Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds

John Y. Campbell, Adi Sunderam, and Luis M. Viceira[1]

First draft: June 2007
This version: July 22, 2010

Abstract

The covariance between US Treasury bond returns and stock returns has moved considerably over time. While it was slightly positive on average in the period 1953–2009, it was unusually high in the early 1980's and negative in the 2000's, particularly in the downturns of 2001–2 and 2008–9. This paper specifies and estimates a model in which the nominal term structure of interest rates is driven by five state variables: the real interest rate, risk aversion, temporary and permanent components of expected inflation, and the "nominal-real covariance" of inflation and the real interest rate with the real economy. The last of these state variables enables the model to fit the changing covariance of bond and stock returns. Log bond yields and term premia are quadratic in these state variables, with term premia determined mainly by the product of risk aversion and the nominal-real covariance. The concavity of the yield curve—the level of intermediate-term bond yields, relative to the average of short- and long-term bond yields—is a good proxy for the level of term premia. The nominal-real covariance has declined since the early 1980's, driving down term premia.

1 Introduction

Are nominal government bonds risky investments, which investors must be rewarded to hold? Or are they safe investments, whose price movements are either inconsequential or even beneficial to investors as hedges against other risks? US Treasury bonds performed well as hedges during the financial crisis of 2008–9, but the opposite was true in the early 1980's. The purpose of this paper is to explore such changes over time in the risks of nominal government bonds.

To understand the phenomenon of interest, consider Figure 1, an update of a similar figure in Viceira (2010). The figure shows the history of the realized beta of 10-year nominal zero-coupon Treasury bonds on an aggregate stock index, calculated using a rolling three-month window of daily data. This beta can also be called the "realized CAPM beta", as its forecast value would be used to calculate the risk premium on Treasury bonds in the Capital Asset Pricing Model (CAPM) that is often used to price individual stocks.

Figure 1 displays considerable high-frequency variation, much of which is attributable to noise in the realized beta. But it also shows interesting low-frequency movements, with values close to zero in the mid-1960's and mid-1970's, much higher values averaging around 0.4 in the 1980's, a spike in the mid-1990's, and negative average values in the 2000's. During the two downturns of 2001–3 and 2008–9, the average realized beta of Treasury bonds was about -0.2. These movements are large enough to cause substantial changes in the risk premium on Treasury bonds that would be implied by the CAPM.

Nominal bond returns respond both to inflation and to real interest rates. A natural question is whether the pattern shown in Figure 1 is due to the changing beta of inflation with the stock market, or of real interest rates with the stock market. Figure 2 summarizes the comovement of inflation shocks with stock returns, using a rolling three-year window of quarterly data and a first-order quarterly vector autoregression for inflation, stock returns, and the three-month Treasury bill yield to calculate inflation shocks. Because inflation is associated with high bond yields and low bond returns, the figure shows the beta of realized deflation shocks (the negative of inflation shocks) which should move in the same manner as the bond return beta reported in Figure 1. Indeed, Figure 2 shows a similar history for the deflation beta as for the nominal bond beta.

1

There is also movement over time in the covariation of long-term real interest rates with the stock market. In the period since 1997, when long-term Treasury inflation-protected securities (TIPS) were first issued, Campbell, Shiller, and Viceira (2009) report that TIPS have had a predominantly negative beta with stocks. Like the nominal bond beta, the TIPS beta was particularly negative in the downturns of 2001–3 and 2008–9. This implies that to explain the time-varying risks of nominal bonds, one needs a model that allows changes over time in the covariances of both inflation and real interest rates with the real economy and the stock market.

In this paper we specify and estimate such a model. Our model allows the covariances of shocks to change over time and potentially switch sign. By specifying stochastic processes for the real interest rate, temporary and permanent components of expected inflation, investor risk aversion, and the covariance of inflation and the real interest rate with the real economy, we can solve for the complete term structure at each point in time and understand the way in which bond market risks have evolved. We find that the covariance of inflation and the real interest rate with the real economy is a key state variable whose movements account for the changing covariance of bonds with stocks and imply that bond risk premia have been much lower in recent years than they were in the early 1980's.

The organization of the paper is as follows. Section 2 reviews the related literature. Section 3 presents our model of the real and nominal term structures of interest rates. Section 4 describes our estimation method and presents parameter estimates and historical fitted values for the unobservable state variables of the model. Section 5 discusses the implications of the model for the shape of the yield curve and the movements of risk premia on nominal bonds. Section 6 concludes. An Appendix to this paper available online (Campbell, Sunderam, and Viceira 2010) presents details of the model solution and additional empirical results.

2   Literature Review

Nominal bond risks can be measured in a number of ways. A first approach is to measure the covariance of nominal bond returns with some measure of the marginal utility of investors. According to the Capital Asset Pricing Model (CAPM), for example, marginal utility can be summarized by the level of aggregate wealth. It follows that the risk of bonds can be measured by the covariance of bond returns with

returns on the market portfolio, often proxied by a broad stock index. Alternatively, the consumption CAPM implies that marginal utility can be summarized by the level of aggregate consumption, so the risk of bonds can be measured by the covariance of bond returns with aggregate consumption growth.

A second approach is to measure the risk premium on nominal bonds, either from average realized excess bond returns or from variables that predict excess bond returns such as the yield spread (Shiller, Campbell, and Schoenholtz 1983, Fama and Bliss 1987, Campbell and Shiller 1991) or a more general linear combination of forward rates (Stambaugh 1988, Cochrane and Piazzesi 2005). If the risk premium is large, then presumably investors regard bonds as risky. This approach can be combined with the first one by estimating an empirical multifactor model that describes the cross-section of both stock and bond returns (Fama and French 1993).

These approaches are appealing because they are straightforward and direct. However, the answers they give depend sensitively on the sample period that is used. The covariance of nominal bond returns with stock returns, for example, is extremely unstable over time and even switches sign (Li 2002, Guidolin and Timmermann 2006, Christiansen and Ranaldo 2007, David and Veronesi 2009, Baele, Bekaert, and Inghelbrecht 2010, Viceira 2010). In some periods, notably the late 1970's and early 1980's, bond and stock returns move closely together, implying that bonds have a high CAPM beta and are relatively risky. In other periods, notably the late 1990's and the 2000's, bond and stock returns are negatively correlated, implying that bonds have a negative beta and can be used to hedge shocks to aggregate wealth.

The average level of the nominal yield spread is also unstable over time as pointed out by Fama (2006) among others. An intriguing fact is that the movements in the average yield spread seem to line up to some degree with the movements in the CAPM beta of bonds. The average yield spread, like the CAPM beta of bonds, was lower in the 1960's and 1970's than in the 1980's and 1990's. Viceira (2010) shows that both the short-term nominal interest rate and the yield spread forecast the CAPM beta of bonds over the period 1962–2007. On the other hand, during the 2000's the CAPM beta of bonds was unusually low while the yield spread was fairly high on average.

A third approach to measuring the risks of nominal bonds is to decompose their returns into several components arising from different underlying shocks. Nominal bond returns are driven by movements in real interest rates, inflation expectations, and the risk premium on nominal bonds over short-term bills. The variances of these components, and their correlations with investor well-being, determine the overall

3

risk of nominal bonds. Campbell and Ammer (1993), for example, estimate that over the period 1952–1987, real interest rate shocks moved stocks and bonds in the same direction but had relatively low volatility; shocks to long-term expected inflation moved stocks and bonds in opposite directions; and shocks to risk premia again moved stocks and bonds in the same direction. The overall effect of these opposing forces was a relatively low correlation between stock and bond returns. However Campbell and Ammer assume that the underlying shocks have constant variances and correlations throughout their sample period, and so their approach fails to explain changes in covariances over time.[2]

Economic theory provides some guidance in modelling the risks of the underlying shocks to bond returns. First, consumption shocks raise real interest rates if consumption growth is positively autocorrelated (Campbell 1986, Piazzesi and Schneider 2006, Gollier 2007); in this case inflation-indexed bonds hedge consumption risk and should have negative risk premia. If the level of consumption is stationary around a trend, however, consumption growth is negatively autocorrelated, inflation-indexed bonds are exposed to consumption risk, and inflation-indexed bond premia should be positive.

Second, inflation shocks are positively correlated with economic growth if demand shocks move the macroeconomy up and down a stable Phillips Curve; but inflation is negatively correlated with economic growth if supply shocks move the Phillips Curve in and out. In the former case, nominal bonds hedge the risk that negative macroeconomic shocks will cause deflation, but in the latter case, they expose investors to the risk of stagflation.

Finally, shocks to risk premia move stocks and bonds in the same direction if bonds are risky, and in opposite directions if bonds are hedges against risk (Connolly, Stivers, and Sun 2005). These shocks may be correlated with shocks to consumption if investors' risk aversion moves with the state of the economy, as in models with habit formation (Campbell and Cochrane 1999).

The term structure model we report in the next section of the paper extends a number of recent term structure models. Dai and Singleton (2002), Bekaert, Engstrom, and Grenadier (2005), Wachter (2006), Buraschi and Jiltsov (2007), and Bekaert, Engstrom, and Xing (2009) specify term structure models in which risk aversion varies over time, influencing the shape of the yield curve. These papers take

---

[2]See also Barsky (1989) and Shiller and Beltratti (1992).

care to remain in the essentially affine class described by Duffee (2002). Bekaert et al. and other recent authors including Mamaysky (2002) and d'Addona and Kind (2006) extend affine term structure models to price stocks as well as bonds. Bansal and Shaliastovich (2010), Eraker (2008), and Hasseltoft (2008) also extend affine term structure models to price stocks and bonds in an economy with long-run consumption risk (Bansal and Yaron 2004). Piazzesi and Schneider (2006) and Rudebusch and Wu (2007) build affine models of the nominal term structure in which a deterministic reduction of inflation uncertainty drives down the risk premia on nominal bonds towards the lower risk premia on inflation-indexed bonds (which can even be negative, as discussed above).[3]

Our introduction of a time-varying covariance between state variables and the stochastic discount factor, which can switch sign, means that we cannot write log bond yields as affine functions of macroeconomic state variables; our model, like those of Beaglehole and Tenney (1991), Constantinides (1992), Ahn, Dittmar and Gallant (2002), and Realdon (2006), is linear-quadratic.[4] To solve our model, we use a general result on the expected value of the exponential of a non-central chi-squared distribution which we take from the Appendix to Campbell, Chan, and Viceira (2003). To estimate the model, we use a nonlinear filtering technique, the unscented Kalman filter, proposed by Julier and Uhlmann (1997), reviewed by Wan and van der Merwe (2001), and recently applied in finance by Binsbergen and Koijen (2008).

3   A Quadratic Bond Pricing Model

We start by formulating a model which, in the spirit of Campbell and Viceira (2001, 2002), describes the term structure of both real interest rates and nominal interest rates. The innovation here is that our model allows for time variation in the covariances between real interest rates, inflation, and the real economy. This results in a term structure where both real and nominal bond yields are linear-quadratic functions

---

[3]In a similar spirit, Backus and Wright (2007) argue that declining uncertainty about inflation explains the low yields on nominal Treasury bonds in the mid-2000's, a phenomenon identified as a "conundrum" by Alan Greenspan in 2005 Congressional testimony.

[4]Duffie and Kan (1996) point out that linear-quadratic models can often be rewritten as affine models if we allow the state variables to be bond yields rather than macroeconomic fundamentals. Buraschi, Cieslak, and Trojani (2008) also expand the state space to obtain an affine model in which correlations can switch sign.

of the vector of state variables and where, consistent with the empirical evidence, the conditional volatilities and covariances of excess returns on real and nominal assets are time varying.

3.1   The SDF and the short-term real interest rate

We start by assuming that the log of the real stochastic discount factor (SDF), $m_{t+1} = \log(M_{t+1})$, follows a linear-quadratic, conditionally heteroskedastic process:

$$-m_{t+1} = x_t + \frac{\sigma_m^2}{2} z_t^2 + z_t \varepsilon_{m,t+1}, \tag{1}$$

whose drift x_t follows an AR(1) process subject to a heteroskedastic shock and a homoskedastic shock,

$$x_{t+1} = \mu_x (1 - \phi_x) + \phi_x x_t + \psi_t \varepsilon_{x,t+1} + \varepsilon_{X,t+1}. \tag{2}$$

It is straightforward to show that the state variable x_t is the short-term log real interest rate. The price of a single-period zero-coupon real bond satisfies

$$P_{1,t} = E_t \left[\exp\{m_{t+1}\} \right],$$

so that its yield $y_{1t} = -\log(P_{1,t})$ equals

$$y_{1t} = -E_t \left[m_{t+1} \right] - \frac{1}{2} \operatorname{Var}_t (m_{t+1}) = x_t. \tag{3}$$

Thus the AR(1) process (2) describes the dynamics of the short-term real interest rate.

The model has two additional state variables, z_t and ψ_t, which govern the time variation in the volatilities of the SDF and the real interest rate respectively. We assume that these state variables both follow standard homoskedastic AR(1) processes:

$$z_{t+1} = \mu_z (1 - \phi_z) + \phi_z z_t + \varepsilon_{z,t+1}, \tag{4}$$

$$\psi_{t+1} = \mu_\psi (1 - \phi_\psi) + \phi_\psi \psi_t + \varepsilon_{\psi,t+1}. \tag{5}$$

In principle both these variables can change sign, but in practice we expect z_t always to be positive (and with negligible effect on likelihood we constrain it to be constant

6

in the empirical results reported later in the paper), while ψ_t changes sign to capture the changing covariances that motivate our model.

The vector of innovations $(\varepsilon_{m,t+1},\ \varepsilon_{x,t+1},\ \varepsilon_{X,t+1},\ \varepsilon_{\psi,t+1},\ \varepsilon_{z,t+1})$ is normally distributed, with zero means and constant variance-covariance matrix. We allow these shocks to be cross-correlated and adopt the notation σ_i^2 to describe the variance of shock ε_i, and σ_{ij} to describe the covariance between shock ε_i and shock ε_j. In this model, σ_m always appears premultiplied by z_t in all pricing equations. This implies that we are unable to identify σ_m separately from z_t. Thus without loss of generality we set σ_m to an arbitrary value of one.[5] This normalization implies that the state variable z_t completely describes the conditional variance of the log real SDF. To reduce the complexity of the equations that follow, we also assume that the shocks to x_t are orthogonal to each other; that is, $\sigma_{xX} = 0$.

