

Does the Failure of the Expectations Hypothesis Matter for Long-Term Investors?

Antonios Sangvinatsos*
New York University

Jessica Wachter†
New York University and NBER

January 21, 2003

Work In Progress

Comments Welcome

*Stern School of Business, New York University, 44 West Fourth Street, Suite 9-190, New York, NY 10012-1126, asangvin@stern.nyu.edu, (212) 998-0718.

†Stern School of Business, New York University, 44 West Fourth Street, Suite 9-190, New York, NY 10012-1126, jwachter@stern.nyu.edu, (212) 998-0799.

Does the Failure of the Expectations Hypothesis Matter for Long-Term Investors?

Abstract

We consider the consumption and portfolio choice problem of a long-run investor when the term structure is affine and when the investor has access to nominal bonds and a stock portfolio. In the presence of unhedgeable inflation risk, there exist multiple pricing kernels that produce the same bond prices, but a unique pricing kernel equal to the marginal utility of the investor. We apply our method to a three-factor Gaussian model with a time-varying price of risk that captures the failure of the expectations hypothesis seen in the data. We extend this model to account for time-varying expected inflation, and estimate the model with both inflation and term structure data. The estimates imply that the bond portfolio for the long-run investor looks very different from the portfolio of a mean-variance optimizer. In particular, the desire to hedge changes in term premia generates large hedging demands for long-term bonds.

1 Introduction

The expectations hypothesis of interest rates states that the premium on long-term bonds over short-term bonds is constant over time. According to this hypothesis, there are no particularly good times to invest in long-term bonds relative to short-term bonds, nor are there particularly bad times. Long-term bonds will always offer the same expected excess return.

While the expectations hypothesis is theoretically appealing, it has consistently failed in U.S. postwar data. Fama and Bliss (1987) and Campbell and Shiller (1991), among others, show that expected excess returns on long-term bonds (term premia) do vary over time, and moreover, it is possible to predict excess returns on bonds using observables such as the forward rate or the term spread. This paper explores the consequences of the failure of the expectations hypothesis for long-term investors.

We estimate a three-factor affine term structure model similar to that proposed in Dai and Singleton (2002a) and Duffee (2002) that accounts for the fact that excess bond returns are predictable. We then solve for the optimal portfolio for an investor taking this term structure as given. Bond market predictability will clearly affect the mean-variance efficient portfolio, but the consequences for long-horizon investors go beyond this. Merton (1971) shows that when investment opportunities are time-varying, a mean-variance efficient portfolio is generally sub-optimal. Long-horizon investors wish to hedge changes in the investment opportunity set; depending on the level of risk aversion, the investor may want more or less wealth when investment opportunities deteriorate than when they improve. As we will show, investors gain by hedging time-variation in the term premia. Thus the investor's bond portfolio looks different from that dictated by mean-variance efficiency.

Despite the obvious importance of bonds to investors, as well as the strength of the empirical findings mentioned above, recent literature on portfolio choice has focused almost exclusively on predictability in stock returns. As shown by Fama and French (1989) and Campbell and Shiller (1988), the price-dividend ratio predicts excess stock returns with a negative sign. Based on this finding, a number of studies (e.g. Balduzzi and Lynch (1999), Barberis (2000), Brandt (1999), Brennan, Schwartz, and Lagnado (1997), Campbell and Viceira (1999), Liu (1999) and Wachter (2002a)) document gains from timing the stock market based on the price-dividend ratio, and from hedging time-variation in expected stock returns. One result of this literature is that when investors have relative risk aversion greater than one, hedging demands dictate that their allocation to stock should increase with the

horizon. A natural question to ask is whether the same mechanism is at work for bond returns. Just as stock prices are negatively correlated with increases in future risk premia on stocks, bond prices are negatively correlated with increases in future risk premia on bonds.¹ This intuition suggests that time-variation in risk premia would cause the optimal portfolio allocation to long-term bonds to increase with horizon.

In the case where the investor allocates wealth between a long and a short-term bond, we show that this intuition holds. Hedging demands induced by time-variation in risk premia more than double the investor's allocation to the long-term bond. Moreover, we find large horizon effects. The investor with a horizon of twenty years holds a much greater percentage of his wealth in long-term bonds than an investor with a horizon of ten years. In the case of multiple long-term bonds, the mean-variance efficient portfolio often consists of a long and short position in long-term bonds. This occurs because of the high positive correlation between bonds of different maturities implied by the model and found in the data. Hedging demands induced by time-varying risk premia generally make the allocation to long-term bonds more extreme. We find that following a myopic strategy and, in particular, failing to hedge time variation in risk premia carries a high cost for the investor in terms of certainty equivalent returns.

