Y_t^*, r_t)$, and the value of the latter contract $V_t^{np}(Y_t^*, r_t)$, where r_t is the contractual mortgage rate. The time- t values are determined after the time- t interest payment is made and after the prepayment decision has been made. We assume that there are no costs to prepay and that the borrower prepays optimally. Under this assumption, prepayment behavior is fully driven by the dynamics of the term structure of interest rates. We do not consider sub-optimal prepayment behavior and the premium associated with it (see Gabaix, Krishnamurthy, and Vigneron (2006)). As in the main text, the face value of the loan is normalized to \$1. Finally, we assume that there is perfect competition in the mortgage market.

Competition implies that the present value of payments must equal the value of the loan at origination. The implied zero-profit mortgage rate at time t , denoted \hat{r}_t^h with $h \in \{p, np\}$, satisfies

$$V_t^p(Y_t^*, \hat{r}_t^p) = V_t^{np}(Y_t^*, \hat{r}_t^{np}) = 1, \quad \forall t = 0, \dots, T. \quad (39)$$

At maturity T , this condition simply states that the principal is paid back in full. The zero-profit rate will be a function of the state Y_t^* and will be different for the FRM with and without prepayment. The contract values can be solved for recursively, working backwards from time T . The recursion at time t satisfies

$$V_t^{np}(Y_t^*, r_t) = \mathbb{E}_t \left[\exp(m_{t+1}^\$) \left(\exp(r_t - 1) + V_{t+1}^{np}(Y_{t+1}^*, r_t) \right) \right], \quad (40)$$

$$V_t^p(Y_t^*, r_t) = \mathbb{E}_t \left[\exp(m_{t+1}^\$) \left(\exp(r_t - 1) + \min \{1, V_{t+1}^{np}(Y_{t+1}^*, r_t)\} \right) \right]. \quad (41)$$

The minimum operator in (41) reflects the prepayment option: the borrower will prepay at time $t + 1$ whenever the rate on a new loan is lower than the rate on the existing loan. This is the case when $V_{t+1}^{np}(Y_{t+1}^*, r_t) > 1$. The prepayment option premium at time 0 is given by the difference in the zero-profit rate between the two contracts, $\hat{r}_0^p(Y_0^*) - \hat{r}_0^{np}(Y_0^*)$. We recalculate this option premium for each month in the sample 1985:1-2006.6.

B.2 Solution Method

We use a Least Squares Monte-Carlo approach to determine the lender's value function backwards in time at different states of the world. See Longstaff and Schwartz (2001) and Stentoft (2004) for a detailed discussion on this approach. We solve the model with a $T = 10$ year horizon and a time step size of one year. We choose a 100-point, equally-spaced grid for the contract rate r_t . For each starting state Y_0^* we wish to evaluate model-implied variables, we use 1000 (500 plus 500 antithetic) simulated paths for the vector of exogenous state variables Y_t^* . We have 258 separate runs, one for each month from 1985.1-2006.6. We reset the pseudo-random number generator to the same seed for all runs. For all points in time, for all simulated paths, and for all contract rates on the grid, we first solve for the realized value for the lender (for both the case with and without prepayment option). This involves knowing the one-period-ahead expected continuation value for the lender. Next we perform a cross-sectional regression to determine the lender's (expected) value in the current period. As regressors we use a complete set of monomials in the exogenous state variables up to degree 2. To improve the accuracy of the numerical procedure, we performed our calculations under the risk neutral measure. Increasing the number of simulated paths beyond 1000 led to no changes in the variables up the reported precision.

After having determined the lender's expected value we determine the zero-profit rate with cubic interpolation. Finally, the expected payments and utility from consumption can be determined by simulating forward, and again using a cross-sectional regression to determine the refinancing rate in case of the FRM with prepayment.

C Monthly ARM Share Data in the UK

We observe the ARM share data for the UK only at a quarterly frequency in the 1993-2001 period. Since all models are specified at monthly frequency, we want to estimate the data points in between. We employ a Kalman filter together with a specification for the ARM dynamics that is motivated by the US ARM share dynamics. The method is discussed in Hansen and Sargent (2004), Chapter 9.13. The goal is to improve upon linear interpolation by postulating reasonable dynamics for the ARM share dynamics. Linear interpolation may be sub-optimal because the average over month 1 to 3 may have very little to do with the average over month 4 to 6. Linear

interpolation introduces dependencies that are not present in the underlying data. The current approach explicitly incorporates the aggregation.

C.1 Kalman Filter

Denote the monthly fraction of ARM mortgages at time t by x_t . We assume that x_t follows an AR(1) model:

$$x_{t+1} = a + bx_t + \epsilon_{t+1}, \epsilon_{t+1} \sim N(0, \sigma^2). \quad (42)$$

The distributional assumption is required to be able to use the Kalman filter. To justify this time series process, we estimate an AR(1) on the ARM share *in the US*. The R^2 of this AR(1) process is 93% over the sample 1985:1-2005:12. The autoregressive parameter equals $\hat{b} = 0.96$.

For the Kalman filter, we need a state transition equation and an observation equation. Towards this end, we introduce the state vector $y_t = (x_t, x_{t-1}, x_{t-2})'$. The state transition equation for the state vector is given by:

$$y_t = \begin{pmatrix} a \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} b & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} y_{t-1} + \begin{pmatrix} \epsilon_t \\ 0 \\ 0 \end{pmatrix} \quad (43)$$

$$\equiv c + Dy_{t-1} + u_t. \quad (44)$$

Since we observe the average mortgage choice over three months, we lag the system for two additional periods:

$$y_t = (I + D + D^2)c + D^3y_{t-3} + u_t + Du_{t-1} + D^2u_{t-2} \quad (45)$$

$$\equiv e + Fy_{t-3} + \xi_t, \quad (46)$$

which results in the transition equation at a quarterly frequency. We observe the average mortgage choice over a quarter, i.e.:

$$z_t = \iota'_{3 \times 1} y_t / 3, \quad (47)$$

which constitutes the observation equation. Finally, we initialize the Kalman filter assuming that the vector of starting values is drawn from the unconditional distribution.

C.2 Empirical Results

We estimate the coefficients a , b , and σ by maximum likelihood. We find $\hat{a} = 2.9828$ (1.7566), $\hat{b} = 0.9457$ (0.0315), and $\sigma = 5.0174$ (0.4751), where the numbers in parentheses are standard errors, computed from the outer product gradient. The solid line in Figure 13 shows the resulting monthly ARM share for the UK.

