

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

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

First draft: June 2007

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. This material is based upon work supported by the National Science Foundation under Grant No. 0214061 to Campbell, and the 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 four state variables: the real interest rate, risk aversion, expected inflation, and the covariance between nominal variables and the real economy. 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.

1 Introduction

Are nominal bonds risky investments, which investors must be rewarded to hold? Or are they safe investments, whose price movements are either inconsequential or possibly even beneficial to investors as a hedge against other risks?

This question can be answered in a number of ways. A first approach is to measure the covariance of nominal bond returns with some measure of investor well-being. According to the Capital Asset Pricing Model (CAPM), for example, investor well-being 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. According to the Consumption CAPM, investor well-being can be summarized by the level of aggregate consumption and 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 returns on bonds or from the average yield spread on long-term bonds over short-term bills. If this risk premium is large, then presumably investors regard bonds as risky.

These approaches are appealing because they are straightforward and direct. Unfortunately, the answers they give appear to depend sensitively on the particular 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 (Guidolin and Timmermann 2004, Baele, Bekaert, and Inghelbrecht 2007, Viceira 2007). 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 early 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 early 1980's and much lower in the late 1990's.

A third approach to measuring the risks of bonds is to decompose bond 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 underlying shocks to bond returns. For example, consumption shocks raise real interest rates if consumption growth is positively autocorrelated (Campbell 1986, Gollier 2005); in this case real 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, real bonds are exposed to consumption risk, and real bond premia should be positive. 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. 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 term structure model that is designed to allow the correlations of shocks, in particular the correlation of inflation with real variables, to change over time. By specifying stochastic processes for real interest rates, expected inflation, and investor risk aversion, we can solve for the complete term structure at each point in time and understand the way in which bond market risks have evolved.

Our approach extends a number of recent term structure models. Bekaert, Engstrom, and Grenadier (BEG, 2004), Wachter (2006), and Buraschi and Jiltsov (2007)

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

all 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 affine class (Dai and Singleton 2002). BEG and other recent authors including Mamaysky (2002) and d’Addona and Kind (2005) extend affine term structure models to price stocks as well as bonds. Our introduction of a time-varying correlation between inflation and real shocks takes us outside the affine class; our model, like those of Constantinides (1992) and Ahn, Dittmar and Gallant (2002), is linear-quadratic. To solve it, 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). We estimate our model using a Kalman filter approach, an extension of the method used in Campbell and Viceira (2001, 2002).

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.

2 A Quadratic Bond Pricing Model

We start by formulating a model which, in the spirit of Campbell and Viceira (2001, 2002), accounts for 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 correlation 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 expected inflation, the volatility of inflation, and the conditional correlation 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.

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

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

Thus the state variable z_t determines time-variation in the volatility of the SDF or, equivalently, in the price of aggregate market risk. In fact, we can interpret our model for the real SDF as a reduced form of a structural model in which aggregate risk aversion changes over time as a function of z_t , as in the habit consumption model of Campbell and Cochrane (1999). While our model does not constrain z_t to remain always positive, our empirical estimates do have this property.

The second 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

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

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

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

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

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

where

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

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 to this paper (Campbell, Sunderam, and Viceira 2007).

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

$$\begin{aligned} r_{n,t+1} - r_{1,t+1} &= p_{n-1,t+1} - p_{n,t} + p_{1,t} \\ &= - \left(\frac{1}{2}B_{x,n-1}^2\sigma_x^2 + \frac{1}{2}B_{z,n-1}^2\sigma_z^2 + B_{x,n-1}B_{z,n-1}\sigma_{xz} \right) \\ &\quad + (B_{x,n-1}\sigma_{mx} + B_{z,n-1}\sigma_{mz})z_t \\ &\quad + B_{x,n-1}\varepsilon_{x,t+1} + B_{z,n-1}\varepsilon_{z,t+1}, \end{aligned} \tag{5}$$

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 (5) that the conditional risk premium on real bonds is

$$E_t[r_{n,t+1} - r_{1,t+1}] + \frac{1}{2} \text{Var}_t(r_{n,t+1} - r_{1,t+1}) = (B_{x,n-1}\sigma_{mx} + B_{z,n-1}\sigma_{mz})z_t, \tag{6}$$

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.³ 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. Following Campbell and Viceira (2001), we model shocks to realized stock returns as a linear combination of shocks to the real rate and shocks to the log stochastic discount factor:

$$r_{e,t+1} - E_t r_{e,t+1} = \beta_{ex}\varepsilon_{x,t+1} + \beta_{em}\varepsilon_{m,t+1} + \varepsilon_{e,t+1}, \quad (7)$$

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, and unpriced because uncorrelated with the SDF.

Substituting (7) into the no-arbitrage condition $E_t [M_{t+1}R_{t+1}] = 1$, the conditional equity risk premium is given by

$$E_t [r_{e,t+1} - r_{1,t+1}] + \frac{1}{2} \text{Var}_t (r_{e,t+1} - r_{1,t+1}) = (\beta_{ex}\sigma_{xm} + \beta_{em}\sigma_m^2) z_t. \quad (8)$$

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

³Note that $\sigma_{mx} = 0$ implies not only $B_{x,n} = 0$, but also $B_{z,n} = 0$, for all n .

