

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

John Y. Campbell, Adi Sunderam, and Luis M. Viceira¹

First draft: June 2007
This version: January 2009

¹Campbell: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, USA, and NBER. Email john_campbell@harvard.edu. Sunderam: Harvard Business School, Boston MA 02163. Email asunderam@hbs.edu. Viceira: Harvard Business School, Boston MA 02163 and NBER. Email lviceira@hbs.edu. We are grateful to Geert Bekaert, Jesus Fernandez-Villaverde, Wayne Ferson, Javier Gil-Bazo, Pablo Guerron, John Heaton, Ravi Jagannathan, Jon Lewellen, Monika Piazzesi, Pedro Santa-Clara, George Tauchen, and seminar participants at the 2009 Annual Meeting of the American Finance Association, Bank of England, European Group of Risk and Insurance Economists 2008 Meeting, Harvard Business School Finance Unit Research Retreat, Imperial College, Marshall School of Business, NBER Fall 2008 Asset Pricing Meeting, Norges Bank, Society for Economic Dynamics 2008 Meeting, Stockholm School of Economics, Tilburg University, Tuck Business School, and Universidad Carlos III in Madrid for helpful comments and suggestions. This material is based upon work supported by the National Science Foundation under Grant No. 0214061 to Campbell, and by Harvard Business School Research Funding.

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 1953–2005, it was particularly high in the early 1980’s and negative in the early 2000’s. 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 covariance between nominal variables and the real economy. The last of these state variables enables the model to fit the changing covariance of bond and stock returns. Log nominal 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 have performed well as hedges during the financial crisis of 2008, but the opposite was true in the late 1970's and early 1980's. The purpose of this paper is to explore such changes over time in the risks of nominal government bonds.

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 2004, Baele, Bekaert, and Inghelbrecht 2007, Christiansen and Rinaldo 2007, Viceira 2007, David and Veronesi 2008). 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 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 was high in the 1980's and much lower in the late 1990's.

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 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, Gollier 2005, Piazzesi and Schneider 2006); 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.

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

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).

In this paper we specify and estimate a model that tracks the economic shocks driving bond returns, and that allows the covariances of shocks, in particular the covariance of inflation with real variables, 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 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 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.

Our approach 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 (2008) specify term structure models in which risk aversion varies over time, influencing the shape of the yield curve. These papers take care to remain in the essentially affine class described by Duffee (2002). Bansal and Shaliastovich (2008) also specify an affine term structure model that explores the implications of the long-run consumption risk model of Bansal and Yaron (2004) for interest rates. Bekaert et al. and other recent authors including Maysky (2002) and d'Addona and Kind (2005) extend affine term structure models to price stocks as well as bonds. Piazzesi and Schneider (2006), Palomino (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 inflation and real shocks, which can switch sign, means that we cannot write log bond yields as affine functions

³In 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.

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 Kojien and van Binsbergen (2008).

The organization of the paper is as follows. Section 2 presents our model of the nominal term structure. Section 3 describes our estimation method and presents parameter estimates and historical fitted values for the unobservable state variables of the model. Section 4 discusses the implications of the model for the shape of the yield curve and the movements of risk premia on nominal bonds. Section 5 concludes. An Appendix to this paper available online (Campbell, Sunderam, and Viceira 2009) presents details of the model solution and additional empirical results.

2 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. However, unlike their model, this model allows for time variation in the risk premia on both real and nominal assets, and for time variation in the covariance between the real economy and inflation and thus between the excess returns on real assets and the returns on nominal assets. The model for the real term structure of interest rates allows for time variation in both real interest rates and risk premia, yet it is simple enough that real bond prices have an exponential affine structure. The nominal side of the model allows for time variation in transitory and persistent components of expected inflation, the volatility of inflation, and the conditional covariance of inflation with the real side of the economy. This results in a nominal term structure where bond yields are linear-quadratic functions of the vector of state variables.

⁴Duffie 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.

2.1 An affine model of the real term structure

We pose a model for the term structure of real interest rates that has a simple linear structure. We assume that the log of the real stochastic discount factor (SDF) $m_{t+1} = \log(M_{t+1})$ follows a linear-quadratic, conditionally heteroskedastic process:

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

where both x_t and z_t follow standard AR(1) processes,

$$x_{t+1} = \mu_x (1 - \phi_x) + \phi_x x_t + \varepsilon_{x,t+1}, \quad (2)$$

$$z_{t+1} = \mu_z (1 - \phi_z) + \phi_z z_t + \varepsilon_{z,t+1}, \quad (3)$$

and $\varepsilon_{m,t+1}$, $\varepsilon_{x,t+1}$, and $\varepsilon_{z,t+1}$ are jointly normally distributed zero-mean shocks with 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 1.

Even though shocks ε are homoskedastic, the log real SDF itself is conditionally heteroskedastic, with

$$\text{Var}_t(m_{t+1}) = z_t^2.$$

The state variable z_t drives the time-varying volatility of the SDF or, equivalently, the price of aggregate market risk or maximum Sharpe ratio in the economy.

This way of modeling time variation in real risk premia is similar to the approach of Lettau and Wachter (2007a,b). 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 it 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.⁵ 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 x_t determines the dynamics of the short-term log real interest rate. The price of a single-period zero-coupon real bond satisfies

$$P_{1,t} = \mathbb{E}_t [\exp \{m_{t+1}\}],$$

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

$$y_{1t} = -\mathbb{E}_t [m_{t+1}] - \frac{1}{2} \text{Var}_t (m_{t+1}) = x_t. \quad (4)$$

Thus the model (1)-(3) allows for time variation in risk premia, yet it preserves simple linear dynamics for the short-term real interest rate.

This model implies that the real term structure of interest rates is affine in the state variables x_t and z_t . Standard calculations (Campbell, Lo, and MacKinlay 1997, Chapter 11) show that the price of a zero-coupon real bond with n periods to maturity is given by

$$P_{n,t} = \exp \{A_n + B_{x,n}x_t + B_{z,n}z_t\}, \quad (5)$$

where

$$\begin{aligned} A_n &= A_{n-1} + B_{x,n-1}\mu_x(1 - \phi_x) + B_{z,n-1}\mu_z(1 - \phi_z) \\ &\quad + \frac{1}{2}B_{x,n-1}^2\sigma_x^2 + \frac{1}{2}B_{z,n-1}^2\sigma_z^2 + B_{x,n-1}B_{z,n-1}\sigma_{xz}, \end{aligned}$$

$$B_{x,n} = -1 + B_{x,n-1}\phi_x,$$

and

$$B_{z,n} = B_{z,n-1}\phi_z - B_{x,n-1}\sigma_{mx} - B_{z,n-1}\sigma_{mz},$$

with $A_1 = 0$, $B_{x,1} = -1$, and $B_{z,1} = 0$. Note that $B_{x,n} < 0$ for all n when $\phi_x > 0$. Details of these calculations are presented in the Appendix (Campbell, Sunderam, and Viceira 2009).