To verify that the Kalman filter does work properly, we verify that the average ARM shares over each three month-period coincides with the quarterly data. We also compare the monthly ARM share that arises from the Kalman filter to the one that arises from a Kalman smoother, and find them to be very similar. Finally, we verify that the monthly data obtained from the Kalman filter are close to the actual monthly data over the 2002.1-2006.6 period, for which we have the monthly data. The red circles in Figure 13 correspond to the actual monthly data; the solid line cuts through these data points.

Table 1: Univariate Regression Analysis of the ARM Share for the US.

This table reports slope coefficients, Newey-West t-statistics (12 lags), and R^2 statistics for univariate regressions of the ARM share on a constant and one regressor, reported in the first column. The regressors are the following variables. The τ -year inflation risk premium $\phi_t^x(\tau)$, the τ -year real rate risk premium $\phi_t^y(\tau)$, the conditional variance of expected inflation V_t^x , and the conditional variance in the real rate V_t^y . The τ -year nominal yield is given by $y_t^{\$}(\tau)$. The τ -one-year yield spread is $y_t^{\$(\tau)} - y_t^{\$(12)}$. $y_t^{\$(FRM)} - y_t^{\$(ARM)}$ denotes the difference between the FRM rate $y_t^{\$(FRM)}$, and the ARM rate $y_t^{\$(ARM)}$. The regressor is lagged by one period, relative to the ARM share. The left panel is for the longest sample from 1985.1-2006.6; the right panel is for the sample over which we have data for both the 5-year and the 10-year treasury inflation-protected security: 1997.7-2006.6. All variables have been multiplied by 100.

	1985.1-2006.6			1997.7-2006.6		
	slope	t-stat	R^2	slope	t-stat	R^2
1. $\phi_t^x(60)$	16.77	5.91	44.12	12.76	8.49	63.52
2. $\phi_t^x(120)$	17.19	4.91	44.59	14.46	6.77	56.24
3. $\phi_t^y(60)$	3.68	0.87	2.19	-8.88	-3.53	28.51
4. $\phi_t^y(120)$	6.88	2.20	12.21	-8.59	-3.01	25.49
5. $\phi_t^{\$(60)} \equiv \phi_t^y(60) + \phi_t^x(60)$	9.99	4.16	32.21	6.47	1.63	11.45
6. $\phi_t^{\$(120)} \equiv \phi_t^y(120) + \phi_t^x(120)$	8.04	3.91	35.13	3.62	0.76	3.33
7. $y_t^{\$(60)} - y_t^{\$(12)}$	0.60	0.21	0.11	0.80	0.35	0.64
8. $y_t^{\$(120)} - y_t^{\$(12)}$	-0.59	-0.32	0.23	0.34	0.22	0.27
9. $y_t^{\$(60)}$	4.46	3.76	37.76	-0.58	-0.29	0.57
10. $y_t^{\$(120)}$	4.98	3.85	39.26	-1.27	-0.41	1.25
11. $y_t^{\$(FRM)} - y_t^{\$(ARM)}$	16.03	3.17	35.31	0.36	0.05	0.01
12. $y_t^{\$(FRM)} - y_t^{\$(ARM)}$ orth.	18.25	3.85	41.23	0.21	0.03	0.00
13. $y_t^{\$(FRM)}$	4.28	3.71	32.87	-2.47	-0.82	5.45
14. V_t^y	-28.91	-2.92	22.84	3.31	0.34	0.76
15. V_t^x	169.20	2.87	25.67	-17.19	-0.25	0.18

Table 2: Multivariate Regression Analysis for ARM Share in US.

This table reports slope coefficients, Newey-West t-statistics, and R^2 statistics for multiple linear regressions of the ARM share on a constant and the variables listed in the first column. The regressors are lagged by one period, relative to the ARM share. The right-hand side variables are demeaned so that the constant gives the average ARM share in the sample. The regressors are the ten-year inflation risk premium $\phi_t^x(120)$, the ten-year real rate risk premium $\phi_t^y(120)$, the conditional variance of expected inflation V_t^x , and the conditional variance in the real rate V_t^y . The ten-one-year yield spread, $y_t^{\$}(120) - y_t^{\$}(12)$, and the utility difference between an ARM contract and an FRM contract with prepayment option, $CEQ^{ARM} - CEQ^{FRM,p}$, are orthogonalized to the other four term structure variables. We run an auxiliary regression of these variables on the first four variables and include the regression residual as a fifth explanatory variable. The left panel is for the longest sample from 1985.1-2006.5; the right panel is for the sample over which we have reliable real yield data: 1997.7-2006.6. Newey-West t-statistics (12 lags) are in parentheses.

Regressors	1985.1-2006.6				1997.7-2006.6			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
constant	28.66 (14.72)	28.66 (16.39)	28.66 (16.53)	28.66 (16.52)	21.61 (14.24)	21.61 (16.10)	21.61 (16.08)	21.61 (16.07)
$\phi_t^x(120)$	15.97 (4.81)	14.27 (5.15)	14.27 (5.14)	14.27 (5.14)	13.38 (5.86)	15.13 (5.09)	15.44 (3.10)	15.18 (5.11)
$\phi_t^y(120)$	2.72 (1.15)	-0.96 (-0.36)	-0.96 (-0.34)	-0.96 (-0.34)	-1.62 (-0.53)	-1.32 (-0.46)	-0.86 (-0.11)	-1.08 (-0.27)
V_t^x		43.89 (0.73)	43.89 (0.76)	43.89 (0.77)		32.50 (0.62)	36.52 (0.40)	34.31 (0.56)
V_t^y		-15.11 (-2.08)	-15.11 (-2.03)	-15.11 (-2.06)		-6.19 (-0.89)	-5.93 (-0.64)	-5.99 (-0.73)
$y_t^{\$}(120) - y_t^{\$}(12)$			3.89 (0.65)				-0.43 (-0.08)	
$CEQ^{ARM} - CEQ^{FRM,p}$				5.82 (0.60)				-1.08 (-0.11)
R^2	46.27	53.93	54.30	54.25	56.82	60.98	60.99	61.00

Table 3: Understanding the Differences Between the UK and the US.

This table reports slope coefficients, Newey-West t-statistics, and R^2 statistics for multiple linear regressions of $\kappa_t(\rho; 120)$ on a constant and the variables listed in the first column. The regressors are the ten-year inflation risk premium $\phi_t^x(120)$, the ten-year real rate risk premium $\phi_t^y(120)$, and the ten-year minus one-year yield spread, orthogonalized to the other two term structure variables. Newey-West t-statistics (12 lags) are in parentheses. The ‘★’ indicates the highest correlation between the rule-of-thumb measure and the ARM share.