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 or, more precisely, for the reciprocal of inflation, which determines the real value of the nominal payments made by the bonds. We assume that log inflation $\pi_t = \log(\Pi_t)$ follows a linear conditionally heteroskedastic process:

$$\pi_{t+1} = \xi_t + \frac{\sigma_\psi^2}{2}\psi_t^2 + \psi_t\varepsilon_{\pi,t+1}, \quad (9)$$

where expected log inflation ξ_t and ψ_t follow

$$\xi_{t+1} = \mu_\xi(1 - \phi_\xi) + \phi_\xi\xi_t + \psi_t\varepsilon_{\xi,t+1}, \quad (10)$$

$$\psi_{t+1} = \mu_\psi(1 - \phi_\psi) + \phi_\psi\psi_t + \varepsilon_{\psi,t+1}, \quad (11)$$

and $\varepsilon_{\pi,t+1}$, $\varepsilon_{\xi,t+1}$, and $\varepsilon_{\psi,t+1}$ are again jointly normally distributed zero-mean shocks with a constant variance-covariance matrix. We allow these shocks to be cross-correlated with the shocks to m_{t+1} , x_{t+1} , and z_{t+1} , and use the same notation as in section 2.1 to denote their variances and covariances.

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 . We assume that this variable drives the volatility of expected inflation as well as the volatility of 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 the real economy, on the other hand.

The process for realized inflation, equation (9), 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 expected log inflation ξ_t . This in turn simplifies the pricing of short-term nominal bonds.

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

$$P_{1,t}^{\$} = \text{E}_t [\exp \{m_{t+1} - \pi_{t+1}\}], \quad (12)$$

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

$$\begin{aligned} y_{1,t+1}^{\$} &= -\text{E}_t [m_{t+1} - \pi_{t+1}] - \frac{1}{2} \text{Var}_t (m_{t+1} - \pi_{t+1}) \\ &= x_t + \xi_t - \sigma_{m\pi} z_t \psi_t, \end{aligned} \quad (13)$$

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

Equation (13) shows that the log of the nominal short rate is the sum of the log real interest rate, expected log inflation, and a nonlinear term that accounts for the correlation between shocks to inflation and shocks to the stochastic discount factor. If inflation is uncorrelated with the SDF ($\sigma_{m\pi} = 0$), the nonlinear term is zero and the Fisher equation holds: that is, the nominal short rate is simply the real short rate plus expected inflation.

It is straightforward to show that the nonlinear term in (13) 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:

$$\text{Cov}_t (m_{t+1}, \pi_{t+1}) = -\sigma_{m\pi} z_t \psi_t. \quad (14)$$

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 (6) or (8). 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 return on a one-period nominal bond and the return on equities. From (7) and (9), 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 of between equities and nominal bonds.

2.4 The nominal term structure

Equation (13) 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:

$$P_{n,t}^{\$} = \exp \left\{ A_n^{\$} + B_{x,n}^{\$} x_t + B_{z,n}^{\$} z_t + B_{\xi,n}^{\$} \xi_t + B_{\psi,n}^{\$} \psi_t + C_{z,n}^{\$} z_t^2 + C_{\psi,n}^{\$} \psi_t^2 + C_{z\psi,n}^{\$} z_t \psi_t \right\}, \quad (15)$$

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 (13), it is immediate to see that $B_{x,1}^{\$} = B_{\xi,1}^{\$} = -1$, $C_{z\psi,1}^{\$} = \sigma_{m\pi}$, and that the remaining coefficients are zero at $n = 1$.

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. We show in the Appendix that the log of the conditional expected gross excess return on an n -period zero-coupon nominal bond is given by

$$\log E_t \left[\frac{P_{n-1,t+1}^{\$}}{P_{n,t}^{\$}} \right] - E_t [r_{1,t+1}^{\$}] = \lambda_{z,n} z_t + \lambda_{\psi,n} \psi_t + \beta_{z,n} z_t^2 + \beta_{\psi,n} \psi_t^2 + \beta_{z\psi,n} z_t \psi_t, \quad (16)$$

where $r_{1,t+1}^{\$} \equiv y_{1,t}^{\$}$ is known at time t , and the coefficients $\lambda_{i,n}$ and $\beta_{i,n}$ are functions of the coefficients $A_n^{\$}$, $B_{i,n}^{\$}$, and $C_{i,n}^{\$}$ and thus are functions of bond maturity and the underlying stochastic processes for real and nominal variables. Explicit expressions for $\lambda_{i,n}$ and $\beta_{i,n}$ are given in Appendix X.

Equation (15) shows that our model implies a nominal term structure of interest rates which 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 expected inflation (ξ_t), and quadratic functions of risk aversion (z_t) and inflation volatility (ψ_t). Thus our model naturally generates four factors that explain bond yields. Equation (16) shows that expected log bond excess returns are time varying. They vary quadratically with risk aversion and inflation volatility, and linearly with the covariance between the log real SDF and inflation ($z_t \psi_t$). In this model, bond risk premia can be either positive or negative as ψ_t switches sign over time.