Our framework generalizes previous studies of portfolio choice when real interest rates vary over time and there is inflation. Brennan and Xia (2002) and Campbell and Viceira (2001) estimate a two-factor Vasicek (1977) term structure model and determine optimal bond portfolios. Both of these studies assume that risk premia on bonds and stocks are constant.² Our study also relates to that of Campbell, Chan, and Viceira (2002) who estimate a vector-autoregression (VAR) including the returns on a long-term bond, a stock index, the dividend yield and the yield spread. Campbell et al. derive an approximate solution to the optimal portfolio choice problem when asset returns are described by the VAR. The advantage of

¹We consider U.S. government bonds that are not subject to default risk. Nonetheless, we use risk premia and term premia interchangeably, as we do not take a stand on the source of the premia.

²Other work on bond returns and portfolio choice includes Brennan and Xia (2000) and Sorensen (1999), who assume that interest rates are Vasicek, and Liu (1999) and Schroder and Skiadas (1999) who assumes general affine dynamics. These studies assume that bonds are indexed, or equivalently, that there is no inflation. Xia (2002) examines the welfare consequences of limited access to nominal bonds under a Vasicek model. Wachter (2002b) shows under general conditions that as risk aversion approaches infinity, the investor's allocation approaches 100% in a long-term indexed bond. None of these papers explore the consequences of bond return predictability.

the VAR approach is that it captures predictability in bond and stock returns in a relatively simple way. The disadvantage is that the term structure is not well-defined; it is necessary to assume that the investor only has access to those bonds included in the VAR. Moreover, estimating bond returns using a VAR gives up the extra information resulting from the no-arbitrage restriction on bonds, namely that bonds have to pay their (nominal) face value when they mature.

Rather than modeling bond return predictability using a VAR, we follow the affine bond pricing literature (e.g. Dai and Singleton (2000, 2002a) and Duffee (2002)) and specify a nominal pricing kernel.³ The drift and diffusion of the pricing kernel is driven by three underlying factors which follow a multivariate Ornstein-Uhlenbeck process. The price of risk is a linear function of the state variables. Thus the model is in the “essentially affine” class proposed by Duffee (2002), and shown by Dai and Singleton (2002a) to capture the pattern of bond predictability in the data.

As a necessary step to showing the implications of affine term structure models for investors, we show how parameters of the inflation process can be jointly estimated with term structure parameters. This joint estimation produces a series for expected inflation that explains a surprisingly high percentage of the variance of realized inflation. This result has implications not only for portfolio choice problems, but for the estimation of term structure models more generally.

The remainder of the paper is organized as follows. Section 2 describes the general form of an economy where nominal bond prices are affine, and there exists equity and unhedgeable inflation. Section 3 derives a closed-form solution for optimal portfolio choice when the investor has utility over terminal wealth and over intermediate consumption. When inflation is introduced, the pricing kernel that determines asset prices is not unique; from the point of view of the investor it is not well-defined. As He and Pearson (1991) show, there is a unique pricing kernel that gives the marginal utility process for the investor.⁴ We derive a closed-form expression for this pricing kernel when incompleteness results from inflation. This expression holds regardless of the form of the term structure. Section 4 uses max-

³For recent surveys of this literature, see Piazzesi (2002) and Dai and Singleton (2002b).

⁴Liu and Pan (2002) and Schroder and Skiadas (1999) also associate the pricing kernel in the economy with the pricing kernel for the investor. In these models, markets are complete, so a unique pricing kernel exists. Schroder and Skiadas (2002) extend the results of Schroder and Skiadas (1999) to cases of incomplete markets. They do not consider the case of incompleteness arising from inflation, which is the focus of our theoretical results.

imum likelihood to estimate the parameters of the process, and demonstrates that the model provides a good fit to term structure data and to inflation. Section 5 discusses the properties of the optimal portfolio for the parameters we have estimated and calculates certainty equivalent losses resulting from sub-optimal strategies.