ρ	US: 1989.12-2006.6					US: 1997.7-2006.6				
	12	24	36	48★	60	12	24	36★	48	60
Constant	0.03 (10.66)	0.04 (10.58)	0.04 (10.69)	0.03 (10.31)	0.02 (7.82)	0.02 (4.57)	0.03 (7.42)	0.03 (14.53)	0.03 (11.74)	0.02 (10.06)
$\phi_t^x(120)$	1.39 (6.90)	1.48 (6.40)	1.29 (5.76)	1.06 (4.98)	0.79 (4.75)	2.37 (12.66)	2.11 (8.35)	1.84 (7.42)	1.65 (10.09)	1.25 (11.91)
$\phi_t^y(120)$	-0.57 (-4.18)	-1.17 (-6.73)	-1.10 (-6.69)	-0.88 (-6.09)	-0.53 (-4.38)	0.47 (1.84)	-0.44 (-1.48)	-0.63 (-5.05)	-0.60 (-5.75)	-0.43 (-4.52)
$y_t^s(120) - y_t^s(12)$	0.74 (16.06)	0.32 (3.92)	-0.06 (-0.68)	-0.29 (-3.64)	-0.36 (-4.73)	0.70 (10.97)	0.30 (3.17)	-0.03 (-0.32)	-0.22 (-3.00)	-0.31 (-4.32)
R^2	84.82	75.35	65.19	59.38	54.48	88.34	83.50	83.28	85.05	84.31
ρ	UK: 1993.1-2006.6					UK: 2002.1-2006.6				
	12	24	36	48★	60	12★	24	36	48	60
Constant	-0.01 (-4.18)	-0.01 (-2.86)	-0.01 (-2.28)	-0.01 (-1.84)	-0.01 (-1.84)	-0.01 (-4.84)	-0.01 (-3.70)	-0.01 (-4.22)	-0.01 (-5.52)	-0.01 (-7.04)
$\phi_t^x(120)$	1.40 (5.78)	1.11 (3.16)	0.73 (1.64)	0.33 (0.67)	-0.04 (-0.10)	2.78 (5.78)	2.65 (3.67)	2.36 (3.72)	1.74 (4.38)	1.95 (4.66)
$\phi_t^y(120)$	1.15 (4.23)	1.21 (2.58)	1.06 (1.79)	0.87 (1.32)	0.54 (0.91)	1.85 (6.85)	1.10 (2.13)	0.24 (0.55)	-0.17 (-0.56)	-0.22 (-0.77)
$y_t^s(120) - y_t^s(12)$	0.41 (2.33)	0.00 (0.00)	-0.05 (-0.16)	-0.21 (-0.58)	-0.27 (-0.85)	0.66 (4.96)	0.17 (0.50)	-0.40 (-1.49)	-0.67 (-2.95)	-0.51 (-2.79)
R^2	79.21	55.81	31.07	12.91	5.14	86.82	46.5	56.36	72.16	78.54

Table 4: Alternative ARM Share Measures in US.

This table reports slope coefficients, Newey-West t-statistics (12 lags) in parentheses, and R^2 statistics for multiple linear regressions of the ARM share on a constant and the variables listed in the first column. The right-hand side variables are demeaned so that the constant gives the average ARM share in the sample. The regressors are the ten-year inflation risk premium $\phi_t^x(120)$, the ten-year real rate risk premium $\phi_t^y(120)$, the conditional variance of expected inflation V_t^x , and the conditional variance in the real rate V_t^y . The first column is our benchmark measure; the other three ARM share measures are defined in Section 4.1. The regressors are lagged by one period, relative to the ARM share. The sample is 1997.7-2006.6.

RHS variables	ARM^1	ARM^2	ARM^3	ARM^4
constant	21.61 (16.10)	18.10 (16.25)	10.92 (15.95)	21.54 (15.76)
$\phi_t^x(120)$	15.13 (5.09)	12.88 (5.27)	6.18 (4.50)	14.98 (5.12)
$\phi_t^y(120)$	-1.32 (-0.46)	-0.39 (-0.16)	0.91 (0.59)	-2.93 (-0.97)
V_t^x	32.50 (0.62)	23.45 (0.56)	5.72 (0.23)	36.86 (0.61)
V_t^y	-6.19 (-0.89)	-3.55 (-0.64)	-5.40 (-2.02)	-4.91 (-0.62)
R^2	60.98	60.51	41.59	63.98

Table 5: Multivariate Regression Analysis for ARM Share in UK.

This table reports slope coefficients, Newey-West t-statistics, and R^2 statistics for multiple linear regressions of the ARM share on a constant and the variables listed in the first column. The regressors are lagged by one period, relative to the ARM share, and are demeaned. The regressors are the ten-year inflation risk premium $\phi_t^x(120)$, the ten-year real rate risk premium $\phi_t^y(120)$, the conditional variance of expected inflation V_t^x , and the conditional variance in the real rate V_t^y . We also consider the ten-year minus one-year yield spread, orthogonalized to the other four term structure variables. The left panel is for the longest sample from 1993.1-2006.6. It uses the monthly ARM share obtained through the Kalman filter, see Appendix C. The right panel is for the sample over which we have monthly data for the ARM share: 2002.1-2006.6. Newey-West t-statistics (12 lags) are in parentheses.

Regressors	1993.1-2006.6			2002.1-2006.6		
	(1)	(2)	(3)	(5)	(6)	(7)
constant	57.53 (19.92)	57.53 (20.27)	57.53 (20.21)	52.52 (27.89)	52.52 (46.25)	52.52 (49.55)
$\phi_t^x(120)$	0.09 (0.40)	-1.77 (-0.30)	-1.77 (-0.29)	26.88 (2.87)	34.38 (2.74)	34.38 (2.56)
$\phi_t^y(120)$	20.47 (2.42)	18.51 (2.04)	18.51 (2.05)	51.57 (8.49)	69.10 (8.39)	69.10 (9.59)
V_t^x		-24.45 (-0.44)	-24.45 (-0.44)		254.89 (3.00)	254.89 (3.48)
V_t^y		27.66 (0.76)	27.66 (0.76)		-24.71 (-0.56)	-24.71 (-0.49)
$y_t^{\$}(120) - y_t^{\$}(12)$			6.87 (0.19)			68.63 (4.01)
R^2	23.05	24.45	24.53	71.69	75.33	76.75

Figure 1: The Share of Adjustable Rate Mortgages in the US.

The figure plots the fraction of all newly originated mortgages that are of the adjustable-rate type. The complementary fraction are fixed-rate mortgages. The data are from the Federal Housing Financing Board and are based on the Monthly Interest Rate Survey sent out to mortgage lenders. It covers all property types: newly constructed homes, and existing homes. ARMs include hybrid mortgages, which may have an initial fixed-interest rate payment period of up to ten years.

