Bond Positions, Expectations, And The Yield Curve*

Monika Piazzesi
Chicago GSB, FRB Minneapolis & NBER

Martin Schneider
NYU, FRB Minneapolis & NBER

October 2007

Abstract

This paper implements a structural model of the yield curve with data on nominal positions and survey forecasts. Bond prices are characterized in terms of investors' current portfolio holdings as well as their subjective beliefs about future bond payoffs. Risk premia measured by an econometrician vary because of changes in investors' subjective risk premia, identified from portfolios and subjective beliefs, but also because subjective beliefs differ from those of the econometrician. The main result is that investors' systematic forecast errors are an important source of business-cycle variation in measured risk premia. By contrast, subjective risk premia move less and more slowly over time.

*Preliminary and incomplete. Comments welcome! We thank Ken Froot for sharing the Goldsmith-Nagan survey data with us, and the NSF for financial support to purchase the Bluechip survey data. We also thank Andy Atkeson, Dimitri Vayanos, seminar participants at the San Francisco Federal Reserve Bank and conference participants at the Atlanta Federal Reserve Bank, UCLA and in Vienna. Email addresses: piazzesi@uchicago.edu, martin.schneider@nyu.edu.

I Introduction

There is a large literature that tries to understand the dynamics of the yield curve through the behavior of optimizing investors. For example, consumption-based asset pricing models start from the fact that, when investors optimize, bond prices can be expressed in terms of investors' beliefs about future asset values and consumption. Model-implied bond prices then consist of expected discounted future bond payoffs, minus a risk premium that depends on the covariance of future bond returns with consumption. Empirical implementation of this idea requires the modeler to postulate a probability distribution that represents investor beliefs. In practice, this distribution is typically supplied by a stochastic process of asset values and consumption, often constrained by cross- equation restrictions implied by the economic model.

This paper proposes an alternative approach to implement a structural model of the yield curve. We start from the fact that, when investors optimize, the prices of bonds can be expressed in terms of the distribution of their future payoffs together with current realized investor asset positions. We construct measures of both objects, using survey expectations of yields as well as data on outstanding nominal assets in the US economy. Model-implied bond prices again depend on expected payoffs minus risk premia, where the latter are identified from beliefs and portfolio holdings.

Our model allows for three sources of time variation in expected excess bond returns measured by an econometrician. First, there is time variation in the subjective volatility of investors' consumption growth (or their return on wealth), which we identify from portfolio data. Second, there is time variation in the conditional covariance of bond returns with investors' continuation utility (or "investment opportunities"), a property of investors' subjective belief. Finally, risk premia measured by an econometrician can vary over time if investors' subjective beliefs do not agree with those of the econometrician. The main result of this paper is that, at least in the model we consider, this third source of time variation in measured expected excess returns is the most important one.

We consider a group of investors who share the same Epstein-Zin preferences and hold the same subjective beliefs about future asset payoffs. Our analysis proceeds in four steps. First, we estimate investors' beliefs about future asset values, combining statistical analysis and survey forecast evidence. Second, we produce a measure of investor asset positions, using quantity data from the Flow of Funds accounts and the CRSP Treasury database. Third, we work out investors' savings and portfolio choice problem given beliefs to derive asset demand, for every period in our sample. Finally, we find equilibrium asset prices by setting asset demand equal to investors' observed asset holdings in the data. We thus arrive at a sequence of model-implied bond prices of the same length as the sample. The model is "successful" if the sequence of model-implied prices is close to actual prices.

In the first step, we document properties of survey forecasts of interest rates over the last four decades. Here we combine evidence from the Blue Chip survey, available since 1982, and its precursor, the Goldsmith-Nagan survey, available since 1970. We compare expected excess returns on bonds implied by predictability regressions that are common in the literature to expected excess returns on bonds perceived by the median survey investor. The main stylized fact from this exercise is that subjective expected excess returns are smaller on average and less countercyclical than conventional measures of expected excess returns. The reason is that predictability regressions do a good job forecasting interest rate drops in recessions, whereas survey forecasters do not. During and after recessions, conventionally measured expected returns thus appear much higher than survey expected excess returns.

To construct investor beliefs that can serve as an input to our asset pricing model, we first estimate a time series model of macro variables and interest rates that nests an affine term structure model. We then assume that investors' subjective belief has the same basic structure and use survey forecast data to estimate the parameters of the Radon-Nikodym derivative of investors' belief with respect to our own "objective" model. We thus obtain a subjective time series model that nests a subjective affine term structure model. The subjective term-structure model has smaller and less variable market prices of risk than its objective counterpart, and does a good job capturing differences in the cyclical properties of subjective and objective expected excess returns.

Since there is a large variety of nominal instruments, an investor's "bond position" is in principle a high-dimensional object. To address this issue, the second step of the analysis uses the subjective term-structure model to replicate positions in many common nominal instruments by portfolios that consist of only three zero coupon bonds. Three bonds work because a two-factor model does a good job describing quarterly movements in the nominal term structure. We use the replicating portfolios to illustrate properties of bonds outstanding in the US credit market. One interesting fact is that the relative supply of longer bonds declined before 1980, as interest rate spreads were falling, but saw a dramatic increase in the 1980s, a time when spreads were extraordinarily high.

We illustrate our asset pricing approach by presenting an exercise where investors are assumed to be "rentiers", that is, they hold only bonds. Rentiers' bond portfolios are taken to be proportional to those of the aggregate US household sector, and we choose preference parameters to best match the mean yield curve. This leads us to consider relatively patient investors with low risk aversion. Our model then allows a decomposition of "objective" risk premia as measured under the objective statistical model of yields into their three sources of time variation. We find that subjective risk premia are small and vary only at low frequencies. This is because both measured bond positions, and the hedging demand for long bonds under investors' subjective belief move slowly over time. In contrast, the difference in subjective and objective forecasts is a source of large time variation in risk premia at business cycle frequencies.

This paper shares the goal of the consumption-based asset pricing literature: to find a model of investor behavior that helps us understand why some bonds have higher returns than others. Indeed, the preferences explored here are the same as in Piazzesi and Schneider (2006; PS). The present paper differs from PS as well as other studies in that it does not claim to directly measure the dynamics of planned consumption. Instead, it only measures beliefs about future asset payoffs. Quantity data enter only in the form of realized asset positions and realized consumption. A second difference is that PS, again following a large literature, restrict beliefs about consumption, inflation, and yields by assuming that agents use the structure of the model itself when forming beliefs.

These differences in approach are minor if consumption is observable and all beliefs are derived from a single stationary probability distribution which also governs the data. Indeed, if the stationary rational-expectations version of the model studied in PS is correct, then the benchmark household sector exercise of this paper will find that model-implied yields have the same properties as yields in the data.[1] A mismatch of model-implied and actual yields would thus indicate that the

[1] Indeed, suppose that the stationary rational expectations version of PS is the "truth" that generated the data. Consider now our benchmark household sector exercise, with beliefs defined as conditionals from a stationary statistical model of asset values. With a long enough sample, returns under the estimated model will be the same as returns

stationary rational expectations version of PS does not fit the data. The approach of this paper can therefore be viewed as an alternative strategy to evaluate stationary rational expectations models.

More generally, the approach of this paper can be used to evaluate models when the rational expectations assumption is not imposed. Our approach has three properties that are particularly helpful in this case. First, it does not assume that agents use the structure of the model to form beliefs. It thus allows for the possibility that a model with recursive preferences is a good model of the risk-return tradeoff faced by investors, but not a good model of belief formation. Second, since beliefs about payoffs are taken as an input to the exercises, we can make use of survey forecasts to discipline the model. Third, since beliefs about planned consumption are not needed, we can derive some pricing implications even when consumption is not observable, as long as we can observe portfolio positions. We make use of this property when we set up the "rentier" exercise.

Related Literature – to be written.

The rest of the paper is structured as follows. Section II introduces the model. Section III describes properties of subjective beliefs measured from surveys. Section IV describes how we model subjective beliefs and estimate them with survey data. Section V explains how we replicate nominal position s by simple portfolios. Section VI reports the asset pricing results.

II Model

A large number of identical investors live forever. Their preferences over consumption plans are represented by Epstein-Zin utility with unitary intertemporal elasticity of substitution. The utility u_t of a consumption plan $(C_\tau)_{\tau=t}^{\infty}$ solves

$$u_t = (1-\beta)\log C_t + \beta \log E_t \left[e^{(1-\gamma)u_{t+1}} \right]^{\frac{1}{1-\gamma}} . \tag{1}$$

under the true PS model. The exercise now derives a sequence of prices from a sequence of optimality conditions. By construction, every such condition also holds under the PS model, where agents solve the same portfolio choice problem. It follows that the statistical properties of the model-implied prices are the same as under the true PS model. Moreover, the distribution of planned consumption would be the same as the distribution of consumption in the data.

Investors' ranking of certain consumption streams is thus given by discounted logarithmic utility. At the same time, their attitude towards atemporal lotteries is determined by the risk aversion coefficient γ. We focus below on the case $\gamma > 1$, which implies an aversion to persistent risks (as discussed in Piazzesi and Schneider 2006).

Investors have access to two types of assets. *Bonds* are nominal instruments that promise dollar-denominated payoffs in the future. In particular, there is a one period bond – from now on, the *short* bond – that pays off one dollar at date $t+1$; it trades at date t at a price e^{-i_t}. Its real payoff is $e^{-\pi_{t+1}}$, where π_t is (log) inflation.[2] In some of our exercises, we also allow investors to trade a *residual asset*, which stands in for all assets other than bonds. The log real return from date t to date $t+1$ is r^{res}_{t+1}, so that its excess return over the short bond is $x^{res}_{t+1} = r^{res}_{t+1} - i_t - \pi_{t+1}$.

In addition to short bonds, investors can buy N other zero-coupon bonds, which—together with the short bond—we refer to as *spanning* bonds. We collect the log nominal prices of these bonds at date t in a vector $\hat{p}_t$, and we collect their log nominal payoffs[3] at date $t+1$ in a vector $\hat{p}^{+1}_{t+1}$. The log excess returns over the short bond from date t to date $t+1$ can thus be written as $\hat{x}_{t+1} = \hat{p}^{+1}_{t+1} - \hat{p}_t - i_t$. Below, the number of long bonds N will correspond to the number of factors in our term structure model: our empirical implementation will use the fact that, under an N-factor model, $N+1$ bonds are sufficient to span the payoffs on all bonds.

Investors start a trading period t with initial wealth $\overline{W}_t$. They decide how to split this initial wealth into consumption as well as investment in the $N+2$ assets. We denote by α^{res}_t the portfolio weight on the residual asset (that is, the fraction of savings invested in that asset), and we collect the portfolio weights on all bonds other than the short bond in a vector $\hat{\alpha}_t$. The household's

[2] This is a simple way to capture that the short (1 period) bond is denominated in dollars. To see why, consider a nominal bond which costs $P^{(1)}_t$ dollars today and pays of \$1 tomorrow, or $1/p^c_{t+1}$ units of numeraire consumption. Now consider a portfolio of p^c_t nominal bonds. The price of the portfolio is $P^{(1)}_t$ units of numeraire and its payoff is $p^c_t/p^c_{t+1} = 1/\pi_{t+1}$ units of numeraire tomorrow. The model thus determines the price $P^{(1)}_t$ of a nominal bond in \$.

[3] This notation is convenient to accomodate the fact that the maturity of zero-coupon bonds changes from one date to the next. For example, assume that there is only one long bond, of maturity n, and let $i^{(n)}_t$ denote its yield to maturity. The long bond trades at date t at a log price $\hat{p}_t = -ni^{(n)}_t$, and it promises a log payoff at date $t+1$ of $\hat{p}^{+1}_{t+1} = -(n-1)\, i^{(n-1)}_{t+1}$.

sequence of budget constraints can then be written as[4]

$$
\begin{aligned}
\bar{W}_{\tau+1} &= R^{w}_{\tau+1}\left(\bar{W}_{\tau}-C_{\tau}\right), \\
R^{w}_{\tau+1} &= e^{i_{\tau}-\pi_{\tau+1}}\left(1+\alpha^{res}_{\tau}e^{x^{res}_{\tau+1}}+\hat{\alpha}^{\top}_{\tau}e^{\hat{x}_{\tau+1}}\right); \quad \tau \geq t.
\end{aligned} \tag{2}
$$

The household problem at date t is to maximize utility (1) subject to (2), given initial wealth $\bar{W}_t$ as well as beliefs about returns. Beliefs about returns are based on current bond prices $\hat{p}_t$, the current short rate i_t, as well as the conditional distribution of the vector $\left(r^{res}_{\tau}, i_{\tau}, \pi_{\tau}, \hat{p}_{\tau}, \hat{p}^{+1}_{\tau}\right)_{\tau>t}$, that is, the return on the residual asset, the short interest rate, the inflation rate and the prices and payoffs on the long spanning bonds. We denote this conditional distribution by G_t.

We now relate bond prices to positions and expectations using investors' optimal policy functions. Since preferences are homothetic and all assets are tradable, optimal consumption and investment plans are proportional to initial wealth. The optimal portfolio weights on long bonds and the residual asset thus depend only on beliefs about returns and can be written as $\hat{\alpha}_t\left(i_t, \hat{p}_t, G_t\right)$ and $\alpha^{res}_t\left(i_t, p_t, G_t\right)$, respectively. Moreover, with an intertemporal elasticity of substitution of one, the optimal consumption rule is $C_t = (1-\beta)\,\bar{W}_t$. Now suppose we observe investors' bond positions: we write B_t for the total dollar amount invested in bonds at date t, and we collect investors' holdings of the two long bonds in the vector $\widehat{B}_t$.

We perform two types of exercises. Consider first a class of investors who invest only in bonds; there is no residual asset. We must then have

$$
\hat{\alpha}_t\left(i_t, \hat{p}_t, G_t\right) = \frac{\hat{B}_t}{B_t}. \tag{3}
$$

These equations can be solved for long bond prices $\hat{p}_t$ as a function of the short rate i_t, bond positions $(B_t, \hat{B}_t)$ and expectations G_t. We can thus characterize yield spreads in terms of these variables.

Second, suppose there is a residual asset. Since investors' total asset holdings are $\beta\bar{W}_t = \frac{\beta}{1-\beta}C_t$,

[4] Here $e^{\hat{x}_t}$ is an N-vector with the jth element equal to $e^{\hat{x}_{t,j}}$.

we must have

$$
\begin{aligned}
\hat{\alpha}_t\left(i_t, \hat{p}_t, G_t\right) &= \frac{1-\beta}{\beta}\frac{\hat{B}_t}{C_t}, \\
\alpha_t^{res}\left(i_t, \hat{p}_t, G_t\right) &= 1-\frac{1-\beta}{\beta}\frac{B_t}{C_t}.
\end{aligned} \tag{4}
$$

These equations can be solved for long bond prices $\hat{p}_t$ and the short rate i_t, as a function of bond positions $(B_t, \hat{B}_t)$, consumption C_t and expectations G_t. This characterizes both short and long yields in terms of positions and expectations.