2 The Economy

As in the affine term structure literature, we specify an exogenous nominal pricing kernel. Because our purpose is modeling predictability in excess bond returns and, as Dai and Singleton (2002a) and Duffee (2002) show, a Gaussian model is best suited for this purpose, we will assume that all variables are homoscedastic.⁵

Let dz denote a $d \times 1$ vector of independent Brownian motions. Let $r(t)$ denote the instantaneous nominal riskfree rate. We assume that

$$r(X(t), t) = \delta_0 + \delta X(t), \quad (1)$$

where $X(t)$ is an $m \times 1$ vector of state variables that follow the process

$$dX(t) = K(\theta - X(t)) dt + \sigma_X dz(t), \quad (2)$$

where σ_X is an $m \times d$ matrix of constants. Suppose there exists a price of risk $\Lambda(t)$ that is linear in $X(t)$:

$$\Lambda(t) = \lambda_1 + \lambda_2 X(t), \quad (3)$$

where λ_1 is $d \times 1$ and λ_2 is $d \times m$. When $\lambda_2 = 0_{d \times m}$, the price of risk is constant and the model is a multifactor version of Vasicek (1977). Given a process for the interest rate r and the price of risk Λ , the pricing kernel is given by:

$$\frac{d\phi(t)}{\phi(t)} = -r(t) dt - \Lambda(t)^\top dz. \quad (4)$$

The pricing kernel determines the price of an asset based on its nominal payoff.

⁵Fisher (1998) shows that a two-factor Gaussian model can partially replicate the failure of the expectations hypothesis, but does not make comparisons across models. Bansal and Zhou (2002) show that a regime-switching is also successful at capturing the failure of the expectations hypothesis in the data. Ahn, Dittmar, and Gallant (2002) discuss an affine-quadratic class of models which, as Brandt and Chapman (2002) show, is also capable of accounting for the expectations hypothesis. Extensions of the results in this paper to quadratic models and models with regime shifts will be considered in future work.

In this economy, bond yields are affine in the state variables $X(t)$. Let $P(X(t), t, s)$ denote the price of such a bond maturing at $s > t$. Then P equals the present discounted value of the bond payoff, namely \$1:

$$P(X(t), t, s) = \phi(t)^{-1} E_t [\phi(s)]$$

As shown by Duffie and Kan (1996) nominal bond prices take the form:

$$P(X(t), t, s) = \exp \{A_2(s-t)X(t) + A_1(s-t)\}, \quad (5)$$

where $A_2(\tau)$ and $A_1(\tau)$ solve a system of ordinary differential equations given in Appendix A. Bond yields are given by

$$\begin{aligned} y(X(t), t, s) &= -\frac{1}{s-t} \log P(X(t), t, s) \\ &= -\frac{1}{s-t} (A_2(s-t)X(t) + A_1(s-t)) \end{aligned} \quad (6)$$

The dynamics of bond prices follow from Ito's lemma:

$$\begin{aligned} \frac{dP(t)}{P(t)} = \left\{ -A_2'(\tau)X(t) - A_1'(\tau) + A_2(\tau)K(\theta - X(t)) + \frac{1}{2}A_2(\tau)\sigma_X\sigma_X^\top A_2(\tau)^\top \right\} dt \\ + A_2(\tau)\sigma_X dz. \end{aligned} \quad (7)$$

The expression for the drift of bond prices can be simplified by applying the expressions for A_2 and A_1 given in Appendix A:

$$\frac{dP(t)}{P(t)} = (A_2(\tau)\sigma_X\Lambda(t) + r(t)) dt + A_2(\tau)\sigma_X dz.$$

Equation (7) shows that bond prices vary with the state variables $X(t)$. The correlation between bond prices and state variables depends on the maturity of the bond through the function $A_2(\tau)$. We will assume (without loss of generality) that δ , K , σ_X , and Λ are such that there are as many non-redundant bonds in the economy as state variables. If this is not the case, then one of the state variables can be removed with no impact on bond prices. With slight abuse of notation, we let $P(t)$ denote a vector of m bond prices, with A_2 the $m \times m$ matrix with rows equal to the corresponding values of $A_2(\tau)$.