⁵Under 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 .

The excess log return on a n -period zero-coupon real bond over a 1-period real bond is given by

$$\begin{aligned}
r_{n,t+1} - r_{1,t+1} &= p_{n-1,t+1} - p_{n,t} + p_{1,t} \\
&= - \left(\frac{1}{2} B_{x,n-1}^2 \sigma_x^2 + \frac{1}{2} B_{z,n-1}^2 \sigma_z^2 + B_{x,n-1} B_{z,n-1} \sigma_{xz} \right) \\
&\quad + (B_{x,n-1} \sigma_{mx} + B_{z,n-1} \sigma_{mz}) z_t \\
&\quad + B_{x,n-1} \varepsilon_{x,t+1} + B_{z,n-1} \varepsilon_{z,t+1}, \tag{6}
\end{aligned}$$

where the first term is a Jensen's inequality correction, the second term describes the log of the expected excess return on real bonds, and the third term describes shocks to realized excess returns. Note that $r_{1,t+1} \equiv y_{1,t}$.

It follows from (6) that the conditional risk premium on real bonds is

$$\text{E}_t [r_{n,t+1} - r_{1,t+1}] + \frac{1}{2} \text{Var}_t (r_{n,t+1} - r_{1,t+1}) = (B_{x,n-1} \sigma_{mx} + B_{z,n-1} \sigma_{mz}) z_t, \tag{7}$$

which is proportional to the state variable z_t . The coefficient of proportionality is $(B_{x,n-1} \sigma_{mx} + B_{z,n-1} \sigma_{mz})$, which can take either sign. It is zero, and thus real bond risk premia are zero, when $\sigma_{mx} = 0$, that is, when shocks to real interest rates are uncorrelated with the stochastic discount factor.⁶ Real bond risk premia are also zero when the state variable z_t is zero, for then the stochastic discount factor is a constant which implies risk-neutral asset pricing.

To gain intuition about the behavior of risk premia on real bonds, consider the simple case where $\sigma_{mz} = 0$ and $\sigma_{mx} > 0$. Since $B_{x,n-1} < 0$, this implies that real bond risk premia are negative. The reason for this is that with positive σ_{mx} , 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. Empirically, however, we estimate a negative σ_{mx} ; this implies procyclical real bond returns that command a positive risk premium, increasing with the level of risk aversion.

2.2 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

⁶Note that $\sigma_{mx} = 0$ implies $B_{z,n} = 0$, for all n .

would be to specify a dividend process and solve for the stock return endogenously in the manner of Bekaert et al. (2005), Mamaysky (2002), and d’Addona and Kind (2005). 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:

$$r_{e,t+1} - E_t r_{e,t+1} = \beta_{ex}\varepsilon_{x,t+1} + \beta_{em}\varepsilon_{m,t+1} + \varepsilon_{e,t+1}, \quad (8)$$

where $\varepsilon_{e,t+1}$ is an identically and independently distributed shock uncorrelated with all other shocks in the model. This shock captures variation in equity returns unrelated to real interest rates, which are not priced because they are uncorrelated with the SDF.

Substituting (8) into the no-arbitrage condition $E_t [M_{t+1}R_{t+1}] = 1$, the conditional equity risk premium is given by

$$E_t [r_{e,t+1} - r_{1,t+1}] + \frac{1}{2} \text{Var}_t (r_{e,t+1} - r_{1,t+1}) = (\beta_{ex}\sigma_{xm} + \beta_{em}\sigma_m^2) z_t. \quad (9)$$

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.

2.3 A model of time-varying inflation risk

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:

$$\pi_{t+1} = \lambda_t + \xi_t + \frac{\sigma_\pi^2}{2}\psi_t^2 + \psi_t\varepsilon_{\pi,t+1}, \quad (10)$$

where expected log inflation is the sum of two components, a permanent component λ_t and a transitory component ξ_t , which follow

$$\lambda_{t+1} = \lambda_t + \varepsilon_{\Lambda,t+1} + \psi_t\varepsilon_{\lambda,t+1}, \quad (11)$$

and

$$\xi_{t+1} = \phi_\xi\xi_t + \psi_t\varepsilon_{\xi,t+1}. \quad (12)$$

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

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 the nominal term structure of interest rates at both ends of the maturity spectrum, the dynamics of realized inflation, and dynamics of the yields on inflation-indexed bonds.

Of course, 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. We have estimated though our model imposing that λ_t follows a stationary process with a highly persistent autoregressive coefficient. In practice this makes no discernible changes to our main empirical conclusions.

The most important innovation in our model is the inclusion of the state variable ψ_t , which multiplies the underlying shocks that drive realized and expected inflation. We specify ψ_t as an AR(1) process with a nonzero mean:

$$\psi_{t+1} = \mu_\psi (1 - \phi_\psi) + \phi_\psi \psi_t + \varepsilon_{\psi,t+1}. \quad (13)$$

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

⁷Without 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.

m_{t+1} , x_{t+1} , and z_{t+1} , and use the same notation as in Section 2.1 to denote their variances and covariances.

This specification implies that the conditional volatility of inflation is 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 a persistent state variable ψ_t which drives the volatility of expected as well as realized inflation. 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 or the covariance between them, 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 real economic variables, on the other hand.

The state variable ψ_t governs the second moments not only of realized inflation, but also of expected inflation. We could assume different processes driving the second moments of realized and expected inflation, but this would increase the complexity of the model considerably. Long-term 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 and realized inflation.

Our specification of the expected inflation process allows for 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. In the absence of a homoskedastic shock to expected inflation, the 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. Our specification avoids imposing proportionality while preserving the parsimony of the model.

Finally, we note that the process for realized inflation, equation (10), 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.

2.4 The short-term nominal interest rate

We now show how to price a single-period nominal bond and derive 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

$$P_{1,t}^{\$} = \mathbf{E}_t [\exp \{m_{t+1} - \pi_{t+1}\}], \quad (14)$$

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

$$\begin{aligned} y_{1,t+1}^{\$} &= -\mathbf{E}_t [m_{t+1} - \pi_{t+1}] - \frac{1}{2} \text{Var}_t (m_{t+1} - \pi_{t+1}) \\ &= x_t + \lambda_t + \xi_t - \sigma_{m\pi} z_t \psi_t, \end{aligned} \quad (15)$$

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

Equation (15) 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. It is straightforward to show that the nonlinear term in (15) is the expected excess 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 SDF:

$$\text{Cov}_t (m_{t+1}, \pi_{t+1}) = -\sigma_{m\pi} z_t \psi_t. \quad (16)$$

When this 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 covariance between inflation and the SDF is determined by the product of two state variables, z_t and ψ_t . Although both variables influence the magnitude of

the covariance, its sign is determined in practice only by ψ_t because, even though we do not constrain z_t to be positive, we estimate it to be so in our sample, consistent with the notion that z_t is a proxy for aggregate risk aversion. Therefore, the state variable ψ_t controls not only the conditional volatility of inflation, but also the sign of the correlation between inflation and the SDF.