Our use of the state variable z_t to model time-varying volatility in the log real SDF, or equivalently time variation in the price of aggregate market risk or maximum Sharpe ratio in the economy, is similar to the approach of Lettau and Wachter (2007, 2010). We can interpret it as a reduced form of a structural model in which aggregate risk aversion changes exogenously over time as in the "moody investor" economy of Bekaert, Engstrom and Grenadier (2005). The model of Campbell and Cochrane (1999), in which movements of aggregate consumption relative to its past history cause temporary movements in risk aversion, is similar in spirit. Such structural models imply a real SDF similar to (1) in which risk aversion is a positive function of z_t. We can also interpret our model as a reduced form of the real SDF generated by the long-run consumption risk model of Bansal and Yaron (2004), in which z_t describes the conditional volatility of log consumption growth.[6] With the first interpretation of our model in mind, we use the terms price of risk or risk aversion interchangeably to refer to z_t.

The state variable ψ_t allows the covariance between the real interest rate and the SDF, and therefore the market price of real interest rate risk, to move over time and even switch sign. In an earlier version of this paper we assumed that the process

---

[5]The same is true with respect to σ_x and ψ_t. However, ψ_t also premultiplies other variables in the model, specifically realized inflation and expected inflation. We choose to normalize to one the volatility of the shocks to realized inflation.

[6]Under such an interpretation our real stochastic discount factor describes the intertemporal marginal rate of substitution of a representative investor with recursive Epstein-Zin preferences facing an exogenous consumption growth process. This process has a persistent drift described by x_t, and it is heteroskedastic, with conditional volatility z_t.

for the real interest rate in (2) was homoskedastic, writing a model in which ψ_t only affects inflation and nominal interest rates. This generates a simpler affine real term structure of interest rates, but is inconsistent with time-variation in the covariance between TIPS returns and the real economy documented by Campbell, Shiller, and Viceira (2009).

We allow for a homoskedastic shock, as well as a heteroskedastic shock, to move the real interest rate because with only a single shock the model would imply a constant Sharpe ratio for real bonds. With only a heteroskedastic shock, the model would also imply that the conditional volatility of the real interest rate would be proportional to the covariance between the real interest rate and the real SDF; equivalently, the conditional correlation of the real rate and the SDF would be constant in absolute value with occasional sign switches. Our specification avoids these implausible implications while remaining reasonably parsimonious.

3.2   The real term structure of interest rates

In the current model, the price of a n-period zero-coupon real bond is an exponential linear-quadratic function of x_t, z_t, and ψ_t. To understand this, consider the standard pricing equation for a two-period bond (Campbell, Lo, and MacKinlay 1997, Chapter 11):

$$P_{2,t} = \mathrm{E}_t\left[\exp\left\{p_{1,t+1} + m_{t+1}\right\}\right], \tag{6}$$

where $p_{n,t} \equiv \log(P_{n,t})$. Since $p_{1,t+1} = -x_{t+1}$, and x_{t+1} and m_{t+1} are jointly conditionally normal, we can write the expectation on the right-hand-side of (6) as

$$
\begin{aligned}
P_{2,t} &= \exp\left\{\mathrm{E}_t\left[-x_{t+1} + m_{t+1}\right] + \frac{1}{2}\mathrm{Var}_t\left(-x_{t+1} + m_{t+1}\right)\right\} \\
&= \exp\left\{-\mu_x\left(1 - \phi_x\right) - \left(1 + \phi_x\right)x_t + \frac{1}{2}\psi_t^2\sigma_x^2 + \frac{1}{2}\sigma_X^2 + \sigma_{xm}z_t\psi_t + \sigma_{Xm}z_t\right\} (7)
\end{aligned}
$$

which depends on x_t, z_t, ψ_t^2 and the product $z_t\psi_t$. Thus a two-period bond is an exponential linear quadratic function of the state variables.

Once we consider bonds with maturity $n > 2$, $P_{n-1,t+1}$ and M_{t+1} are no longer jointly lognormal because $P_{n-1,t+1}$ is an exponential-quadratic function of normally distributed variables. However, the Appendix shows that we can still derive a closed-

form solution for the price of the bond that takes the form

$$P_{n,t} = \exp\left\{A_n + B_{x,n}x_t + B_{z,n}z_t + B_{\psi,n}\psi_t + C_{z,n}z_t^2 + C_{\psi,n}\psi_t^2 + C_{z\psi,n}z_t\psi_t\right\}. \quad (8)$$

The coefficients A_n, $B_{i,n}$, and $C_{i,n}$ solve a set of recursive equations given in the Appendix. These coefficients are functions of the maturity of the bond (n) and the coefficients that determine the stochastic processes for the real SDF and state variables x_t, z_t, and ψ_t. From equation (3), it is immediate to see that $B_{x,1} = -1$, and that the remaining coefficients are zero at $n = 1$, and from equation (7) that $A_2 = -\mu_x(1 - \phi_x) + \sigma_X^2/2$, $B_{x,2} = -(1 + \phi_x)$, $B_{z,2} = \sigma_{Xm}$, $C_{\psi,2} = \sigma_x^2/2$, $C_{z\psi,2} = \sigma_{xm}$, and $B_{\psi,2} = C_{z,2} = 0$.

This model of the real term structure of interest rates generates time-varying real bond risk premia that depend on z_t and the product $z_t\psi_t$ (see Appendix). Once again, the 2-period bond is helpful to understand this result. The excess log return on a 2-period zero-coupon real bond over a 1-period real bond is given by

$$
\begin{aligned}
r_{2,t+1} - r_{1,t+1} &= p_{1,t+1} - p_{2,t} + p_{1,t} \\
&= -\frac{1}{2}\psi_t^2\sigma_x^2 - \frac{1}{2}\sigma_X^2 - \sigma_{xm}z_t\psi_t - \sigma_{Xm}z_t - \psi_t\varepsilon_{x,t+1} - \varepsilon_{X,t+1}, \quad (9)
\end{aligned}
$$

where the first two terms are a Jensen's inequality correction (the form of which depends on our simplifying assumption that $\sigma_{xX} = 0$), the middle two terms describe the log of the expected excess return on real bonds, and the last two terms describe shocks to realized excess returns.

It follows from (9) that the conditional risk premium on the 2-period real bond is

$$\mathrm{E}_t\left[r_{2,t+1} - r_{1,t+1}\right] + \frac{1}{2}\mathrm{Var}_t\left(r_{2,t+1} - r_{1,t+1}\right) = -\left(\sigma_{Xm} + \sigma_{xm}\psi_t\right)z_t, \quad (10)$$

which is proportional to z_t. The coefficient of proportionality can take either sign and varies over time with the state variable ψ_t. To gain intuition about the 2-period real bond risk premium, consider the simple case where $\sigma_{Xm} = 0$ and $\sigma_{xm}\psi_t > 0$. This implies that real bond risk premia are negative. The reason for this is that with positive $\sigma_{xm}\psi_t$, the real interest rate tends to rise in good times and fall in bad times. Since real bond returns move opposite the real interest rate, real bonds are countercyclical assets that hedge against economic downturns and command a negative risk premium.

3.3   Pricing equities

We want our model to fit the changing covariance of bonds and stocks, and so we must specify a process for the equity return within the model. One modelling strategy would be to specify a dividend process and solve for the stock return endogenously in the manner of Mamaysky (2002), Bekaert et al. (2005), and d'Addona and Kind (2006). However we adopt a simpler approach. Following Campbell and Viceira (2001), we model shocks to realized stock returns as a linear combination of shocks to the real interest rate and shocks to the log stochastic discount factor:

$$r_{e,t+1} - \mathrm{E}_t\, r_{e,t+1} = \beta_{ex}\varepsilon_{x,t+1} + \beta_{eX}\varepsilon_{X,t+1} + \beta_{em}\varepsilon_{m,t+1} + \varepsilon_{e,t+1}, \tag{11}$$

where $\varepsilon_{e,t+1}$ is an identically and independently distributed shock uncorrelated with all other shocks in the model. This shock captures movements in equity returns that are both unrelated to real interest rates and carry no risk premium because they are uncorrelated with the SDF.

Substituting (11) into the no-arbitrage condition $E_t\left[M_{t+1}R_{t+1}\right] = 1$, the Appendix shows that the conditional equity risk premium is given by

$$\mathrm{E}_t\left[r_{e,t+1} - r_{1,t+1}\right] + \frac{1}{2}\mathrm{Var}_t\left(r_{e,t+1} - r_{1,t+1}\right) = \left(\beta_{ex}\sigma_{xm} + \beta_{eX}\sigma_{Xm} + \beta_{em}\sigma_m^2\right)z_t. \tag{12}$$

The equity premium, like all risk premia in our model, is proportional to risk aversion z_t. It depends not only on the direct sensitivity of stock returns to the SDF, but also on the sensitivity of stock returns to the real interest rate and the covariance of the real interest rate with the SDF.

Equation (11) does not attempt to capture heteroskedasticity in stock returns. Although such heteroskedasticity is of first-order importance for understanding stock prices, we abstract from it here in order to maintain the parsimony of our term structure model.

3.4   Modelling inflation

To price nominal bonds, we need to specify a model for inflation. We assume that log inflation $\pi_t = \log(\Pi_t)$ follows a linear-quadratic conditionally heteroskedastic process:

$$\pi_{t+1} = \lambda_t + \xi_t + \frac{\sigma_\pi^2}{2}\psi_t^2 + \psi_t\varepsilon_{\pi,t+1}, \tag{13}$$

where ψ_t is given in (5) and expected log inflation is the sum of two components, a permanent component λ_t and a transitory component ξ_t.

The dynamics of the components of expected inflation are given by

$$\lambda_{t+1} = \lambda_t + \varepsilon_{\Lambda,t+1} + \psi_t \varepsilon_{\lambda,t+1}, \tag{14}$$

and

$$\xi_{t+1} = \phi_\xi \xi_t + \psi_t \varepsilon_{\xi,t+1}. \tag{15}$$

The presence of an integrated component in expected inflation removes the need to include a nonzero mean in the stationary component of expected inflation.

We assume that the underlying shocks to realized inflation, the components of expected inflation, and conditional inflation volatility—$\varepsilon_{\pi,t+1}$, $\varepsilon_{\lambda,t+1}$, $\varepsilon_{\Lambda,t+1}$, $\varepsilon_{\xi,t+1}$, and $\varepsilon_{\psi,t+1}$—are again jointly normally distributed zero-mean shocks with a constant variance-covariance matrix.[7] We allow these shocks to be cross-correlated with the shocks to m_{t+1}, x_{t+1}, and z_{t+1}, and use the same notation as in Section 3.1 to denote their variances and covariances.

Our inclusion of two components of expected inflation gives our model the flexibility it needs to fit simultaneously persistent shocks to both real interest rates and expected inflation. This flexibility is necessary because both realized inflation and the yields of long-dated inflation-indexed bonds move persistently, which suggests that both expected inflation and the real interest rate follow highly persistent processes. At the same time, short-term nominal interest rates exhibit more variability than long-term nominal interest rates, which suggests that a rapidly mean-reverting state variable must also drive the dynamics of nominal interest rates. By allowing for a permanent component and a transitory component in expected inflation, our model can capture parsimoniously the dynamics of short- and long-term nominal bond yields, realized inflation, and the yields on inflation-indexed bonds.[8]

---

[7]Without loss of generality we set σ_π to an arbitrary value of 1, for reasons similar to those we use to set σ_m to an arbitrary value of 1.

[8]It might be objected that in the very long run a unit-root process for expected inflation has unreasonable implications for inflation and nominal interest rates. Regime-switching models have been proposed as an alternative way to reconcile persistent fluctuations with stationary long-run behavior of interest rates (Garcia and Perron 1996, Gray 1996, Bansal and Zhou 2002, Ang, Bekaert, and Wei 2008). We do not pursue this idea further here, but in principle there is no reason why our model could not be rewritten using discrete regimes to capture persistent movements in expected inflation. As a robustness check, we have estimated though our model imposing that λ_t follows

Because the state variable ψ_t multiplies shocks in all the nominal equations, the conditional volatility of both inflation and expected inflation are both time varying. A large empirical literature in macroeconomics has documented changing volatility in inflation. In fact, the popular ARCH model of conditional heteroskedasticity (Engle 1982) was first applied to inflation. Our model captures this heteroskedasticity using the persistent state variable ψ_t which drives the volatility of expected as well as realized inflation.

The state variable ψ_t governs not only the second moments of realized inflation and expected inflation, but also the volatility of the real interest rate. We could assume different processes driving the second moments of realized and expected inflation and the real interest rate, but this would increase the complexity of the model considerably. Long-term nominal bond yields depend primarily on the persistent component of expected inflation; therefore the state variable that governs the second moments of this state variable is the most important one for the behavior of the nominal term structure. We keep our model parsimonious by assuming that the same state variable drives the second moments of transitory expected inflation, realized inflation, and the real interest rate. This is consistent with evidence that the volatility of returns on inflation indexed bonds is positively correlated with the volatility of returns on nominal bonds (Campbell, Shiller, and Viceira 2009).[9]

Since we model ψ_t as an AR(1) process, it can change sign. The sign of ψ_t does not affect the variances of expected or realized inflation, the covariance between them, or their covariance with the real interest rate, because these moments depend on the square ψ_t^2. However the sign of ψ_t does determine the sign of the covariance between expected and realized inflation, on the one hand, and the log real SDF, on the other hand. For this reason we will refer to ψ_t as the nominal-real covariance, although it also determines the covariance of the real interest rate with the real SDF and thus real bond risk premia.

We allow both a homoskedastic shock $\varepsilon_{\Lambda,t+1}$ and a heteroskedastic shock $\psi_t \varepsilon_{\lambda,t+1}$ to impact the permanent component of expected inflation. The reasons for this assumption are similar to those that lead us to assume two shocks for the real interest rate process. In the absence of a homoskedastic shock to expected inflation, the

---

a stationary process with a highly persistent autoregressive coefficient. In practice this makes no discernible changes to our main empirical conclusions.

[9]Although not reported in the article, the correlation in their data between the volatility of nominal US Treasury bond returns and the volatility of TIPS returns is slightly greater than 0.7.

conditional volatility of expected inflation would be proportional to the conditional covariance between expected inflation and real economic variables. There is no economic reason to expect that these two second moments should be proportional to one another, and the data suggest that the conditional covariance can be close to zero even when the conditional volatility remains positive. Put another way, the presence of two shocks allows the conditional correlation between real and nominal variables to vary smoothly rather than being fixed in absolute value with occasional sign switches. Since long-term expected inflation is the main determinant of long-term nominal interest rates, we allow two shocks to this process but for parsimony allow only heteroskedastic shocks to transitory expected and realized inflation.