2.5 Special cases

Our general quadratic term structure model nests three important constrained models. First, if we constrain z_t and ψ_t to be constant, 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 (x_t) and expected inflation (ξ_t). Second, if we constrain only ψ_t to be constant over time, our model becomes a three-factor affine yield model where both real bond risk premia and nominal bond risk premia vary in proportion to aggregate risk aversion (z_t). This model captures the spirit of recent work on the term structure of interest rates by Bekaert, Engstrom, and Grenadier (2004), Buraschi and Jiltsov (2006), Wachter (2006) and others in which time-varying risk aversion is the only cause of time variation in bond risk premia. Finally, if we constrain only z_t to be constant 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 inflation volatility.

Since the coefficients of the nominal bond pricing function are complicated functions of the parameters of the model, we now present estimates of these parameters, and discuss the properties of bond prices and bond returns given our estimates.

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 estimates of the model based on a standard Kalman filter approach. Given the nonlinear structure of the model, this inherently linear approach does not produce maximum likelihood estimates of the parameters of the model, but rather quasi-maximum likelihood estimates. Although these estimates are not efficient, they are still consistent and asymptotically normal. They also provide us with a reasonable check of the ability of our data to explain important aspects of the time series and cross-sectional behavior of interest rates. Moreover, these

estimates provide useful initial values for the Efficient Method of Moments of Gallant and Tauchen (1996), which we plan to implement in the future to obtain efficient estimates of the parameters of the model.

The Kalman filter approach starts with the specification of 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.

Our first set of measurement equations relates observable nominal bond yields to the vector of state variables. Specifically, we use the relation between nominal zero-coupon bond log yields $y_{n,t}^s = -\log(P_{n,t}^s)/n$ and the vector of state variables implied by equation (15). We use monthly yields on constant maturity 3-month, 1-year, 3-year and 10-year zero-coupon nominal bonds 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.

To this set of equations we add a second set of four measurement equations. The first equation in this set is given by equation (9), which relates observed inflation rates to expected inflation and inflation volatility, plus a measurement error term. The second is the equation for realized log equity returns $r_{e,t+1}$ implied by (4), (7), and (8).

The third additional measurement equation uses the dividend yield on equities $D_{e,t}/P_{e,t}$ to identify z_t as

$$\frac{D_{e,t}}{P_{e,t}} = d_0 + d_1 z_t + \varepsilon_{D/P,t+1}, \quad (17)$$

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 is known to forecast future equity returns, and that in our model expected equity excess returns are proportional to z_t , as shown in (8). Thus we are effectively proxying aggregate risk aversion with a linear transformation of the aggregate dividend yield on equities.

The fourth additional measurement equation uses the implication of our model that the conditional covariance between equity returns and nominal bond returns is time varying. The Appendix derives an expression for this conditional covariance, a linear function of z_t and ψ_t . Following Viceira (2007), we construct the realized covariance between daily stock returns and bond returns using a 1-year rolling window, and assume that this covariance measures the true conditional covariance with error. Given that equation (17) identifies z_t , this final measurement equation helps us identify ψ_t .

To implement our additional measurement equations, we use monthly observations of CPI inflation, monthly total returns and dividend yields on the value-weighted portfolio comprising the stocks traded in the NYSE, AMEX and NASDAQ, and daily total returns on bonds and equities extracted from CRSP. To compute dividend yields, we use the standard procedure of using a 1-year backward-looking average of dividends to deal with intra-year seasonal effects in dividends.

The Kalman filter uses the system of measurement equations we have just formulated, together with the set of transition equations (2), (3), (10), and (11) 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.

3.2 Parameter estimates

Table 1 presents monthly estimates of our general model over the period January 1953-December 2005. The table also estimates the three constrained models described in Section 2.5. These are the models that constrain z_t , or ψ_t , or both variables to be constant over time. The estimates of the general model and the constant- z_t model are quasi-maximum likelihood estimates, since in those models log prices are non linear functions of the underlying state variables. The estimates of the constant- ψ_t model and the constant- z_t -and- ψ_t model are true maximum likelihood estimates, since these models fall within the affine yield class. The table reports parameter estimates in natural units, together with their asymptotic standard errors.

Table 1 shows that the state variables in the model are all highly persistent. They all have autoregressive coefficients above .95 and, in the case of log expected inflation (ξ_t) and aggregate risk aversion (z_t), the point estimate of the autoregressive

coefficient is exactly one, although the standard errors around the estimates are fairly large. The model estimates expected inflation to be much more persistent than real interest rates in the postwar period. This result is consistent with the estimates of the model with constant z_t and ψ_t in Campbell and Viceira (2001, 2002) using data through 1999.⁴ The estimated persistence of risk aversion z_t is not surprising in light of observation equation (17), 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).

The high persistence of the processes for the state variables makes it difficult to estimate their unconditional means. Accordingly, we have estimated our model requiring that the unconditional mean of the log real interest rate x_t equals the average ex-post log real interest rate in our sample period. This average is 1.54% per annum, which implies a value for μ_x of 0.001283, as reported in Table 1. We also require that the unconditional mean of the log inflation rate π_t equals the average log inflation rate in our period, which is 3.63% per annum. This in turn implies a value for μ_ξ of 0.0029.⁵

Table 1 shows large differences in the volatility of shocks to the state variables. The estimated one-month conditional volatility of the annualized real interest rate is about 14 basis points, and the average one-month conditional volatility of annualized expected inflation is about 5 basis points. Both of these estimates are statistically significant. The unconditional standard deviations of the real interest rate and expected inflation are of course much larger because of the high persistence of these processes; in fact, the population unconditional standard deviation of expected inflation is undefined because this process is estimated to have a unit root. We estimate the average conditional volatility of realized inflation to be about 1.88%.⁶ Finally,

⁴Campbell 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.