This property of the single-period nominal risk premium carries over to the entire nominal term structure. In our model the risk premium on real assets varies over time and increases or decreases as a function of aggregate risk aversion, as shown in (7) or (9). 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.

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 (8) and (10), this covariance is given by

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

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.

2.5 A quadratic model of the nominal term structure

Equation (15) shows that the log nominal short rate is a linear-quadratic function of the state variables in our model. We show in the Appendix that this property carries over to the entire zero-coupon nominal term structure. The price of a n -period zero-coupon nominal bond is an exponential linear-quadratic function of the vector of state

variables:

$$P_{n,t}^{\$} = \exp \left\{ \begin{aligned} &A_n^{\$} + B_{x,n}^{\$}x_t + B_{z,n}^{\$}z_t + B_{\lambda,n}^{\$}\lambda_t + B_{\xi,n}^{\$}\xi_t + B_{\psi,n}^{\$}\psi_t \\ &+ C_{z,n}^{\$}z_t^2 + C_{\psi,n}^{\$}\psi_t^2 + C_{z\psi,n}^{\$}z_t\psi_t \end{aligned} \right\}, \quad (17)$$

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 (15), 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 $n = 1$.

Equation (17) shows that the nominal term structure of interest rates is a linear-quadratic function of the vector of state variables. 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. In our model, expected bond excess returns are time varying in risk aversion (z_t) and the covariance between the log real SDF and inflation ($z_t\psi_t$).

Specifically, 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 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 as ψ_t switches sign over time.

2.6 Special cases

Our quadratic term structure model nests three constrained models of particular interest. First, if we constrain z_t to be constant but allow ψ_t to vary over time, 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 report estimates of this constrained model, which has many of the same properties as our unconstrained model.

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. We report estimates of this constrained model also.

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.

3 Model Estimation

3.1 Data and estimation methodology

The term structure model presented in Section 2 generates nominal bond yields which are linear-quadratic functions of a vector of latent state variables. We now take this model to the data, and present maximum likelihood estimates of the model based on 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 Koijen and van Binsbergen (2008) apply the method

to a prediction problem in finance.⁸

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 ten measurement equations in total.

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 (17). We use yields on constant maturity 3-month, 1-year, 3-year and 10-year zero-coupon nominal bonds sampled at a quarterly frequency from a monthly dataset for the period January 1953-December 2005. This dataset is spliced together from two sources. From January 1953 through July 1971 we use data from McCulloch and Kwon (1993) and from August 1971 through December 2005, we use data from the Federal Reserve Board constructed by Gürkaynak, Sack, and Wright (2006). 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 1 illustrates our nominal interest rate data by plotting the 3-month and 10-year nominal yields, and the spread between them, over the period 1953-2005. 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

⁸Koijen and van Binsbergen'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 stable 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.

movements in the covariance of nominal and real variables.

Our fifth measurement equation is given by equation (10), 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.

The sixth 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. Because the history of TIPS is relatively short—we have data only for the period January 1998 through December 2007—we use data on constant maturity UK inflation-linked gilts to construct a hypothetical sample of TIPS yields back to January 1985. Specifically, we splice together 10-year UK inflation-linked gilt yields for the period January 1985-December 1999 with 10-year TIPS yields for the period Jan 2000-Dec 2005.⁹ Before 1985, 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 2 illustrates our real bond yield series. The decline in UK inflation-indexed yields since the mid-1990's, and in US TIPS yields since the year 2000, are clearly visible in this figure. The divergence of the two inflation-indexed series around the turn of the millennium is a puzzle that may in part be explained by the immaturity of the TIPS market in this period.

Our seventh and eighth measurement equations use data on an equity index, the CRSP value-weighted portfolio comprising the stocks traded in the NYSE, AMEX and NASDAQ. The seventh equation describes realized log equity returns $r_{e,t+1}$ using equations (4), (8), and (9). The eighth uses the dividend yield on equities $D_{e,t}/P_{e,t}$, measured with a one-year backward moving average of dividends, relating

⁹We take historical yield series for TIPS and inflation-indexed gilts from the Global Financial Database. We have also estimated our model using a time series of fitted TIPS data from January 1985 through December 2005. There, we estimate a regression of 10-year TIPS yields on a constant and 10-year UK inflation-linked gilt yields for the period January 1999-December 2005. We then use the fitted values of the TIPS yield as our observed time series of TIPS yields for the period January 1985-December 1999, and observed TIPS yields for the period Jan 2000-Dec 2005. This modification makes little difference to our model estimates. In the Appendix we also report results when we drop the TIPS measurement equation altogether. Without the evidence of variable long-term TIPS yields, we estimate a less persistent real interest rate, but other properties of our model are little changed.

the dividend yield to z_t as

$$\frac{D_{e,t}}{P_{e,t}} = d_0 + d_1 z_t + \varepsilon_{D/P,t+1}, \quad (18)$$

where $\varepsilon_{D/P,t+1}$ is a measurement error term uncorrelated with the fundamental shocks of the model. This measurement equation is motivated by the fact that the dividend yield appears to forecast future equity returns, and that in our model expected equity excess returns are proportional to z_t , as shown in (9). Thus we are effectively proxying aggregate risk aversion with a linear transformation of the aggregate dividend yield on equities. In additional empirical exercises described in Section 4.4 below, we replace (18) with alternative specifications that help identify z_t from the time series of bond excess returns.

Figure 3 plots the history of the dividend yield since 1953. The increase in the 1970's, followed by the long decline from the early 1980's to the year 2000, is interpreted by our model to mean that risk premia increased in the middle of our sample period and declined at the end.

Finally, our ninth and tenth measurement equations use the implication of our model that the conditional covariance between equity returns and nominal bond returns and the conditional volatility of nominal bond returns are time varying. The Appendix derives an expression for these conditional second moments, which are linear functions of z_t and ψ_t . Following Viceira (2007), we construct the realized covariance between daily stock returns and bond returns and the realized variance of daily nominal bond returns using a 1-quarter rolling window of daily stock returns and Treasury fixed-term bond returns from CRSP from 1964 onwards; before that we use a trailing 12-month window of monthly observations, as CRSP does not have daily observations of bond returns before 1964. We assume that these realized second moments measures the true conditional second moments with error. Given that equation (18) identifies z_t , these final measurement equations help us identify ψ_t .