Portfolio choice when beliefs are driven by a normal VAR

We now restrict beliefs to obtain tractable approximate formulas for investors' portfolio policies. Let z_t denote a vector of exogenous state variables that follows a vector autoregression with homoskedastic normally distributed shocks. We assume that log short interest rates i_t and log inflation π_t are linear functions of z_t and that log excess returns $x_t = (x_t^{res}, \hat{x}_t')'$ are linear functions of z_t and z_{t-1}. Households' belief about future returns, interest rates and inflation at date t is now defined as the conditional implied by the VAR. As a result, the state vector for the household problem is $\left(\bar{W}_t, z_t\right)$.

We use the approximation method proposed by Campbell, Chan and Viceira (2003). The basic idea is that the log return on a portfolio in a discrete time problem is well approximated by a discretized version of its continuous-time counterpart. In our setup, the log return on wealth is approximated by

$$
\log R_{t+1}^w \approx i_t - \pi_{t+1} + \alpha_t^\top x_{t+1} + \frac{1}{2}\alpha_t^\top\left(\operatorname{diag}\left(\Sigma_{xx}\right) - \Sigma_{xx}\alpha_t\right), \tag{5}
$$

where Σ_{xx} is the one-step-ahead conditional covariance matrix of excess returns x_{t+1}, and α_t denotes the vector of portfolio weights $\left(\alpha_t^{res}, \hat{\alpha}_t^\top\right)^\top$.

If this approximation is used for the return on wealth, the investor's value function can be written as $v_t\left(\bar{W}_t, z_t\right) = \log \bar{W}_t + \tilde{v}_t$, where $\tilde{v}_t$ is linear-quadratic in the state vector z_t. Moreover, the optimal portfolio is

$$\text{(6)} \qquad \alpha_t \approx \frac{1}{\gamma}\Sigma_{xx}^{-1}\left(E_t\left[x_{t+1}\right]+\frac{1}{2}\text{diag}\left(\Sigma_{xx}\right)\right)+\left(1-\frac{1}{\gamma}\right)\Sigma_{xx}^{-1}cov_t\left(x_{t+1},\pi_{t+1}\right)$$
$$-\left(1-\frac{1}{\gamma}\right)\Sigma_{xx}^{-1}cov_t\left(x_{t+1},\tilde{v}_{t+1}\right).$$

If $\gamma = 1$ – the case of separable logarithmic utility – the household behaves "myopically", that is, the portfolio composition depends only on the one-step-ahead distribution of returns. More generally, the first line in (6) represents the myopic demand of an investor with one-period horizon and risk aversion coefficient γ. To obtain intuition, consider the case of independent returns, so that Σ_{xx} is diagonal. The first term then says that the myopic investor puts more weight on assets with high expected returns and low variance, and more so when risk aversion is lower. The second term says that, if $\gamma > 1$, the investor also likes assets that provide insurance against inflation, and buys more such insurance assets if risk aversion is higher. For general Σ_{xx}, these statements must be modified to take into account correlation patterns among the individual assets.

For a long-lived household with $\gamma \neq 1$, asset demand also depends on the covariance of excess returns and future continuation utility $\tilde{v}(z_{t+1})$. Continuation utility is driven by changes in investment opportunities: a realization of z_{t+1} that increases $\tilde{v}$ is one that signals high returns on wealth ("good investment opportunities") in the future. Agents with $\gamma > 1$ prefer relatively more asset payoff in states of the world where investment opportunities are bad. As a result, an asset that pays off when investment opportunities are bad is attractive for a high-γ agent. He will thus demand more of it than a myopic agent.

Explicit price formulas

Consider the case without a residual asset. Using the portfolio policy (6), equation (3) can be rearranged to provide an explicit formula for long bond prices:

$$\hat{p}_t = -i_t + E_t\hat{p}_{t+1} + \frac{1}{2}\text{diag}\left(\Sigma_{\hat{x}\hat{x}}\right)$$
$$\text{(7)} \qquad -\gamma\Sigma_{\hat{x}\hat{x}}\widehat{B}_t/B_t + (\gamma-1)\,cov_t\left(\hat{x}_{t+1},\pi_{t+1}\right) - (\gamma-1)\,cov_t\left(\hat{x}_{t+1},\tilde{v}_{t+1}\right).$$

The first line is (log) discounted expected future price, where the variance term appears because of Jensen's inequality. The second line is the risk premium, which consists of three parts. The first is proportional to $\Sigma_{\hat{x}\hat{x}}\widehat{B}_t/B_t$, the covariance of excess returns with the excess return on wealth $\hat{x}_{t+1}^{\top}\widehat{B}_t/B_t$. In the case of log utility ($\gamma = 1$) this covariance represents the entire risk premium. More generally, a higher risk aversion coefficient drives up the compensation required for covariance with the return on wealth.[5] The second term is an inflation risk premium. For $\gamma > 1$, this premium is negative: households have to be compensated less to hold an asset that provides insurance against inflation. Finally, the third term is a premium for covariance with future investment opportunities. An asset that insures households against bad future investment opportunities – by paying off less when continuation utility $\tilde{v}_{t+1}$ is high – commands a lower premium.

Explicit price formulas are also available when investors have access to a residual asset. Let $\alpha_t^W = \left(\hat{B}_t^{\top}/C_t, 1 - B_t/C_t\right)$ denote the investor's wealth portfolio. Using equation (6), we can rearrange (4) as

$$\begin{aligned}
\hat{p}_t &= -i_t + E_t\left[\hat{p}_{t+1}\right] + \frac{1}{2}\mathrm{diag}\left(\Sigma_{\hat{x}\hat{x}}\right) \\
&\quad -\gamma\delta\Sigma_{\hat{x}x}\alpha_t^W + (\gamma-1)\,cov_t\left(\hat{x}_{t+1},\pi_{t+1}\right) - (\gamma-1)\,cov_t\left(\hat{x}_{t+1},\tilde{v}_{t+1}\right), \\
i_t &= E_t\left[r_{t+1}^{res}\right] + E_t\left[\pi_{t+1}\right] + \frac{1}{2}var_t\left(x_{t+1}^{res}\right) + (\gamma-1)\,cov_t\left(x_{t+1}^{res},\pi_{t+1}\right) \\
&\quad -(\gamma-1)\,cov_t\left(x_{t+1}^{res},\tilde{v}_{t+1}\right) - \gamma\delta\Sigma_{x^{res}x}\alpha_t^W.
\end{aligned} \tag{8}$$

where $\delta = \beta^{-1} - 1$. The risk premium on long bonds now also depends on the covariance between excess bond returns and the excess return on the residual asset (through the expression $\Sigma_{\hat{x}x}\alpha_t^W$). The short rate depends on moments of the residual assets as well as expected inflation. Expectations about the real return on the residual asset and the perceived risk premium on that asset fix the real interest rate.

If the risk premium is constant, a version of the expectations hypothesis holds: on average up to a constant, buying a long bond at t and holding it to maturity should cost the same as buying a short bond at t, earning interest i_t on it and then buying the long bond only at $t+1$. Our model distinguishes three reasons for changes in risk premia. First, since the composition of wealth changes over time, for example with changes in the relative amount of different bonds in $\hat{B}_t$, there

[5] Since consumption is proportional to savings, or wealth, the first term in the risk premium also represents the covariance of returns with consumption growth, multiplied by risk aversion.

can be time variation in risk premia. Second, the strength of hedging demand may vary over time. For our numerical results below, the function $\tilde{v}$ will be approximately linear-quadratic in z_{t+1}, and so the need for insurance against bad states will indeed vary over time. Third, investors may have expectations of future prices that are not rational. This implies that even if their subjective risk premia are constant, the modeler may be able to predict excess returns on long bonds with some variable known at time t. This predictability reflects the systematic forecast errors by investors.

III Survey forecasts

We measure subjective expectations of interest rates with survey data from two sources. Both sources conduct comparable surveys that ask approximately 40 financial market professionals for their interest-rate expectations at the end of each quarter and record the median survey response. Our first source are the Goldsmith-Nagan surveys that were started in mid-1969 and continued until the end of 1986. These surveys ask participants about their one-quarter ahead and two-quarter ahead expectations of various interest rates, including the 3-month Treasury bill, the 12-month Treasury bill rate, and a mortgage rate. Our second source are Bluechip Financial Forecasts, a survey that was started in 1983 and continues until today. This survey asks participants for a wider range of expectation horizons (from one to six quarters ahead) and about a larger set of interest rates. The most recent surveys always include 3-month, 6-month and 1-year Treasury bills, the 2-year, 5-year, 10-year and 30-year Treasury bonds, and a mortgage rate.[6]

Deviations of subjective expectations from objective expectations of interest rates have consequences for expected excess returns on bonds. We define the (log) excess return on an n-period bond for a h-period holding period as the log-return from t to $t+h$ on the bond in excess of the h-period interest rate:

$$rx_{t+h}^{(n)} = p_{t+h}^{(n-h)} - p_t^{(n)} - i_t^{(h)}.$$

[6] The survey questions ask for constant-maturity Treasury yield expectations. To construct zero-coupon yield expectations implied by the surveys, we use the following approximation. We compute the expected change in the n-year constant-maturity yield. We then add the expected change to the current n-year zero-coupon yield.

The objective expectation E of an excess returns can be decomposed as follows:

$$(9) \qquad \begin{aligned} \underbrace{E_t\left[rx_{t+h}^{(n)}\right]}_{\text{objective premium}} &= E_t^*\left[rx_{t+h}^{(n)}\right] + E_t\left[p_{t+h}^{(n-h)}\right] - E_t^*\left[p_{t+h}^{(n-h)}\right] \\ &= \underbrace{E_t^*\left[rx_{t+h}^{(n)}\right]}_{\text{subjective premium}} + \underbrace{(n-h)\left(E_t^*\left[i_{t+h}^{(n-h)}\right] - E_t\left[i_{t+h}^{(n-h)}\right]\right)}_{\text{subj. - obj. interest-rate expectation}} \end{aligned}$$

This expression shows that, if subjective expectations E^* of interest rates deviate from their objective expectations E, the objective premium is different from the subjective premium. In particular, if the difference between objective and subjective beliefs changes in systematic ways over time, the objective premium may change over time even if the subjective premium is constant.

We can evaluate equation (9) based on our survey measures of subjective interest-rate expectations $E_t^*\left[i_{t+h}^{(n-h)}\right]$ for different maturities n and different horizons h. To measure objective interest-rate expectations $E_t\left[i_{t+h}^{(n-h)}\right]$, we estimate unrestricted VAR dynamics for a vector of interest rates with quarterly data over the sample 1952:2-2007:1 and compute their implied forecasts. Later, in Section IV, we will impose more structure on the VAR by assuming the absence of arbitrage and using a lower number of variables in the VAR, and thereby check the robustness of the empirical findings we document here. The vector of interest rates Y includes the 1-year, 2-year, 3-year, 4-year, 5-year, 10-year and 20-year zero-coupon yields. We use data on nominal zero-coupon bond yields with longer maturities from the McCulloch file available from the website http://www.econ.ohio-state.edu/jhm/ts/mcckwon/mccull.htm. The sample for these data is 1952:2 - 1990:4. We augment these data with the new Gurkaynak, Sack, and Wright (2006) data. We compute the forecasts by running OLS directly on the system $Y_{t+h} = \mu + \phi Y_t + u_{t+h}$, so that we can compute the $h-$horizon forecast simply as $\mu + \phi Y_t$.

Figure 1 plots the left-hand side of equation (9), expected excess returns under objective beliefs as a black line, and the second term on the right-hand side of the equation, the difference between subjective and objective interest-rate expectations, as a gray line. For the short post-1983 sample for which we have Bluechip data, we have data for many maturities n and many forecasting horizons h. The lower two panels of Figure 1 use maturities $n = 3$ years and 11 years and a horizon of $h = 1$ year, so that we deal with expectations of the $n - h = 2$ year and 10 year interest rate. These combinations of n and h are in the Bluechip survey, and the VAR includes these two maturities as well so that the computation of objective expectations is easy. For the long post-1970 sample, we need to combine data from the Goldsmith-Nagan and Bluechip surveys. The upper left panel shows the $n = 1.5$ year bond and $h = 6$ month holding period. from the estimated VAR (which includes the $n - h = 1$ year yield.) This works, because both surveys include the $n - h = 1$ year interest rate and a $h = 6$-month horizon. The VAR delivers an objective 6-month ahead expectation of the 1-year interest rate. For long bonds, we do not have consistent survey data over this long sample. To get a rough idea of long-rate expectations during the Great Inflation, we take the Goldsmith-Nagan data on expected mortgage-rate changes and the Bluechip data on expected 30-year Treasury-yield over the next $h = 2$ quarters and add them to the current 20-year zero-coupon yield. The VAR produces a $h = 2$ quarter ahead forecast of the 20-year yield.

Figure 1 also shows NBER recessions as shaded areas. The plots indicate that expected excess returns under objective beliefs and the difference between subjective and objective interest-rate expectations have common business-cycle movements. The patterns appear more clearly in the lower panels which use longer (1 year) horizons. This is not surprising in light of the existing predictability literature which documents that expected excess returns on bonds and other assets are countercyclical when we look at longer holding periods, such as one year (e.g., Cochrane and Piazzesi 2005.) In particular, expected excess returns are high right after recession troughs. The lower panels show indeed high values for both series around and after the 1991 and 2001 recessions. The series are also high in 1984 and 1996, which are years of slower growth (as indicated, for example, by employment numbers) although they were not classified as recessions. For shorter holding periods, the patterns are also there in the data but they are much weaker. However, the upper panels show additional recessions where similar patterns appear. For example, the two series in both panels are high in the 1970, 1974, 1980 and 1982 recessions or shortly afterwards.