Our framework allows for the existence of other assets besides bonds. For concreteness, we assume there exists a stock portfolio with price dynamics

$$\frac{dS(t)}{S(t)} = (\sigma_S\Lambda + r) dt + \sigma_S dz, \quad (8)$$

The row vector σ_S is assumed to be linearly independent of the rows of σ_X , so that the stock is not spanned by bonds. We can then group the existing assets into the vector process:

$$\begin{pmatrix} dP(t) \\ dS(t) \end{pmatrix} = \text{diag} \begin{pmatrix} P \\ S \end{pmatrix} (\mu(t) dt + \sigma dz), \quad (9)$$

where

$$\sigma = \begin{pmatrix} A_2 \sigma_X \\ \sigma_S \end{pmatrix}, \quad (10)$$

and μ is such that

$$(\mu - \iota r) = \sigma \Lambda \quad (11)$$

with ι equal to an $(m+1) \times 1$ vector of ones. Because we have assumed there exist m non-redundant bonds, and because the stock is not redundant, the variance-covariance matrix of the assets, $\sigma \sigma^\top$ is invertible.

Equation (11) shows why this specification allows for predictable excess returns. Because Λ is a function of the state variables $X(t)$, the instantaneous expected excess return $\mu - r$ will also be a function of $X(t)$. The structure of λ_2 will determine how quantities that are correlated with the state variables, such as the yield spread, predict asset returns.

So far, we have described the nominal economy. Because we are interested in the strategies for an investor who cares about real wealth, it is necessary to define a process for the price level. Define a stochastic price level $\Pi(t)$ such that

$$\frac{d\Pi(t)}{\Pi(t)} = \pi(X(t), t) dt + \sigma_\Pi dz. \quad (12)$$

It is assumed that $\pi(t)$ is affine in the state variables. In particular:⁶

$$\pi(t) = \zeta_0 + \zeta X(t). \quad (13)$$

It may at first seem unnatural to require that expected inflation be a linear function of the state variables. However, it is no different than defining the underlying variables as the real interest rate and an expected inflation process. For example, Brennan, Wang, and Xia (2002) and Campbell and Viceira (2001) consider complete-market economies where a real riskfree rate exists, and where the real riskfree rate and the expected inflation rate jointly follow a multivariate Ornstein-Uhlenbeck

⁶It is sufficient for the portfolio choice results to require that $r(t) - \pi(t)$ is an affine function. However, (1) is required to achieve affine bond prices.

process. We could redefine our state variables so that the first state variable is given by π , the expected inflation rate, and the second state variable by $r - \pi$, which in a model of complete markets and constant risk premia, is a constant plus the real riskfree rate. The model would be the same as what we have now, only the notation would be different. The formulation (13) has the econometric advantage that it allows us to estimate these processes as a sum of orthogonal state variables. It is also convenient in the case of incomplete markets where, as we discuss further below, a real riskfree rate does not exist.

While we started by defining a pricing kernel for nominal assets, we could have equivalently defined payoffs in real terms, and defined a pricing kernel for real assets. In fact, any nominal pricing kernel $\phi(t)$ is associated with a “real” pricing kernel. For an asset with nominal value $V(s)$ at time s , the price at time t (assuming the asset pays no dividends between t and s) equals

$$V(t) = E_t \left[\frac{\phi(s)}{\phi(t)} V(s) \right]. \quad (14)$$

It follows directly from (14) that for the real payoff $V(s)/\Pi(s)$,

$$\frac{V(t)}{\Pi(t)} = E_t \left[\frac{\phi(s)\Pi(s)}{\phi(t)\Pi(t)} \left(\frac{V(s)}{\Pi(s)} \right) \right]. \quad (15)$$

Therefore $\phi(t)\Pi(t)$ is a valid pricing kernel when asset prices are expressed in real terms. This also follows from the interpretation of $\phi(t)$ as a system of Arrow-Debreu state prices. Normalizing $\phi(0) = 1$ and $\Pi(0) = 1$, $\phi(t)$ is a ratio of units of consumption at time 0 to dollars at time t . Then $\phi(t)\Pi(t)$ is a ratio of consumption at time 0 to consumption at time t . We choose to model prices in nominal rather than real terms for ease of comparison to the affine term structure literature.

Given that nominal prices can be transformed into real prices, one may ask whether the inflation process (12) plays a substantive role in the analysis. It does, as long as we assume, realistically, that the price level cannot be perfectly hedged by trading in the underlying assets. When Π cannot be perfectly hedged, there exists an asset that is riskless in nominal terms but not in real terms. Because markets are incomplete, inflation matters.