⁵Equation (9) implies that

$$\mu_\xi \approx E[\pi_{t+1}] - \frac{\sigma_\psi^2}{2} (\mu_\psi^2 + \sigma_\psi^2),$$

from which we can extract μ_ξ after replacing $E[\pi_{t+1}]$ with its sample mean, and the moments of ψ_t with their estimated values.

⁶The average volatility of expected inflation is computed as $(\mu_\psi^2 + \sigma_\psi^2)^{1/2} \sigma_\xi$, and the average

shocks to risk aversion have an annualized conditional volatility of about 1.84%.

Table 1 also reports the covariance structure of the shocks. We estimate σ_{xm} to be negative, which implies that the real interest rate is countercyclical, real bond returns are procyclical, and term premia on real bonds are positive. We also estimate $\sigma_{m\pi}$ to be negative, which according to equation (14) implies a positive average correlation between the SDF and inflation, since both z_t and ψ_t have positive means. Thus on average, we estimate inflation to be countercyclical, which implies a positive inflation risk premium in the nominal term structure.

In the equity market, we estimate positive 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}). The first loading would imply a negative equity premium but the second implies a positive equity premium, and this effect dominates.

The constrained models in Table 1 produce estimates for their unconstrained parameters which are generally in line with those of the unconstrained model. In particular, the constrained models also produce highly persistent processes, fit a more persistent process for expected inflation than for the real rate, and deliver negative estimates of $\sigma_{m\pi}$ and, with the exception of the constant- z_t model, of σ_{xm} . Finally, Table 1 reports the losses in log likelihood from imposing the constraints. These losses are extremely large, essentially because of our auxiliary measurement equations for the dividend yield and the conditional covariance of bond and stock returns. The loss in log likelihood is considerably larger when we constrain inflation volatility to be constant than when we constrain risk aversion to be constant.

3.3 Fitted state variables

Figure 1 plots the fitted time series of the latent state variables implied by our model estimates. Panel A in Figure 1 plots the time series of the short-term real interest rate. Panel A shows that short-term real rates were higher on average in the first half of our sample than in the second half, and reached a maximum of about 9.5% in the early 1980's. However, they appear to be more volatile in the second half of the sample where, interestingly, we estimate the real interest rate to be significantly negative at different points in time, particularly in the recessions of the early 1990's

volatility of realized inflation as $\left(\mu_\psi^2 + \sigma_\psi^2\right)^{1/2} \sigma_\pi$.

and early 2000's. This increased real interest rate volatility is in contrast with recent estimates of volatility in macroeconomic variables showing that growth, investment, and inflation volatility declined in the 1980's and 1990's (Stock and Watson, 2002).

Panel B in Figure 1 plots the time series of z_t . This is effectively a scaled version of the dividend yield given our assumption in equation (17) that the dividend yield is a multiple of z_t plus white noise measurement error. Consistent with our interpretation of z_t as aggregate risk aversion, the model chooses the scale factor in (17) so that z_t is positive everywhere.

Panel C in Figure 1 plots the time series of expected inflation. Expected inflation exhibits a familiar hump shape over the postwar period. It was low, even negative, in the 1950's and 1960's, increased during the 1970's and reached maximum values of about 10% in the first half of the 1980's. Since then, it has experienced a secular decline to about 1% at the end of the sample. The volatility of expected inflation also shows a hump shape; its decline in the second half of the sample is consistent with the results in Stock and Watson (2002).

Finally, Panel D in Figure 1 shows the time series of ψ_t . As we have noted, this variable is identified through the covariance of stock returns and bond returns. Panel D is to a close approximation a scaled version of the realized covariance of stock returns and bond returns, whose time series behavior we discuss below. It is important to note though that ψ_t can and does switch sign over time. Sign switches tend to be persistent, though there are also transitory changes in sign that appear to be related to "flight-to-quality" events in the bond and stock markets. In particular, we estimate ψ_t to be slightly negative on average for most of the 1950's and 1960's, positive or highly positive on average for most of 1970's, 1980's and the first half of the 1990's, and negative on average afterwards. The volatility of ψ_t appears to have increased significantly in the period that started in the 1970's. During this period, ψ_t has experimented brief periods in which it was either highly positive, such as in the early 1980's, or highly significantly negative, such as 1987.

Figure 2 provides a sense of the fit of the model in the time series dimension. These figures plot the observed and model-fitted time series of the covariance between stock returns and bond returns, the equity dividend yield, the 3-month nominal bond yield, and the 10-year nominal bond yield. The fitted time series of these variables almost perfectly overlap with the observed time series, reflecting the fact that the estimation algorithm achieves a good fit to the time series of stock and bond yields and the stock-bond covariance.

Panel A in Figure 2 plots the covariance between stock returns and bond returns. Consistent with the results in Viceira (2007), this covariance was negative in the 1950's, relatively stable around zero in the 1960's, and much more volatile since the 1970's. The stock-bond covariance was highly positive in the early 1980's, and since then it appears to have undergone a secular decline.