Finally, we note that the process for realized inflation, equation (13), is formally similar to the process for the log SDF (1), in the sense that it includes a Jensen's inequality correction term. The inclusion of this term simplifies the process for the reciprocal of inflation by making the log of the conditional mean of $1/\Pi_{t+1}$ the negative of the sum of the two state variables λ_t and ξ_t. This in turn simplifies the pricing of short-term nominal bonds.

3.5 The short-term nominal interest rate

The real cash flow on a single-period nominal bond is simply $1/\Pi_{t+1}$. Thus the price of the bond is given by

$$P_{1,t}^{\$} = \mathrm{E}_t \left[\exp \left\{ m_{t+1} - \pi_{t+1} \right\} \right], \tag{16}$$

so the log short-term nominal rate $y_{1,t+1}^{\$} = -\log\left(P_{1,t}^{\$}\right)$ is

$$
\begin{aligned}
y_{1,t+1}^{\$} &= -\mathrm{E}_t \left[m_{t+1} - \pi_{t+1} \right] - \frac{1}{2} \mathrm{Var}_t \left(m_{t+1} - \pi_{t+1} \right) \\
&= x_t + \lambda_t + \xi_t - \sigma_{m\pi} z_t \psi_t,
\end{aligned}
\tag{17}
$$

where we have used the fact that $\exp\left\{m_{t+1} - \pi_{t+1}\right\}$ is conditionally lognormally distributed given our assumptions.

Equation (17) shows that the log of the nominal short rate is the sum of the log real interest rate, the two state variables that drive expected log inflation, and a nonlinear term that accounts for the correlation between shocks to inflation and shocks to the stochastic discount factor. This nonlinear term is the expected excess

13

return on a single-period nominal bond over a single-period real bond. Thus it measures the inflation risk premium at the short end of the term structure. It equals the conditional covariance between realized inflation and the log of the real SDF:

$$\text{Cov}_t\left(m_{t+1}, \pi_{t+1}\right) = -\sigma_{m\pi} z_t \psi_t. \tag{18}$$

Given positive risk aversion z_t, the sign of this covariance depends on the sign of the state variable ψ_t. When the covariance is positive, short-term nominal bonds are risky assets that have a positive risk premium because they tend to have unexpectedly low real payoffs in bad times. Of course, this premium increases with risk aversion z_t. When the covariance is negative, short-term nominal bonds hedge real risk; they command a negative risk premium which becomes even more negative as aggregate risk aversion increases.

The conditional covariance between the SDF and inflation also determines the covariance between the excess returns on real and nominal assets. Consider for example the conditional covariance between the real return on a one-period nominal bond and the real return on equities, both in excess of the return on a one-period real bond. From (11) and (13), this covariance is given by

$$\text{Cov}_t\left(r_{e,t+1} - r_{1,t+1}, y^{\$}_{1,t+1} - \pi_{t+1} - r_{1,t+1}\right) = -\left(\beta_{ex}\sigma_{x\pi} + \beta_{em}\sigma_{m\pi}\right)\psi_t,$$

which moves over time and can change sign. This implies that we can identify the dynamics of the state variable ψ_t from the dynamics of the conditional covariance between equities and nominal bonds.[10]

3.6 The nominal term structure of interest rates

Equation (17) writes the log nominal short rate as a linear-quadratic function of the state variables. We show in the Appendix that this property carries over to the entire zero-coupon nominal term structure. Just like the price of a n-period zero-coupon real bond, the price of a n-period zero-coupon nominal bond is an exponential linear-quadratic function of the vector of state variables:

$$P^{\$}_{n,t} = \exp\left\{ \begin{array}{c} A^{\$}_n + B^{\$}_{x,n}x_t + B^{\$}_{z,n}z_t + B^{\$}_{\lambda,n}\lambda_t + B^{\$}_{\xi,n}\xi_t + B^{\$}_{\psi,n}\psi_t \\ +C^{\$}_{z,n}z^2_t + C^{\$}_{\psi,n}\psi^2_t + C^{\$}_{z\psi,n}z_t\psi_t \end{array} \right\}, \tag{19}$$

---

[10]We can also identify ψ_t from the covariance between equities and real bonds, and we do so in our estimation.

where the coefficients $A_n^\$$, $B_{i,n}^\$$, and $C_{i,n}^\$$ solve a set of recursive equations given in the Appendix. These coefficients are functions of the maturity of the bond (n) and the coefficients that determine the stochastic processes for real and nominal variables. From equation (17), it is immediate to see that $B_{x,1}^\$ = B_{\xi,1}^\$ = B_{\lambda,1}^\$ = -1$, $C_{z\psi,1}^\$ = \sigma_{m\pi}$, and that the remaining coefficients are zero at $n = 1$.

In equation (19), log bond prices are affine functions of the short-term real interest rate (x_t) and the two components of expected inflation (λ_t and ξ_t), and quadratic functions of risk aversion (z_t) and inflation volatility (ψ_t). Thus our model naturally generates five factors that explain bond yields.

We can now characterize the log return on long-term nominal zero-coupon bonds in excess of the short-term nominal interest rate. Since bond prices are not exponential linear functions of the state variables, their returns are not conditionally lognormally distributed. But we can still find an analytical expression for their conditional expected returns. The Appendix derives an expression for the log of the conditional expected gross excess return on an n-period zero-coupon nominal bond which varies quadratically with risk aversion z_t and linearly with the covariance between the log real SDF and inflation ($z_t\psi_t$). Thus in this model, bond risk premia can be either positive or negative depending on the value of ψ_t.

Intuitively, the risk premium on nominal bonds varies over time as a function of both aggregate risk aversion and the covariance between inflation and the real side of the economy. If this covariance switches sign, so will the risk premium on nominal bonds. At times when inflation is procyclical—as will be the case if the macroeconomy moves along a stable Phillips Curve—nominal bond returns are countercyclical, making nominal bonds desirable hedges against business cycle risk. At times when inflation is countercyclical—as will be the case if the economy is affected by supply shocks or changing inflation expectations that shift the Phillips Curve in or out— nominal bond returns are procyclical and investors demand a positive risk premium to hold them.

3.7 Special cases

Our quadratic term structure model nests four constrained models of particular interest. First, if we constrain z_t to be constant and the real interest rate to be homoskedastic, our model reduces to a single-factor affine yield model for the term

structure of real interest rates, and a linear-quadratic model for the term structure of nominal interest rates. In this constrained model, real bond risk premia are constant, but nominal bond risk premia vary with the covariance between inflation and the real economy. We estimated this model in an earlier version of our paper.

Second, if we constrain ψ_t to be constant but allow z_t to vary over time, our model becomes a four-factor affine yield model where both real bond risk premia and nominal bond risk premia vary in proportion to aggregate risk aversion. This model captures the spirit of recent work on the term structure of interest rates by Bekaert, Engstrom, and Grenadier (2005), Buraschi and Jiltsov (2007), Wachter (2006) and others in which time-varying risk aversion is the only cause of time variation in bond risk premia.

Third, if we constrain z_t to be constant but allow ψ_t to vary over time, our model still remains in the class of exponential quadratic term structure models for both the real term structure and the nominal term structure. Specifically, the Appendix shows that equation (8) for the price of a real zero-coupon bond with maturity n becomes

$$P_{n,t} = \exp\left\{A_n + B_{x,n}x_t + B_{\psi,n}\psi_t + C_{\psi,n}\psi_t^2\right\},\tag{20}$$

while equation (19) for the price of a nominal zero-coupon bond with maturity n becomes

$$P_{n,t}^\$ = \exp\left\{A_n^\$ + B_{x,n}^\$ x_t + B_{\lambda,n}^\$ \lambda_t + B_{\xi,n}^\$ \xi_t + B_{\psi,n}^\$ \psi_t + C_{\psi,n}^\$ \psi_t^2\right\}.\tag{21}$$

In this model, time variation in bond risk premia is driven exclusively by the changing real-nominal covariance, i.e., by changes in the quantity of risk. This model is the one that we estimate in the empirical work of the next section.

Finally, if we constrain both z_t and ψ_t to be constant over time, and we allow expected inflation to have only the transitory component ξ_t, our model reduces to the two-factor affine yield model of Campbell and Viceira (2001, 2002), where both real bond risk premia and nominal bond risk premia are constant, and the factors are the short-term real interest rate and expected inflation. Allowing expected inflation to have a permanent component λ_t results in an expanded version of this affine yield model with permanent and transitory shocks to expected inflation.

16

4 Model Estimation

4.1 Data and estimation methodology

The term structure model presented in Section 3 generates real and nominal bond yields which are linear-quadratic functions of a vector of latent state variables. We now use this model as a laboratory to study the joint behavior of observed yields on nominal and inflation-indexed bonds, realized inflation, survey-based measures of expected inflation, stock returns, inflation-indexed bond returns, and nominal bond returns, and their second moments.

We start our exploration by presenting the data we use, and the corresponding maximum likelihood estimates of our model. Since our state variables are not observable, and the observable series have a nonlinear dependence on the latent state variables, we obtain maximum likelihood estimates via a nonlinear Kalman filter. Specifically, we use the unscented Kalman filter estimation procedure of Julier and Uhlmann (1997).

The unscented Kalman filter is a nonlinear Kalman filter which works through deterministic sampling of points in the distribution of the innovations to the state variables, does not require the explicit computation of Jacobians and Hessians, and captures the conditional mean and variance-covariance matrix of the state variables accurately up to a second-order approximation for any type of nonlinearity, and up to a third-order approximation when innovations to the state variables are Gaussian. Wan and van der Merwe (2001) describe in detail the properties of the filter and its practical implementation, and Binsbergen and Koijen (2008) apply the method to a prediction problem in finance.[11]

To use the unscented Kalman filter, we must specify a system of measurement equations that relate observable variables to the vector of state variables. The filter uses these equations to infer the behavior of the latent state variables of the model. We use twelve measurement equations in total.

---

[11]Binsbergen and Koijen's application has linear measurement equations and nonlinear transition equations, whereas ours has linear transition equations and nonlinear measurement equations. The unscented Kalman filter can handle either case. We have also checked the robustness of our estimates by re-estimating our model using the "square root" variant of the filter, which has been shown to be more estable when some of the state variables follow heteroskedastic processes. This variant produces estimates which are extremely similar to the ones we report in the paper.

Our first four measurement equations relate observable nominal bond yields to the vector of state variables. Specifically, we use the relation between nominal zero-coupon bond log yields $y_{n,t}^{\$} = -\log(P_{n,t}^{\$})/n$ and the vector of state variables implied by equation (19). We use yields on constant maturity 3-month, 1-year, 3-year and 10-year zero-coupon nominal bonds sampled at a quarterly frequency for the period 1953.Q1-2009Q3. These data are spliced together from two sources. From the first quarter of 1953 through the first quarter of 1961 we sample quarterly data from the monthly dataset developed by McCulloch and Kwon (1993), and from the second quarter of 1961 through the last quarter of 2009 we sample quarterly data from the daily dataset constructed by Gürkaynak, Sack, and Wright (2006, updated through 2009, GSW henceforth). We assume that bond yields are measured with errors, which are uncorrelated with each other and with the structural shocks of the model.

We sample the data at a quarterly frequency in order to minimize the impact of high-frequency noise in the measurement of some of our key variables—such as realized inflation—while keeping the frequency of observation reasonably high (Campbell and Viceira 2001, 2002). By not having to fit all the high-frequency monthly variation in the data, our estimation procedure can concentrate on uncovering the low-frequency movements in interest rates which our model is designed to capture.

Figure 3 illustrates our nominal interest rate data by plotting the 3-month and 10-year nominal yields, and the spread between them, over the period 1953-2009. Some well-known properties of the nominal term structure are visible in this figure, notably the greater smoothness and higher average level of the 10-year nominal interest rate. The yield spread shows large variations in response to temporary movements in the 3-month bill rate, but also a tendency to be larger since the early 1980's than it was in the first part of our sample. Our model will explain this tendency as the result of movements in the real interest rate, transitory expected inflation, and the covariance of nominal and real variables.

Our fifth measurement equation is given by equation (13), which relates the observed inflation rate to expected inflation and inflation volatility, plus a measurement error term. We use the CPI as our observed price index in this measurement equation. We complement this measurement equation with another one that uses data on inflation expectations from the Survey of Professional Forecasters for the period 1968.Q4-2009.Q3. Specifically, we use the median forecast of growth in the GDP price index over the next quarter. We relate this observed measure of expected inflation to the sum of equations (14) and (15) in our model plus a measurement error term. Fig-

18

ure 4 plots the history of realized inflation and our survey based measure of expected inflation. Average inflation was higher in the first half of our sample, peaked in the late 1970's and early 1980's, and declined afterwards; inflation was essentially zero or even negative at both ends of our sample period, i.e., the 1950s and the 2000's, when it was also volatile. Expected inflation exhibits a pattern similar to realized inflation, albeit smoother. This figure implies that the long-term decline in short and long nominal interest rates that started in the early 1980's was at least partly caused by declining inflation expectations.

The seventh measurement equation relates the observed yield on constant maturity Treasury inflation protected securities (TIPS) to the vector of state variables, via the pricing equation for real bonds generated by our model. We obtain data on constant maturity zero-coupon 10-year TIPS dating back to the first quarter of 1999 from GSW (2008). Before 1999, we treat the TIPS yield as missing, which can easily be handled by the Kalman filter estimation procedure. As with nominal bond yields, we sample real bond yields at a quarterly frequency, and we assume that they are measured with errors, which are uncorrelated with each other and with the structural shocks of the model.

Figure 5 illustrates our real bond yield series. The decline in the TIPS yield since the year 2000, and the spike in the fall of 2008, are clearly visible in this figure. Campbell, Shiller, and Viceira (2009) document that this decline in the long-term real interest rate, and the subsequent sudden increase during the financial crisis, occurred in inflation-indexed bond markets around the world. In earlier data from the UK, long-term real interest rates were much higher on average during the 1980's and 1990's. Our model will explain such large and persistent variation in the TIPS yield primarily as the result of persistent movements in the short-term real interest rate.

Our eighth measurement equation uses data on an equity index, the CRSP value-weighted portfolio comprising the stocks traded in the NYSE, AMEX and NASDAQ. This equation describes realized log equity returns $r_{e,t+1}$ using equations (3), (11), and (12).