The connection between market incompleteness and the lack of a real riskfree rate can also be seen from the real pricing kernel associated with the price of risk Λ . From Ito’s lemma, it follows that

$$\frac{d(\phi(t)\Pi(t))}{\phi(t)\Pi(t)} = (-r(t) + \pi(t) - \sigma_{\Pi}\Lambda(t)) dt + (\sigma_{\Pi} - \Lambda(t)) dz \quad (16)$$

If a real riskfree rate were to exist, its real return must equal $-r + \pi(t) - \sigma_{\Pi}\Lambda(t)$, the drift of the real pricing kernel. Note however that $\sigma_{\Pi}\Lambda(t)$ is not well-defined if markets are incomplete. In particular, we could replace Λ by some other price of risk $\tilde{\Lambda}$. As long as $\sigma_X\tilde{\Lambda}$ and $\sigma_S\tilde{\Lambda}$ were the same as $\sigma_X\Lambda$ and $\sigma_S\Lambda$, then Λ and $\tilde{\Lambda}$ would result in the same asset prices. However, they would in general lead to different values of $\sigma_{\Pi}\Lambda$, and thus different real riskfree rates.

In what follows, it will be useful to distinguish the unique price of risk that both prices, and is spanned by, the underlying assets:

$$\Lambda^* = \sigma^\top (\sigma\sigma^\top)^{-1} \sigma\Lambda = \sigma^\top (\sigma\sigma^\top)^{-1} (\mu - r\iota). \quad (17)$$

The last equality shows that Λ^* is not dependent on which pricing kernel Λ , is chosen, as long as Λ correctly prices the underlying assets. One reason Λ^* is useful is that its norm is equal to the maximal Sharpe ratio:

$$\max_{\sigma} \frac{\sigma\Lambda^*}{\sqrt{\sigma\sigma^\top}} = \frac{(\Lambda^*)^\top \Lambda^*}{\sqrt{(\Lambda^*)^\top \Lambda^*}} = \sqrt{(\Lambda^*)^\top \Lambda^*},$$

which follows from the Cauchy inequality. The maximum Sharpe ratio is always positive, even if Λ^* is not; this is because an investor can take both short and long positions in any asset. Because we have assumed homoscedasticity, Λ^* has the same functional form as Λ , with

$$\lambda_1^* = \sigma^\top (\sigma\sigma^\top)^{-1} \sigma\lambda_1 \quad (18)$$

$$\lambda_2^* = \sigma^\top (\sigma\sigma^\top)^{-1} \sigma\lambda_2. \quad (19)$$

replacing λ_1 and λ_2 in (3).

The investment opportunity set can be summarized as follows. The investor has access to an asset with riskless nominal return r , and $m + 1$ risky assets whose nominal price dynamics are described by (9), (10), (11). Nominal markets are complete in that there exists a full term structure of nominal bonds.⁷ However, real markets are incomplete, because no asset spans unexpected inflation. Equivalently, there is no asset that is riskless in real terms.

3 Optimal portfolio choice

In this section, we derive the optimal portfolio allocation for an investor who takes bond and stock prices as given. Section 3.1 describes the general form of the solution

⁷Below we will also consider cases where the investor has access to only a subset of the bonds (incomplete nominal markets).

when there is unexpected inflation. Section 3.2 specializes to the case of an affine term structure.

3.1 Complete nominal markets: General results

We first solve the portfolio choice problem for an investor with power utility over terminal wealth at date T , and then generalize to the case of consumption withdrawal. We assume that the investor solves:

$$\max_{W(T) > 0} E_t \left[\frac{(W(T)/\Pi(T))^{1-\gamma}}{1-\gamma} \right], \quad (20)$$

such that $W(T)$ can be achieved by taking positions in the underlying assets with initial wealth $W(0)$:

$$\frac{dW(t)}{W(t)} = w(t)^\top (\mu(t) - r(t)\iota) dt + r(t) dt + w(t)^\top \sigma(t) dz \quad (21)$$

where $w(t)$ is an $(m+1) \times 1$ vector of portfolio weights that satisfies integrability conditions. To disallow doubling strategies, we require that $W(t) > 0$ for all t (see Dybvig and Huang (1988)).