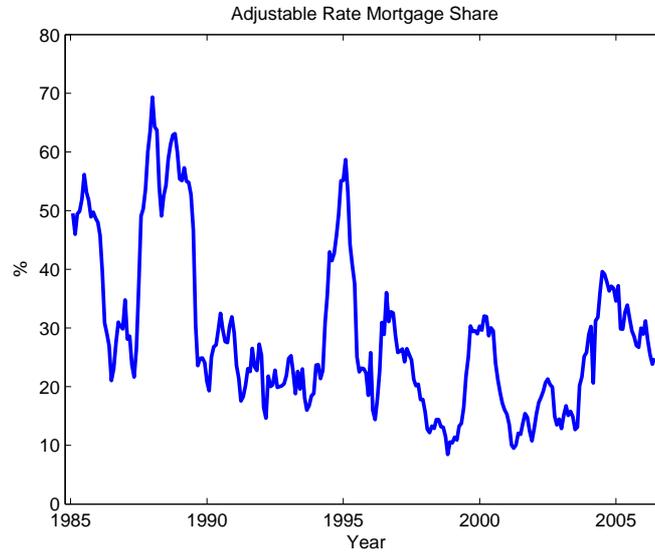


Figure 2: The Inflation Risk Premium and the ARM Share in the US.

The figure plots the fraction of all mortgages that are of the adjustable-rate type against the left axis, and the inflation risk premium 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 nominal and real 5-year bond yields are from McCulloch and start in January 1997. The inflation expectation is the median long-term (10-year) inflation forecast from the Survey of Professional Forecasters (SPF).

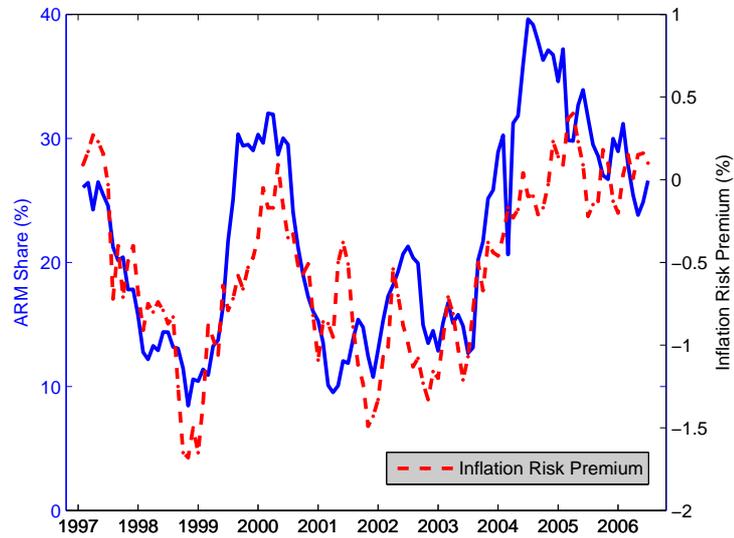


Figure 3: VAR Estimation for the US.

The figure plots long-term risk premia. The 5-year risk premia for expected inflation risk and for real rate risk are plotted in the left panel. The 10-year risk premia, formed by subtracting the 10-year real rate yield data from the VAR-implied 10-year real rate, are plotted in the right panel.

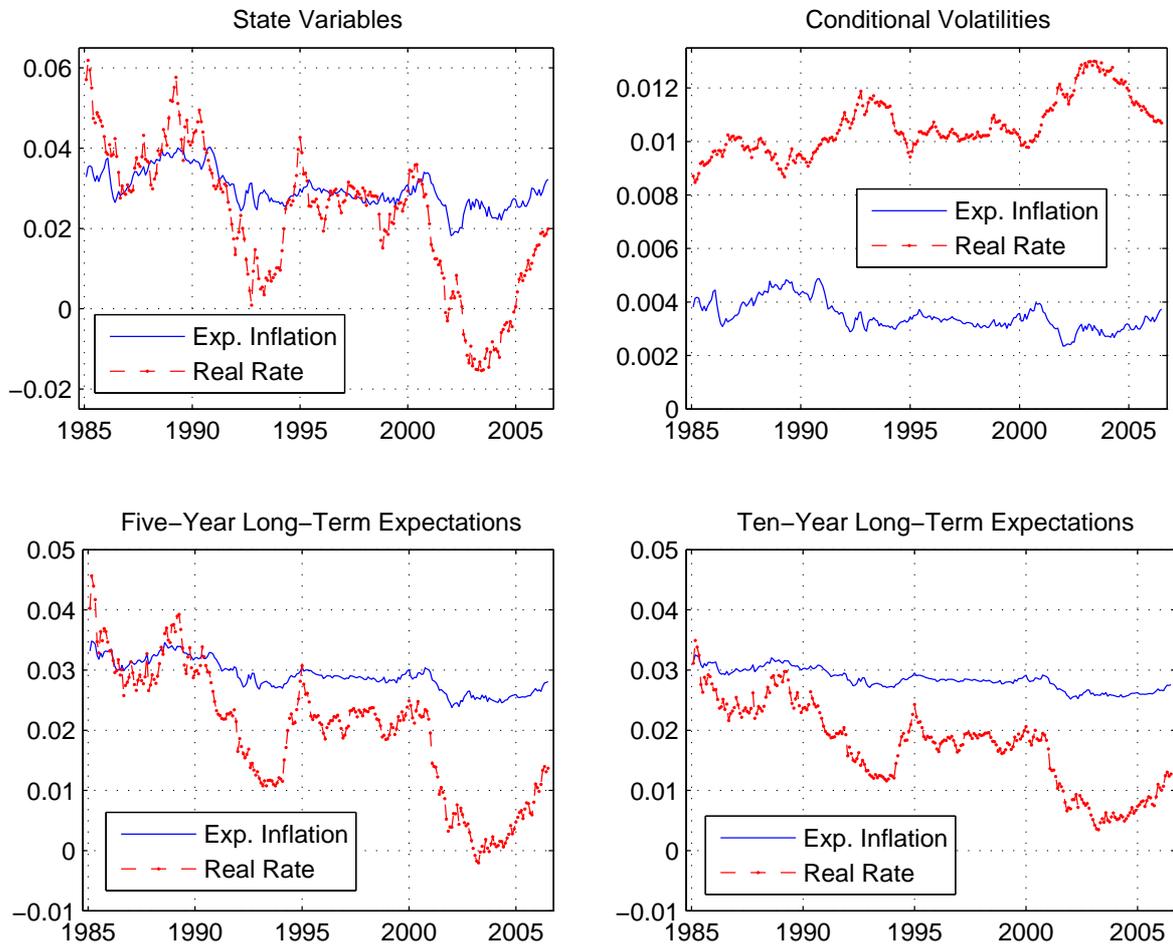


Figure 4: Inflation and Real Rate Risk Premia in the US.