The last four measurement equations use the implications of our model for: i) the conditional covariance between equity returns and real bond returns, (ii) the conditional covariance between equity returns and nominal bond returns, (iii) the conditional volatility of real bond returns, and (iv) the conditional volatility of nominal bond returns. The Appendix derives expressions for these time-varying conditional

19

second moments, which are functions of z_t and ψ_t. Following Viceira (2010), we construct the analogous realized second moments using high-frequency data. We obtain daily stock returns from CRSP. We calculate daily nominal bond returns from daily GSW nominal yields from 1961.Q2 onwards, and daily real bond returns from daily GSW real yields from 1999.Q1 onwards.[12] We then compute the variances and covariances realized over quarter t and treat these as the conditional (expected) moments at quarter $t-1$ plus measurement error. These measurement equations help us identify z_t and ψ_t.

The data used in these measurement equations are plotted in Figure 6 for real bonds and in Figure 7 for nominal bonds. The left panel of each figure shows the realized covariance between daily stock and bond returns, while the right panel shows the realized variance of daily bond returns. The thick lines in each panel show a smoothed version of the raw data.

Figure 7 shows that both the stock-nominal bond covariance series and the nominal bond variance series increase in the early 1970's and, most dramatically, in the early 1980's. In the early 1960's, the early 2000's, and the late 2000's the covariance spikes downward while the variance increases; the spikes are particularly pronounced in the financial crisis of 2008–2009. Our model will interpret these as times when the nominal-real covariance was negative.[13] Figure 6 shows that the stock-real bond covariance series and the real bond variance series follow patterns very similar to those of nominal bonds for the overlapping sample period.

The unscented Kalman filter uses the system of measurement equations we have just formulated, together with the set of transition equations (2), (4), (5), (13), (14), and (15) that describe the dynamics of the state variables, to construct a pseudo-likelihood function. We then use numerical methods to find the set of parameter values that maximize this function and the asymptotic standard errors of the parameter estimates. Specifically, we use the outer product method to compute maximum likelihood asymptotic standard errors.

---

[12]We calculate daily returns on the n year bond from daily yields as $r_{n,t+1} = ny_{n,t} - (n-1/264)\,y_{n,t+1}$. We assume there are 264 trading days in the year, or 22 trading days per month. Prior to 1961.Q2, we calculate monthly returns from monthly McKullock-Kwon nominal yields, and calculate variances and covariances using a rolling 12-month return window.

[13]Figure 7 also shows a brief downward spike in the realized bond-stock covariance around the stock market crash of October 1987. However this movement is so short-lived that it does not cause our estimated nominal-real covariance to switch sign.

4.2   Model constraints for estimation

In exploratory data analysis, we have found that our model estimates generate negligible variation in z_t (risk aversion) unless we use extraneous information to pin down this state variable. For example, in the first version of this paper we added a measurement equation linking risk aversion to the dividend-price ratio of the aggregate stock market; we have also estimated versions of the model that effectively identify z_t through a measurement equation linking the expected excess return on nominal bonds to empirically successful predictor variables such as the yield spread (Shiller, Campbell, and Schoenholtz 1983, Fama and Bliss 1987, Campbell and Shiller 1991) or the linear combination of forward rates suggested by Cochrane and Piazzesi (2005). However all versions of our model, with or without such additional measurement equations, have similar implications for the nominal-real covariance ψ_t and its impact on bond prices.

For simplicity, and given our primary interest in the state variable ψ_t, we report results for a model with no additional measurement equations that constrains z_t to be constant. Thus bond pricing equations are given by (20) and (21). Of course in this constrained model σ_m is no longer unidentified, and we estimate it freely.

Even this constrained model has a large number of shocks, and we have found that it is difficult to estimate the model allowing an unconstrained variance-covariance matrix for the shocks. Therefore we also constrain many of the covariances in the model to be zero. The unconstrained parameters are the covariances of the heteroskedastic shocks to the real interest rate and permanent expected inflation, and the shocks to transitory expected inflation and realized inflation, with the stochastic discount factor; the covariances of the transitory component of expected inflation with realized inflation and the heteroskedastic shock to the real interest rate; and the covariance of realized inflation with the heteroskedastic shock to the real interest rate. The homoskedastic shocks to the real interest rate and permanent expected inflation are assumed to be uncorrelated with other shocks.

With these constraints on the variance-covariance matrix, we allow freely estimated risk premia on all the state variables except the nominal-real covariance, as well as a risk premium for realized inflation that affects the level of the short-term nominal interest rate. We allow correlations among real interest rates, realized inflation, and the transitory component of expected inflation, while imposing that the permanent component of expected inflation is uncorrelated with movements in the

21

transitory state variables. This constraint is natural if one believes that long-run expected inflation is determined by central bank credibility, which is moved by political developments rather than business-cycle fluctuations in the economy. A likelihood ratio test of the constrained model cannot reject it against the fully parameterized model.

4.3   Parameter estimates

Table 1 presents quarterly parameter estimates over the period 1953-2009 and their asymptotic standard errors for the model with constant z_t. It shows that the real interest rate is the most persistent state variable, with an autoregressive coefficient of 0.94. This coefficient implies that shocks to the real interest rate have a half life of 11 quarters. This persistence reflects the observed variability and persistence of TIPS yields. The nominal-real covariance and the transitory component of expected inflation are persistent processes in our model, with half-lives of about 3 and 6 quarters respectively. Of course the model also includes a permanent component of expected inflation. If we model expected inflation as a single stationary AR(1) process, as we did in the first version of this paper, we find expected inflation to be more persistent than the real interest rate.[14]All persistence coefficients are precisely estimated, with very small asymptotic standard errors.

Table 1 shows large differences in the volatility of shocks to the state variables. The one-quarter conditional volatility of the homoskedastic shock to the annualized real interest rate is estimated to be about 45 basis points, and the average one-quarter conditional volatility of the heteroskedastic shock to the annualized real interest rate is estimated to be 97 basis points. The average one-quarter conditional volatility of the transitory component of annualized expected inflation is about 72 basis points, and the average one-quarter conditional volatility of annualized realized inflation is about 334 basis points.[15] By contrast, the average one-quarter conditional volatilities of the shocks to the permanent component of expected inflation are very small. Of

---

[14]Campbell and Viceira (2001, 2002) also estimate expected inflation to be more persistent than the real interest rate in a model with constant z_t and ψ_t and a stationary AR(1) process for expected inflation. Campbell and Viceira do find that when the estimation period includes only the years after 1982, real interest rates appear to be more persistent than expected inflation, reflecting the change in monetary policy that started in the early 1980's under Federal Reserve chairman Paul Volcker. We have not yet estimated our quadratic term structure model over this subsample.

[15]We compute the average conditional volatilities of the heteroskedastic shock to the real interest

course, the unconditional standard deviations of the real interest rate and the two components of expected inflation are much larger because of the high persistence of the processes; in fact, the unconditional standard deviation of the permanent component of expected inflation is undefined because this process has a unit root. With the exception of the volatility of the heteroskedastic shock to the permanent component of expected inflation, the volatility parameters for the real interest process and inflation are all precisely estimated with very small asymptotic standard errors.

Table 1 also reports the unrestricted correlations among the shocks and their asymptotic standard errors. We report correlations instead of covariances to facilitate interpretation. We compute their standard errors from those of the primitive parameters of the model using the delta method. The Appendix reports covariances and their asymptotic standard errors.

There is a correlation of almost -0.24 between ξ_t and $-m_t$ shocks. Although the correlation coefficient is marginally significant at conventional significance levels, the Appendix shows that the covariance is more precisely estimated. This negative correlation implies that the transitory component of expected inflation is countercyclical, generating a positive risk premium in the nominal term structure, when the state variable ψ_t is positive; but transitory expected inflation is procyclical, generating a negative risk premium, when ψ_t is negative. The absolute magnitude of the correlation between λ_t and $-m_t$ shocks is larger at around -0.80, implying that the risk premium for permanent shocks to expected inflation is larger than the risk premium for transitory shocks to expected inflation. However, this covariance has a very large standard error.

We also estimate a statistically insignificant and economically very small positive correlation between π_t and $-m_t$ shocks. The point estimate implies that short-term inflation risk is very small, and that nominal Treasury bills have a very small or zero inflation risk premium. Finally, we estimate a statistically significant negative correlation of -0.32 between x_t and $-m_t$ shocks, implying a time-varying term premium on real bonds that is positive when ψ_t is positive.

In the equity market, we estimate statistically insignificant small loadings of stock

---

rate, the components of expected inflation, and realized inflation as $\left(\mu_\psi^2 + \sigma_\psi^2\right)^{1/2}$ times the volatility of the underlying shocks. For example, we compute the average conditional volatility of realized inflation as $\left(\mu_\psi^2 + \sigma_\psi^2\right)^{1/2} \sigma_\pi$.

returns on shocks to the real interest rate (β_{ex} and β_{eX}), and a much larger and statistically significant positive loading on shocks to the negative of the log SDF (β_{em}). Naturally this estimate implies a positive equity risk premium.

4.4   Fitted state variables

How does our model interpret the economic history of the last 55 years? That is, what time series does it estimate for the underlying state variables that drive bond and stock prices? Figure 8 shows our estimates of the real interest rate x_t.[16] The model estimates a process for the real interest rate that is high on average, with a spike in the early 1980's, and becomes more volatile and declining in the second half of the sample. Higher-frequency movements in the real interest rate were often countercyclical in this period, as we see the real rate falling in the recessions of the early 1970's, early 1990's, early 2000's, and at the end of our sample period in 2008–2009. The real interest rate also falls around the stock market crash of 1987. However there are important exceptions to this pattern, notably the very high real interest rate in the early 1980's, during Paul Volcker's campaign against inflation, and a short-lived spike in the fall of 2008. This spike and generally the history of the real interest rate since the late 1990's follow the history of TIPS during this period, shown in Figure 5. Thus the model attributes the history of long-dated TIPS yields mostly to changes in the short-term real rate x_t. While the state variable ψ_t is also relevant for TIPS yields, it plays a secondary role.

Figure 9 plots the components of expected inflation. The permanent component of expected inflation, in the left panel, exhibits a familiar hump shape over the postwar period. It was low, even negative, in the 1950's and 1960's, increased during the 1970's and reached a maximum value of about 10% in the first half of the 1980's. Since then, it has experienced a secular decline and remained close to 2% throughout the 2000's.

The transitory component of expected inflation, in the right panel, was particularly high in the late 1970's and 1980, indicating that investors expected inflation to decline gradually from a temporarily high level. The transitory component has been predominantly negative since then till almost the end of our sample period, implying

---

[16]As we have noted already, we constrain z_t to be constant. Maximum likelihood estimation of the model that allows z_t to vary over time produces an estimate of z_t which is almost constant.

that our model attributes the generally high levels of yield spreads during the second half of our sample period at least partly to investor pessimism about increases in future inflation. By estimating a generally negative transitory component of expected inflation, the model is also able to explain simultaneously the low average nominal short-term interest rate and the high average real short-term interest rate in the latter part of our sample period.

Finally, Figure 10 shows the time series of ψ_t. As we have noted, this variable is identified primarily through the covariance of stock returns and bond returns and the volatility of bond returns—both nominal and real. The state variable ψ_t exhibits low volatility and an average close to zero in the period leading up to the late 1970's, with briefly negative values in the late 1950's, and an upward spike in the early 1970's. It becomes much more volatile starting in the late 1970's through the end of our sample period. It rises to large positive values in the early 1980's and stays predominantly positive through the 1980's and 1990's. However, in the late 1990's it switches sign and turns predominantly negative, with particularly large downward spikes in the period immediately following the recession of 2001 and in the fall of 2008, at the height of the financial crisis of 2008–2009. Thus ψ_t not only can switch sign, it has done so during the past ten years. Overall, the in-sample average for ψ_t is positive.

The state variables we have estimated can be used to calculate fitted values for observed variables such as the nominal term structure, real term structure, realized inflation, analysts' median inflation forecast, and the realized second moments of bond and equity returns. We do not plot the histories of these fitted values to save space. They track the actual observed yields on nominal bonds, inflation forecasts, and the realized stock-nominal bond covariance very closely, and closely the yields on TIPS, realized inflation, and the rest of the realized second moments included in the estimation. In general, our model is rich enough that it does not require measurement errors with high volatility to fit the observed data on stock and bond prices.

5 Term Structure Implications

5.1 Moments of bond yields and returns

Although our model fits the observed history of real and nominal bond yields, an important question is whether it must do so by inferring an unusual history of shocks, or whether the observed properties of interest rates emerge naturally from the properties of the model at the estimated parameter values. In order to assess this, Tables 2 and 3 report some important moments of bond yields and returns.

The tables compare the sample moments in our historical data with moments calculated by simulating our model 1,000 times along a path that is 250 quarters (or 62 and a half years) long, and averaging time-series moments across simulations. In each table, sample moments are shown in the first column and model-implied moments in the second column. The third column reports the fraction of simulations for which the simulated time-series moment is larger than the corresponding sample moment in the data. These numbers can be used as informal tests of the ability of the model to fit each sample moment. Although our model is estimated using maximum likelihood, these diagnostic statistics capture the spirit of the method of simulated moments (Duffie and Singleton 1993, Gallant and Tauchen 1996), which minimizes a quadratic form in the distance between simulated model-implied moments and sample moments.

In Table 2 the short-term interest rate is a three-month rate and moments are computed using a three-month holding period, while in Table 3 the short-term interest rate is a one-year rate and the holding period is one year. The use of a longer short rate and holding period in Table 3 follows Cochrane and Piazzesi (2005), and shows us how our model fits lower frequency movements at the longer end of the yield curve.

The first two rows of Tables 2 and 3 report the sample and simulated means for nominal bond yield spreads, calculated using 3 and 10 year maturities, and the third and fourth rows look at the volatilities of these spreads. In all cases our model provides a fairly good fit to average yield spreads, slightly understating the average 3-year spread and understating the average 10-year spread (that is, slightly overstating the average concavity of the yield curve). However the model systematically overstates the volatility of yield spreads, a problem that appears in almost all our 1,000 simulations.

26

In each table, the next four rows show how our models fit the means and standard deviations of realized excess returns on 3-year and 10-year nominal bonds. In order to calculate three-month realized returns from constant-maturity bond yields, we interpolate yields between the constant maturities we observe, doing this in the same manner for our historical data and for simulated data from our models. Most of our annual realized returns do not require interpolation, but in the early part of our sample, before 1971, we must also interpolate the 9-year bond yield to calculate the annual realized return on 10-year bonds. Just as with yield spreads, the model provides a good fit to mean excess returns, but systematically somewhat overstates the volatility of excess returns on 3-year bonds, although it provides a statistically better fit for the excess return on 10-year bonds.