The data used in these measurement equations are plotted in Figure 4. The left panel of the 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. Both series increase in the early 1970's and, most dramatically, in the early 1980's. In the early 1960's and the early 2000's, the covariance spikes downward while the variance increases. Our

model will interpret these as times when the nominal-real covariance changes sign.¹⁰

The unscented Kalman filter uses the system of measurement equations we have just formulated, together with the set of transition equations (2), (3), (11), (12), and (13) 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.

Despite the parsimony of our term structure model, the number of parameters to estimate is fairly large relative to the data series available for their estimation, and we find it difficult to estimate precisely all the elements of the variance-covariance matrix of shocks. Consequently, we estimate our model constraining many of these covariances to be zero. The unconstrained parameters are the covariances of the first four state variables and realized inflation with the stochastic discount factor, the covariances of the transitory component of expected inflation with the real interest rate and realized inflation, and the covariance of the real interest rate with realized inflation.

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 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.

¹⁰Figure 4 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.

3.2 Parameter estimates

Table 1 presents quarterly parameter estimates over the period 1953-2005. We estimate the full model and two constrained models described in Section 2.6, with constant z_t and ψ_t respectively.¹¹ All the models constrain certain shock covariances as described above. We discuss full-model parameter estimates first, and then parameter variation in the constrained models.

Table 1 shows that risk aversion is a persistent process, with an autoregressive coefficient of 0.957 implying a half-life of about 3.9 years. This result is unsurprising in light of the measurement equation (18), which links z_t to the equity dividend yield, since the dividend yield is known to be highly persistent and possibly even nonstationary (Stambaugh 1999, Lewellen 2004, Campbell and Yogo 2006). Our estimate of the autoregressive coefficient for z_t inherits the estimated persistence in the quarterly dividend yield.

The real interest rate is also a persistent process, with shocks that have a half-life slightly above 3.5 years. This persistence reflects the observed variability and persistence of TIPS yields; in the Appendix we show that the half-life of real interest rate shocks declines to about 2.5 years when we exclude the TIPS measurement equation.

The nominal-real covariance and the transitory component of expected inflation are the least persistent processes in our model, with half-lives of about 5.5 and 4.5 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.¹²

¹¹In practice, we constrain z_t to be constant by setting a large value for d_1 in (18). This results in a time series of z_t which has an extremely low volatility. We find that setting $d_1 = 10$ makes z_t constant for all practical purposes. The right hand columns of Table 1 report estimates for two additional models that we discuss in Section 4.4 below.

¹²Campbell 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.

Table 1 shows large differences in the volatility of shocks to the state variables. The one-quarter conditional volatility of the annualized real interest rate is estimated to be about 37 basis points, the average one-quarter conditional volatility of the transitory component of annualized expected inflation is about 114 basis points, and the average one-quarter conditional volatility of annualized realized inflation is about 178 basis points.¹³ By contrast, the average one-quarter conditional volatilities of the shocks to the permanent component of expected inflation are very small. Of 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.

Table 1 also reports the unrestricted correlations among the shocks. Two correlations stand out as particularly significant, both statistically and economically. First, there is a correlation of almost -0.5 between ξ_t and $-m_t$ shocks. This 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 correlation between λ_t and $-m_t$ shocks is much smaller, implying that the risk premium for permanent shocks to expected inflation is close to zero.

Second, there is a correlation of about -0.1 between shocks to the real interest rate x_t and shocks to realized inflation π_t . This implies that inflation shocks have driven real interest rates down when ψ_t is positive, but real interest rates have increased in response to realized inflation when ψ_t is negative. Given the high levels of ψ_t that we estimate in the late 1970's, and the lower levels that we estimate more recently, the changing response of real interest rates to realized inflation is qualitatively consistent with a shift towards more strongly anti-inflationary monetary policy discussed in Clarida, Gali, and Gertler (2000).

We also estimate a statistically insignificant and economically small negative correlation between π_t and $-m_t$ shocks. The point estimate implies that realized inflation is countercyclical, and nominal Treasury bills have a small positive risk premium,

¹³We compute the average conditional volatilities of the components of expected inflation and realized inflation as $(\mu_\psi^2 + \sigma_\psi^2)^{1/2}$ times the volatility of the underlying shocks. For example, we compute the average conditional volatility of realized inflation as $(\mu_\psi^2 + \sigma_\psi^2)^{1/2} \sigma_\pi$.

when ψ_t is positive. Finally, we estimate a small negative correlation between x_t and $-m_t$ shocks, implying that the real interest rate is countercyclical, real bond returns are procyclical, and term premia on real bonds are positive.

In the equity market, we estimate positive and statistically significant loadings of stock returns on both shocks to the real interest rate (β_{ex}) and shocks to the negative of the log SDF (β_{em}). Since the real interest rate is slightly countercyclical, the former effect modestly reduces the equity risk premium, but overall it is positive and increasing in aggregate risk aversion z_t .

Parameter estimates for the model with z_t constrained to be constant, in the second column of Table 1, are quite similar to those in our full model. However when we constrain ψ_t to be constant, in the third column of Table 1, we find that the volatility of the temporary component of expected inflation increases by a factor of four, and the correlation between shocks to this state variable and shocks to the stochastic discount factor declines by a factor of six. In the absence of variation in ψ_t , the model must explain more of the variation in intermediate-term bond yields by transitory variation in expected inflation; and in order to match the average slope of the intermediate-term yield curve, this requires a lower risk premium on transitory expected inflation.

At the bottom of Table 1, we report the change in log likelihood from imposing constraints on our general model. The restriction that z_t is constant is rejected at the 1% level. The restriction that ψ_t is constant is much more strongly rejected, because the model has no way to fit the movements of bond volatility and bond-stock covariation with constant ψ_t .

3.3 Fitted state variables

How does our model interpret the economic history of the last 50 years? That is, what time series does it estimate for the underlying state variables that drive bond and stock prices? Figure 5 shows our estimates of the real state variables, the real interest rate x_t in the left panel and risk aversion z_t in the right panel. The real interest rate is estimated to be unusually low for much of the 1970's and towards the end of our sample period. Higher-frequency movements in the real interest rate are generally countercyclical, as we see the real rate falling in the recessions of the early 1970's, early 1990's, and early 2000's. The real interest rate also falls around the

stock market crash of 1987. The major exception to this pattern is the very high real interest rate in the early 1980's during Paul Volcker's campaign against inflation. The average level of the short-term real interest rate is fairly high at 3.17%; this is partly due to our need to fit the average yields on TIPS, as the average real interest rate falls by about 130 basis points when we drop the TIPS measurement equation from the model. Risk aversion z_t displays modest variation whose pattern closely matches the history of the dividend-price ratio on US equities, an outcome that is almost guaranteed by our assumption in equation (18) that the dividend-price ratio is a constant, plus risk aversion, plus serially uncorrelated noise.