The remaining panels of the figure plot relatively familiar financial time series. The equity dividend yield, shown in Panel B, trended down throughout the 1950's and the 1960's. This trend reversed during the 1970's, when the dividend yield moved up and reached a maximum in the early 1980's. Since then, it has experienced a steady decline, with only a small reversal in the early 2000's.

Panels C and D show the time series of the observed and fitted short-term and long-term nominal interest rates. Both short-term and nominal long-term interest rates exhibit a pronounced hump shape, similar to the pattern we fit for expected inflation, and shown in Panel B of Figure 1. In fact, visual inspection of the plot of the time series of expected inflation and the time series of the 10-year nominal yield shows that they are extremely similar. The short-term nominal rate is considerably more volatile than the long-term nominal rate.

Table 2 reports fitted sample moments derived from these estimates, not only for our full model but also for our three constrained models. The mean ex post real interest rate is slightly lower than the mean ex ante real interest rate, because the mean ex post inflation rate is slightly higher than the mean ex ante inflation rate; our sample period had a slight preponderance of positive inflation surprises. The model does a good job of matching the historical moments of the real interest rate and inflation. The full model implies modest positive term premia, which are generally lower than those implied by our restricted models, and fits the realized returns on three-year bonds better than those models.

4 Implications for the Nominal Term Structure

4.1 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 Figure 3 the zero-coupon log nominal yield curves 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. Panels A through Panel D illustrate the yield curves that obtain when we vary x_t , z_t , ξ_t , and ψ_t respectively. We plot maturities up to 10 years, or 120 months.

The central line in each of panels in Figure 3 describes the yield curve generated by our model when all state variables are evaluated at their in-sample mean. This yield curve has a positive slope, with a spread between the 10-year rate and the 1-month rate of about 110 basis points. This spread is similar to the historical average spread in our sample period. The curve is more concave at maturities up to five years, and considerably flatter at longer maturities. The intercept of the curve implies a short-term nominal interest rate of about 5.4%, in line with the average short-term nominal interest rate in our sample.

Panel A in Figure 3 shows that changes in the real interest rate move the short end of the nominal yield curve but have almost no effect on the long end of the yield curve; thus they alter the slope of the curve. This effect is intuitive given that we have estimated a mean-reverting real interest rate process with a half-life of about 13 quarters. Such a process should not have a large effect on a 10-year zero-coupon bond yield.

Panel B shows that changes in z_t have almost no effect on the intercept of the nominal yield curve, but have 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 much 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.

Panel C shows that changes in expected inflation affect short- and long-term nominal yields almost equally, causing parallel shifts in the level of the nominal yield curve. This effect reflects the extremely high persistence that we have estimated for expected inflation.

The most interesting results are those shown in Panel D of Figure 3. Here we see that changes in ψ_t have almost no effect on the short end of the yield curve, but they have strong effects on both the middle of the curve and the long end. When ψ_t moves from its sample mean to its sample maximum, intermediate-term bond yields rise but long-term bond yields do not. This reflects two opposing effects of ψ_t on yields. On the one hand, when ψ_t increases nominal bonds have higher return volatility, 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. On the other hand, when ψ_t increases, nominal bonds become more systematically risky and investors demand a higher risk premium. As ψ_t moves from its sample mean to its sample maximum, the two effects roughly cancel at the long end of the yield curve but the greater risk premium dominates in the middle of the yield curve, driving intermediate yields up relative to both short and long-term yields.

As ψ_t moves from its sample mean to its sample minimum, however, it moves from slightly positive to slightly negative and there is relatively little change in the volatility of bond returns. Thus the convexity effect is small relative to the risk premium effect, and in panel D we see that the long end of the yield curve falls when ψ_t approaches its sample minimum.

Figure 3 allows 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. Traditional 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. Figure 3 suggests that in our model, expected inflation is the level factor; the short-term real interest rate and risk aversion both contribute to the slope factor; and the covariance of nominal and real variables drives the curvature factor and, when it is not too high, the slope factor.

4.2 The determinants of bond risk premia

In the previous section we saw that both risk aversion z_t and the nominal-real covariance ψ_t are important determinants of long-term nominal interest rates in our model. The reason for this is that these variables have powerful effects on risk premia. In fact, the main determinant of nominal bond risk premia is the product $z_t\psi_t$.

Figure 4 illustrates this by plotting the time series of the monthly risk premium on a 10-year nominal zero-coupon bond, $\log E_t[P_{119,t+1}^{\$}/P_{120,t}^{\$}] - E_t[r_{1,t+1}^{\$}]$, together with the time series of $z_t\psi_t$ scaled to have approximately the same standard deviation. In principle, we know from equation (16) that the nominal-bond risk premium in our model is a linear combination of z_t , ψ_t , their squares, and their cross-product. The figure shows that in practice, the cross-product $z_t\psi_t$ generates most of the variation in the risk premium..

Our model fits the time series of postwar US bond risk premia with three periods that broadly coincide with three distinct periods for capital market and macroeconomic conditions. The first period includes most of the 1950's and 1960's. This was a period of a stable tradeoff between growth and inflation and sharply declining risk aversion, and our estimated bond risk premia reflect that; they are positive early in the period, negative for most of the 1950's, and close to zero in the 1960's. The second period includes the 1970's and the first half of the 1980's. This was a period of an unstable relation between inflation and growth, a declining stock market, and increasing risk aversion. Our estimated bond risk premia during this period are positive on average, with considerable volatility. The third period runs from the mid-1980's until the end of our sample period. This period has been characterized by a return to stable growth and inflation, a rising stock market, and declining risk aversion. Estimated bond risk premia show a declining trend in both their mean and their volatility, and become negative at the end of the sample.