Table 1 shows summary statistics of subjective beliefs measured from surveys. During the short Bluechip sample, the average difference between realized interest rates and their one-quarter ahead subjective expectation is negative for short maturities and close to zero, or slightly positive for longer maturities. The average forecast error is -15 basis points for the 3-quarter interest rate and -45 basis points for the 6-quarter interest rates. These two mean errors are the only ones

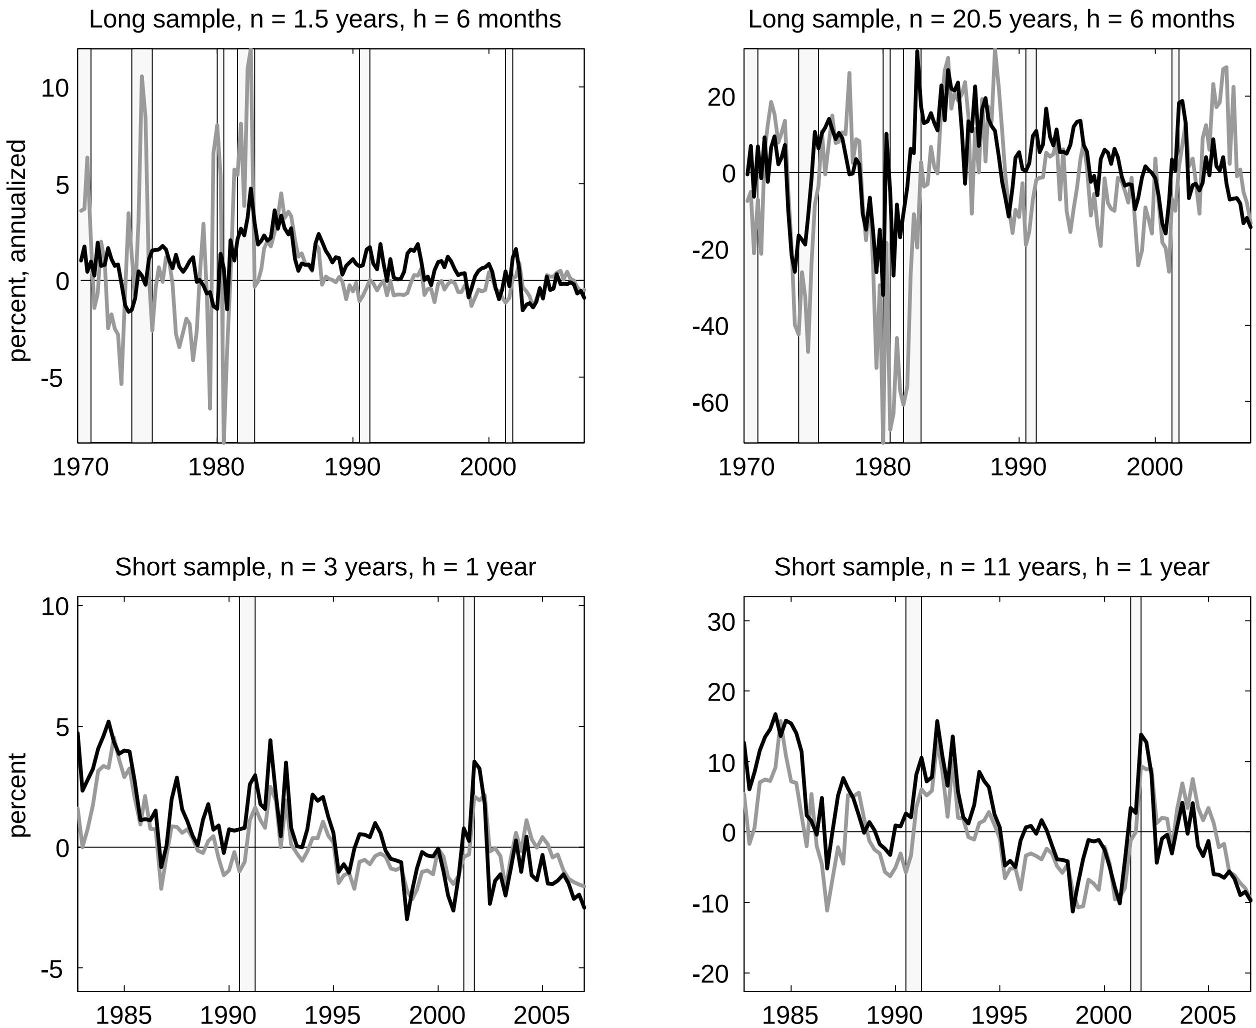

Figure 1: Each panel shows objective expectations of excess returns in black (the left-hand side of equation (9)) and the difference between subjective and objective interest-rate expectations in gray (the second term on the right-hand side of the equation) for the indicated bond maturity n and holding period/forecast horizon h. Shaded areas indicate NBER recessions. The numbers are annualized and in percent. The upper panels show data over a longer sample than the lower panels.

that are statistically significant, considering the sample size of 98 quarters (which means that the ratio of mean to standard deviation needs to be multiplied by roughly 10 to arrive at the relevant t-statistic.) There is stronger evidence of bias at the 1-year horizon, where on average subjective interest-rate expectations are above subsequent realizations for all maturities.

The upward bias in subjective expectations may partly explain why we observe positive average excess returns on bonds. The right-hand side of equation (9) shows why: if objective expectations are unbiased, then $E_t^* \left[i_{t+h}^{(n-h)} \right] > E_t \left[i_{t+h}^{(n-h)} \right]$ on average, which raises the value of the left-hand side of the equation. The magnitude of the bias is also economically significant. For example, the -56 basis-point bias in subjective expectations of the 1-year interest rate can easily account for the 39 basis-point objective premium of the 2-year bond. For higher maturities, we need to multiply the subjective bias by $n-1$ as in equation (9). For example, the -52 basis point bias in 2-year interest rate expectations multiplied by $n-h=2$ more than accounts for the 57 basis point objective premium.

When we match up these numbers, it is important to keep in mind that subjective biases and objective premia are measured imprecisely, because they are computed with small data samples. In particular, over most of the Bluechip sample, interest rates were declining.

To sum up, the evidence presented in this section suggests that subjective interest-rate expectations deviate from the objective expectations that we commonly measure from statistical models. Table 1 suggests that these deviations may account for average objective premia. Figure 1 suggest that these deviations may also be responsible for the time-variation in objective bond premia.

Table 1: Subjective Biases And Objective Bond Premia

	horizon	maturity n								
subj. bias	h	3 qtr	6 qtr	1 year	2 year	3 year	5 year	7 year	10 year	30 year
average	1 qrt	−0.15	−0.45	−0.15	−0.11	−0.07	0.00	0.05	0.12	0.01
	1 year	−0.57	−0.86	−0.56	−0.52	−0.47	−0.38	−0.32	−0.24	−0.30
stdev										
	1 qrt	0.57	0.83	0.79	0.76	0.76	0.73	0.71	0.67	0.60
	1 year	1.41	1.64	1.66	1.55	1.47	1.35	1.28	1.22	1.12
obj. premium	1 year				0.39	0.57	0.76	0.85	0.83	−0.31

Note: The table reports summary statistics of subjective expectational errors computed as $i_{t+h}^{n-h} - E_t^* \left[i_{t+h}^{(n-h)} \right]$ for the indicated horizon h and maturity n. The data are quarterly Bluechip Financial Forecasts from 1983:1-2007:1, 98 quarters. The numbers are annualized and in percent. The last two rows are average excess returns computed as sample average of $rx_{t+h}^{(n-h)} = p_{t+h}^{(n-h)} - p_t^{(n)} - i_t^{(h)}$ for the indicated holding period h and maturity n. The quarterly zero-coupon yield data for the years 1952:2 - 1990:4 are from the McCulloch files and for the years 1991:1-2007:1 from the new Gurkaynak, Sack, and Wright (2006) dataset. The numbers are annualized and in percent.

IV Modeling investor beliefs

The previous section has documented some properties of survey forecasts of interest rates. In order to implement our asset pricing model, we need investors' subjective conditional distributions over future asset returns. Subsection A. describes a general setup to construct such distributions. In subsection B., we report estimation results for a specific model of beliefs.

A. Setup

The basic idea is to start from an objective probability, provided by a statistical model of macro variables and yields that fits the data well from our (the modelers') perspective. A second step then uses survey forecasts to estimate the Radon-Nikodym derivative of investors' subjective probability, denoted P^*, with respect to the objective probability P.

Objective probabilities

In order to choose portfolios, investors in our model form beliefs about interest rates, inflation, and possibly the return on a residual asset. We describe the joint distribution of these variables by a large state space system that nests in particular an affine term structure model for yields. Let h_t denote an S-vector of observables that includes all relevant macro variables, and may also include some interest rates. As before, let $i_t^{(n)}$ denote the yield to maturity on an n-period zero coupon bond. We represent the joint distribution of h_t and interest rates under the objective probability

P by

$$
\begin{aligned}
(10)\qquad h_t &= \mu + \eta_h s_{t-1} + e_t \\
(11)\qquad s_t &= \phi_s s_{t-1} + \sigma_s e_t, \\
(12)\qquad f_t &= \eta_f s_t \\
(13)\qquad i_t^{(n)} &= a_n + b_n' f_t, \qquad n = 1, 2, ...
\end{aligned}
$$

Here s_t and e_t are S-vectors of state variables and i.i.d. zero-mean normal shocks with $E e_t e_t' = \Omega$, respectively. Moreover, f_t is an F-vector of term-structure factors which are in turn linear combinations (for example, selections) of the state variables. The term-structure model implies coefficients a_n and b_n that describe yields as affine functions of the factors. Cross-equation restrictions need to be imposed on the matrices in (10)-(13) to ensure that the term-structure factors are Markov and that yields in h_t are consistent with the term-structure model.

We distinguish two types of state variables and observables. The first Y state variables are term structure factors s_t^y that are each identified (up to a constant) with a particular yield or yield spread, with the latter collected in the first Y components h_t^y of h_t. In particular, the first component of h_t is always the short interest rate $i_t^{(1)}$ and the first state variable is the demeaned short interest rate, that is, $s_{t,1} = i_t^{(1)} - \mu_1$. The other $S - Y$ state variables z_t^o are expected values of macro variables h_t^o; they drive the remaining $F - Y$ term structure factors. We can rewrite the first three equations of (10)-(13) as

$$
\begin{aligned}
\begin{pmatrix} h_t^y \\ h_t^o \end{pmatrix} &= \begin{pmatrix} \mu^y \\ \mu^o \end{pmatrix} + \begin{pmatrix} & \phi_y \\ 0 & I_{(S-Y)} \end{pmatrix} \begin{pmatrix} s_{t-1}^y \\ s_{t-1}^o \end{pmatrix} + e_t, \\
\begin{pmatrix} s_t^y \\ s_t^o \end{pmatrix} &= \begin{pmatrix} \phi_y \\ \phi_o \end{pmatrix} s_{t-1} + \begin{pmatrix} I_Y & 0 \\ & \sigma_o \end{pmatrix} e_t, \\
f_t &= \eta_f s_t = \begin{pmatrix} I_Y & 0 \\ 0 & \eta_o \end{pmatrix} \begin{pmatrix} s_t^y \\ s_t^o \end{pmatrix}
\end{aligned}
$$

where I_N is an identity matrix of size N. The first Y state equations are copies of the first Y

observation equations, up to the constant vector μ^y. In addition, the restrictions imply that e_t is the forecast error on a forecast of the observables h_t given all past observables $(h_\tau)_{\tau<t}$.

To ensure that the term-structure factors are Markov, we assume that there exists an $F \times F$ matrix ϕ_f, such that $\eta_f \phi_s = \phi_f \eta_f$. The vector f_t can then be represented as an AR(1) process even if $S > F$:

$$\begin{aligned} f_t &= \eta_f s_t = \eta_f \phi_s s_{t-1} + \eta_f \sigma_s e_t \\ &= \phi_f \eta_f s_{t-1} + \eta_f \sigma_s e_t \\ &= \phi_f f_{t-1} + \sigma_f e_t. \end{aligned} \tag{14}$$

The general structure allows for $F - Y$ term structure factors that are linear combinations of forecasts of macro variables. For example, the matrix η_o could be a selection matrix that makes expected inflation a term-structure factor. Importantly, expected inflation can be a term-structure factor even if inflation itself cannot be represented as a component of the AR(1) process z_t.

The general structure nests two useful special cases. The first assumes that all term-structure factors can be identified with yields or spreads, that is, $F = Y$ and $\eta_f = (I_Y, 0)$. The Markov restriction is then $\phi_y = (\phi_f, 0)$ for some $Y \times Y$ matrix ϕ_f. In other words, macro variables are assumed to not help forecasts yields, given the information in the factors z_t^y. We also have $\sigma_f = (I_Y, 0)$, so that σ_f simply picks out the first two components of e_t. The second special case assumes that all forecasts of macro variables included in the system are themselves term structure factors, that is, $F = S$ and $\eta_f = I_S$. The Markov restriction is then simply $\phi_s = \phi_f$ for some $S \times S$ matrix ϕ_f, which is always satisfied.

Term-structure coefficients

We assume that there are no arbitrage opportunities in bond markets. As a result, there exists a "risk neutral" probability Q under which bond prices are discounted present values of bond payoffs. In particular, the prices $P^{(n)}$ of zero-coupon bonds with maturity n satisfy the recursion

$$P_t^{(n)} = e^{-i_t} E_t^Q \left[P_{t+1}^{(n-1)} \right]$$

with terminal condition $P_t^{(0)} = 1$.

We specify the Radon-Nikodym derivative ξ_t^Q of the risk neutral probability Q with respect to the objective probability P by $\xi_1^Q = 1$ and

$$\frac{\xi_{t+1}^Q}{\xi_t^Q} = \exp\left(-\frac{1}{2}\lambda_t' \sigma_f \Omega \sigma_f' \lambda_t - \lambda_t' \sigma_f e_{t+1}\right),$$

where λ_t is an F-vector. Since the innovations to the factors $\sigma_f e_t$ are normal with variance $\sigma_f \Omega \sigma_f'$ under the objective probability P, ξ_t^Q is a martingale under P. The vector λ_t contains the "market prices of risk" associated with the innovations $\sigma_f e_{t+1}$ to the term-structure factors. Indeed, the (log) expected excess returns at date t on a set of assets with payoffs proportional to $\exp(\sigma_f e_{t+1})$ is equal to $\sigma_f \Omega \sigma_f' \lambda_t$. For the purpose of pricing bonds, it is sufficient to specify market prices of risk for shocks to term-structure factors. At the same time, we are not ruling out that agents worry about other shocks as well.

We assume that risk premia are linear in the term-structure factors, that is,

$$\lambda_t = l_0 + l_1 f_t,$$

for some $F \times 1$ vector l_0 and some $F \times F$ matrix l_1. Standard calculations then deliver that bond prices are exponential linear functions of the factors

$$P_t^{(n)} = e^{-i_t} E_t^Q \left[P_{t+1}^{(n-1)}\right] = e^{-i_t} E_t \left[\frac{\xi_{t+1}^Q}{\xi_t^Q} P_{t+1}^{(n-1)}\right] = \exp\left(A_n + B_n^\top f_t\right)$$

where A_n is a scalar and B_n is an $F \times 1$ vector of coefficients that depend on maturity n.

The recursion for bond prices implies that the coefficients are computed from the difference equations

$$\begin{aligned} A_{n+1} &= A_n - B_n' \sigma_f \Omega \sigma_f' \; l_0 + \frac{1}{2} B_n' \sigma_f \Omega \sigma_f' B_n - \mu_1 \\ B_{n+1} &= \left(\phi_f - \sigma_f \Omega \sigma_f' l_1\right)' B_n - e_1 \end{aligned}$$

where e_1 is the first unit vector of length F and initial conditions are given by $A_0 = 0$ and $B_0 = 0_{F \times 1}$.

The coefficients for the short (one-period) bond are thus $A_1 = -\mu_1$ and $B_1 = -e_1$. Given these formulas for bond prices, interest rates $i_t^{(n)} = -\ln P_t^{(n)}/n$ are also linear functions of the factors with the coefficients $a_n = -A_n/n$ and $b_n = -B_n/n$ that appear in equation (13).

From objective to subjective beliefs

We assume that investors' belief has the same basic structure as our time series model. Investors also have in mind a state space representation of h_t and an affine term structure model for the yields $y_t^{(n)}$. Moreover, they recognize the deterministic relationship between term structure factors and yields; in other words, their model of yields also involves the risk neutral measure Q used to price bonds above. However, investors' subjective distribution of the state variables need not be the same as the distribution of these variables under the objective probability P.