The figure plots long-term risk premia. The five-year risk premia for expected inflation risk and for real rate risk are plotted in the left panel. The ten-year risk premia, formed by subtracting the ten-year real rate yield data from the VAR-implied five-year real rate, are plotted in the right panel.

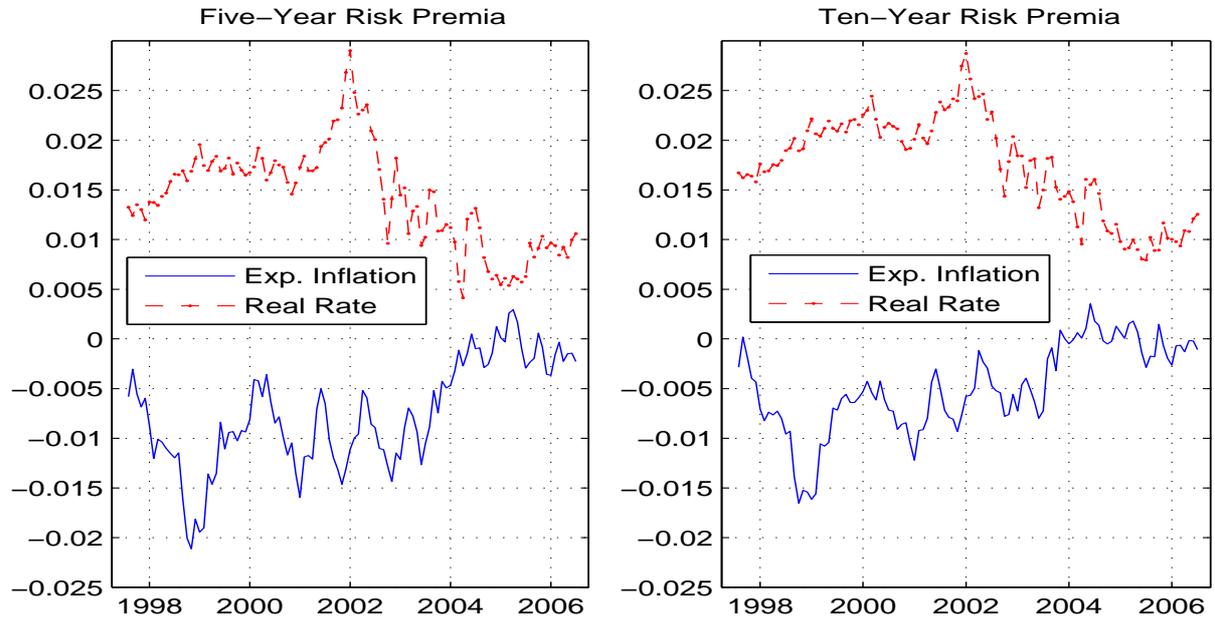


Figure 5: Extending the Sample of Inflation and Real Rate Risk Premia for the US.

The figure plots ten-year risk premia. The solid black line is the ten-year nominal risk premium. It is computed as the difference between the observed nominal ten-year yield $\{y_t^{\$}(120)\}$ and the average expected nominal ten-year yield. These expectations are readily obtained from the VAR for the entire 1985.1-2006.6 period. The circled purple line is the projected ten-year real rate risk premium, $\phi_t^y(120)$, formed as the product of the state variables z and the regression coefficients in Equation (22). These loadings are estimated on the 1997.7-2006.6 sub-sample. The circled turquoise line is the inflation risk premium, $\phi_t^x(120)$. It is formed as the difference between the nominal risk premium and the inflation risk premium according to Equation (3). For the 1997.7-2006.6 sample, we overlay the actual risk premia (reported earlier in Figure 4) on the projected risk premia.

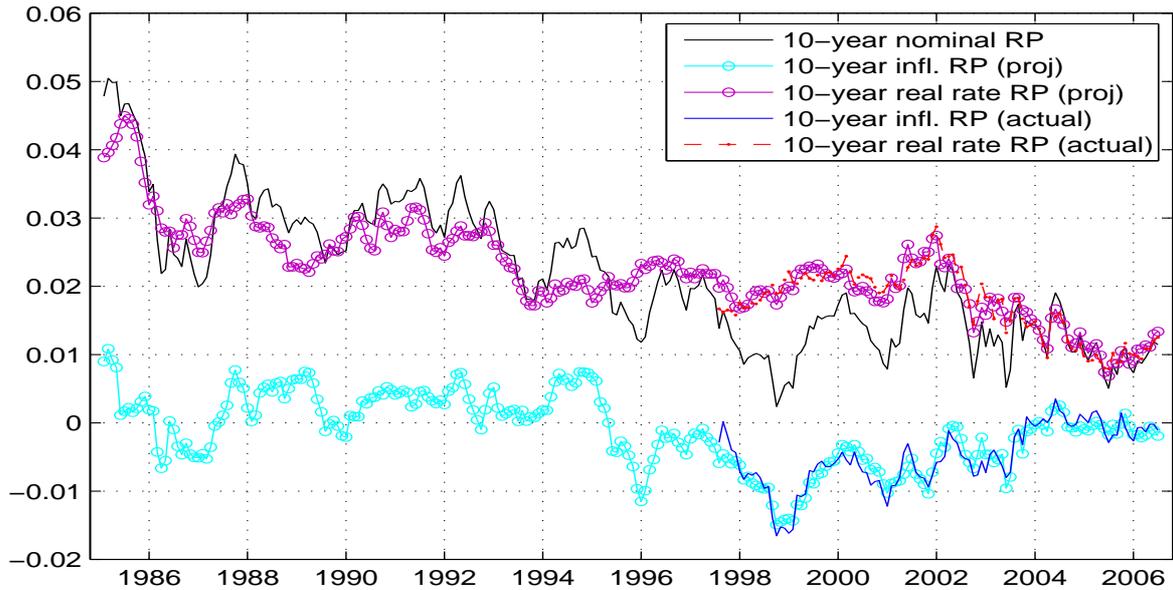


Figure 6: Explaining Bond Risk Premia in the US.