The next four rows of each table summarize our model description of TIPS yields. The model generates an average TIPS yield that is higher than the observed average. We do not believe this is an extremely serious problem, as our estimates imply higher real interest rates earlier in our sample period, before TIPS were issued, than in the period since 1997 over which we measure the average TIPS yield. Thus the discrepancy may result in part from the short and unrepresentative period over which we measure the average TIPS yield in the data.

The model implies a zero or slightly negative average yield spread and positive average realized excess return. The difference between these two statistics reflects the effect of Jensen's Inequality; equivalently, it is the result of convexity in long-term bonds. The slightly positive average risk premium results from our negative estimate of ρ_{xm} in Table 1, which implies that the real interest rate is countercyclical on average.

5.2   State variables and the yield curve

Given our estimated term structure model, we can now analyze the impact of each of our four state variables on the nominal yield curve, and thus get a sense of which components of the curve they affect the most. To this end, we plot in Figures 11 through 14 the zero-coupon log nominal yield curve and, when appropriate, the zero-coupon log real yield curve generated by our model when one of the state variables is at its in-sample mean, maximum, and minimum, while all other state variables are at their in-sample means. Thus the central line describes the yield curve—real or nominal—generated by our model when all state variables are evaluated at their

in-sample mean. For simplicity we will refer to this curve as the "mean log yield curve."[17] We plot maturities up to 10 years, or 40 quarters.

Figure 11 plots the zero-coupon log real yield curve in the left panel, and the zero-coupon log nominal yield curve in the right panel, that obtain when we vary the short-term real rate x_t. The left panel shows that the mean log real yield curve generated by our model is gently upward sloping, with an intercept of about 2.3% and a 10-year yield spread of about 80 basis points.

The right panel of Figure 11 shows that the mean log nominal yield curve has a somewhat greater positive slope, with a spread between the 10-year rate and the 1-quarter rate of just under 110 basis points. This spread is slightly lower than the 115 basis point historical average spread in our sample period. The yield curve is concave, flattening out at maturities beyond five years. The intercept of the curve implies a short-term nominal interest rate of about 5.3%, in line with the average short-term nominal interest rate in our sample.

Figure 11 shows that changes in the real interest rate alter both the level and the slope of the real and nominal yield curves. However, the slope effects are modest because the real interest rate is so persistent in our model.

Figure 12 plots the effect of changes in the components of expected inflation, λ_t and ξ_t, on the nominal yield curve. The left panel shows that changes in the permanent component λ_t of expected inflation affect short- and long-term nominal yields almost equally, causing parallel shifts in the level of the nominal yield curve. The right panel shows that, by contrast, changes in the transitory component ξ_t of expected inflation have a much stronger effect on the short end of the curve than on the long end, causing changes in the slope of the curve. These effects reflect the fact that the shocks to the transitory component of expected inflation have a relatively short half-life of about 18 months.

The most interesting results are shown in Figure 13. This figure illustrates the nominal and real yield curves that obtain when we vary ψ_t. In both panels, increases in ψ_t from the sample mean to the sample maximum increase intermediate-term yields and lower long-term yields, while decreases in ψ_t to the sample minimum lower both

---

[17]Strictly speaking this is a misnomer in the case of the nominal yield curve, for two reasons. First, the log nominal yield curve is a non-linear function of the vector of state variables. Second, its unconditional mean is not even defined, since one of the state variables follows a random walk. Thus at most we can compute a mean yield curve conditional on initial values for the state variables.

intermediate-term and long-term yields. Thus ψ_t alters the concavity of both the real yield curve and the nominal yield curve.

The impact of ψ_t on the concavity of the nominal yield curve results from two features of our model. First, nominal bond risk premia increase with maturity rapidly at intermediate maturities and slowly at short maturities because intermediate maturities are exposed both to transitory and permanent shocks to expected inflation. When ψ_t is positive, this generates a steep yield curve at shorter maturities, and a flatter one at longer maturities. When ψ_t changes sign, however, the difference in risk prices pulls intermediate-term yields down more strongly than long-term yields.

Second, when ψ_t is far from zero bond returns are unusually volatile, and through Jensen's Inequality this lowers the bond yield that is needed to deliver any given expected simple return. This effect is much stronger for long-term bonds; in the terminology of the fixed-income literature, these bonds have much greater "convexity" than short- or intermediate-term bonds. Therefore extreme values of ψ_t tend to lower long-term bond yields relative to intermediate-term yields.

Figures 11 through 13 allow us to relate our model to traditional factor models of the term structure of interest rates, and to provide an economic identification of those factors. Following Litterman and Scheinkman (1991), many term structure analyses distinguish a "level" factor, a "slope" factor, and a "curvature" factor. The first of these moves the yield curve in parallel; the second moves the short end relative to the long end; and the third moves intermediate-term yields relative to short and long yields.

Figures 11 and 12 suggest that in our model, the permanent component of expected inflation is the main contributor to the level factor. The short-term real interest rate and particularly the transitory component of expected inflation both contribute to the slope factor by moving short-term nominal yields more than long-term nominal yields. Finally, Figure 13 shows that the covariance of nominal and real variables drives the curvature factor and also has some effect on the slope factor. Putting these results together, the curvature factor is likely to be the best proxy for the nominal-real covariance.

5.3 Time-variation in bond risk premia

In the previous section we saw that in our model, the nominal-real covariance ψ_t is an important determinant of medium- and long-term nominal interest rates. The reason for this is that these variables have powerful effects on risk premia. In fact, nominal bond risk premia are almost perfectly proportional to the product ψ_t.

Figure 14 illustrates this fact. The left panel plots the simulated expected excess return on 3-year and 10-year nominal bonds over 3-month Treasury bills against ψ_t.[18] The right panel of the figure shows the term structure of risk premia as ψ_t varies from its sample mean to its sample minimum and maximum. Risk premia spread out rapidly as maturity increases, and 10-year risk premia vary from -75 to 140 basis points. The reason for this asymmetry is that we observe large negative values of ψ_t towards the end of our sample period.

The full history of our model's 10-year term premium is illustrated in Figure 15. The figure shows fairly stable risk premia of about 0.2% during the 1950's and 1960's, and a run up later in the 1970's to a peak above 1.5% in the early 1980's. A long decline in risk premia later in the sample period was accentuated around the recession of the early 2000's and during the financial crisis of 2008–2009, bringing the risk premium to its sample minimum of -0.75%. This time series reflects the shape in the nominal-real covariance ψ_t illustrated in Figure 10.

We saw in Figure 13 that the nominal-real covariance ψ_t influences the curvature of the yield curve as well as its slope. Other factors in our model, such as the real interest rate, also influence the slope of the yield curve but do not have much effect on its curvature. Given the dominant influence of ψ_t on bond risk premia, the curvature of the yield curve should be a good empirical proxy for risk premia on nominal bonds.

In fact, an empirical result of this sort has been reported by Cochrane and Piazzesi (CP, 2005). Using econometric methods originally developed by Hansen and Hodrick (1983), and implemented in the term structure context by Stambaugh (1988), CP show that a single linear combination of forward rates is a good predictor of excess bond returns at a wide range of maturities. CP work with a 1-year holding period

---

[18] In the general model with time varying z_t, the nominal-bond risk premium in our model is a linear combination of z_t, z_t^2, and $z_t\psi_t$. But the maximum likelihood estimate of the general model produces an almost constant z_t. Thus in practice, ψ_t generates almost all the variation in the risk premium.

and a 1-year short rate. They find that the combination of forward rates that predicts excess bond returns is tent-shaped, with a peak at 3 or 4 years, implying that bond risk premia are high when intermediate-term interest rates are high relative to both shorter-term and longer-term rates; that is, they are high when the yield curve is strongly concave.

Our model interprets this phenomenon as the result of changes in the nominal-real covariance ψ_t. As ψ_t increases, it raises the risk premium for the transitory component of expected inflation and strongly increases the intermediate-term yield, but it has a damped or even perverse effect on long-term yields because the permanent component of expected inflation has little systematic risk and the convexity of long bonds causes their yields to fall with volatility. Thus the best predictor of excess bond returns is the intermediate-term yield relative to the average of short- and long-term yields.

Figure 16 illustrates the estimated coefficients in a CP regression of annual excess bond returns over the 1-year short rate, averaged across maturities from 2 to 5 years in the manner of CP, onto 1-year, 3-year, and 5-year forward rates. The fitted value in the data is tent-shaped, as reported by CP; the fitted value implied by our model has a similar shape but a much smaller magnitude. A caveat is that when we add 2- and 4-year forward rates to the regression, we do not reliably recover the tent shape either in the data or in the model, as the regressors are highly collinear and so the regression coefficients become unstable.

Despite these promising qualitative results, the predictability of bond returns is small in our model. The bottom panels of Tables 2 and 3 illustrate this point. In the first three rows of Table 2 we report the standard deviations of true expected 3-month excess returns within our model. Our model implies an annualized standard deviation for the expected excess return on 3-year bonds of about 9 basis points, and for the expected excess return on 10-year bonds of about 15 basis points.[19] This variation is an order of magnitude smaller than the annualized standard deviations of realized excess bond returns, implying that the true explanatory power of predictive regressions in our model is tiny. There is also modest variability of about 11 basis points in the true expected excess returns on TIPS.

---

[19]Yield interpolation for 3-month returns in Table 2, and for annual returns on 10-year bonds in the early part of our sample in Table 3, may exaggerate the evidence for predictability; however the same yield interpolation is used for simulated data from our models, so the comparison of results across columns is legitimate. We have used our simulations to examine the effect of interpolation. We find that interpolation does slightly increase measured bond return predictability, but the effect is modest.

The next three rows report the standard deviations of fitted values of Campbell-Shiller (1991, CS) predictability regressions of annualized nominal bond excess returns onto yield spreads of the same maturity at the beginning of the holding period. At the 3-month horizon shown in Table 2, the standard deviations in the data are 104 basis points for 3-year bonds, and 251 basis points for 10-year bonds. These numbers are considerably larger than the true variability of expected excess returns in our model, implying that our model cannot match the behavior of these predictive regressions.

We also report the standard deviations of fitted values generated by CS regressions on simulated data from our various models. For our full model, the regressions deliver fitted values that are considerably more volatile than the true expected excess returns. The reason for this counterintuitive behavior is that there is important finite-sample bias in the CS regression coefficients of the sort described by Stambaugh (1999). In the case of regressions of excess bond returns on yield spreads, by contrast with the better known case of regressions of excess stock returns on dividend yields, the Stambaugh bias is negative (Bekaert, Hodrick, and Marshall 1997). In our full model, where the true regression coefficient is positive but close to zero, the Stambaugh bias increases the standard deviation of fitted values by generating spurious negative coefficients. Nonetheless, the standard deviation of fitted values in the model is still considerably smaller than in the data.

We obtain somewhat more promising results using a procedure that approximates the approach of Cochrane and Piazzesi (2005, CP). We regress excess bond returns on 1-, 3-, and 5-year forward rates at the beginning of the holding period, and report the standard deviations of fitted values.[20] This procedure generates comparable standard deviations of fitted values in the model and in the data, at least for predicting excess 3-year bond returns. Once again, however, this finding is largely driven by small-sample bias as the fitted values in the model have a much higher standard deviation than the true expected excess returns.

Finally, Table 4 asks whether our model generates proxies for bond risk premia, constrained to be linear combinations of 1-year, 3-year, and 5-year forward rates, that perform well in the historical data. The table compares the empirical R^2 statistics for unconstrained CP regressions with the empirical R^2 statistics that result from regressing bond returns on the combinations of forward rates that, in simulated data generated by our term structure model, best predict bond returns. In the top panel

---

[20]Cochrane and Piazzesi impose proportionality restrictions across the regressions at different maturities, but we do not do this here.

we allow free regression coefficients on these forward rate combinations, while in the bottom panel we restrict them to have a unit coefficient as implied by our model simulations. Our model generates the type of predictability for expected excess bond returns shown by CP, particularly at the 1-year horizon. However, it cannot match the explanatory power of unrestricted CP regressions.

These results show that although our model does generate time-varying bond risk premia, the implied variation in risk premia is smaller and has a different time-series pattern from that implied by CS and CP regressions. In the CS case, the difference in time-series behavior can be understood visually by comparing the history of the yield spread shown in Figure 3 with the history of the model-implied bond risk premium shown in Figure 15. The former has a great deal of business-cycle variation, while the latter has a hump shape with a long secular decline from the early 1980's through the late 2000's. The fitted value from a CP regression lines up somewhat better with the model-implied bond risk premium, but it too spikes up in the recessions of the early 1990's and early 2000's in a way that has no counterpart in Figure 15.

6   Conclusion

We have argued that term structure models must confront the fact that the covariances between nominal and real bond returns, on the one hand, and stock returns, on the other, have varied substantially over time. Analyses of asset allocation traditionally assume that broad asset classes have a stable structure of risk over time; our empirical results imply that in the cases of nominal and inflation-indexed government bonds, at least, this assumption is seriously misleading.

We have proposed a term structure model in which real and nominal bond returns are driven by five factors: the real interest rate, the level of risk aversion or volatility of the stochastic discount factor, transitory and permanent components of expected inflation, and a state variable that governs the covariances of inflation and the real interest rate with the stochastic discount factor. We have estimated a version of the model allowing for time-variation in all these variables except for the level of risk aversion. The model implies that the risk premia of nominal bonds have changed over the decades because of changes in the covariance between inflation and the real economy. Nominal bond risk premia were positive in the early 1980's, when bonds covaried strongly with stocks, and negative in the 2000's, particularly during the

downturn of 2008–9, when bonds hedged equity risk.

Our model explains the qualitative finding of Cochrane and Piazzesi (2005) that a tent-shaped linear combination of forward rates, with a peak at about 3 years, predicts excess bond returns at all maturities better than maturity-specific yield spreads. In our model, the covariance between inflation and the real economy has its largest effect on the risk premium for a transitory component of expected inflation, which moves the 3-year nominal yield. There is a more modest effect on the risk premium for a permanent component of expected inflation, which is important for the 10-year nominal yield. In addition, there is a Jensen's Inequality effect of increasing volatility on yields. In the language of fixed-income investors, longer-term bonds have "convexity" which becomes more valuable when volatility is high, driving down bond yields. At the long end of the yield curve, the convexity effect is slightly stronger for high levels of the nominal-real covariance, whereas at the intermediate portion of the curve, the risk premium effect dominates. Hence, the level of intermediate yields relative to short- and long-term yields is a good proxy for the nominal-real covariance and thus for the risk premium on nominal bonds.