To define investors' subjective beliefs, we represent the Radon-Nikodym derivative of investors' subjective belief P^* with respect to the objective probability P by a stochastic process ξ_t^*, with $\xi_1^* = 1$ and

$$\frac{\xi_{t+1}^*}{\xi_t^*} = \exp\left(-\frac{1}{2}\kappa_t'\Omega\kappa_t - \kappa_t' e_{t+1}\right). \tag{15}$$

Since e_t is normal with variance Ω under the objective probability P, ξ_t^* is a martingale under P. Since e_t is the error in forecasting h_t, the process κ_t can be interpreted as investors' bias in their forecast of h_t. Like the risk premia λ_t above, the forecast bias is affine in state variables, that is $\kappa_t = k_0 + k_1 s_t$.

Standard calculations now deliver that the dynamics of h_t and yields under P^* can be represented by

$$\begin{aligned}
h_t &= \mu - \Omega k_0 + (\eta_h - \Omega k_1)s_{t-1} + e_t^* && (16)\\
s_t &= -\sigma_s \Omega k_0 + (\phi_s - \sigma_s' \Omega k_1)s_{t-1} + \sigma_s e_t^*, && (17)\\
f_t &= \eta_f s_t && (18)\\
i_t^{(n)} &= a_n + b_n' f_t, \qquad n = 1, 2, ..., && (19)
\end{aligned}$$

where e_t^* is i.i.d. mean-zero normal with covariance matrix Ω. The vector k_0 thus affects investors' subjective mean of h_t and also the state variables s_t, whereas the matrix k_1 determines how their forecasts of h deviate from the objective forecasts as a function of the state s_t. Since investors use the same risk neutral probability Q to prices bonds, the equations for interest rates (19) involve the same coefficients as in (13).

We further impose restrictions such that the term-structure factors can be represented as an AR(1) process under P^*:

$$\begin{aligned} f_t &= -\eta_f \sigma_s \Omega k_0 + \eta_f (\phi_s - \sigma_s' \Omega k_1) s_{t-1} + \eta_f \sigma_s e_t^* \\ &= -\eta_f \sigma_s \Omega k_0 + (\phi_f - k_f) \eta_f s_{t-1} + \eta_f \sigma_s e_t^* \\ &= \mu_f^* + \phi_f^* f_{t-1} + \sigma_f e_t^*. \end{aligned}$$

Since investors price assets under the risk-neutral measure Q, but their belief is P^* rather than P, their subjective market prices of risk are in general not equal to λ_t. Instead, we impose restrictions such that there is a market price of risk process $\lambda_t^* = \lambda_t + \kappa_t^f$, so that the bond prices computed earlier are also risk-adjusted present discounted values of bond payoffs under the subjective belief P^*:

$$P_t^{(n)} = e^{-i_t} E_t^Q \left[P_{t+1}^{(n-1)} \right] = e^{-i_t} E^* \left[\exp \left(-\frac{1}{2} \lambda_t^{*\prime} \sigma_f \Omega \sigma_f' \lambda_t^* - \lambda_t^{*\prime} \sigma_f e_{t+1} \right) P_{t+1}^{(n-1)} \right].$$

B. Results

We have empirically implemented two versions of the general model (10)-(13). In both cases, a period is a quarter, and there are $S = 4$ observables h_t: the 1 quarter interest rate, the spread between the 5 year and the 1 quarter rate, inflation, and consumption growth. The "2 factor system" uses only the short rate and the spread as term structure factors, that is, $F = Y = 2$. Under the "4-factor system", expected inflation and expected consumption growth are also factors ($S = F = 4$). In what follows, we report detailed results for the two factor model. The main message is similar for the four factor model.

Data

We use the data on zero-coupon interest rates and survey forecasts described in detail in section

III. Moreover, we use quarterly data on nondurables and service consumption and inflation measured by the Personal Consumption deflator obtained from the NIPA tables. The sample consists of end-of-quarter observations for 1952:2 - 2004:3 (which will be the sample for our positions data.)

Estimation

Estimation of the 2-factor model proceeds in four steps. First, we set μ_y equal to the sample mean of h_t^y and find the parameters ϕ_f and $var\left(\sigma_f e_t\right)$ that govern the VAR for the term structure factors (14) using standard SUR. Second, we estimate the parameters l_0 and l_1 that describe the objective risk premia, given the VAR estimates from the first stage. This is done by minimizing the sum of squared fitting errors for a set of yields, subject to the constraint that the 1 quarter and 5 year rates are matched exactly.

The third step is to estimate the full system. Here we use information already gained from the term structure estimation in the first step: since $Y = F = 2$, we have that $\phi_y = (\phi_f,\, 0)$ as well as $\sigma_f = (I_2,\, 0)$, which implies that $var\left(\sigma_f e_t\right)$ is the top left 2×2 submatrix of Ω. The first step thus already delivers estimates for the first two rows of the matrices ϕ_s and η_s, and for three elements in Ω. We estimate the remaining 23 parameters of the full system (10)-(11) by maximum likelihood holding the term-structure parameters already estimated in step 1 fixed at their estimated values. This step also produces a sequence of estimates $(\hat{s}_t)$ for the realized values of the state variables s_t.

The fourth step is to estimate the parameters k_0 and k_1 that govern the Radon-Nikodym derivative (15) of the subjective belief. The current implementation of the model uses only interest rate forecasts (and not yet survey forecast data on inflation and growth.) It thus restricts attention to biases to the term structure factors. In particular, we let $k_0' = (k_0^{f\prime}\ 0)$ for a 2×1 vector k_0^f and let k_1 consists of all zeros, except a 2×2 matrix k_1^f in the top left corner. We estimate these six parameters by minimizing an objective function that penalizes differences between model-implied subjective forecasts and survey forecasts. This step is thus technically similar to the estimation of risk premia l_0 and l_1 under the objective probability in step 2.

In particular, we select forecast horizons of 1 quarter and 1 year, as well as yields of maturity 1 quarter, 1 year and 10 years. Consider some date where we have a survey forecast of some yield over some horizon. Given parameters k_0 and k_1, as well as the parameter and state variable estimates from step 3, we can find the forecast of that yield, for that date and horizon, under the subjective belief P^*. The objective function now sums up differences between survey forecasts and model-implied subjective forecasts for every date, yield and horizon. It also adds, for every date and for the same forecast horizons, differences between the model-implied subjective forecasts of inflation and consumption growth and the model-implied objective forecasts of inflation and consumption growth. The presence of these last terms implies that the estimation cannot distort the macro forecasts "too much" in order to fit the interest rate forecasts.

Term structure dynamics

Panel A in Table 2 reports parameter estimates.[7] The estimated dynamics of the factors are highly persistent; the eigenvalues of the matrix ϕ_f are 0.96 and 0.75. The two factors are contemporaneously negatively correlated and the spread is negatively correlated with the short rate lagged less than year, and positively correlated with longer lags of the short rate. The short rate is negatively correlated with the spread lagged less than three years, with weak correlation for longer lags.

The parameter estimates of l_0 and l_1 govern the behavior of the conditional Sharpe ratio $-\lambda_t$ (mean excess return divided by standard deviation) on long bonds. Since the standard deviation of excess returns in this model is constant, and the factors f_t are mean zero, the large negative estimate of the first l_0 component indicates that expected excess returns on long bonds are positive. The entries in l_1 are negative and indicate that expected excess returns on long bonds are high in periods with high short rate or high spreads. The dependence of expected excess returns on spreads captures that model-implied expected excess returns are countercyclical.

[7] This draft does not yet report standard errors. Standard errors can be computed by GMM, taking into account the multi-step nature of the estimation.

TABLE 2: ESTIMATION OF TERM STRUCTURE MODEL

Panel A: Parameter Estimates

ϕ_f		$\sigma_f \Omega^{\frac{1}{2}} \times 100$		$\sigma_f \Omega^{\frac{1}{2}} l_0$	$\sigma_f \Omega^{\frac{1}{2}} l_1$	
0.952	0.108	0.248	0	-0.218	-21.2	-55.8
0.016	0.758	-0.126	0.114	0.033	-4.4	-115.9

Panel B: Fitting errors for bond yields (annualized)

	maturity (in quarters)					
	1 qtr	4 qrts	20 qrts	40 qrts	60 qrts	80 qrts
mean absolute errors (in %)	0	0.26	0	0.23	0.37	.46

Panel B reports by how much the model-implied yields differ from observed yields on average. By construction, the model hits the 1-quarter and 5-year interest rates exactly, because these rates are included as factors. For intermediate maturities, the error lies within the $.23 - .46$ percent range. We will see below that these errors are sufficiently small for our purposes.

Subjective vs. objective dynamics

Table 3 reports estimation results for the change of measure from the objective to the subjective belief. For the two factor model, we estimate 6 parameters, two in k_0 and four in k_1. Rather than report these estimates directly, Panel A of the table shows the implied factor dynamics and market prices of risk of the investor's subjective term structure model. The market price of risk $\lambda_t^* = \lambda_t - \kappa_t$ have been premultiplied by the volatility matrix $\Omega^{\frac{1}{2}} \sigma_f'$ from Table 2 to make them comparable to the market prices of risk in the earlier table. Panel B reports mean absolute distances between the survey forecasts and model-implied forecasts, for both the subjective belief and the objective statistical model. Comparison of these errors provides a measure of how well the change of measure works to capture the deviation of survey forecasts from statistical forecasts.

The results show that the improvement is small for short-horizon forecasts of short yields. However, there is a marked reduction of errors for 1-year forecasts, especially for the 10-year bond. Figure 2 shows where the improvements in matching the long-bond forecasts come from. The top

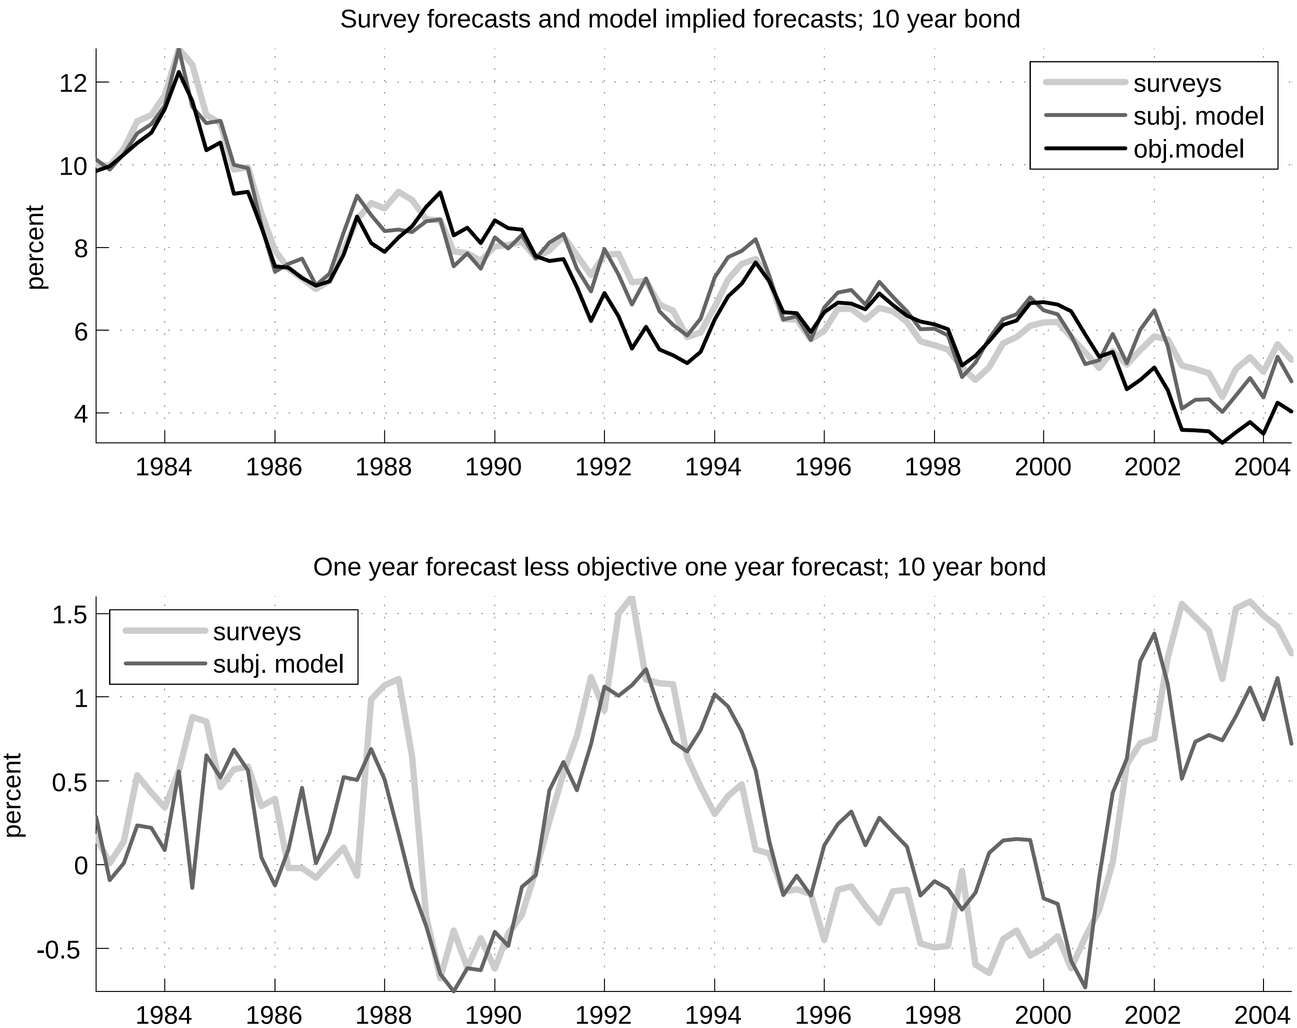

Figure 2: The top panel shows one-year ahead forecasts of the 10-year zero coupon rate constructed from survey data in Section III, together with the corresponding forecasts from our objective and subjective models. The bottom panel shows the difference between the survey forecast and the objective model forecast, as well as the difference between the subjective and objective model forecasts.

panel shows one-year ahead forecasts of the 10-year zero coupon rate constructed from survey data in Section III, together with the corresponding forecasts from our objective and subjective models, for the sam,ple 1982:4-2004:3. All forecasts track the actual 10-year rate over this period, which is natural given the persistence of interest rates. The largest discrepancies between the survey forecasts and the subjective model on the one hand, and the objective model on the other hand, occur during and after the recessions of 1990 and 2001. In both periods, the objective model quickly forecasts a drop in the interest rate, whereas investors did not actually expect such a drop. The subjective model captures this property.

For our asset pricing application, we are particularly interested in how well the subjective model captures deviations of survey forecasts of long interest rates from their statistical forecasts over the business cycle. As discussed in Section III, this forecast difference is closely related to measured expected excess returns. The bottom panel of the figure focuses again on forecasting a 10-year rate over one year, and plots the difference between the survey forecast and the objective model forecast, as well as the difference between the subjective and objective model forecasts. It is apparent that both forecast differences move closely together at business cycle frequencies, increasing during and after recessions. We thus conclude that the subjective model is useful to capture this key fact about subjective forecasts that matters for asset pricing.