Figure 6 plots the components of expected inflation. The permanent component of expected inflation, in the left panel, exhibits a familiar hump shape over the post-war period. It was low, even negative, in the 1950's and 1960's, increased during the 1970's and reached a maximum value of almost 8% in the first half of the 1980's. Since then, it has experienced a secular decline to just below 2% at the end of the sample. 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, implying that our model attributes the generally high levels of yield spreads in the second half of our sample period at least partly to investor pessimism about increases in future inflation.

Finally, Figure 7 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. The state variable ψ_t has a hump shape, with upward spikes in the early 1970's and early 1980's, and downward spikes in the early 1990's and early 2000's, that matches those time series. Although ψ_t is predominantly positive, it can switch sign and is estimated to have done so early in the sample and in the the period immediately following the recession of 2001. Although the 2007-08 period is not included in our sample, casual observation suggests a negative ψ_t for that period as well.

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, and equity dividend yield. We do not plot the histories of these fitted values as they track the actual observed variables extremely closely. That is, 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.

4 Term Structure Implications

4.1 Properties of the term structure

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 thereafter: moments implied by our full model in the second column, our constrained model with constant z_t in the third column, and our constrained model with constant ψ_t in the fourth column. The remaining columns report results for additional models discussed in Section 4.4 below. 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.

Below each model-implied moment in Tables 2 and 3, we report in square brackets 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.

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. All our models tend to understate the average yield spreads in the data, particularly the average 10-year

spread, and overstate the volatility of yield spreads. The latter is a more serious problem as almost all of our 1,000 simulations generate excessively volatile yield spreads.

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. All the models somewhat understate mean excess returns, but the differences are not statistically significant. The models overstate the volatility of three-year realized excess returns, and understate the volatility of ten-year realized excess returns.

The next four rows of each table summarize our models' descriptions of TIPS yields. The models generate average TIPS yields that are higher than those in the data. All our term structure models imply extremely small risk premia on inflation-indexed bonds; thus the risk premia on nominal bonds are primarily determined by the covariance of inflation with the real economy.

The bottom panels of Tables 2 and 3 look at evidence on bond return predictability. In the first three rows we report the standard deviations of true expected excess returns within each of our estimated models. Our full model, and the constrained model with constant risk aversion, imply an annualized standard deviation for the expected excess return on 3-year bonds of about 30-35 basis points, and for the expected excess return on 10-year bonds of 30-50 basis points.¹⁴ While this variation is economically meaningful, it 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 these models is on the order of 1%. There is almost no variability in true expected excess returns on TIPS in any of the models we estimate. Finally, the model with a constant nominal-real covariance ψ delivers

¹⁴Yield 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.

only trivial variability in expected excess returns.

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 81 basis points for 3-year bonds, and 255 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 slightly less volatile than the true expected excess returns. For the models with constant ψ , however, the simulated fitted values are much more volatile than the true expected excess returns. In fact, there is little variation across these models in the standard deviations of regression fitted values.

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, the Stambaugh bias diminishes the standard deviation of the fitted values; in the constant- ψ model, where the true regression coefficient is almost zero, the bias increases the standard deviation of fitted values.

We also examine predictability 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.¹⁵ Results are broadly comparable to those reported for the CS regressions.

¹⁵Cochrane and Piazzesi impose proportionality restrictions across the regressions at different maturities, but we do not do this here.

4.2 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 8 through 11 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.”¹⁶ We plot maturities up to 10 years, or 40 quarters.

Figure 8 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 very slightly upward sloping, with an intercept just above 3%. This slight upward slope is consistent with our estimates of the full model, which imply that the real rate is countercyclical and thus that real bond risk premia are positive on average.

The right panel of Figure 8 shows that the mean log nominal yield curve also exhibits a positive slope, with a spread between the 10-year rate and the 1-month rate of about 100 basis points. This spread is slightly lower than the 113 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 8 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 9 examines the effect of changes in aggregate risk aversion z_t on the log real

¹⁶Strictly 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.

yield curve in the left panel, and the log nominal yield curve in the right panel. The left panel shows that changes in z_t have no noticeable effects on the real yield curve: the “maximum,” “minimum,” and “mean” real yield curves overlap each other. In fact, we estimate $B_{z,n}$ in (5) to be very close to zero. Thus real bond risk premia are approximately constant in our model.

The right panel of Figure 9 shows that changes in z_t have a somewhat stronger but still very modest effect on the nominal yield curve. This effect is asymmetric across the maturity spectrum. Changes in z_t have almost no effect on the intercept of the nominal yield curve, but have more noticeable effects on the long end of the curve. When other state variables are at their in-sample means, nominal bonds are moderately risky and thus their yields increase when risk aversion z_t increases. This effect is more powerful for long-term bonds than for short-term bonds. Thus risk aversion, like the real interest rate, alters the slope of the nominal yield curve; but it does so by moving the long end of the curve rather than the short end.

Figure 10 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 are qualitatively similar to those of changes in the real rate, and 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 11. This figure illustrates the yield curves that obtain when we vary ψ_t . Changes in ψ_t have almost no effect on the short end of the yield curve, and have a stronger effect on the middle of the curve than on the long end. Thus ψ_t can change the shape of the yield curve from concave to convex.

The impact of ψ_t on the concavity of the yield curve results from two features of our model. First, we have estimated the price of risk to be much higher for the transitory component of expected inflation than for the permanent component. This difference in risk prices normally generates a steep yield curve at shorter maturities, and a flatter one at longer maturities, since long-term nominal bond returns are driven primarily by the permanent component of expected inflation. 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 nominal bonds; in the terminology of the fixed-income literature, these bonds have much greater "convexity" than short- or intermediate-term bonds. Therefore unusually high positive values of ψ_t tend to lower long-term bond yields relative to intermediate-term yields. (Unusually low negative values of ψ_t do so as well, but these are less common because ψ_t has a positive mean.)

Figures 8 through 11 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 8 and 10 suggest that in our model, the permanent component of expected inflation is the main contributor to the level factor, with the real rate and the transitory component of expected inflation also driving changes in the average level of interest rates. The short-term real interest rate and particularly the transitory component of expected inflation both contribute to the slope factor by moving short-term yields more than long-term yields. In principle risk aversion also contributes to the slope factor by moving long-term yields, but Figure 9 shows that its effect is extremely small. Finally, Figure 11 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.

4.3 The determinants of bond risk premia

In the previous section we saw that in our model, both risk aversion z_t and the nominal-real covariance ψ_t are important determinants of 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 $z_t\psi_t$.

Figure 12 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 the product $z_t\psi_t$. In principle, the nominal-bond risk premium in our model is a linear combination of z_t , z_t^2 , and $z_t\psi_t$. The figure shows that in practice, the cross-product $z_t\psi_t$ generates almost all the variation in the risk premium, because the pure risk premium on real bonds is close to zero.