Although our results are qualitatively consistent with empirical findings of predictability in excess bond returns, our model does not replicate the high explanatory power of regressions that predict excess US Treasury bond returns from yield spreads and forward rates. Our estimates of variation in the nominal-real covariance and the level of risk aversion deliver risk premia whose standard deviation is an order of magnitude smaller than the standard deviation of realized excess bond returns. Related to this, the risk premia implied by our model have trended strongly downwards since the early 1980's, in line with the declining covariance between bond and stock returns, whereas the fitted values of predictive regressions for bond returns have not trended in this way.

The results we have presented can be extended in a number of directions. First, it would be of interest to allow both persistent and transitory variation in the nominal-real covariance, as we have done for expected inflation. This might allow our model to better fit both the secular trends and cyclical variation in the realized covariance between bonds and stocks.

Second, we can estimate our model using data from other countries, for example the UK, where inflation-indexed bonds have been actively traded since the mid-1980's. This will be particularly interesting since the evidence of bond return predictability is considerably weaker outside the US (Bekaert, Hodrick, and Marshall 2001, Campbell

2003) and may better fit the predictability generated by our model.

Third, it would be possible to derive stock returns from primitive assumptions on the dividend process, as in the recent literature on affine models of stock and bond pricing (Mamaysky 2002, Bekaert, Engstrom, and Grenadier 2005, d'Addona and Kind 2006, Bekaert, Engstrom, and Xing 2009).

Fourth, we can consider other theoretically motivated proxies for the stochastic discount factor. An obvious possibility is to look at realized or expected future consumption growth, as in recent papers on consumption-based bond pricing by Piazzesi and Schneider (2006), Abhyankar and Lee (2008), Eraker (2008), Hasseltoft (2008), Lettau and Wachter (2010), and Bansal and Shaliastovich (2010). A disadvantage of this approach is that consumption is not measured at high frequency, so one cannot use high-frequency data to track a changing covariance between bond returns and consumption growth.

Finally, we can explore the relation between our covariance state variable ψ_t and the state of monetary policy and the macroeconomy. We have suggested that a positive ψ_t corresponds to an environment in which the Phillips Curve is unstable, perhaps because supply shocks are hitting the economy or the central bank lacks anti-inflationary credibility, while a negative ψ_t reflects a stable Phillips Curve. It would be desirable to use data on inflation and output, and a structural macroeconomic model, to explore this interpretation. The connection between the bond-stock covariance and the state of the macroeconomy should be of special interest to central banks. Many central banks use the breakeven inflation rate, the yield spread between nominal and inflation-indexed bonds, as an indicator of their credibility. The bond-stock covariance may be appealing as an additional source of macroeconomic information.

References

Abhyankar, Abhay and Soyeon Lee, 2008, "Are Government Bonds Risky Assets?", unpublished paper, University of Edinburgh.

Ahn, Dong-Hyun, Robert F. Dittmar, and A. Ronald Gallant, 2002, "Quadratic Term Structure Models: Theory and Evidence", *Review of Financial Studies* 15, 243–288.

Ang, Andrew, Geert Bekaert, and Min Wei, 2008, "The Term Structure of Real Rates and Expected Inflation", *Journal of Finance* 63, 797–849.

Backus, David K. and Jonathan H. Wright, 2007, "Cracking the Conundrum", *Brookings Papers on Economic Activity* 1, 293–316.

Baele, Lieven, Geert Bekaert, and Koen Inghelbrecht, 2010, "The Determinants of Stock and Bond Return Comovements", *Review of Financial Studies* 23, 2374–2428.

Bansal, Ravi and Ivan Shaliastovich, 2010, "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets", unpublished paper, Duke University.

Bansal, Ravi, and Amir Yaron, 2004, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance* 59, 1481-1509.

Bansal, Ravi and Hao Zhou, 2002, "Term Structure of Interest Rates with Regime Shifts", *Journal of Finance* 57, 1997–2043.

Barsky, Robert B., 1989, "Why Don't the Prices of Stocks and Bonds Move Together?", *American Economic Review* 79, 1132–1145.

Beaglehole, David R. and Mark S. Tenney, 1991, "General Solutions of Some Interest Rate-Contingent Claim Pricing Equations", *Journal of Fixed Income*, September, 69–83.

Bekaert, Geert, Eric Engstrom, and Steve Grenadier, 2005, "Stock and Bond Returns with Moody Investors", unpublished paper, Columbia University, University of Michigan, and Stanford University.

Bekaert, Geert, Eric Engstrom, and Yuhang Xing, 2009, "Risk, Uncertainty, and Asset Prices", *Journal of Financial Economics* 91, 59–82.

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, "On Biases in Tests of the Expectations Hypothesis of the Term Structure of Interest Rates", *Journal of Financial Economics* 44, 309–348.

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 2001, "Peso Problem Explanations for Term Structure Anomalies", *Journal of Monetary Economics* 48, 241–270.

Binsbergen, Jules H. van and Ralph S.J. Koijen, 2008, "The Impact of Non-Linearities in Present-Value Models," unpublished paper, Stanford University and University of Chicago.

Buraschi, Andrea, Anna Cieslak, and Fabio Trojani, 2008, "Correlation Risk and the Term Structure of Interest Rates", unpublished paper, Imperial College London and University of St. Gallen.

Buraschi, Andrea and Alexei Jiltsov, 2007, "Habit Formation and Macroeconomic Models of the Term Structure of Interest Rates", *Journal of Finance* 62, 3009–3063.

Campbell, John Y., 1986, "Bond and Stock Returns in a Simple Exchange Model", *Quarterly Journal of Economics* 101, 785–804.

Campbell, John Y., 2003, "Consumption-Based Asset Pricing", Chapter 13 in George Constantinides, Milton Harris, and Rene Stulz eds. *Handbook of the Economics of Finance Vol. IB*, North-Holland, Amsterdam, 803–887.

Campbell, John Y. and John Ammer, 1993, "What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns", *Journal of Finance* 48, 3–37.

Campbell, John Y., Y. Lewis Chan, and Luis M. Viceira, 2003, "A Multivariate Model of Strategic Asset Allocation", *Journal of Financial Economics* 67, 41–80.

Campbell, John Y. and John H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy* 107, 205–251.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.

Campbell, John Y. and Robert J. Shiller, 1991, "Yield Spreads and Interest Rate Movements: A Bird's Eye View", *Review of Economic Studies* 58, 495–514.

Campbell, John Y., Robert J. Shiller, and Luis M. Viceira, 2009, "Understanding Inflation-Indexed Bond Markets", *Brookings Papers on Economic Activity* 79-120, Spring 2009.

Campbell, John Y., Adi Sunderam, and Luis M. Viceira, 2010, "Appendix to Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds", available online at http://kuznets.fas.harvard.edu/~campbell/papers.html.

Campbell, John Y. and Luis M. Viceira, 2001, "Who Should Buy Long-Term Bonds?", *American Economic Review* 91, 99–127.

Campbell, John Y. and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, New York, NY.

Campbell, John Y. and Motohiro Yogo, 2006, "Efficient Tests of Predictability in Stock Returns", *Journal of Financial Economics* 81, 27–60.

Christiansen, Charlotte and Angelo Ranaldo, 2007, "Realized Bond-Stock Correlation: Macroeconomic Announcement Effects", *Journal of Futures Markets* 27, 439–469.

Clarida, Richard, Jordi Gali, and Mark Gertler, 2000, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory", *Quarterly Journal of Economics* 115, 147–180.

Cochrane, John H. and Monika Piazzesi, 2005, "Bond Risk Premia", *American Economic Review* 95, 138–160.

Connolly, Robert, Chris Stivers, and Licheng Sun, 2005, "Stock Market Uncertainty and the Stock-Bond Return Relation", *Journal of Financial and Quantitative Analysis* 40, 161–194.

Constantinides, George M., 1992, "A Theory of the Nominal Term Structure of Interest Rates", *Review of Financial Studies* 5, 531–552.

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, 1985, "An Intertemporal General Equilibrium Model of Asset Prices", *Econometrica* 53, 363–384.

d'Addona, Stefano and Axel H. Kind, 2006, "International Stock-Bond Correlations in a Simple Affine Asset Pricing Model", *Journal of Banking and Finance* 30, 2747–2765.

Dai, Qiang and Kenneth Singleton, 2002, "Expectations Puzzles, Time-Varying Risk Premia, and Affine Models of the Term Structure", *Journal of Financial Economics* 63, 415–442.

David, Alexander and Pietro Veronesi, 2009, "What Ties Return Volatilities to Price Valuations and Fundamentals?", unpublished paper, University of Calgary and University of Chicago.

Duffee, Greg, 2002, "Term Premia and Interest Rate Forecasts in Affine Models", *Journal of Finance* 57, 405–443.

Duffie, Darrell and Rui Kan, 1996, "A Yield-Factor Model of Interest Rates", *Mathematical Finance* 6, 379–406.

Duffie, Darrell and Kenneth J. Singleton, 1993, "Simulated Moments Estimation of Markov Models of Asset Prices", *Econometrica* 61, 929–952.

Engle, Robert F., 1982, "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation", *Econometrica* 50, 987–1008.

Eraker, Bjørn, 2008, "Affine General Equilibrium Models," *Management Science* 54, 2068–2080.

Fama, Eugene F., 2006, "The Behavior of Interest Rates" *Review of Financial Studies* 19, 359-379.

Fama, Eugene F. and Robert R. Bliss, 1987, "The Information in Long-Maturity Forward Rates", *American Economic Review* 77, 680-692.

Fama, Eugene F. and Kenneth R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds", *Journal of Financial Economics* 25, 23–49.

Fama, Eugene F. and Kenneth R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds", *Journal of Financial Economics* 33, 3–56.

Gallant, Ronald and George Tauchen, 1996, "Which Moments to Match", *Econometric Theory* 12, 657-681.

Garcia, Rene and Pierre Perron, 1996, "An Analysis of the Real Interest Rate under Regime Shifts", *Review of Economics and Statistics* 78, 111–125.

Gollier, Christian, 2007, "The Consumption-Based Determinants of the Term Structure of Discount Rates", *Mathematics and Financial Economics* 1, 81–101.

Gray, Stephen F., 1996, "Modelling the Conditional Distribution of Interest Rates as a Regime-Switching Process", *Journal of Financial Economics* 42, 27–62.

Guidolin, Massimo and Allan Timmermann, 2006, "An Econometric Model of Nonlinear Dynamics in the Joint Distribution of Stock and Bond Returns", *Journal of Applied Econometrics* 21, 1–22.

Gurkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2006, "The U.S. Treasury Yield Curve: 1961 to the Present", FEDS working paper 2006-28, available online at http://www.federalreserve.gov/pubs/feds/2006/200628.

Gurkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2008, "The TIPS Yield Curve and Inflation Compensation", FEDS working paper 2008-05, available online at http://www.federalreserve.gov/pubs/feds/2008/200805.

Hansen, Lars P. and Robert J. Hodrick, 1983, "Risk Averse Speculation in the Forward Foreign Exchange Market: An Econometric Analysis of Linear Models", in J. Frenkel (ed.) *Exchange Rates and International Macroeconomics*, University of Chicago Press, Chicago.

Hasseltoft, Henrik, 2008, "The 'Fed-Model' and the Changing Correlation of Stock and Bond Returns: An Equilibrium Approach," unpublished paper, Stockholm School of Economics.

Julier, S. and J. Uhlmann, 1997, "A New Extension of the Kalman Filter to Nonlinear Systems", in *Proceedings SPIE Signal Processing, Sensor Fusion, and Target Recognition VI,* 3068, 182–193.

Lewellen, Jonathan, 2004, "Predicting Returns with Financial Ratios", *Journal of Financial Economics* 74, 209–235.

Lettau, Martin, and Jessica Wachter, 2007, "Why is Long-horizon Equity Less Risky? A Duration-based Explanation of the Value Premium," *Journal of Finance*, 62(1), 55-92.

Lettau, Martin and Jessica Wachter, 2010, "The Term Structures of Equity and Interest Rates", forthcoming *Journal of Financial Economics.*

Li, Lingfeng, 2002, "Macroeconomic Factors and the Correlation of Stock and Bond Returns", Yale ICF Working Paper 02-46.

Litterman, Robert and Jose Scheinkman, 1991, "Common Factors Affecting Bond Returns", *Journal of Fixed Income*, June.

Mamaysky, Harry, 2002, "Market Prices of Risk and Return Predictability in a Joint Stock-Bond Pricing Model", Yale ICF Working Paper 02-25.

McCulloch, J. Huston and H. Kwon, 1993, "US Term Structure Data, 1947–1991", Working Paper 93-6, Ohio State University.

Piazzesi, Monika and Martin Schneider, 2006, "Equilibrium Yield Curves", in D. Acemoglu, K. Rogoff, and M. Woodford (eds.) *NBER Macroeconomics Annual 2006*, 317-363, MIT Press, Cambridge, MA.

Realdon, Marco, 2006, "Quadratic Term Structure Models in Discrete Time", *Finance Research Letters* 3, 277–289.

Rudebusch, Glenn and Tao Wu, 2007, "Accounting for a Shift in Term Structure Behavior with No-Arbitrage and Macro-Finance Models", *Journal of Money, Credit, and Banking* 39, 395–422.

Shiller, Robert J. and Andrea Beltratti, 1992, "Stock Prices and Bond Yields: Can Their Comovements Be Explained in Terms of Present Value Models?", *Journal of Monetary Economics* 30, 25–46.

Stambaugh, Robert F., 1988, "The Information in Forward Rates: Implications for Models of the Term Structure", *Journal of Financial Economics* 21, 41–70..

Stambaugh, Robert F., 1999, "Predictive Regressions", *Journal of Financial Economics* 54:375-421, 1999.

Vasicek, Oldrich, 1977, "An Equilibrium Characterization of the Term Structure", *Journal of Financial Economics* 5, 177–188.

Viceira, Luis M., 2010, "Bond Risk, Bond Return Volatility, and the Term Structure of Interest Rates", forthcoming *International Journal of Forecasting*.

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

Wan, Eric A. and Rudolph van der Merwe, 2001, "The Unscented Kalman Filter", Chapter 7 in Simon Haykin ed., *Kalman Filtering and Neural Networks*, Wiley, New York, NY.

Table 1: Parameter estimates.