Table 3: Estimation of Subjective Belief

Panel A: Parameter Estimates

μ_f^*	ϕ_f^*		$\Omega^{\frac{1}{2}}\sigma'_f\lambda_0^*$	$\Omega^{\frac{1}{2}}\sigma'_f l_1^*$	
.002	.996	.129	-0.145	-3.4	-47.1
-.001	-.044	.887	0.043	-37.4	5.8

Panel B: Mean absolute fitting errors for yield forecasts (% p.a.)

	subjective model			objective model		
	maturity of forecasted yield in quarters			in quarters		
forecast horizon	1 qtr	4 qrts	80 qrts	1 qtr	4 qtrs	80 qtr
1 quarter	0.19	0.37	0.35	0.19	0.40	0.36
1 year	0.25	0.36	0.34	0.29	0.47	0.58

The parameter estimates in Table 3 also show that investors' forecasts are on average very close to the forecasts from our statistical model: the unconditional means of the factors under P^* are very small. At the same time, the factors are more persistent under P^*; in particular, other things equal, a one-percent increase in the spread increases the forecast of the spread next period by .88%, as opposed to .75% under the statistical model. The estimated market prices of risk show that the investor requires less compensation for short rate risk than what is suggested by the statistical model: the Sharpe ratio on an asset with payoff proportional to the future price of a short bond drops to .14 from .21 under the statistical model. Moreover, there is less time variation of risk premia than under the statistical model. In particular, as the spread plays a larger role in forecasting future spreads, it plays a smaller role in moving around risk premia for spread shocks.

V Bond positions

In this section, we use the (subjective) term structure model estimated in the previous section to represent the universe of bonds available to investors in terms of a small number of "spanning bonds". In subsection A., we construct, for every zero-coupon bond, a portfolio of three bonds – a short bond and two long bonds – that replicates closely the return on the given zero-coupon bond. In subsection B., we then show how statistics on bond positions in the US economy can be converted into a time series of positions in the three bonds.

A. Replicating Zero Coupon Bonds

According to the term structure model, the price $P_t^{(n)}$ of a zero coupon bond of maturity n at date t is well described by $\exp\left(A_n + B_n' f_t\right)$. We now select N long bonds, zero-coupon bonds with maturity greater than one, and stack their coefficients in a vector $\widehat{a}$ and a matrix $\widehat{b}$. Our goal is to construct a portfolio containing the long bonds and the short bond such that the return on the portfolio replicates closely the return on any other zero-coupon bond with maturity n. We use a discretization of continuous time returns, similar to those used above for the return on wealth. We approximate the change in price on an n-period bond by

$$\begin{aligned}
P_{t+1}^{(n-1)} - P_t^{(n)} &\approx P_t^{(n)} \left(A_{n-1} - A_n + B_{n-1}'(f_{t+1} - f_t) + (B_{n-1} - B_n)' f_t + \frac{1}{2} B_{n-1}' \sigma_f \sigma_f' B_{n-1}' \right) \\
&= P_t^{(n)} \left(A_{n-1} - A_n + B_{n-1}' \mu_f^* + B_{n-1}'(\phi_f^* - I) f_t + (B_{n-1} - B_n)' f_t + \frac{1}{2} B_{n-1}' \sigma_f \sigma_f' B_{n-1}' \right) \\
&\quad + P_t^{(n)} B_{n-1}' \sigma_f \varepsilon_{t+1}^* \\
(20) \qquad &=: a_t^{(n)} + b_t^{(n)} \sigma_f \varepsilon_{t+1}^*
\end{aligned}$$

Conditional on date t, we thus view the change in value of the bond as an affine function in the shocks to the factors $\sigma_f \varepsilon^*_{t+1}$. Its distribution is described by $N+1$ time-dependent coefficients: the constant $a_t^{(n)}$ and the loadings $b_t^{(n)}$ on the N shocks. In particular, we can calculate coefficients $\left(a_t^{(1)}, b_t^{(1)}\right)$ for the short bond, and we can arrange coefficients for the N long bonds in a vector $\hat{a}_t$ and a matrix $\hat{b}_t$. Now consider a portfolio that contains θ_1 units of the short bond and $\hat{\theta}_i$ units of the ith long bond. The change in value of this portfolio is also an affine function in the factor shocks and we can set it equal to the change in value of any n-period bond:

$$(21) \qquad \begin{pmatrix} \theta_1 & \hat{\theta}' \end{pmatrix} \begin{pmatrix} a_t^{(1)} & b_t^{(1)} \\ \hat{a}_t & \hat{b}_t \end{pmatrix} \begin{pmatrix} 1 \\ \sigma_f \varepsilon^*_{t+1} \end{pmatrix} = \begin{pmatrix} a_t^{(n)} & b_t^{(n)} \end{pmatrix} \begin{pmatrix} 1 \\ \sigma_f \varepsilon^*_{t+1} \end{pmatrix}.$$

Since the $(N+1)\times(N+1)$ matrix of coefficient on the left hand side is invertible for a nondegenerate term structure model, we can select the portfolio $\begin{pmatrix} \theta_1 & \hat{\theta}' \end{pmatrix}$ to make the conditional distribution of the value change in the portfolio equal to that of the bond.

Replicating portfolios based on a two-factor model

When stated in terms of units of bonds $\left(\theta_1, \hat{\theta}\right)$, the replicating portfolio for the n-period bond answers the question: how many spanning bonds are equivalent to one n-period bond? For our work below it is more convenient to define portfolio weights that answer the question: how many dollars worth of spanning bonds are equivalent to one dollar worth of invested in the n-period bonds? The answer to this question can be computed using the units $\left(\theta_1, \hat{\theta}\right)$ and the prices of spanning bonds. Figure 3 provides the answer computed from the two-factor term structure model estimated above. Since the term structure model is stationary, these weights do not depend on calendar time.

The maturity of the n-period bond to be replicated is measured along the horizontal axis. The three lines are the portfolio weights $\theta_i P_t^{(i)}/P_t^{(n)}$ on the different spanning bonds i; they sum to one for every maturity n. As spanning bonds i, we have selected the 1-, 8- and 20-quarter bonds. For simplicity, we refer to the long spanning bonds as the middle and the long bond, respectively. The figure shows that the spanning bonds are replicated exactly by portfolio weights of one on themselves. More generally, the replicating portfolios for the most average neighboring bonds. For example, most of the bonds with maturities in between the 1-quarter and 8-quarter bond are

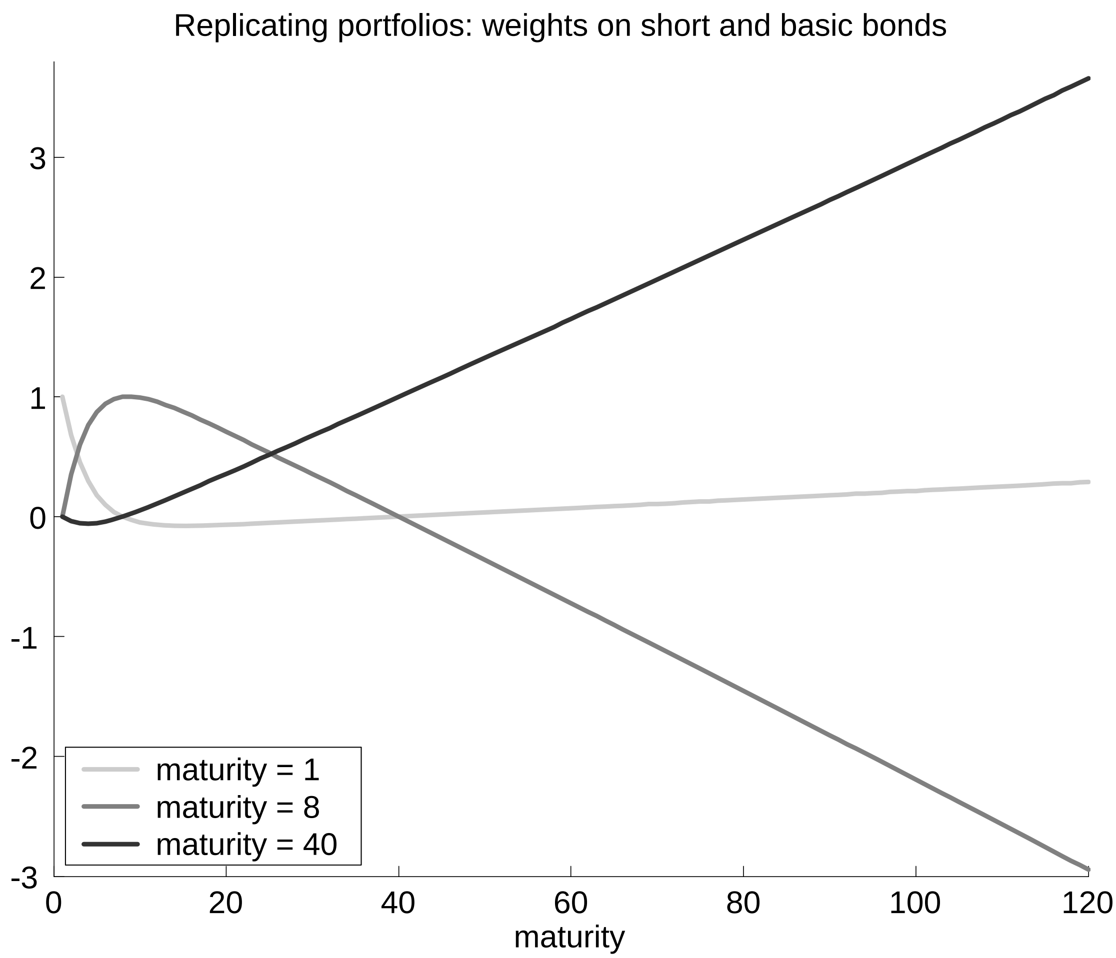

Figure 3: Replicating portfolios; the maturity in quarters of the bond to be replicated is measured along the x-axis.

generating by simply mixing these two bonds, although there is also a small short position in the long bond. Similarly, most of the bonds with maturities in between the middle and long bond are generating by mixing those two bonds. Intuitively, mixing of two bonds will lead to expected returns that are linear in maturity, whereas adding a third bond helps generate curvature.

Quality of the approximation

We now provide some evidence that the approximation of a zero-coupon bond by a portfolio of spanning bonds is decent for our purposes. There are two dimensions along which we would like to obtain a good approximation. First, we would like the value of the approximating portfolio to be the same as the value of the zero-coupon bond. This is relevant for measuring the supply of bonds:

below we will take existing measures of the quantity of zero-coupon-bonds held by households and convert them into portfolios of the small number of spanning bonds that are tradable by agents in our model. Along this dimension, the approximation is essentially as good as the term structure model itself. For a replicating portfolio defined by (21), the portfolio value $e^{-i_t}\theta_1 + \hat{P}_t'\hat{\theta}$ differs from the bond value $P_t^{(n)}$ only to the extent that the term structure model does not fit bonds of maturity n. The additional approximation error introduced by the matching procedure is less than .0001 basis points.

Second, we would like the conditional distribution of bond returns to be the same as that of the portfolio return. This is important because we would like agents in our model to have bond investment opportunity sets that are similar to those of actual households who trade bonds of many more maturities. Figure 4 gives an idea about the goodness of the approximation (21) by comparing statistics of the actual return implied by the term structure model and the approximating return. All statistics are unconditional moments computed from our sample, using the realizations of the term structure factors. For example, to obtain the difference in means in the top panel, we compute (i) quarterly returns from the term structure model using the formula $\exp(A_n + B_n' f_t)$ for prices, and (ii) quarterly returns based on the approximate formula for price changes (20) and subtract the mean of (i) from the mean of (ii).

On average, the two return distributions are quite similar for all maturities. The mean returns differ by less than 10 basis points for all bonds shorter than 30 years. The difference in variance is at most 30 basis points. The approximation error increases with maturity, as do the mean return and the variance of returns themselves. The approximate mean return on bonds is always within 5% of the true mean return, while the approximate variance is within 5-15% of the true variance. Larger errors tend to arise for longer bonds. In addition to the univariate distribution of a return, it is also of interest how it covaries with other returns. If the term structure model is correct, then a parsimonious way to check this is to consider the correlation with the two factor innovations. The lower panel of the figure reports the difference between the correlation coefficients of the true and approximate returns with the factor innovations. These differences are very small.

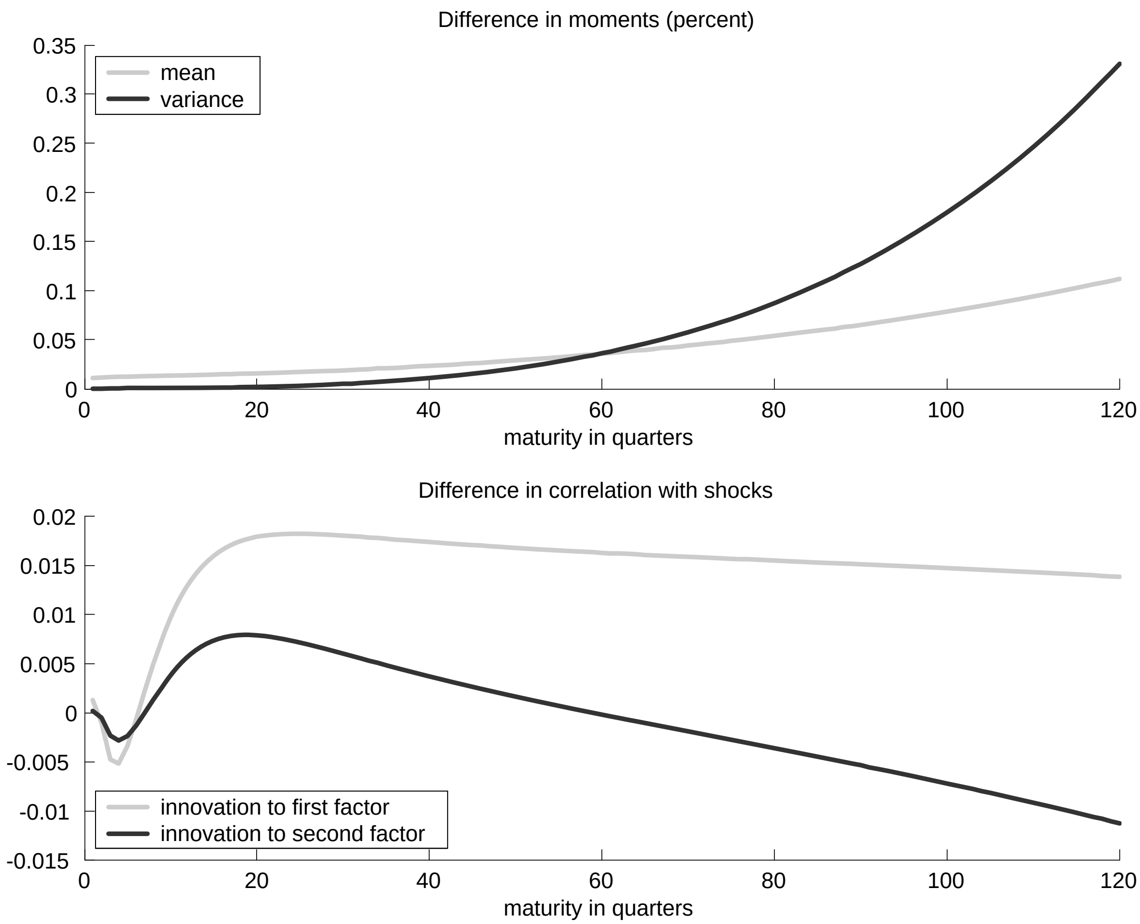

Figure 4: Approximation errors

B. Replicating nominal instruments in the US economy

We now turn to more complicated fixed-income instruments. The Flow-of-Funds (FFA) provides data on book value for many different types of nominal instruments. Doepke and Schneider (2006; DS) sort these instrument into several broad classes, and then use data on interest rates, maturities and contract structure to construct, for every asset class and every date t, a certain net payment stream that the holders of the asset expect to receive in the future. Their procedure takes into account credit risk in instruments such as corporate bonds and mortgages. They use these payment streams to restate FFA positions at market value and assess the effect of changes in inflation expectations on wealth.