The right panel of the figure shows the term structure of risk premia as $z_t\psi_t$ 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 -250 to 350 basis points. The reason for this asymmetry is that in our sample period, large positive values of ψ_t coincided with large positive values of z_t , whereas large negative values of ψ_t coincided with smaller values of z_t .

The full history of our model's 10-year term premium is illustrated in Figure 13. The figure shows fairly stable risk premia below 1% during the 1950's and 1960's, then an upward spike to about 2% in the early 1970's and a major run up later in the 1970's to a peak of 3.5% in the early 1980's. A long decline in risk premia later in the sample period was accentuated briefly around the stock market crash of 1987 and the recession of the early 2000's, bringing the risk premium to its sample minimum of -2.5%. This time series reflects both the hump shape in the nominal-real covariance, illustrated in Figure 7, and the generally declining level of risk aversion, illustrated in Figure 5.

We saw in Figure 11 that the nominal-real covariance ψ_t and the product $z_t\psi_t$ influence 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 $z_t\psi_t$ 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

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 14 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 very similar shape but a somewhat 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.

Table 4 asks whether our models generate 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 each of our term structure models, best predict bond returns. In the top panel 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. All of our models, including the model with constant ψ_t but variable z_t , generate reasonably successful proxies for expected excess bond returns, particularly at the 3-year horizon. However, none of the models fully match the explanatory power of unrestricted CP regressions.

4.4 Alternative model specifications

As a final empirical exercise, we ask whether it is possible to increase the predictability of excess bond returns within our model by altering the set of measurement equations. Specifically, we drop the measurement equation for the equity dividend yield, and replace it with a measurement equation for a regression-based estimate of the bond risk premium. We first estimate a Campbell-Shiller (CS) regression of the excess 10-year bond return on the 10-year yield spread, with a 3-month holding period. We treat the fitted value of this regression as a noisy measurement of the true expected excess return on a 10-year bond. The resulting estimates are reported in the columns of Tables 1-4 labelled “CS EXR”. Alternatively, we estimate a Cochrane-Piazzesi (CP) regression of the excess 10-year bond return on 1-, 3-, and 5-year forward rates, again with a 3-month holding period. We treat the fitted value of the CP regression as a noisy measurement of the true expected excess return on a 10-year bond and report the results in the columns of Tables 1-4 labelled “CP EXR”.

Table 1 shows that these two alternative estimates have considerably greater volatility of risk aversion. In the CS EXR case, risk aversion is much less persistent, because the yield spread has more short-term variation than does the CP combination of forward rates. The CS EXR case also generates a high risk premium for shocks to the transitory component of expected inflation.

The bottom panels of Tables 2 and 3 show that the alternative estimates have considerably greater volatility of true expected excess bond returns. They also generate much greater volatility of fitted values in simulated CS and CP regressions. With these estimates, the CS and CP empirical results are no longer problematic for the model except for 10-year CS regressions, which still have anomalously volatile fitted values in the historical data.

Finally, Table 4 reports the ability of the alternative estimates to generate linear combinations of 1-year, 3-year, and 5-year forward rates that predict excess bond returns in the historical data. The CS EXR case fails at this task, because the model has fit the excess bond return to the 10-year yield spread. The CP EXR case performs better, but still no better than the base case reported in the second column of the table.

5 Conclusion

In this paper we have argued that a changing covariance between nominal and real variables is of central importance in understanding the term structure of nominal interest rates. Analyses of asset allocation traditionally assume that broad asset classes have a stable structure of risk over time; our empirical results suggest that in the case of nominal bonds, at least, this assumption is seriously misleading.

Our term structure model implies that the risk premia of nominal bonds have changed over the decades, in part with movements in risk aversion that are proxied by changes in the equity dividend yield, and in part with changes in the covariance between inflation and the real economy. Nominal bond risk premia were particularly high in the early 1980's, when bonds covaried strongly with stocks and risk aversion was high; they were negative in the early 2000's, when bonds covaried negatively with stocks, but at this time risk aversion was somewhat lower, so negative bond risk premia were smaller in magnitude.

Our model explains the 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 risk premium and convexity effects almost cancel 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, in the base case our model does not replicate the high explanatory power of regressions predicting excess 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. In addition, the Stambaugh (1999) finite-sample bias in predictive regressions reduces the explanatory power of yield spreads for excess returns in data simulated from our model. We can ameliorate this problem by adding measurement equations to force our model to fit the time series behavior of regression-based estimates of bond risk premia, but the predictive power of the 10-year yield spread for excess 10-year bond returns remains largely unexplained by our model.

The results we have presented are preliminary, and can be extended in a number of directions. First, 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.

Second, we can 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 2005, Bekaert, Engstrom, and Xing 2008).

Third, we can consider other theoretically motivated proxies for the stochastic discount factor in the economy. 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), Bansal and Shaliastovich (2007), Lettau and Wachter (2007b), and Abhyankar and Lee (2008).

Fourth, we can ask our model to fit a wider range of conditional second moments for asset returns, for example the volatility of returns on inflation-indexed bonds and the covariance of these bonds with stocks.

Fifth, we can explore alternative models for risk aversion, or equivalently, the conditional volatility of the stochastic discount factor. We have assumed that risk aversion moves closely with the dividend yield on stocks, but term structure models of changing risk premia with fewer restrictions on the stochastic discount factor, such as Dai and Singleton (2002), have been more successful in matching the explanatory power of predictive regressions.

Sixth, we can enrich our description of the real interest rate. This is desirable in itself and may be necessary to fit the behavior of inflation-indexed bond returns. Since the short-term real interest rate is controlled by the Federal Reserve, its covariance with the stochastic discount factor and the stock market reflects the policy rule of

the monetary authority. To the extent that the Federal Reserve cuts the real interest rate when the stock market is weak, and raises it when the stock market is strong, the covariance between x_t and $-m_t$ is negative. If such policy behavior has altered over time, then this covariance too would be time-varying rather than constant.

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, while a negative ψ_t reflects a stable Phillips Curve. It would be desirable to use data on inflation and output to explore this interpretation.

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, 2007, “The Determinants of Stock and Bond Return Comovements”, unpublished paper, Tilburg University, Columbia University, and Ghent University.
- Bansal, Ravi and Ivan Shaliastovich, 2007, “Risk and Return in Bond, Currency, and Equity 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, 2008, “Risk, Uncertainty, and Asset Prices”, forthcoming *Journal of Finance*.

- 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.
- 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. 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., Adi Sunderam, and Luis M. Viceira, 2009, “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, 2005, “International Stock-Bond Correlations in a Simple Affine Asset Pricing Model”, unpublished paper, Columbia University and University of St. Gallen.
- 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, 2008, “Inflation and Earnings Uncertainty and Volatility Forecasts: A Structural Form Approach”, 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.
- 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, 2005, “The Consumption-Based Determinants of the Term Structure of Discount Rates”, unpublished paper, University of Toulouse.
- 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, 2004, “An Econometric Model of Non-linear Dynamics in the Joint Distribution of Stock and Bond Returns”, unpublished paper.
- 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>.
- 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.