Our fitted bond risk premia also exhibit short episodes where they reach extreme positive or negative values. These episodes are related to the occurrence of extreme economic or financial events, such as the large increases in interest rates in the early 1980's during the Volcker period, which drove bond risk premia sharply higher, or the stock market crash of October 1987, which produced a "flight to quality" and sharply lower bond risk premia.

Figure 5 explores the impact of changes in z_t and ψ_t in more detail. Panel A in the

figure plots bond risk premia as a function of maturity n when all state variables are at their sample mean, and z_t is at its mean, minimum, and maximum values. Panel B is identical in structure to Panel A, except that it varies ψ_t instead of z_t . Panel C varies the product $z_t\psi_t$.

Consistent with the analysis of the impact of state variables on bond yields shown in Figure 3, Panel A shows that z_t always increases bond risk premia, and that the impact is increasing in the maturity of the bond. However, even at a maturity of 20 years, the conditional bond risk premium evaluated at the maximum value of z_t is only slightly larger than 2% per annum. By contrast, Panel B shows that ψ_t has a very pronounced effect on bond risk premia: The conditional risk premium on a 20-year bond evaluated at the maximum value of ψ_t is about 15% per annum. Moreover, conditional bond risk premia are increasingly negative as a function of maturity when ψ_t is at its sample minimum. Panel C shows a similar pattern to Panel B, but the positive risk premia at the maximum are accentuated while the negative risk premia at the minimum are less extreme. The reason 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 much smaller values of z_t .

We saw in Figure 3 that the nominal-real correlation ψ_t influences the curvature of the yield curve as well as its slope. Other factors in our model, such as the real interest rate, also influence the slope of the yield curve but do not have much effect on its curvature. Given the dominant influence of ψ_t on bond risk premia, illustrated in Figure 5, the curvature of the yield curve may 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 (2005). Using econometric methods originally developed by Hansen and Hodrick (1983), and implemented in the term structure context by Stambaugh (1988), Cochrane and Piazzesi show that a single linear combination of forward rates is a good predictor of excess bond returns at a wide range of maturities. Interestingly, this combination of forward rates 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.

Table 3 reports a similar exercise to that of Cochrane and Piazzesi. Using both our raw data and the fitted yield curves from our estimated models, we regress both ten-year and three-year realized excess returns onto the yield spread and onto our

estimated optimal combination of forward rates, which we call “Tent”. These regressions are fairly similar in the data and in all of our models, since all the models track the in-sample behavior of the yield curve fairly well. Like Cochrane and Piazzesi, we find that the tent variable produces a higher R^2 statistic and stronger statistical significance than does the yield spread.

Each of our models implies a different history for the expected (as opposed to the realized) excess bond return. We regress our model-implied expected returns onto the yield spread and the tent variable, and again find that the tent variable delivers a better fit. The difference in fit is particularly pronounced in our full model where both z_t and ψ_t move over time.

Finally, we regress realized returns onto the expected excess returns implied by each of our models. We obtain coefficients close to one for our full model, but smaller and even in one case negative coefficients for our restricted models. This finding strongly suggests that our full model is necessary to explain the predictability of excess bond returns in our postwar US dataset.

5 Conclusion

In this paper we have argued that changing covariances between nominal and real variables are 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 and stocks, 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 relatively low, so negative bond risk premia were modest 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 3 or 4 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 opposing effects on longer-term bond yields. It raises them by increasing the risk premium, but it lowers them through a Jensen’s Inequality effect of increasing volatility. In the language of fixed-income investors, longer-term bonds have “convexity” which becomes more valuable when volatility is high. At the long end of the yield curve, these two effects roughly 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 hence for the risk premium on nominal bonds.

The results we have presented are preliminary and this research can be extended in a number of directions. First, we can use alternative estimation methods, such as the Efficient Method of Moments of Gallant and Tauchen (1996), to properly handle econometric difficulties caused by the nonlinearity of our term structure model.

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 2004, d’Addona and Kind 2005).

Third, we can ask our model to fit a wider range of conditional second moments for asset returns; this may require us to generalize the model to allow heteroskedasticity in real as well as nominal variables.

Fourth, we can ask our model to fit data from the last ten years on the yields of TIPS (Treasury inflation-protected securities) as well as the longer time series for nominal bond yields. This additional source of information will allow us to ask, for example, whether real bond returns have stable covariances with stock returns as implied by our model.

Fifth, 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 more directly to explore this interpretation.

Sixth, we can enrich our description of the real interest rate. To the extent that 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. For example, the hypothesis 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 so-called “Greenspan put”) would imply a negative covariance between x_t and m_t . If such policy behavior has altered over time, then this covariance too would be time-varying rather than constant.

Finally, we can apply our model to other countries with different inflation histories. One particularly interesting country is the UK, where inflation-indexed bonds have been actively traded since the mid-1980’s.