Here we determine, for every broad asset class, a replicating portfolio that consists of spanning bonds. For every asset class and every date t, DS provide a certain payment stream, which we can view as a portfolio of zero-coupon bonds. By applying equation (21) to every zero-coupon bond, and then summing up the resulting replicating portfolios across maturities, we obtain a replicating portfolio for the asset class at date t. Figure 5 illustrates replicating portfolios for Treasury bonds and mortgages. The top panel shows how the weights on the spanning bonds in the replicating portfolio for Treasury bonds have changed over the postwar period.

The reduction of government debt after the war went along with a shortening of maturities: the weight on the longest bond declined from over 60% in 1952 to less than 20% in 1980. This development has been somewhat reversed since 1980.[8] The bottom panel shows that the effective maturity composition of mortgages was very stable before the 1980s, with a high weight on long bonds. The changes that apparent since the 1980s are driven by the increased use of adjustable rate mortgages.

We do not show replicating portfolios for Treasury bills, municipal bonds and corporate bonds, since the portfolio weights exhibit few interesting changes over time. All three instruments are represented by essentially constant portfolios of only two bonds: T-bills correspond to about 80% short bonds and 20% middle bonds, that for corporate bonds corresponds to about 60% middle bonds and 40% long bonds, and the replicating portfolio for municipal bonds has 70% long bonds and 30% middle bonds. The final asset class is a mopup group of short instruments, which we replicate by a short bond.

Replicating aggregate FFA household sector positions

We compute measure aggregate household holdings in the FFA at date t. To derive their positions in spanning bonds, we compute replicating portfolios for household positions in the FFA. One important issue is how to deal with indirect bond positions, such as bonds held in a pension plan or bonds held by a mutual fund, the shares in which are owned by the household sector. Here we make use of the calculations in DS who consolidate investment intermediaries in the FFA to arrive at effective bond positions.

Applying the replicating portfolios for the broad asset classes to FFA household sector positions

[8] The figure shows only the portfolios corresponding to outstanding Treasury *bonds*, not including bills. DS use data from the CRSP Treasury data base to construct a separate series for bills.

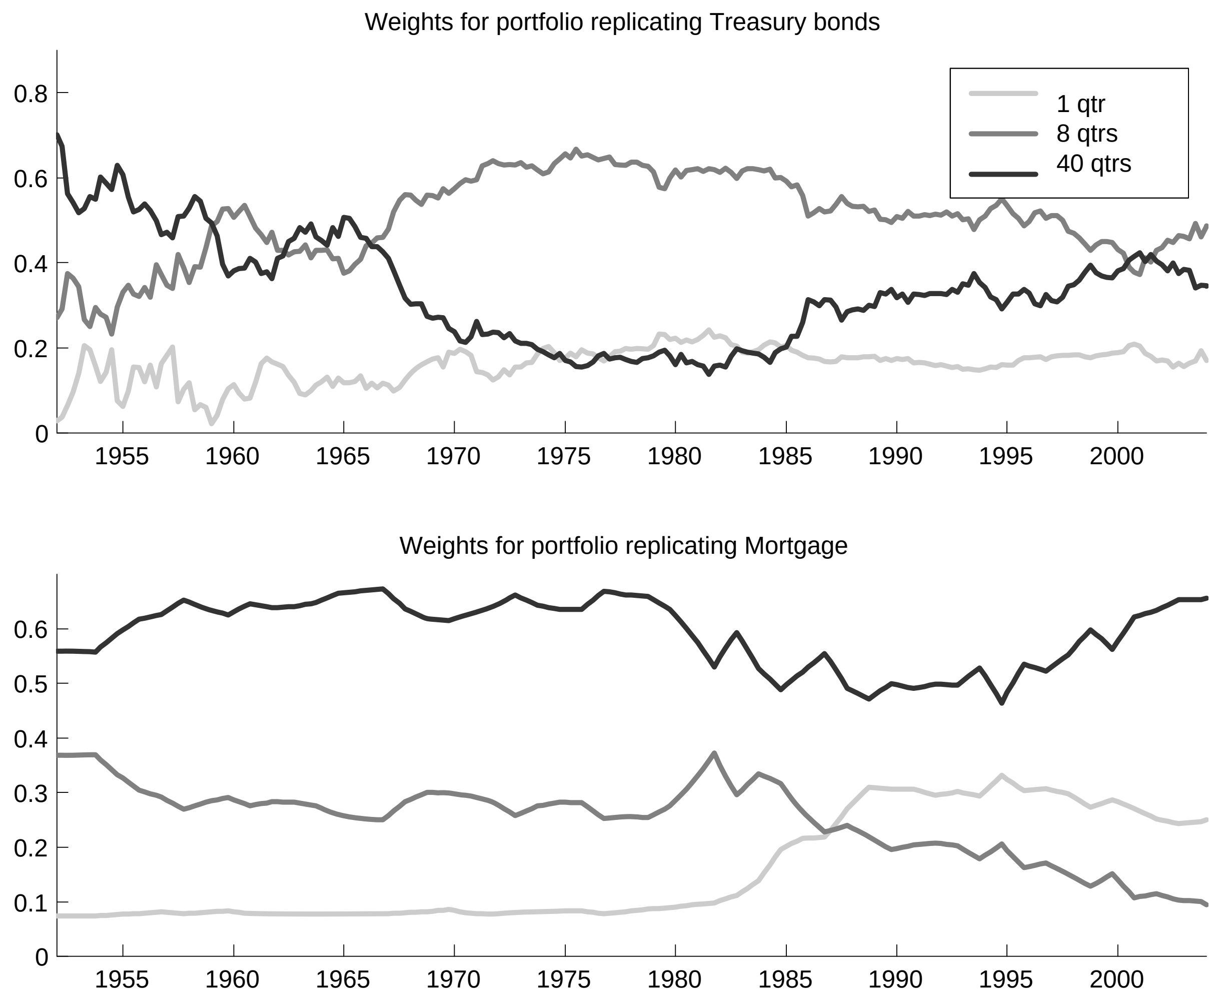

Figure 5: Replicating weights for Treasury bonds and mortgages

delivers three time series for holdings in spanning bonds, which are plotted in Figure 6. It is apparent that the early 1980s brought about dramatic changes in US bond portfolios. Until then, the positions in short bonds had been trending slightly upwards, whereas the positions in long and middle bonds had been declining. This pattern was reversed during the 1980s and early 1990s. The increase in the share of long bonds was partly due to changes in the composition of Treasury debt, as seen in Figure 5.

VI Interest Rates

In this section, we report numerical results for model implied yields for the rentier model.

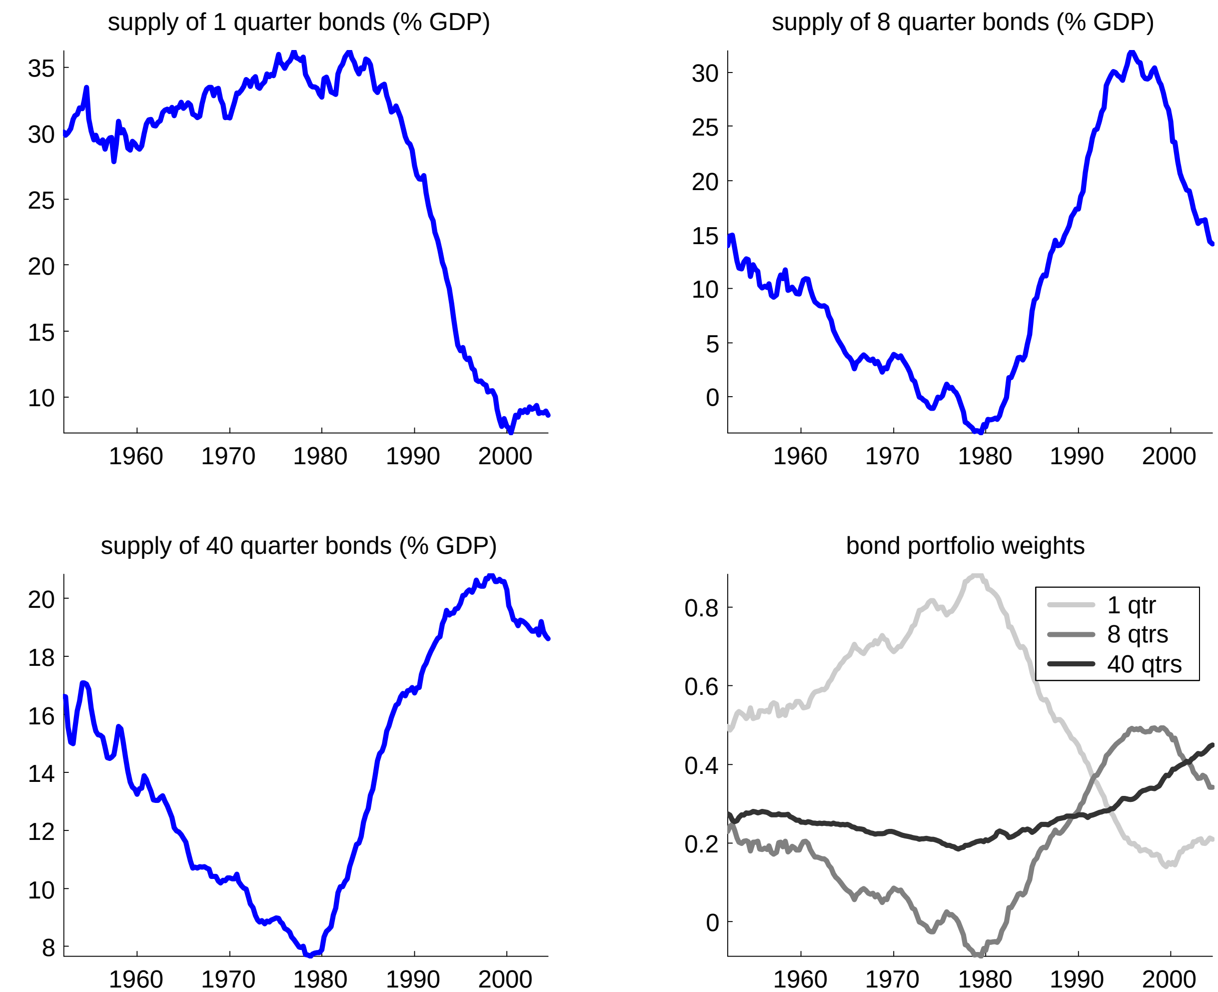

Figure 6: US household sector positions in the bond market

A. The rentier model

The starting point for the rentier exercise is a series of bond portfolio holdings. We consider an investor who invests only in bonds, and whose holdings are proportional to the given series of holdings. Since we do not observe the consumption of these bond investors, we do not allow for a residual asset, take the short rate as exogenous and only derive model-implied long bond prices. The idea behind the exercise is to explore what risk premia look like if there is a subset of investors for whom bonds are very important. As for the series of holdings, we have used US households' holdings of Treasury bonds, their holdings of all nominal assets except deposits, as well as US households' net position of nominal assets. The main lessons have been similar for all cases. In what follows, we report only the results for the total net position.

Formally, the input to the rentier exercise is a sequence of portfolio weights $\hat{\alpha}_t$ that represents the class of bonds considered in terms of our spanning bonds, a sequence of observed short-term interest rates, as well as a sequence of beliefs G_t about future interest rates and inflation that is needed to solve the portfolio choice problem. The sequence of beliefs G_t about future interest rates and inflation is a sequence of conditionals from the subjective probability distribution over yields and inflation estimated in subsection B..

Returns and preference parameters

The subjective state space system (16)-(19) together with the formulas for yields from the term structure model implies a VAR in (i) the real return on the short bond, (ii) excess returns on the middle and long bonds and (iii) the state variables z_t. This VAR is the basis for our portfolio choice problem, and allows us to use the approximation method of Campbell et al., as discussed in Section II. Given a solution for the portfolio demand, we then use equation (7) to derive a sequence of model-implied long bond prices.

Table 4: Subjective Moments of Excess Returns (% p.a.)

excess returns (%)	middle bond	long bond
mean	.64	.48
cond. standard deviation	2.68	12.72
cond. correlation matrix	1	.83
	.	1
correlation with short bond	-.17	.07

The short bond has a subjective mean real return 1.76% and one-quarter-ahead conditional volatility 0.56% in annualized percentage terms. Table 4 summarizes relevant moments of the subjective distribution of quarterly excess returns implied by the model for the middle and long spanning bonds (which 2 and 10 year maturities.) In terms of conditional Sharpe ratios (mean excess return divided by standard deviation), the middle bonds looks subjectively more attractive than the long bond. It is also notable that the long and middle bond returns are highly positively correlated. No longer bond covaries much with the short bond return.