The figure plots the R-squared of a regression of the nominal bond risk premium $\phi_t^S(\tau)$ on a constant and both the 10-year real rate premium ($\phi_t^y(120)$) and inflation risk premium ($\phi_t^x(120)$), see (23). The maturity τ ranges from two to 10 years and is depicted on the horizontal axis. The regressions are estimated over the period 1985.1-2006.6.

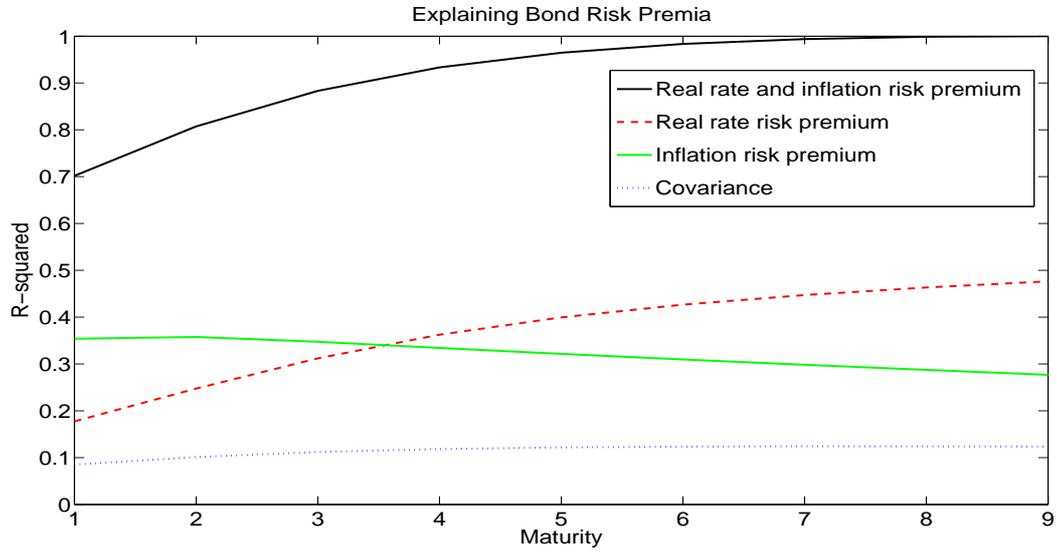


Figure 7: The ARM Share for the US in More Detail.

The figure plots the fraction of all newly originated mortgages that are of the adjustable-rate type between 1992.1 and 2006.6. The first three series are from the Federal Housing Financing Board and are based on the Monthly Interest Rate Survey sent out to mortgage lenders. The first series includes all hybrid mortgages. The second series excludes hybrids with an initial fixed-rate period of more than five years, and the third series excludes hybrids with an initial fixed-rate period of more than three years. The last series is from Freddie Mac and is based on the Primary Mortgage Market Survey. Like the first measure, it contains all ARM originations.

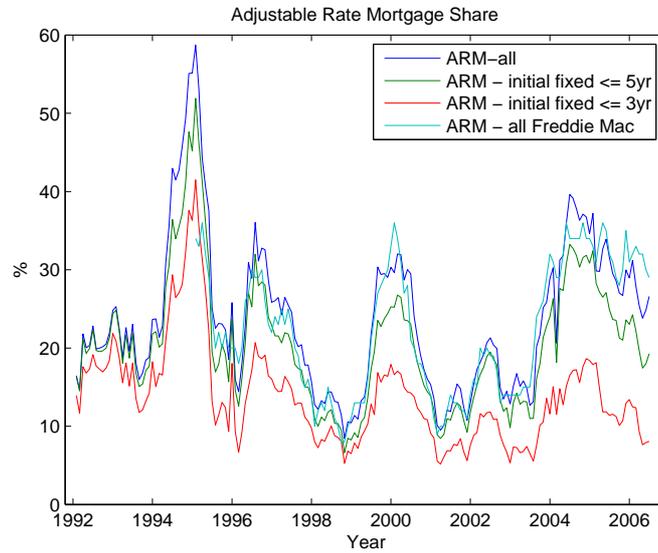


Figure 8: Price Sensitivity to Changes in the Real Rate and Expected Inflation for the US.

The figure plots the price sensitivities of the FRM contract with and without prepayment to the real interest rate y (top panel) and expected inflation x (bottom panel). The mortgage values are determined within the model of Appendix A. The analogous fixed-income securities are a regular bond (FRM without prepayment) and a callable bond (FRM with prepayment).

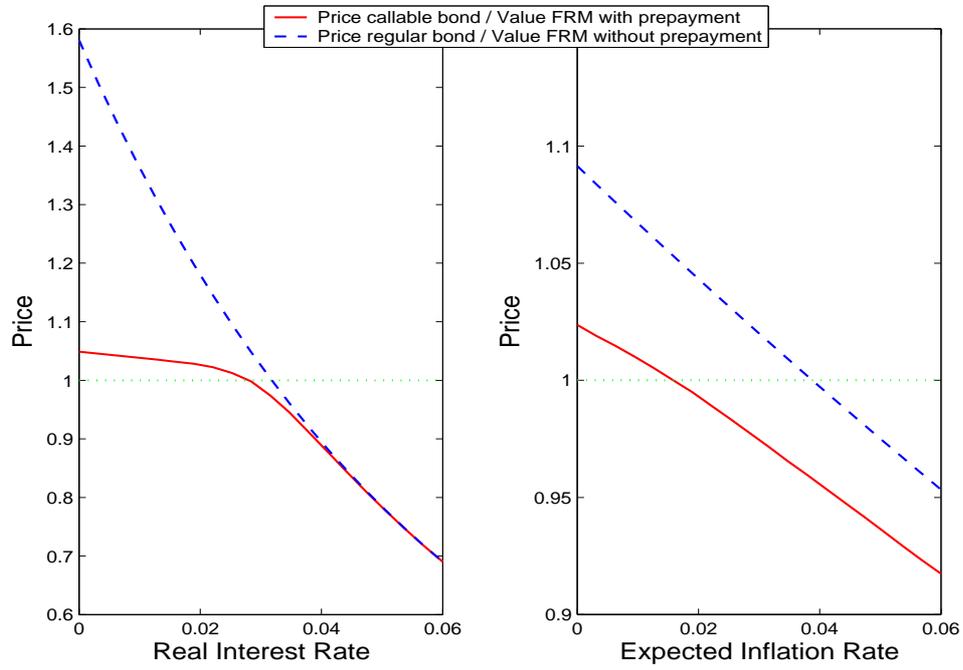


Figure 9: Alternative Model of Bond Risk Premia and the ARM Share.