References

- 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.
- 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.
- Barsky, Robert B., 1989, “Why Don’t the Prices of Stocks and Bonds Move Together?”, *American Economic Review* 79, 1132–1145.
- Bekaert, Geert, Eric Engstrom, and Steve Grenadier, 2004, “Stock and Bond Returns with Moody Investors”, unpublished paper, Columbia University, University of Michigan, and Stanford University.
- Buraschi, Andrea and Alexei Jiltsov, 2007, “Habit Formation and the Term Structure of Interest Rates”, *Journal of Finance*, forthcoming.
- 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, 2007, “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.
- 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 Finance* 63, 415–442.
- David, Alexander and Pietro Veronesi, 2004, “Inflation and Earnings Uncertainty and Volatility Forecasts”, unpublished paper, Washington University and University of Chicago.
- 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.
- Gallant, Ronald and George Tauchen, 1996, "Which Moments to Match", *Econometric Theory* 12, 657-681.
- Gollier, Christian, 2005, "The Consumption-Based Determinants of the Term Structure of Discount Rates", unpublished paper, University of Toulouse.
- 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, Sack, and Wright, 2006, 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.
- Lewellen, Jonathan, 2004, "Predicting Returns with Financial Ratios", *Journal of Financial Economics* 74, 209-235.
- 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, 2007, "Equilibrium Yield Curves", in D. Acemoglu, K. Rogoff, and M. Woodford (eds.) *NBER Macroeconomics Annual 2006*, 317-363, MIT Press, Cambridge, MA.

- 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.
- Stock, James H. and Mark Watson, 2002, “Has the Business Cycle Change and-Why?,” in M. Gertler and K. Rogoff (eds.) *NBER Macroeconomics Annual: 2002*, MIT Press, Cambridge, MA.
- 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.




Figure 1: Time series of state variables. Panel A (top left) shows the real interest rate x_t , panel B (top right) shows risk aversion z_t , panel C (bottom left) shows expected inflation ξ_t , and panel D (bottom right) shows the covariance of real and nominal variables ψ_t .




Figure 2: Fitted and actual variables. Panel A (top left) shows the fitted and actual covariance of bond and stock returns. Panel B (top right) shows the fitted and actual dividend-price ratio on the aggregate stock market. Panel C (bottom left) shows the fitted and actual short-term nominal interest rate. Panel D (bottom right) shows the fitted and actual long-term nominal interest rate.




Figure 3: Yield curves in relation to state variables. Each panel shows the yield curves that prevail when one state variable is at its sample minimum, mean, and maximum, holding all other state variables at their mean. Panel A (top left) varies the real interest rate x_t , Panel B (top right) varies risk aversion z_t , Panel C (bottom left) varies expected inflation ξ_t , and Panel D (bottom right) varies the covariance of nominal and real variables ψ_t .




Figure 4: The risk premium on 10-year nominal zero-coupon bonds. The figure plots the expected excess return on 10-year nominal zero-coupon bonds, together with the cross-product of state variables $z_t \psi_t$.




Figure 5: Term premia and state variables. Each panel shows the expected excess return on nominal bonds as a function of maturity, when one state variable is at its sample minimum, mean, and maximum, and all other state variables are at their sample means. Panel A (top left) varies risk aversion z_t , Panel B (top right) varies the covariance of nominal and real variables ψ_t , and Panel C (bottom left) varies the cross-product $z_t\psi_t$.

Table 1: Parameter Estimates, 1953:1-2005:12

Standard errors in parentheses

Covariances and their s.e.'s are multiplied by 1e7

	Full model	Constant z	Constant Ψ	Constant z and Ψ
μ_x	0.003 (0.001)	0.002 (0.002)	0.002 (0.000)	0.002 (0.000)
μ_ξ	0.002 (0.001)	0.000 (0.011)	0.000 (0.010)	0.003 (0.039)
μ_z	0.791 (0.029)	0.204 (0.012)	0.197 (0.011)	0.158 (0.003)
μ_Ψ	0.025 (0.004)	0.264 (0.009)	0.357 (0.008)	0.034 (0.001)
φ_x	0.964 (0.033)	0.962 (0.031)	0.956 (0.031)	0.948 (0.037)
φ_ξ	1.000 (0.065)	1.000 (0.063)	0.999 (0.049)	1.000 (0.083)
φ_z	1.000 (0.259)		0.991 (0.241)	
φ_Ψ	0.957 (0.003)	0.917 (0.002)		
σ_x	0.001 (0.000)	0.001 (0.001)	0.004 (0.000)	0.000 (0.000)
σ_ξ	0.005 (0.001)	0.004 (0.000)	0.005 (0.000)	0.007 (0.000)
σ_z	0.005 (0.009)		0.038 (0.000)	
σ_Ψ	0.022 (0.077)	0.175 (0.052)		
σ_π	0.164 (0.004)	0.027 (0.002)	0.038 (0.004)	0.087 (0.002)
$\sigma(x,\xi)$	-0.114 (0.013)	-0.185 (14.23)	-48.651 (15.42)	-6.180 (165.43)
$\sigma(x,z)$	-0.866 (0.022)		-1385.186 (394.31)	
$\sigma(x,\Psi)$	2.721 (128.87)	61.744 (331.30)		
$\sigma(x,m)$	-1.385 (245.05)	373.608 (29613)	-1828.266 (10.68)	-2426.206 (29181)
$\sigma(x,\pi)$	2.039 (44.21)	32.712 (245.05)	1641.075 (100944)	16.340 (221.97)
$\sigma(\xi,z)$	0.137 (0.012)		35.066 (251177)	
$\sigma(\xi,\Psi)$	-29.260 (83883)	20.215 (62.58)		
$\sigma(x,m)$	-20010.460 (297.37)	1249.589 (0.013)	767.449 (6802)	-4693.761 (141.29)
$\sigma(x,\pi)$	0.007 (12.40)	-41.765 (5008)	-13029.359 (10949)	240.900 (245.05)
$\sigma(z,\Psi)$	521.408 (61.19)			
$\sigma(z,m)$	2.317 (0.012)		227.706 (140.68)	