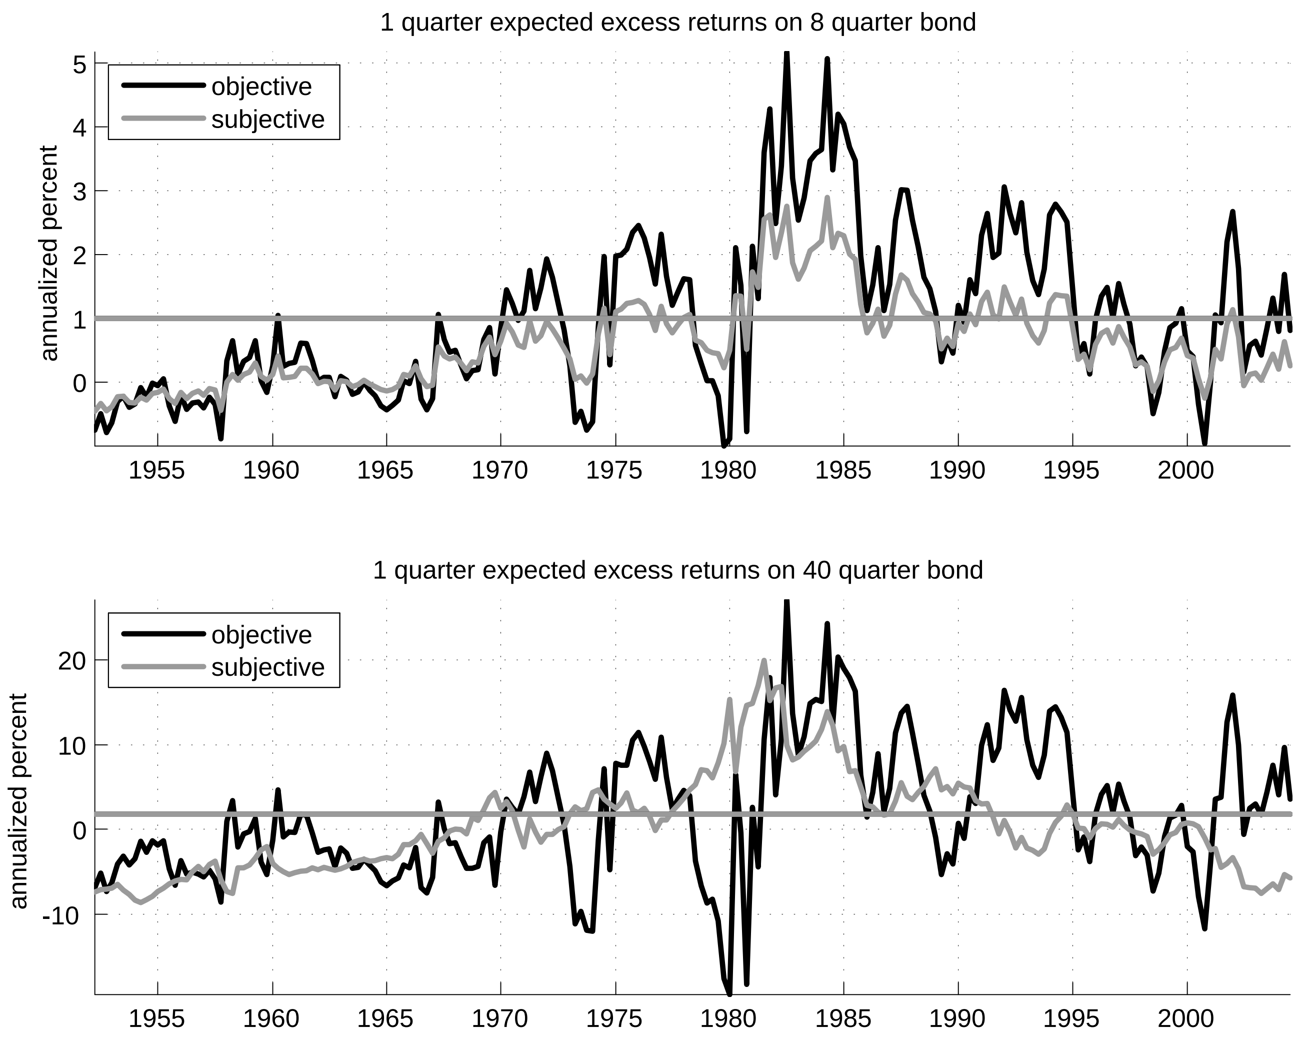

Figure 7: Expected excess holding period returns over one quarter, for the 2 year bond (top panel) and the 10 year bond (bottom panel), under the objective anmd subjective models, in annualized percent.

There is a role for market timing because of the predictability of excess long bond returns. This is illustrated in Figure 7. The light gray lines in both panels are the subjective expected excess returns under the subjective model. They exhibit both a low frequency movement – long bonds were less attractive in the beginning of the sample and again recently – and a business cycle component – expected excess returns tend to tick up during and after recessions. For comparison, the figure also plots expected excess returns under the objective model. Interestingly, business cycle swings in expected excess returns are much smaller under the subjective model. This reflects the difference in market prices of risk for the two models discussed above. Under the subjective model, the yield is more persistent, and does not drive risk premia to the extent it does under the objective model. As a result, investors perceive less opportunities for bond market timing than if their belief was given by the objective.

Mean term premia and yield volatility

Table 5 reports means and standard deviations for three sets of parameters. As a benchmark, we include a log investor ($\gamma = 1$). In this case, the mean long rate is matched up to 4 basis points, while the mean short rate is matched up to 21 basis points. By equation (7), the log investor model shuts down any time variation in subjective risk premia due to changes in changes in future investment opportunities. In addition, risk aversion is small, so that subjective consumption risk premia measured via bond portfolio weights are essentially zero. The model nevertheless generates positive term premia, because prices reflect investors' subjective expectations of future payoffs. For the same reason, model implied yield are almost as volatile as those in the data.

Table 5: Moments of Equilibrium Yields

nominal yields (% p.a.)		short bond	middle bond	long bond
Data	mean	5.27	5.90	6.53
	std. dev.	1.47	1.45	1.33
Log utility	mean		5.69	6.57
	std. dev.		1.37	1.28
$\gamma = 2$, $\beta = .97$	mean		5.67	6.53
	std. dev.		1.34	1.25
$\gamma = 20$, $\beta = .8$	mean		5.67	6.53
	std. dev.		1.32	1.23

We have searched over a grid of values for β and γ in order to find values that match exactly the mean excess returns on the two yields. It turns out that such a pair of values does not exist for $\gamma < 200$. This is somewhat surprising, since the parameters β and γ do have different effects on portfolio demand. In particular, risk aversion γ lowers the myopic demand for bonds, but increases the hedging demand for assets that provide insurance against bad future investment opportunities.

In contrast, the discount factor affects only the size of the hedging demand. However, it appears that the myopic and hedging demands for the two bonds react to both parameters in a similar way, so that they cannot be identified from the two average premia. In fact, the mean spread between the long and middle bond reacts very little to changes in β and γ. The reason is that it is due in part to systematic differences in payoff expectations between the two bonds (induced by the subjective model), which affect prices independently of the risk premium terms that are responsive to γ and β.

We report results for two other values of the risk aversion coefficient, in each case choosing the discount factor to match the mean long rate. Our leading example, to be explored further below, is the case $\gamma = 2$, a common choice for risk aversion in macro models; matching the long rate then requires $\beta = .97$. The resulting summary statistics are quite close to the log investor case. Relative to the log case, the mean long rate falls, even though risk aversion has gone up. This is because a patient investor can hold the long bond to hedge against bad investment opportunities, that is, low future short interest rates. This reduces the required compensation for risk. The table also provides a high risk aversion example: for $\gamma = 20$, we must pick $\beta = .8$. As risk aversion increases, the desire to hedge becomes stronger which by itself would increase the demand for the long bond and lower the long rate. The reduction in β weakens the hedging motive to keep the mean long rate at its observed value. This tradeoff says that if high risk aversion of rentiers is to be consistent with yield spreads in the data, then the effective planning horizon of the rentier must be shorter than that of the typical RBC agent.

Model-implied yield spreads

Figure 8 plots yield spreads from the data together with yield spreads implied by the model for the case $\gamma = 2$, $\beta = .97$. The middle panel shows that the model matches the 10-year spread quite well. The top and bottom panel show the two-year spread, as well as the spread of the ten year rate over the two year rate, respectively. Both panels reflects the fact that the mean 2-year rate is lower in the model than in the data. At the same time, the changes in model-implied and observed spreads track each other rather closely. The dynamics of spreads is driven in part by expectations of future yields. It is interesting to ask whether the model implied yields exhibit the

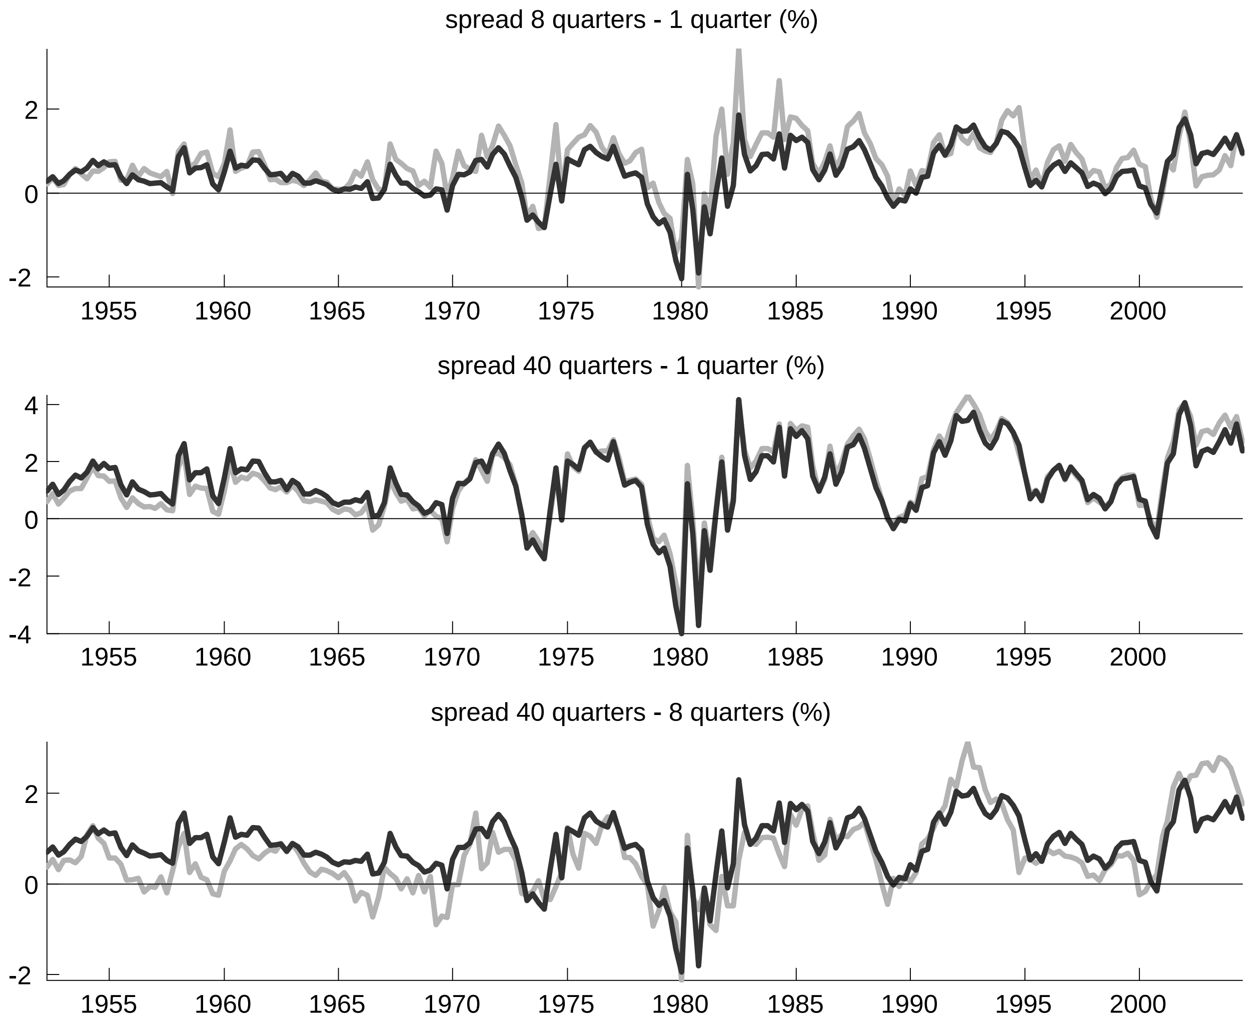

Figure 8: Yield spreads (as % p.a.) for the data sample (light lines) and the model-implied sample (dark lines).

same movements relative to statistical expected future yields as do their counterparts in the data. This is done in Figure 9, which reports risk premia relative to the expectations hypothesis for the 10 year yield.

All three lines in the figure represent the difference between a 10-year yield and its expectations hypothesis counterpart from the first line of (7):

$$i_t^{(n)} - \frac{1}{n} i_t - \frac{(n-1)}{n} E_t i_{t+1}^{(n-1)} - \frac{1}{2} var_t \left(x_{t+1}^{(n)} \right),$$

where $x_{t+1}^{(n)}$ is the 1-quarter excess holding period return on the n-period bond. In other words, the lines show the difference between the forward rate on a contract that promises an $(n-1)$-period bond one period from date t. The three lines differ in what forward rate is used, and how the expectation is formed. The light gray line labelled "data – obj. EH" shows the difference between the forward rate from the data and the expected rate under the objective probability. It is thus proportional to measured expected excess returns, which tend to be high during and after recessions.

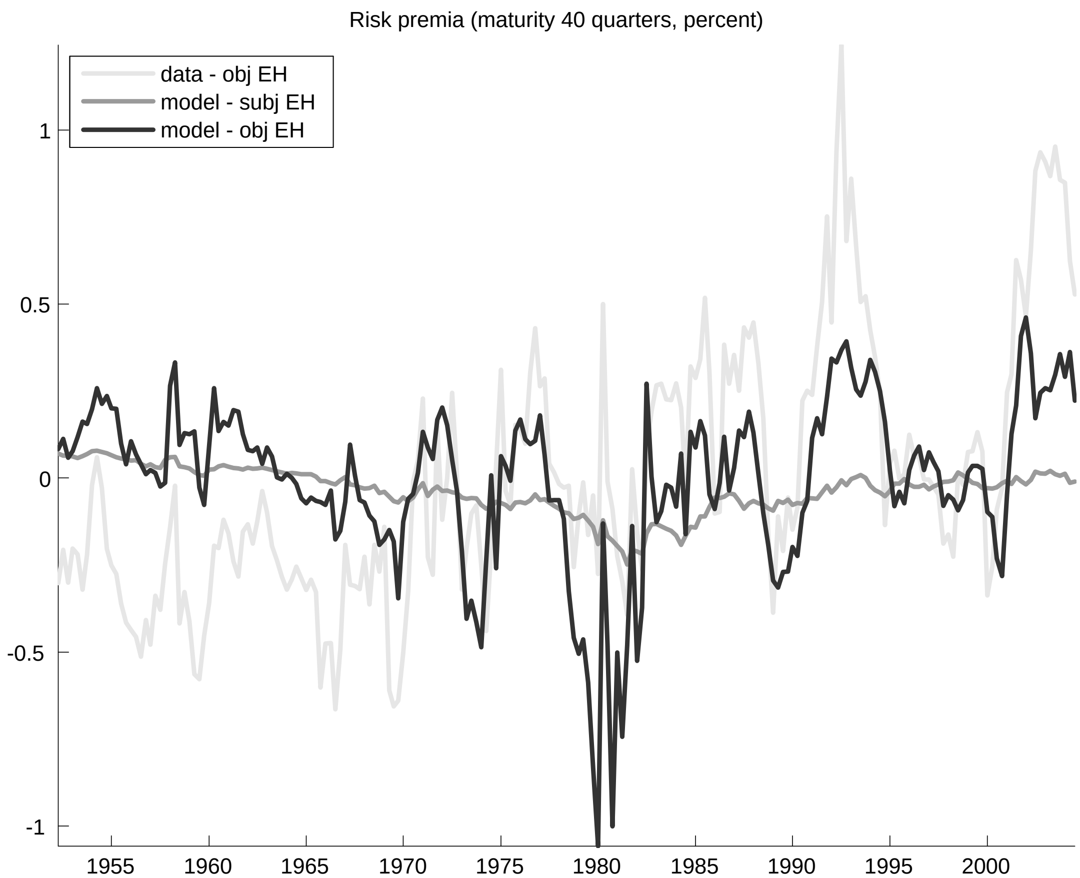

Figure 9: Risk premia (actual yield minus yield predicted under the expectations hypothesis) for (i) the data sample using predictions from the objective modol (light gray), (ii) the model implied sample using predictions from the subjective model (black), and (iii) the model implied sample using predicitions from the subjective model.