The figure plots the time series of bond risk premia (red dashed line) following from the model in Section 5 for the US. Bond risk premia are computed as the difference between the 10-year yield and the 3-year moving average of short rates, i.e., $\kappa_t(36;120)$. The left axis displays the magnitude of this measure of bond risk premia, scaled by a factor 100. The blue line corresponds to the ARM share in the US, and its values are depicted on the right axis. The time series runs from 1987.12 to 2006.6.

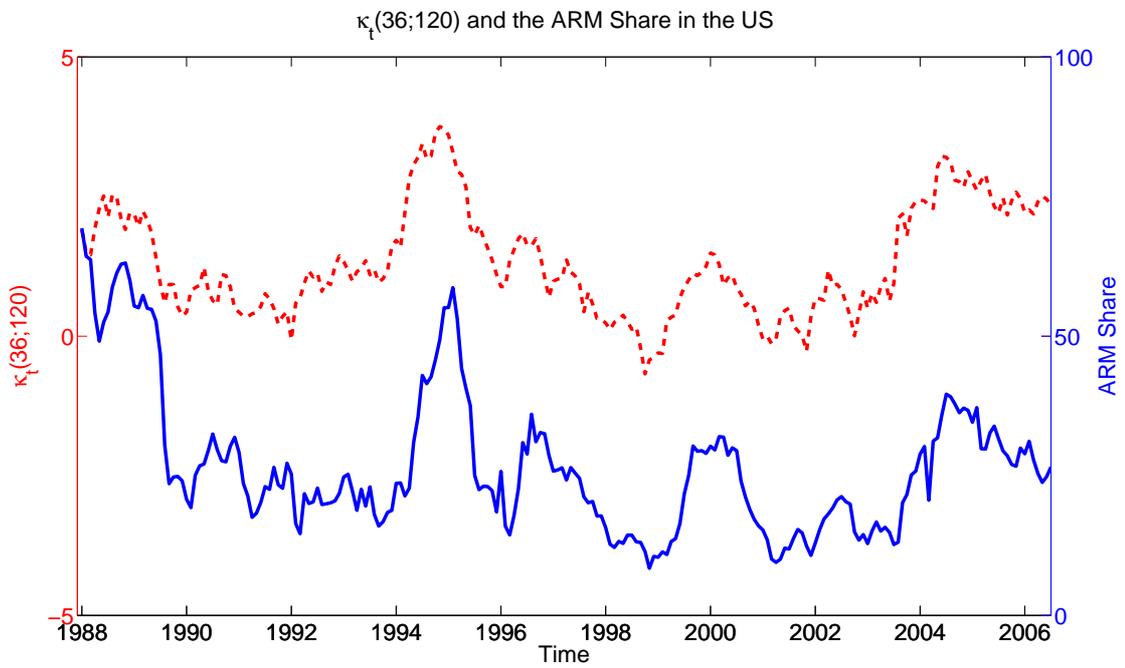


Figure 10: Alternative Model of Bond Risk Premia and the ARM Share.

The figure plots the correlation of bond risk premia (red dashed line) following from the model in Section 5 with the ARM share. Bond risk premia are computed as the difference between the 10-year yield and the ρ -month moving average of short rates, i.e., $\kappa_t(\rho; 120)$. The blue bars correspond to $\rho = 24, 36, 48, 60$. The red line corresponds to the correlation between the yield spread (i.e., $\rho = 1$) and the ARM share. The red dashed line depicts the correlation between the 10-year yield and the ARM share (i.e., $\rho = \infty$). The left panel uses treasury yields as yield variable, and the right panel the ARM and FRM mortgages rates. The time series runs from 1985.1 to 2006.6. Since we require a ρ -month history to compute the average of short rates, the time series effectively used to compute the correlations are reduced by the number of months to compute this average.

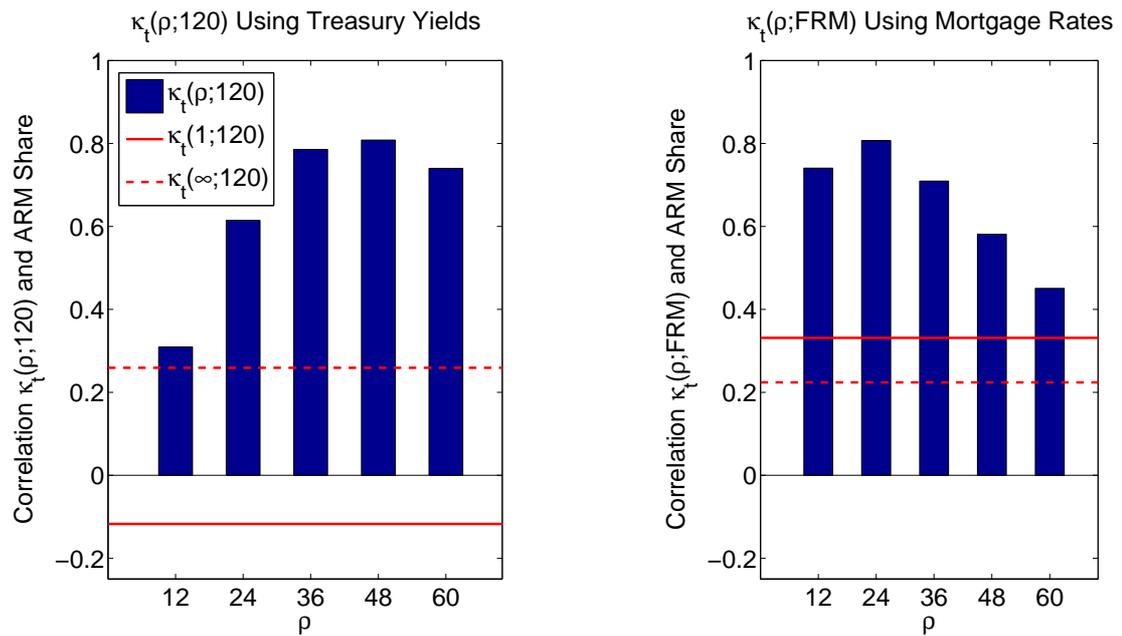


Figure 11: VAR Estimation for the UK.

The figure plots long-term risk premia. The 5-year risk premia for expected inflation risk and for real rate risk are plotted in the left panel. The 10-year risk premia, formed by subtracting the 10-year real rate yield data from the VAR-implied 10-year real rate.