Parameter	Estimate	Std Err
Parameter Estimates		
μ_x x 10^3	8.889	0.338
μ_ψ x 10^3	7.419	1.165
ϕ_x	0.939	0.004
ϕ_ξ	0.885	0.008
ϕ_ψ	0.751	0.030
σ_m x 10^2	6.815	1.003
σ_X x 10^3	1.133	0.092
σ_x x 10^1	2.894	0.373
σ_λ x 10^4	3.052	4.905
σ_Λ x 10^4	5.725	1.160
σ_ξ x 10^1	2.142	0.413
σ_ψ x 10^3	3.843	0.460
β_{eX}	1.136	5.658
β_{ex} x 10^2	-0.005	4.736
β_{em} x 10^2	8.748	4.007
$\rho_{x\xi}$	0.000	0.121
ρ_{xm}	-0.323	0.105
$\rho_{x\pi}$	0.000	0.549
$\rho_{\lambda m}$	-0.797	1.302
$\rho_{\xi m}$	-0.237	0.156
$\rho_{\xi\pi}$	-0.136	1.168
$\rho_{m\pi}$	0.011	0.060

Table 2: Sample and Implied Moments. Yield spreads (YS) are calculated over the 3mo yield. Realized excess returns (RXR) are calculated over a 3mo holding period, in excess of the 1yr yield. Units are annualized percentage points. Simulation columns report means across 1000 replications, each of which simulates a time-series of 250 quarters. The $\sigma(\widehat{CP})$ row reports the standard deviation of the fitted values from a Cochrane-Piazzesi style regression of RXR on the 1-, 3-, and 5-yr forward rates at the beginning of the holding period. The $\sigma(\widehat{CS})$ row reports the standard deviation of the fitted values from a Campbell-Shiller style regression of RXR on the same-maturity YS at the beginning of the holding period. In the rightmost column we report the fraction of simulation runs where the simulated value exceeds the data value. [†] Data moments for the 10yr return require 117mo yields. We interpolate the 117mo yield linearly between the 5yr and the 10yr [‡] TIPS entries refer to the 10yr spliced TIPS yield. We have this data 1/1999-9/2009.

Sample and Implied Moments			
Moment	Actual Data	Model	Above
3yr YS mean	0.62	0.54	0.44
10yr YS mean	1.15	0.83	0.36
3yr YS stdev	0.45	0.57	0.89
10yr YS stdev	0.70	1.18	1.00
3yr RXR mean	1.17	1.11	0.49
10yr RXR mean	2.21	1.71	0.39
3yr RXR stdev	4.37	5.64	1.00
10yr RXR stdev	11.16	10.57	0.24
10yr TIPS yield mean	2.58[‡]	3.51	0.98
10yr TIPS YS mean		0.02	
10yr TIPS RXR mean		0.36	
10yr TIPS RXR stdev		8.82	

Predictive Regressions			
Moment	Actual Data	Model	Above
3yr EXR stdev		0.09	
10yr EXR stdev		0.15	
10yr TIPS EXR stdev		0.11	
3yr RXR $\sigma(\widehat{CS})$	1.04	0.34	0.01
10yr RXR $\sigma(\widehat{CS})$	2.51[†]	0.61	0.00
10yr TIPS RXR $\sigma(\widehat{CS})$		0.53	
3yr RXR $\sigma(\widehat{CP})$	0.79	0.70	0.36
10yr RXR $\sigma(\widehat{CP})$	2.06[†]	1.33	0.11

Table 3: Sample and Implied Moments. Yield spreads (YS) are calculated over the 1yr yield. Realized excess returns (RXR) are calculated over a 1yr holding period, in excess of the 1yr yield. Units are annualized percentage points. Simulation columns report means across 1000 replications, each of which simulates a time-series of 250 quarters. The $\sigma(\widehat{CP})$ row reports the standard deviation of the fitted values from a Cochrane-Piazzesi style regression of RXR on the 1-, 3-, and 5-yr forward rates at the beginning of the holding period. The $\sigma(\widehat{CS})$ row reports the standard deviation of the fitted values from a Campbell-Shiller style regression of RXR on the same-maturity YS at the beginning of the holding period. In the rightmost column we report the fraction of simulation runs where the simulated value exceeds the data value. [†] Data moments for the 10yr return require 9yr yields. We interpolate the 9yr yield linearly between the 5yr and the 10yr. [‡] TIPS entries refer to the 10yr spliced TIPS yield. We have this data 1/1999-9/2009.

Sample and Implied Moments			
Moment	Actual Data	Model	Above
3yr YS mean	0.37	0.31	0.42
10yr YS mean	0.90	0.57	0.28
3yr YS stdev	0.51	0.75	0.99
10yr YS stdev	1.11	1.96	1.00
3yr RXR mean	0.75	0.66	0.43
10yr RXR mean	1.84[†]	1.25	0.30
3yr RXR stdev	3.17	4.00	0.97
10yr RXR stdev	10.32[†]	9.48	0.18
10yr TIPS yield mean	2.58[‡]	3.54	0.99
10yr TIPS YS mean		-0.05	
10yr TIPS RXR mean		0.17	
10yr TIPS RXR stdev		7.80	

Predictive Regressions			
Moment	Actual Data	Model	Above
3yr EXR stdev		0.08	
10yr EXR stdev		0.21	
10yr TIPS EXR stdev		0.19	
3yr RXR $\sigma(\widehat{CS})$	0.91	0.41	0.08
10yr RXR $\sigma(\widehat{CS})$	3.54[†]	0.95	0.00
10yr TIPS RXR $\sigma(\widehat{CS})$		0.83	
3yr RXR $\sigma(\widehat{CP})$	1.16	0.87	0.23
10yr RXR $\sigma(\widehat{CP})$	4.51[†]	2.09	0.00

Table 4: Forecasting excess returns. The table below reports the R^2 for regressions in our data of actual data RXR on linear combinations of the actual data 1-, 3-, and 5-yr forward rates at the beginning of the holding period. The unconstrained column estimates the best combination in the data, and thus corresponds to the first stage of the Cochrane-Piazzesi procedure. In the other columns, the combination is restricted to be the one estimated in long-sample simulation regressions of simulated RXR on simulated forward rates. In the first panel, we allow this simulation-generated combination to be scaled up. In the second panel, we do not allow scaling. Realized excess returns (RXR) are calculated over 3mo and 1yr holding periods. [†] Data moments for the 10yr return require 9yr yields. These yields are in our dataset 8/1971-9/2009. For the earlier part of the sample we interpolate the 9yr yield linearly between the 5yr and the 10yr.

Forecasting Excess Returns			
Holding Period	Moment	Unconstrained	Model
3-month	3yr RXR	0.032	0.022
3-month	10yr RXR	0.031	0.013
1-yr	3yr RXR	0.132	0.093
1-yr	10yr RXR	0.181[†]	0.080

Forecasting Excess Returns: No scaling			
Holding Period	Moment	Unconstrained	Model
3-month	3yr RXR	0.032	0.001
3-month	10yr RXR	0.022	0.000
1-yr	3yr RXR	0.127	0.007
1-yr	10yr RXR	0.097[†]	0.002
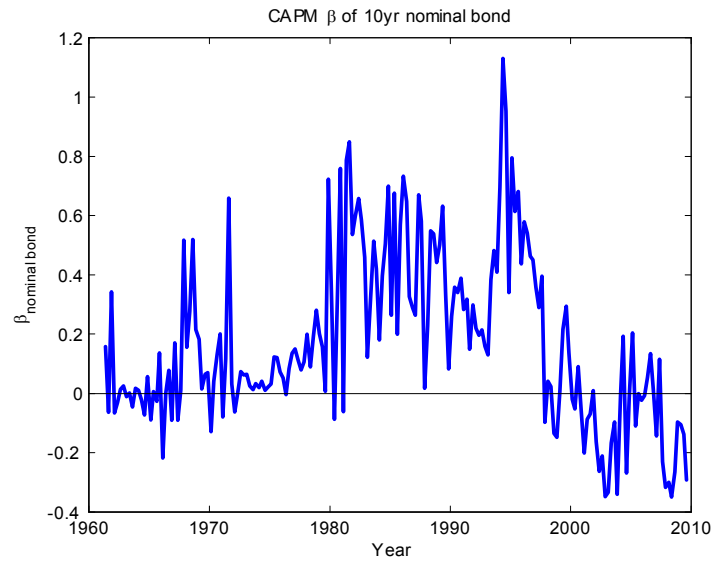
Figure 1: Time series of the CAPM β of the 10-year nominal bond.
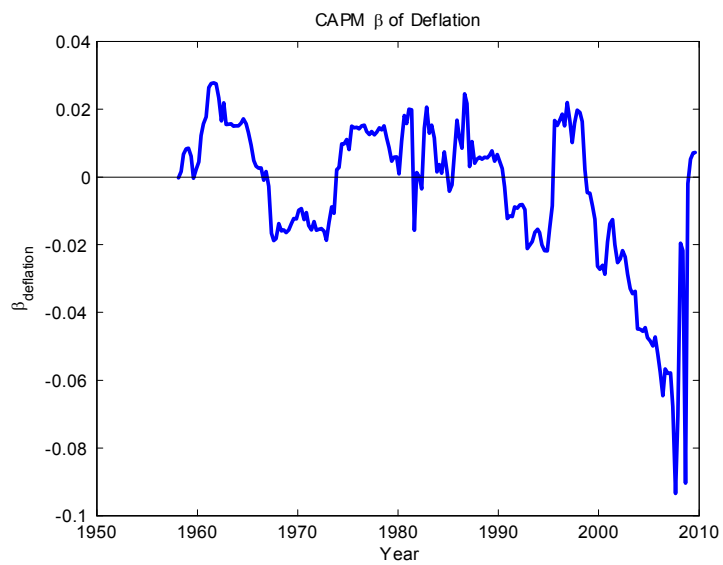
Figure 2: Time series of the CAPM β of deflation.
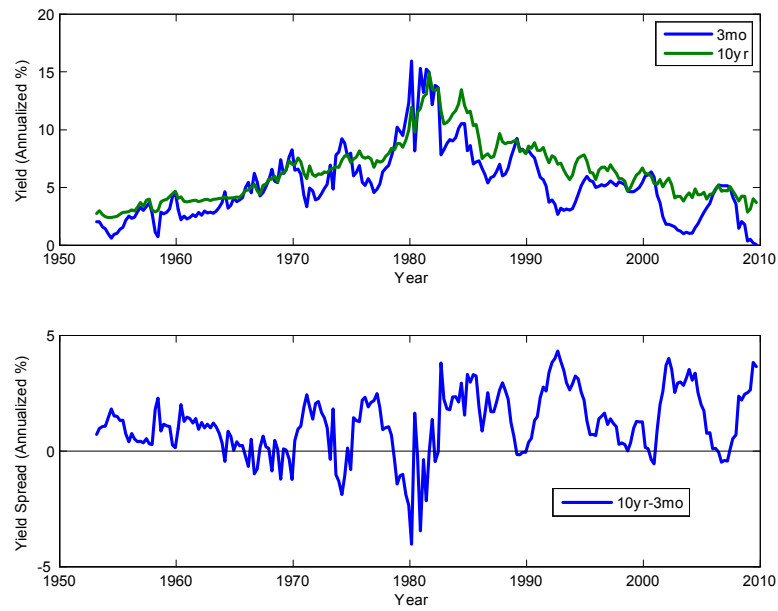
Figure 3: Time series of 3-month and 10-year nominal yields and yield spread.
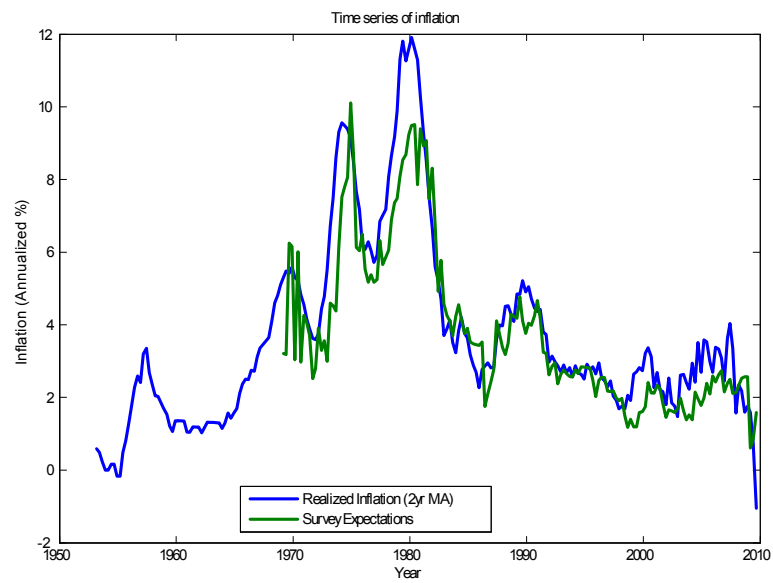
Figure 4: Time series of inflation.
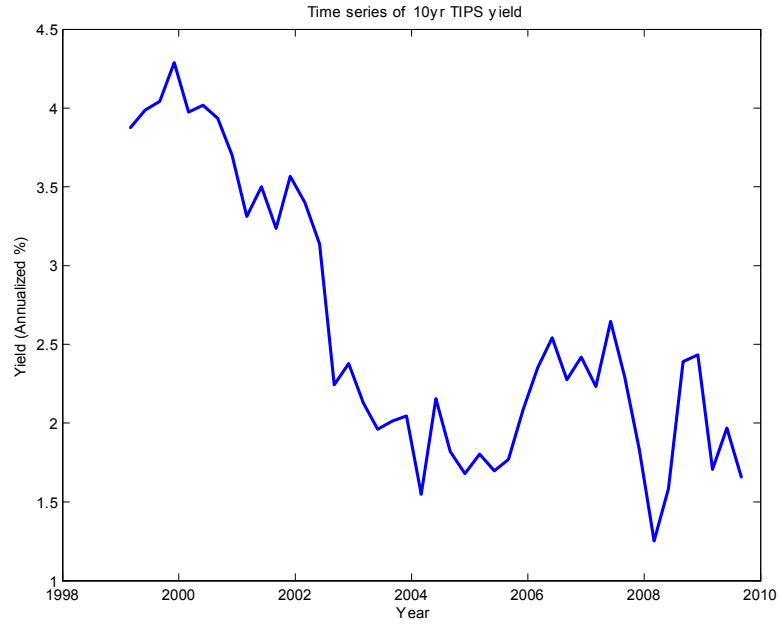
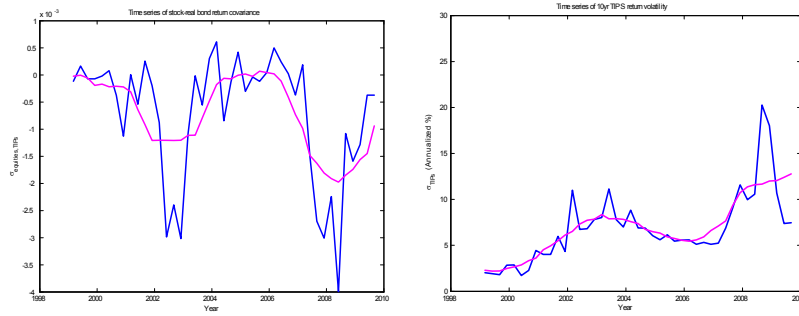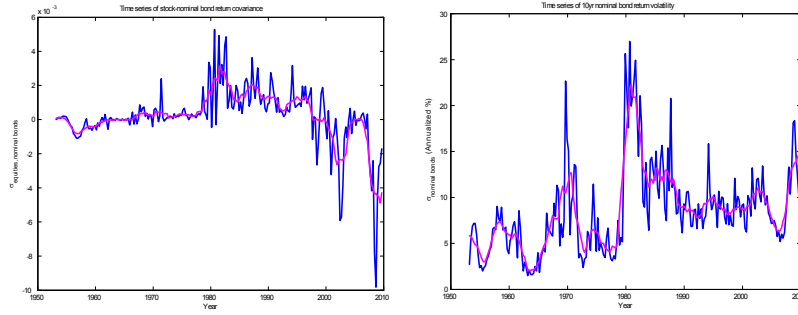**Figure 5: Time series of US 10-year inflation-indexed yields.**

Figure 6: Time series of real bond second moments. The figure on the left shows the covariance between stock and real bond returns. The figure on the right shows variance of real bond returns. The smoothed line in each figure is a 2-year equal-weighted moving average.