The black line represents a model implied objective risk premium – the difference between the forward rate implied by the model and the expected rate under the objective probability. It exhibits both a low frequency component and a business cycle component that comoves with the risk premium from the data. To illustrate the source of these movements, the dark gray line shows the subjective risk premium that is, the difference between the forward rate implied by the model and the expected rate under the subjective probability. The subjective premium looks like a smoothed version of the objective premium. It follows that the business cycle frequency movements in the objective risk premium have little to do with the subjective risk premium as perceived by investors. Instead, those movements are due to the differences between subjective and objective forecasts documented in Sections III.

Investors' subjective perception of risk is reflected in the subjective risk premium, and is also responsible for the low frequency components in the model implied objective risk premium. By equation (7), changes in the subjective premium must be due to changes in hedging opportunities or changes in bond positions. At the current parameter values, risk aversion is low and changes in positions do not have large effects. Instead, the V-shaped pattern in the subjective risk premium is due to changes in the covariance between long bond returns and future continuation utility. Intuitively, the insurance that long bonds provide to long horizon investors against drops in future short rates became more valuable in the 1980s when short rates were high. As a result, the compensation required for holding such bonds declined.

To explore the role of changes in bond positions, the dark gray line in Figure 10 shows the parts $\gamma\Sigma_{\hat{x}\hat{x}}\widehat{B}_t/B_t$ of the model-implied subjective risk premia for the 2-year and 10-year bonds that are due to changes in positions. This compensation for risk is small and also moves slowly, at similar frequencies as the portfolio weights derived in Section 6. For comparison, the figure also shows, at different scales, low frequency components of the respective term spreads that have been extracted using a band-pass filter. The increase in the spread in the 1980s thus went along with a shift towards longer bond positions that contributed positively to risk premia. However, the drop in subjective risk premia before 1980s did not go along with movements in the spread.

The quantitative lessons of this section are confirmed when we start the exercise from other classes of bond portfolios. Matching the long interest rate under the subjective model point us to small risk aversion coefficients, which in turn make the contribution of subjective risk premia, and especially the part due to bond positions, is quantitatively small. A qualitative difference is that when we eliminate deposits and liabilities from the portfolio, the middle and long bonds become relatively more important, and there is some evidence of business cycle variation in the model implied subjective risk premia. This suggests that a different model of conditional variance might lead to a larger role for subjective risk premia measured via bond positions.

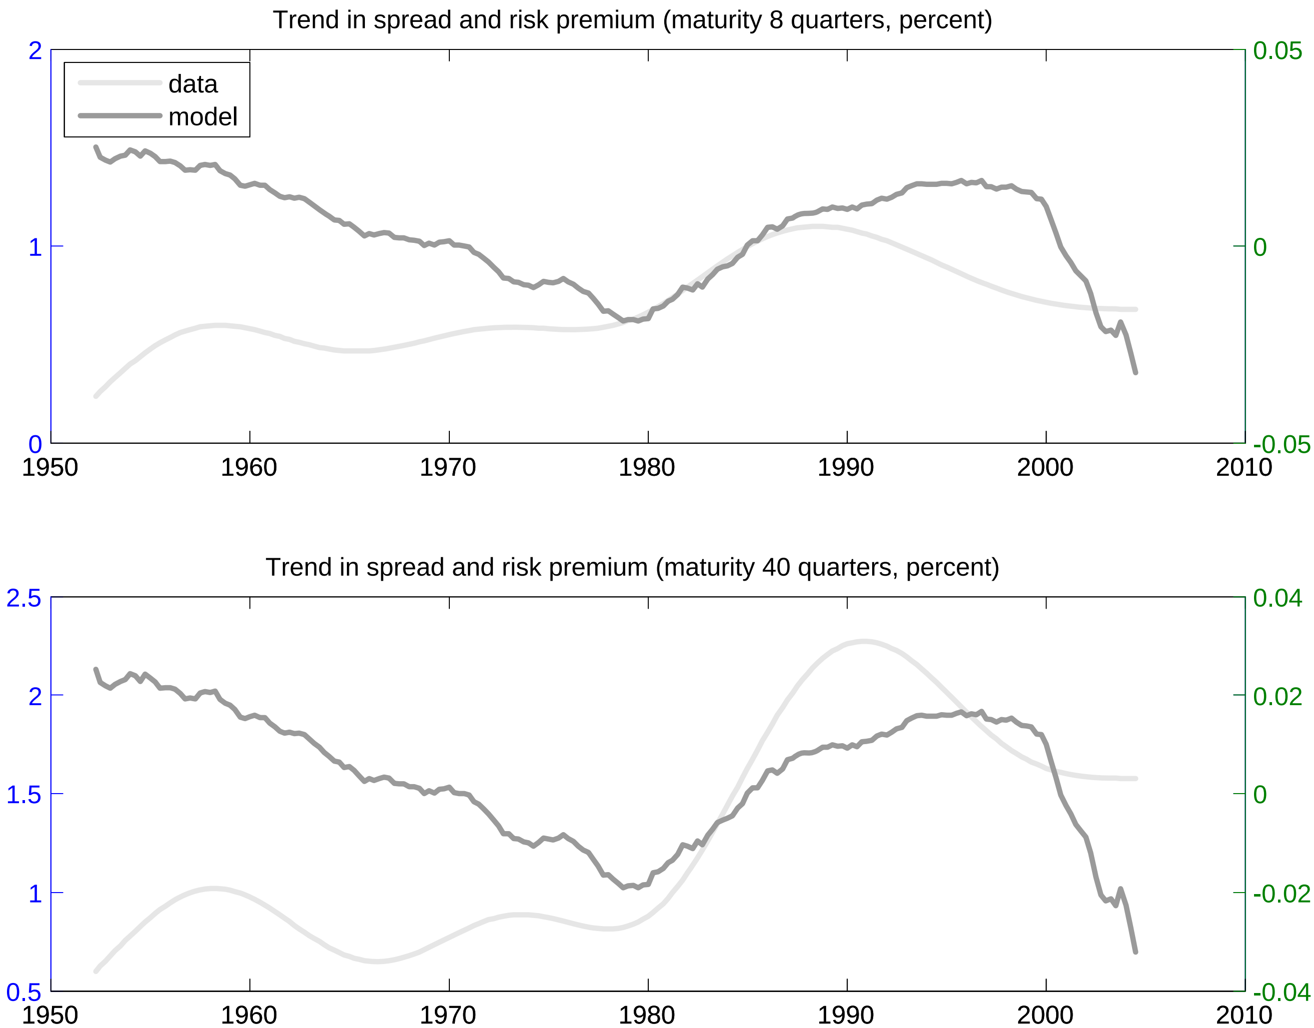

Figure 10: Trend in the term spread constructed with bandpass filter (light line, left scale) and model-implied risk premium (dark line, right scale).

References

Backus, D., Gregory, A., Zin, S. 1989. Risk Premiums in the Term Structure: Evidence from Artificial Economies. *Journal of Monetary Economics* 24, 371-399.

Bansal, R., A. Yaron 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481-1509.

Campbell, J.Y., Chan Y.L., Viceira, L.M. 2003. A Multivariate Model of Strategic Asset Allocation. *Journal of Financial Economics* 67(1), 41-80.

Campbell, J.Y., Viceira, L.M. 2001. Who should buy long-term bonds? *American Economic Review* 91(1), 99-127.

Chun, A.L. 2006. Expectations, Bond Yields and Monetary Policy. Manuscript, University of Montreal.

Cochrane, J.H., M. Piazzesi 2005. Bond Risk Premia. *American Economic Review* 95:1, 138-160.

Dai, Q., K.J. Singleton 2002. Expectation Puzzles, Time-varying Risk Premia, and Affine Models of the Term Structure *Journal of Financial Economics* 63, 415-441.

Dunn, K., K.J. Singleton, 1986. Modeling the Term Structure of Interest Rates Under Non-separable Utility and Durability of Goods. *Journal of Financial Economics* 17, 27-55.

Doepke, M., Schneider M. 2006. Inflation and the Redistribution of Nominal Wealth. *Journal of Political Economy* 114(6), 1069-1097.

Epstein, L., Zin, S. 1989. Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57, 937-69

Fama, E.F., Bliss, R.R.. The Information in Long-Maturity Forward Rates. *American Economic Review*, 1987, (77), 680-92.

Greenwood, J. and D. Vayanos 2007, Bond Supply and Expected Excess Returns. Manuscript, HBS and LSE.

Gürkaynak, R.S., Sack, B., Wright, J.H., 2006. The U.S. Treasury Yield Curve: 1961 to the Present. Manuscript, Federal Reserve Board of Governors.

Hansen, L.P., J.C. Heaton, and N. Li 2005. Consumption strikes back? Measuring Long Run Risk. Manuscript, University of Chicago.

Kim, D., A. Orphanides 2007. Term Structure Estimation with Survey Data on Interest Rate Forecasts. Manuscript, Federal Reserve Board.

Krishnamurthy, A., Vissing-Jorgensen, A. 2007. The Demand for Treasury Debt. *NBER Working Paper* #12881.

Lucas, R.E. Jr. 1978. Asset Prices in an Exchange Economy. *Econometrica* 46, 1429-1445.

Piazzesi, M., Schneider, M. 2006. Equilibrium Yield Curves. *NBER Macroannual 2006.*

Wachter, J. 2006. A Consumption-Based Model of the Term Structure of Interest Rates. *Journal of Financial Economics* 79, 365-399.

A Appendix

The goal of this appendix is to describe an equilibrium model of the US economy. The equations about optimal behavior and model-implied prices described in the body of the paper (for the case with a residual asset) hold in this equilibrium model. The model goes beyond these equations because it (i) derives initial wealth of households as the (endogenous) value of exogenous quantities of asset endowments and (ii) and specifies trades between US households and other sectors of the economy, such as foreigners. The households in the model optimize for some given exogenous expectations. This specification allows expectations to be consistent with survey evidence or with some learning mechanism.

The model describes a single trading period t. In each period, the US household sector trades assets with another sector of the economy. We call the other sector the "rest of the economy" (ROE), which stands in for the government, business, and foreign sectors. Households enter the trading period t endowed with one unit of the residual asset. Their initial wealth $\overline{W}_t$ also comprises dividends from the residual asset as well as the value of all bonds written on the rest of the economy that households bought at date $t-1$, denoted $\bar{B}_t$:

$$\overline{W}_t =: P_t^{res} + D_t + \bar{B}_t. \tag{A-1}$$

Households decide how to split this initial wealth into consumption as well as investment in the $N+2$ assets. More specifically, the household problem at date t is to maximize utility (1) subject to the budget constraint (2), given initial wealth $\bar{W}_t$ as well as beliefs about (a conditional distribution of) the relevant future price variables $\left(x_\tau^{res}, i_\tau, \pi_\tau, \hat{p}_\tau, \hat{p}_\tau^{+1}\right)_{\tau>t}$: the excess return on the residual asset, the short interest rate, the inflation rate and the prices and payoffs on the long spanning bonds.

The trading strategy of the rest of the economy is exogenous. In particular, at date t the ROE sector sells residual assets worth $P^{res} f_t$. This trade captures, for example, the construction of new houses, and the net issuance of new equity. The ROE also trades in the bond market. It is convenient to summarize bond trades in terms of the value of outstanding bonds written on the

ROE. At date t, the ROE redeems all short bonds issued at date $t-1$, and it also buys back all outstanding long bonds, at a total cost of $\bar{B}$. Moreover, the ROE issues new bonds worth B_t, so that its net sale of bonds is $B_t - \bar{B}_t$, which could be positive or negative.

To define new issues of individual bonds, we collect the values of new long bonds in a vector $\widehat{B}_t$. The value of outstanding short bonds is then $B_t - \iota_t' \widehat{B}$, where ι is an N-vector of ones. The ROE trading strategy is thus set up so that the same set of $N+1$ types of bonds – namely the short bond and the long bonds – are held by households at the end of every trading period. For example, suppose that the only long bond is a zero-coupon bond with a maturity of n periods. Between dates $t-1$ and t, households can then hold short (1-period) and long (n-period) bonds. At date t, the ROE buys back all long bonds (which now have maturity $n-1$), and again issues new 1-period short and n-period long bonds, and so on.

Equilibrium

We solve for a sequence of temporary equilibria. For each trading date t, we take as given (i) the strategy of the rest of the economy, summarized by its asset trades $\left(P^{res} f_t, \bar{B}_t, B_t, \widehat{B}_t\right)$, (ii) dividends D_t on the residual asset earned by households and (iii) household expectations about (that is, the conditional distribution of) the relevant price variables $\left(r_\tau^{res}, i_\tau, \pi_\tau, \hat{p}_\tau, \hat{p}_\tau^{+1}\right)_{\tau>t}$, which comprise the return on the residual asset, the short interest rate, the inflation rate and the prices and payoffs on the long bonds for all future periods. We characterize equilibrium prices as functions of these three inputs by equating household asset demand to the net asset supply provided by the trading strategy of the ROE.

Formally, an equilibrium consists of sequences of short interest rates and (log) long bond prices $(i_t, \hat{p}_t)$ as well as optimal choices by households $(C_t, \alpha_t^{res}, \hat{\alpha}_t)$ such that, at every date t, all four asset markets clear:

$$\begin{aligned} \alpha_t^{res} \left(\overline{W}_t - C_t\right) &= P_t^{res} + P^{res} f_t; \\ \hat{\alpha}_t \left(\overline{W}_t - C_t\right) &= \widehat{B}_t; \\ \overline{W}_t - C_t &= B_t + P_t^{res} + P^{res} f_t \end{aligned} \tag{A-2}$$

Here the first equation clears the market for the residual asset, the second equation clears the markets for the long bonds and the last equation ensures that total savings equals the total value of outstanding assets, which implies that the market for short bonds also clears. This system of $N+2$ equations determines the $N+2$ asset prices $(P_t^{res}, i_t, \hat{p}_t)$. While the price of the residual asset P_t^{res} appears directly in (A-2), bond prices enter via the effect of bond returns on portfolio demand.

A sequence of temporary equilibria imposes weaker restrictions on allocations and prices than a standard rational expectations equilibrium. In particular, the definition above does not directly connect what happens at different trading periods. On the one hand, we do not require that the initial wealth of households is derived from its choices in the previous period. For example, the sequence $\left(\bar{B}_t\right)$ of payoffs from bonds bought earlier is an exogenous input to the model. On the other hand, we do not impose conditions relating return expectations at date t to model-implied realized (or expected) returns in future periods, as one would do when imposing rational expectations. At the same time, if there is a rational expectations equilibrium of our model that accounts for observed asset prices and household sector choices, then it also gives rise to a sequence of temporary equilibria.

The fact that our model allows for trades between the household sector and the rest of the economy distinguishes it from the endowment economies frequently studies in the asset pricing literature. In particular, our model accommodates nonzero personal savings. Combining (A-1) and the last equation in (A-2), we obtain the flow-of-funds identity

$$C_t + P^{res} f_t + \left(B_t - \bar{B}_t\right) = D_t. \tag{A-3}$$

The dividend on the residual asset D_t corresponds to personal income less net personal interest income. As a result, $D_t - C_t$ is personal savings less net interest. It consists of $P^{res} f_t$ – net purchases of all assets other than bonds – and $B_t - \bar{B}_t$ – net purchases of bonds less net interest. A typical endowment economy model of the type studied by Lucas (1978) instead assumes that bonds are in zero net supply and that household net wealth is a claim on future consumption, so that $C_t = D_t$ and $P^{res} f_t = 0$.