$\sigma(z,\pi)$	-2.055 (1501)		11616.570 (271.40)	
$\sigma(\Psi,m)$	17.742 (90.79)	3588.713 (140.68)		
$\sigma(\Psi,\pi)$	-0.449 (768.62)	2841.709 (271.40)		
$\sigma(m,\pi)$	-0.021 (143.25)	-3425.235 (105.54)	-6080.120 (119.05)	-9955.811 (155.30)
β_{ex}	16.389 (3.905)	25.858 (3.151)	24.101 (3.823)	32.573 (2.925)
β_{em}	0.043 (0.001)	0.056 (0.002)	0.059 (0.001)	0.050 (0.002)
d0	0.007 (0.000)	0.010 (0.000)	0.004 (0.008)	0.010 (0.000)
d1	0.313 (0.002)	0.308 (0.002)	0.305 (0.002)	0.126 (0.002)
Change in Log-Likelihood	0	-3531	-5641	-8179

Table 2: Fitted Sample Moments

	Full model	Constant z	Constant Ψ	Constant z and Ψ	Data
Ex ante real rate, mean	1.850	0.828	1.604	2.291	
Ex ante real rate, stdev	0.561	0.558	0.545	0.543	
Ex post real rate, mean	1.513	0.719	1.496	2.169	1.540
Ex post real rate, stdev	0.730	0.519	0.620	0.698	0.895
Expected inflation, mean	3.517	4.596	3.773	3.083	
Expected inflation, stdev	0.706	0.703	0.757	0.754	
Realized inflation, mean	3.626	4.593	3.816	3.144	3.772
Realized inflation, stdev	0.918	0.704	0.810	0.861	1.053
10 yr expected excess return, mean	0.639	1.623	0.995	1.918	
10 yr expected excess return, stdev	0.474	0.838	0.079	0.000	
Conditional stdev of 10 yr excess return	9.219				
3 yr expected excess return, mean	0.212	0.136	0.871	1.225	
3 yr expected excess return, stdev	0.157	0.165	0.070	0.000	
Conditional stdev of 3 yr excess return	4.019				
10yr realized return, mean	0.910	0.685	0.848	0.812	1.040
10yr realized return, stdev	9.576	10.101	9.976	9.964	11.59
3yr realized return, mean	0.640	1.461	0.998	0.998	0.646
3yr realized return, stdev	2.981	3.851	3.848	3.839	3.112

Table 3: Predictability of Excess Bond Returns
T-statistics in parentheses

	Full model	Constant z	Constant Ψ	Constant z and Ψ	Data
10yr RXR = a+bSPR	2.617 (9.303)	4.544 (9.642)	2.671 (9.542)	2.649 (9.879)	2.631 (7.659)
R2	0.130	0.130	0.132	0.130	0.136
10yr RXR = a+bEXR	0.870 (3.651)	0.367 (2.647)	-2.951 (-2.024)	0 (0)	
R2	0.025	0.012	0.014	0	
10yr RXR = a+bTent	0.027 (12.800)	0.029 (13.064)	0.029 (13.877)	0.030 (12.064)	0.0318 (11.264)
R2	0.214	0.215	0.241	0.244	0.266
10yr EXR = a+bSPR	0.183 (3.744)	0.012 (0.136)	0.057 (7.052)	0 (0)	
R2	0.151	0.239	0.934	0	
10yr EXR = a+bTent	0.003 (10.303)	0.001 (12.282)	0.000 (3.929)	0 (0)	
R2	0.356	0.385	0.931	0	
3yr RXR = a+bSPR	1.049 (6.634)	0.592 (6.865)	0.873 (5.366)	0.846 (5.201)	1.186 8.701
R2	0.105	0.144	0.093	0.092	0.145
3yr RXR = a+bEXR	1.183 (5.566)	0.633 (3.052)	0.325 (0.654)	0 (0)	
R2	0.098	0.093	0.052	0	
3yr RXR = a+bTent	0.007 (9.677)	0.007 (10.412)	0.010 (9.652)	0.006 (9.341)	0.008 (12.356)
R2	0.170	0.216	0.175	0.168	0.230
3yr EXR = a+bSPR	0.061 (3.744)	0.028 (0.990)	0.066 (5.009)	0 (0)	
R2	0.151	0.055	0.932	0	
3yr EXR = a+bTent	0.001 (10.302)	0.000 (9.064)	0.000 (3.929)	0 (0)	
R2	0.256	0.162	0.931	0	
Tent = a+b*psi_t	14.198 (11.502)	0.951 (12.068)			
R2	0.288	0.299			