- Julier, S. and J. Uhlmann, 1997, “A New Extension of the Kalman Filter to Non-linear Systems”, in *Proceedings SPIE Signal Processing, Sensor Fusion, and Target Recognition VI*, 3068, 182–193.
- Koijen, Ralph S.J. and Jules H. van Binsbergen, 2008, “Predictive Regressions: A Present-Value Approach”, unpublished paper, Duke University and New York University.
- Lewellen, Jonathan, 2004, “Predicting Returns with Financial Ratios”, *Journal of Financial Economics* 74, 209–235.
- Lettau, Martin, and Jessica Wachter, 2007a, “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, 2007b, “The Term Structures of Equity and Interest Rates”, unpublished paper, Columbia University and University of Pennsylvania.
- 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”, unpublished paper, University of York.
- 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., 2007, “Bond Risk, Bond Return Volatility, and the Term Structure of Interest Rates”, unpublished paper, Harvard Business School.
- 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 Estimates					
Parameter	Full Model	Constant z	Constant ψ	CS EXR	CP EXR
$\mu_x \times 10^2$	1.02 (0.02)	1.01 (0.02)	1.00 (0.02)	0.87 (0.01)	1.07 (0.01)
$\mu_z \times 10^1$	2.36 (0.32)	2.37 (0.33)	2.27 (0.69)	3.09 (0.09)	5.80 (0.12)
$\mu_\psi \times 10^3$	3.97 (1.24)	4.45 (1.43)	1.58 (0.09)	0.89 (0.21)	4.00 (0.81)
ϕ_x	0.95 (0.02)	0.95 (0.02)	0.93 (0.02)	0.87 (0.01)	0.92 (0.01)
ϕ_ξ	0.86 (0.02)	0.86 (0.02)	0.89 (0.02)	0.96 (0.01)	0.90 (0.02)
ϕ_z	0.96 (0.08)		0.96 (0.08)	0.39 (0.02)	0.91 (0.01)
ϕ_ψ	0.88 (0.08)	0.89 (0.09)		0.75 (0.01)	0.83 (0.04)
$\sigma_x \times 10^4$	9.17 (0.73)	9.52 (0.77)	5.57 (0.34)	24.15 (0.42)	10.18 (0.51)
$\sigma_\lambda \times 10^4$	3.94 (36.16)	5.84 (43.40)	7.26 (40.85)	8.12 (13.54)	3.75 (22.96)
$\sigma_\Lambda \times 10^4$	7.01 (0.28)	6.99 (0.28)	9.26 (0.22)	11.15 (0.19)	7.53 (0.22)
σ_ξ	0.64 (0.09)	0.63 (0.08)	2.51 (0.01)	0.52 (0.02)	0.58 (0.03)
$\sigma_z \times 10^3$	6.52 (3.77)		5.22 (2.24)	84.49 (30.21)	26.53 (9.47)
$\sigma_\psi \times 10^3$	1.99 (0.32)	1.98 (0.28)		1.11 (0.07)	1.97 (0.23)
β_{ex}	22.77 (0.01)	21.69 (0.01)	-56.70 (0.91)	4.86 (0.01)	23.94 (0.01)
$\beta_{em} \times 10^2$	8.42 (0.05)	8.40 (0.05)	7.31 (0.44)	3.81 (0.08)	7.36 (0.06)
b_d	0.44 (0.01)		0.59 (0.07)		
$a_d \times 10^1$	-0.72 (0.14)		-1.01 (0.44)		

Parameter Estimates					
Parameter	Full Model	Constant z	Constant ψ	CS EXR	CP EXR
$\rho_{x\xi}$	-0.00 (0.01)	-0.00 (0.01)	0.01 (0.03)	0.00 (0.01)	-0.00 (0.01)
ρ_{xm}	-0.03 (0.01)	-0.03 (0.01)	-0.05 (0.01)	-0.00 (0.01)	-0.03 (0.01)
$\rho_{x\pi}$	-0.12 (0.03)	-0.12 (0.03)	-0.13 (0.01)	-0.02 (0.03)	-0.08 (0.02)
$\rho_{\lambda m}$	0.01 (1.36)	0.01 (1.09)	0.02 (0.31)	0.01 (0.23)	0.01 (0.69)
$\rho_{\xi m}$	-0.48 (0.04)	-0.48 (0.03)	-0.08 (0.01)	-0.89 (0.01)	-0.35 (0.01)
$\rho_{\xi\pi}$	0.16 (0.07)	0.16 (0.06)	0.03 (0.01)	0.29 (0.03)	0.23 (0.03)
ρ_{zm}	0.03 (0.37)		0.04 (0.27)	-0.01 (0.31)	0.01 (0.49)
$\rho_{\pi m}$	-0.03 (0.04)	-0.03 (0.07)	-0.06 (0.02)	0.01 (0.03)	-0.03 (0.02)

Log Likelihoods					
	Full Model	Constant z	Constant ψ	CS EXR	CP EXR
$\Delta\text{Log-likelihood}$	0	-5	-303	N/A	N/A
p-value		0.01	0.000	N/A	N/A

Table 2. Sample and implied moments for 3mo excess returns. Yield spreads (YS) are calculated over the 3mo yield. Realized excess returns (RXR) are calculated over a 3mo holding period, in excess of the 3mo 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. Below each entry we report in brackets 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 1/1985-12/2005.

Sample and Implied Moments						
Moment	Actual Data	Full model	Constant z	Constant ψ	CS EXR	CP EXR
3yr YS mean	.674	.355 [.201]	.400 [.239]	.159 [.047]	.216 [.035]	.731 [.558]
10yr YS mean	1.13	.478 [.154]	.557 [.185]	.233 [.036]	.413 [.061]	1.20 [.553]
3yr YS stdev	.401	.714 [.996]	.748 [1.00]	.725 [1.00]	.515 [.990]	.662 [.997]
10yr YS stdev	.642	1.16 [.998]	1.23 [1.00]	1.29 [1.00]	.830 [.984]	1.09 [.998]
3yr RXR mean	1.06	.694 [.308]	.781 [.357]	.225 [.128]	.508 [.160]	1.41 [.707]
10yr RXR mean	1.79	.910 [.238]	1.05 [.289]	.330 [.126]	1.36 [.371]	2.19 [.622]
3yr RXR stdev	4.01	4.85 [.928]	5.12 [.969]	5.58 [1.00]	4.96 [1.00]	5.42 [1.00]
10yr RXR stdev	10.00	8.40 [.011]	8.69 [.035]	10.13 [.631]	12.45 [1.00]	10.66 [.782]
10yr TIPS yield mean	3.37 [‡]	4.07 [.999]	4.05 [1.00]	4.04 [1.00]	3.46 [.868]	4.34 [1.00]
10yr TIPS YS mean		.001	.000	.014	-.035	.048
10yr TIPS RXR mean		.039	.041	.027	.013	.095
10yr TIPS RXR stdev		3.27	3.35	1.43	3.60	2.58