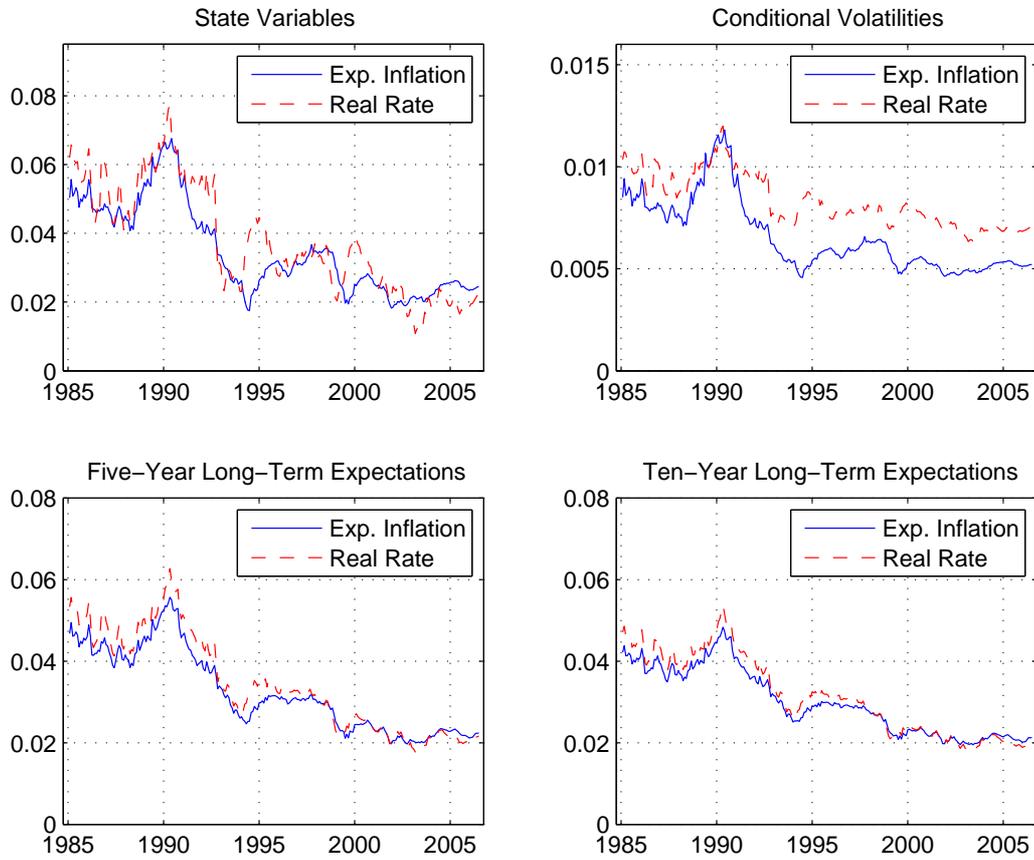


Figure 12: Inflation and Real Rate Risk Premia in the UK.

The figure plots long-term risk premia. The 5-year risk premia for expected inflation risk and for real rate risk are plotted in the left panel. The 10-year risk premia, formed by subtracting the 10-year real rate yield data from the VAR-implied 10-year real rate.

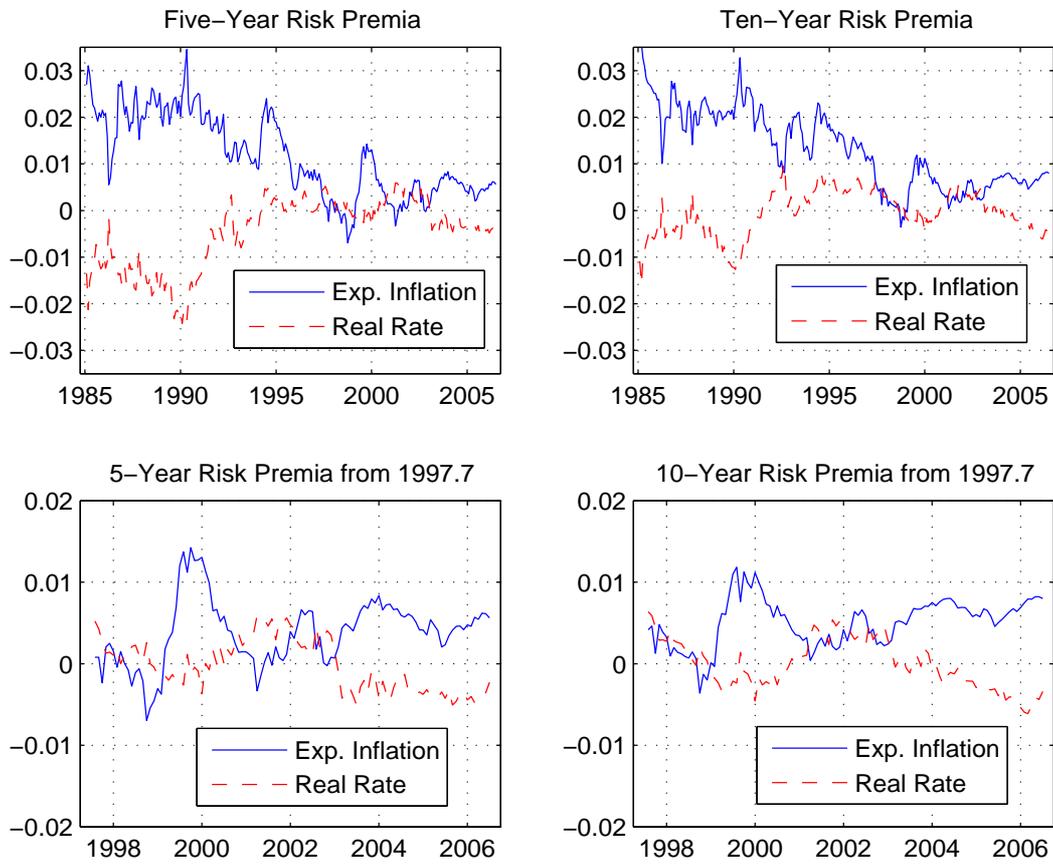


Figure 13: The Share of Adjustable Rate Mortgages in the UK.

The figure plots the fraction of all newly originated mortgages that are of the adjustable-rate type. The complementary fraction are fixed-rate mortgages. The data are from the Council of Mortgage Lenders (sheet ML5) and are based on the Survey of Mortgage Lenders before April 2005 and based on Product Sales Data reported to the CML after April 2005. It covers both house purchases and remortgages. Adjustable rate mortgages are the sum of standard variable rate contracts (SVR), discounted (variable rate) mortgages and trackers. Fixed rate mortgages are the sum of fixed contracts and capped contracts. The red-circled line plots the monthly data, which are only available from January 2002 onwards. The solid blue line is a monthly time-series, which we generate from quarterly temporally aggregated data that start in 1993.I. See Appendix C for details on this procedure.