Figure 7: Time series of nominal bond second moments. The figure on the left shows the covariance between stock and nominal bond returns. The figure on the right shows variance of nominal bond returns. The smoothed line in each figure is a 2-year equal-weighted moving average.
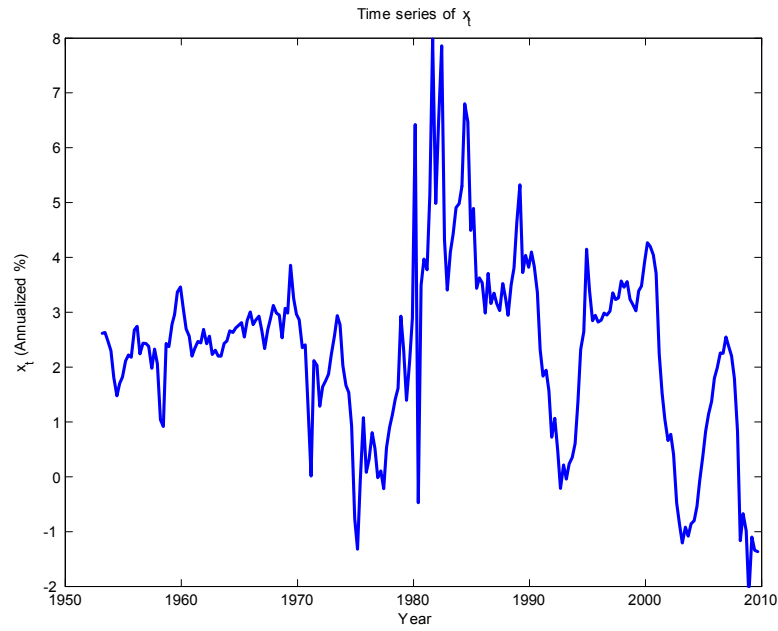
Figure 8: Time series of real rate. The figure on the left plots the time series of x_t, the real interest rate.
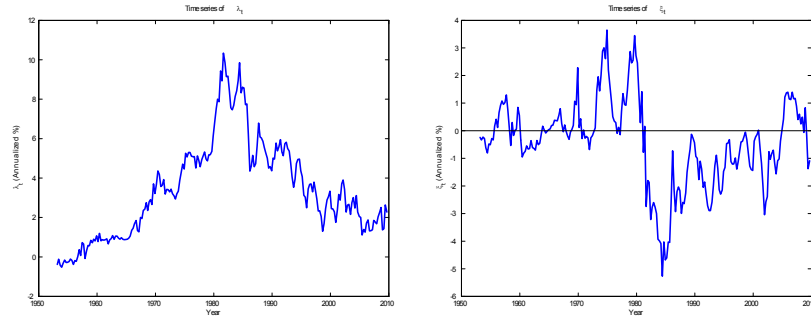
Figure 9: Time series of permanent and transitory components of expected inflation. The figure on the left plots the time series of λ_t, the permanent component of expected inflation. The figure on the right plots the time series of ξ_t, the temporary component of expected inflation.
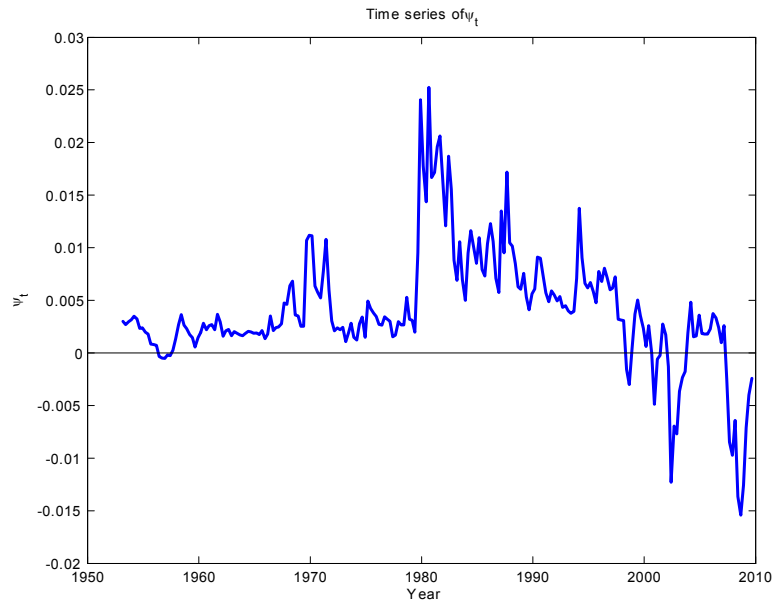
Figure 10: Time series of ψ_t.
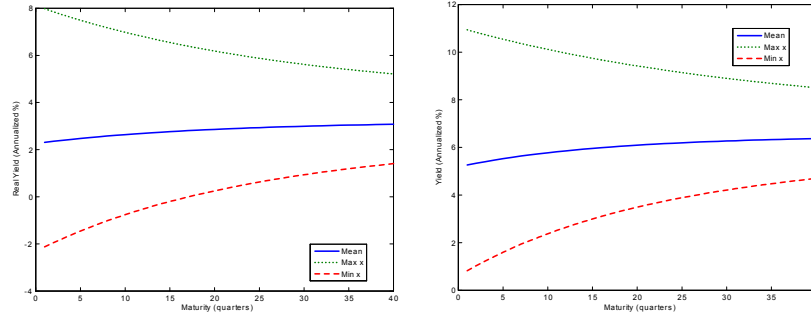
Figure 11: Responses of yield curves to x_t. The left hand figure shows the response of the real yield curve, and the right hand figure shows the response of the nominal yield curve, as x_t is varied between its sample minimum and maximum while all other state variables are held fixed at their sample means.
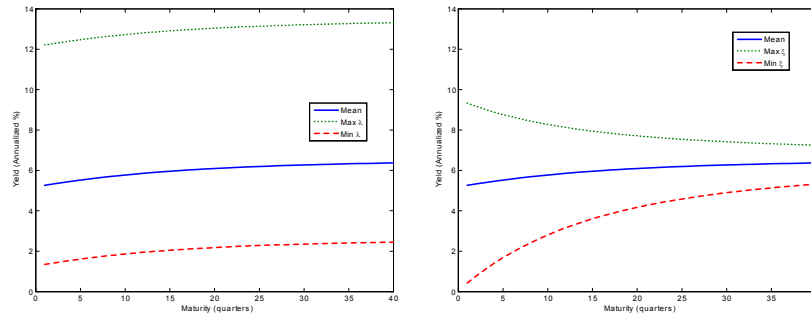
Figure 12: Responses of yield curves to λ_t **and** ξ_t. The left hand figure shows the response of the nominal yield curve to the permanent component of expected inflation λ_t, and the right hand figure shows the response to the transitory component of expected inflation ξ_t, as these state variables are varied between their sample minima and maxima while all other state variables are held fixed at their sample means.
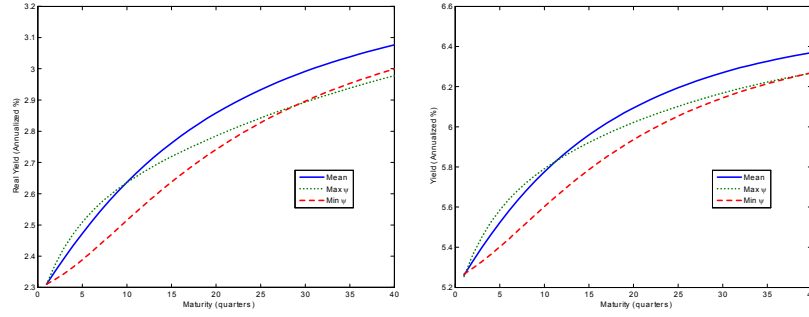
Figure 13: Responses of yield curves to ψ_t . The left hand figure shows the response of the real yield curve, and the right hand figure shows the response of the nominal yield curve, to ψ_t as it is varied between its sample minima and maxima while all other state variables are held fixed at their sample means.
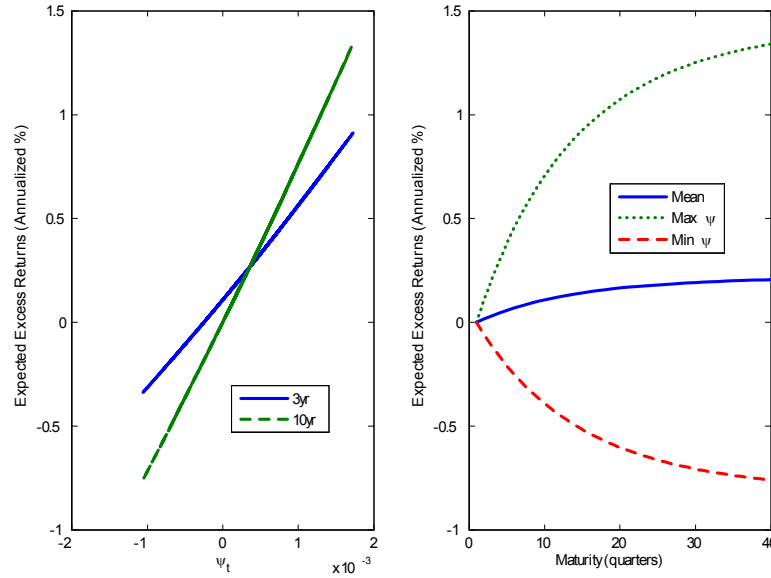

Figure 14: Responses of nominal expected excess returns to ψ_t**.** The left hand figure shows the expected excess returns on 3-year and 10-year nominal bonds over 3-month Treasury bills, as functions of ψ_t. The right hand figure shows the term structure of expected excess nominal bond returns as ψ_t is varied between its sample minimum and maximum while all other state variables are held fixed at their sample means.
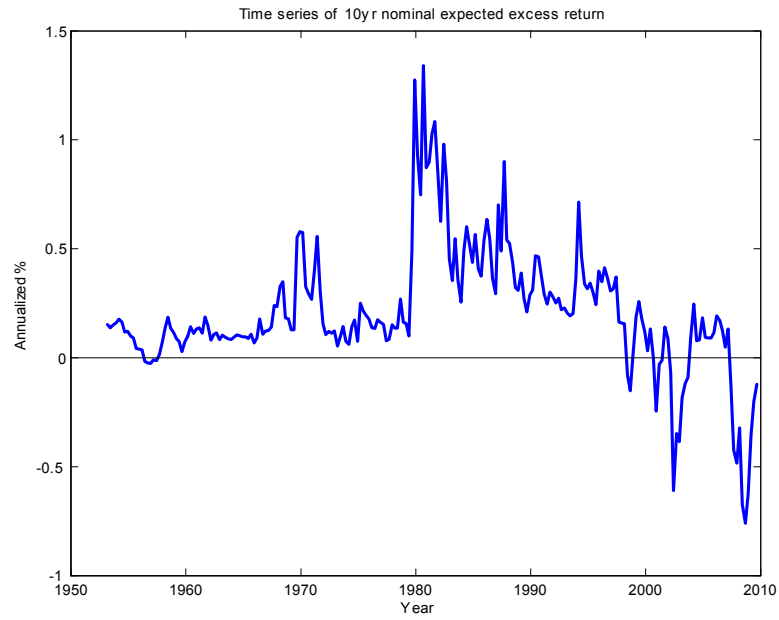
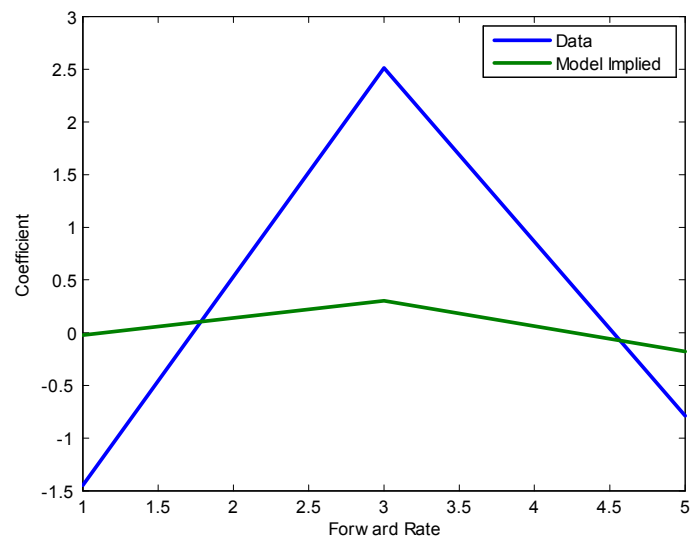**Figure 15: Time series of expected excess returns for 10-year nominal bonds.**

Figure 16: Data- and model- implied Cochrane-Piazzesi relationships.