Predictive Regressions						
Moment	Actual Data	Full model	Constant z	Constant ψ	CS EXR	CP EXR
3yr EXR stdev		.334	.339	.007	1.60	1.04
10yr EXR stdev		.408	.420	.010	3.81	1.51
10yr TIPS EXR stdev		.003	.000	.001	.001	.004
3yr RXR $\sigma(\widehat{CS})$.810	.300 [.037]	.314 [.056]	.285 [.023]	.629 [.259]	.575 [.255]
10yr RXR $\sigma(\widehat{CS})$	2.55 [†]	.488 [.000]	.493 [.000]	.510 [.000]	.843 [.003]	.649 [.000]
10yr TIPS RXR $\sigma(\widehat{CS})$.179	.076	.184	.188	.132
3yr RXR $\sigma(\widehat{CP})$	1.00	.653 [.118]	.678 [.142]	.647 [.093]	1.29 [.785]	1.28 [.760]
10yr RXR $\sigma(\widehat{CP})$	2.23 [†]	1.10 [.013]	1.13 [.022]	1.27 [.033]	3.11 [.835]	2.02 [.369]

Table 3. Sample and implied moments for 1yr excess returns. 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. Below each entry we report in brackets 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 1/1985-12/2005.

Sample and Implied Moments						
Moment	Actual Data	Full model	Constant z	Constant ψ	CS EXR	CP EXR
3yr YS mean	.360	.210 [.233]	.234 [.294]	.123 [.079]	.150 [.091]	.456 [.643]
10yr YS mean	.821	.348 [.128]	.393 [.196]	.169 [.054]	.361 [.108]	.924 [.571]
3yr YS stdev	.486	.878 [.998]	.925 [.999]	.928 [1.00]	.600 [.951]	.805 [.999]
10yr YS stdev	1.05	1.80 [.999]	1.91 [.997]	2.06 [1.00]	1.28 [.936]	1.74 [.996]
3yr RXR mean	.647	.438 [.292]	.482 [.356]	.110 [.137]	.335 [.223]	.916 [.705]
10yr RXR mean	1.49 [†]	.718 [.191]	.797 [.255]	.204 [.140]	1.08 [.383]	1.79 [.598]
3yr RXR stdev	2.86	3.40 [.876]	3.59 [.922]	3.96 [1.00]	3.25 [.941]	3.81 [.982]
10yr RXR stdev	9.10 [†]	7.44 [.010]	7.68 [.038]	9.06 [.507]	10.34 [.953]	9.43 [.622]
10yr TIPS yield mean	3.37 [‡]	4.07 [.999]	4.05 [1.00]	4.04 [1.00]	3.46 [.857]	4.34 [1.00]
10yr TIPS YS mean		-.003	-.003	.011	-.031	.042
10yr TIPS RXR mean		.026	.028	.022	.001	.082
10yr TIPS RXR stdev		2.94	3.02	1.26	2.95	2.27

Predictive Regressions						
Moment	Actual Data	Full model	Constant z	Constant ψ	CS EXR	CP EXR
3yr EXR stdev		.359	.367	.009	.899	1.16
10yr EXR stdev		.504	.525	.015	2.68	2.01
10yr TIPS EXR stdev		.003	.000	.002	.001	.013
3yr RXR $\sigma(\widehat{CS})$	0.88	.363 [.051]	.386 [.069]	.280 [.072]	.292 [.016]	.496 [.149]
10yr RXR $\sigma(\widehat{CS})$	3.44 [†]	.717 [.010]	.745 [.000]	.722 [.002]	.909 [.007]	.901 [.005]
10yr TIPS RXR $\sigma(\widehat{CS})$.314	.322	.123	.289	.230
3yr RXR $\sigma(\widehat{CP})$	1.23	.805 [.095]	.843 [.147]	.762 [.149]	1.02 [.254]	1.39 [.620]
10yr RXR $\sigma(\widehat{CP})$	4.49 [†]	1.79 [.000]	1.83 [.000]	2.08 [.006]	3.27 [.090]	2.91 [.062]

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 8/1971-12/2005. For the earlier part of the sample we interpolate the 9yr yield linearly between the 5yr and the 10yr.

Forecasting Excess Returns							
	Moment	Unconstrained	Full model	Constant z	Constant ψ	CS EXR	CP EXR
3-month holding period	3yr RXR	.049	.030	.030	.029	.014	.024
	10yr RXR	.050	.031	.031	.023	.003	.014
1-year holding period	3yr RXR	.184	.146	.146	.134	.059	.127
	10yr RXR	.194 [†]	.172	.171	.129	.020	.113

Forecasting Excess Returns: No scaling							
	Moment	Unconstrained	Full model	Constant z	Constant ψ	CS EXR	CP EXR
3-month holding period	3yr RXR	.049	.029	.029	.028	.002	.024
	10yr RXR	.050	.017	.017	.017	.001	.012
1-year holding period	3yr RXR	.180	.124	.124	.134	.002	.121
	10yr RXR	.214 [†]	.056	.055	.071	.014	.051




Figure 1: Time series of 3-month and 10-year nominal yields and yield spread.




Figure 2: Time series of US and UK 10-year inflation-indexed yields.




Figure 3: Time series of dividend-price ratio. The dividend price ratio is constructed using the CRSP value-weighted index and a one-year backward moving average of dividends.




Figure 4: Time series of 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 5: Time series of real state variables. The figure on the left plots the time series of x_t , the real interest rate. The figure on the right plots the time series of z_t , risk aversion.




Figure 6: 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 7: Time series of nominal-real covariance.




Figure 8: 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 9: Responses of yield curves to z_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 z_t is varied between its sample minimum and maximum while all other state variables are held fixed at their sample means.




Figure 10: Responses of nominal yield curve 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 11: Responses of nominal yield curve to ψ_t . The left hand figure shows the response of the nominal yield curve to the nominal-real covariance 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 12: Responses of nominal expected excess returns to $z_t \psi_t$. 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 the product $z_t \psi_t$. The right hand figure shows the term structure of expected excess nominal bond returns as the product $z_t \psi_t$ is varied between its sample minimum and maximum while all other state variables are held fixed at their sample means.




Figure 13: Time series of expected excess returns for 10-year nominal bonds.




Figure 14: Data- and model- implied Cochrane-Piazzesi relationships.