To solve this problem, it is convenient to use the martingale technique of Cox and Huang (1989), Karatzas, Lehoczky, and Shreve (1987) and Pliska (1986) generalized to the case of incomplete markets by He and Pearson (1991). He and Pearson show that, for some endogenous pricing kernel, the dynamic budget constraint (21) can be replaced by a static budget constraint. Let $\phi_\nu(t)$ denote this endogenous pricing kernel. Given the pricing kernel $\phi_\nu(t)$, the static budget constraint equals:

$$E[\phi_\nu(T)W(T)] = W(0). \quad (22)$$

Therefore, for some Lagrange multiplier l , the investor's first-order condition equals

$$\frac{W(T)^{-\gamma}}{\Pi(T)^{1-\gamma}} = l\phi_\nu(T),$$

and the optimal terminal wealth policy is given by

$$W(T) = (l\phi_\nu(T)\Pi(T)^{1-\gamma})^{-\frac{1}{\gamma}}. \quad (23)$$

Substituting back into (22) gives the expression for l .⁸

⁸Solving (22) for l implies

$$l = W(0)^{-\gamma} \left(E \left(\phi_\nu(T)^{1-\frac{1}{\gamma}} \Pi(T)^{1-\frac{1}{\gamma}} \right) \right)^\gamma.$$

The investor's terminal wealth policy has an economic interpretation. Rearranging,

$$\frac{W(T)}{\Pi(T)} = (l\phi_\nu(T)\Pi(T))^{-\frac{1}{\gamma}}. \quad (24)$$

The left hand side is equal to real wealth. The term in the parenthesis on the right hand side is proportional to $\phi_\nu(T)\Pi(T)$. This equals the real pricing kernel corresponding to the nominal kernel ϕ_ν . Thus (24) states that the greater the price of a given state, the less the agent will consume in that state. The lower the risk aversion (γ), the more the agent adjusts terminal wealth in response to changes in the state-price density. Note however, that ϕ_ν is also implicitly a function of γ .

The optimal portfolio allocation is derived using (23). Following Cox and Huang (1989), define a new state variable equal to the real wealth of the log utility investor. In our environment with inflation, this state variable equals:

$$Z(t) = (l\phi_\nu(t)\Pi(t))^{-1}. \quad (25)$$

No-arbitrage implies that wealth at time t must equal the present discounted value of wealth at time T , where the discounting is accomplished by the state-price density:

$$\begin{aligned} W(t) &= \phi_\nu(t)^{-1} E_t \left[\phi_\nu(T)\Pi(T)Z(T)^{\frac{1}{\gamma}} \right] \\ &= \Pi(t)Z(t) E_t \left[Z(T)^{\frac{1}{\gamma}-1} \right]. \end{aligned} \quad (26)$$

The next theorem characterizes the optimal wealth and portfolio weights.

Theorem 1 *Assume that the investor has utility over terminal wealth with coefficient of relative risk aversion γ . At time t , optimal wealth takes the form*

$$W(t) = \Pi(t)Z(t)^{\frac{1}{\gamma}} F(X(t), t, T), \quad (27)$$

where $Z(t)$ is given by (25). The minmax pricing kernel equals

$$\frac{d\phi_\nu}{\phi_\nu} = -r dt - (\Lambda^* + \nu)^\top dz,$$

with

$$\nu = (1 - \gamma) \left(\sigma_\Pi - (\sigma_\Pi \sigma^\top) (\sigma \sigma^\top)^{-1} \sigma \right)^\top. \quad (28)$$

The function F satisfies the partial differential equation

$$\begin{aligned} \frac{1-\gamma}{\gamma}(r-\pi)F + F_X \left(K(\theta - X) + \frac{1}{\gamma}\sigma_X(\Lambda^* + \nu) + \frac{\gamma-1}{\gamma}\sigma_X\sigma_\Pi^\top \right) + F_t + \\ \frac{1}{2} \left(\frac{1-\gamma}{\gamma}((\Lambda^* + \nu)^\top(\Lambda^* + \nu) + \sigma_\Pi^\top\sigma_\Pi)F + \text{tr} \left(F_{XX}\sigma_X\sigma_X^\top \right) \right) = \\ \frac{\gamma-1}{\gamma}\sigma_\Pi(\Lambda^* + \nu)F + F_X\sigma_X(\Lambda^* + \nu), \quad (29) \end{aligned}$$

with boundary condition $F(X(T), T, T) = 1$.⁹ The optimal portfolio allocation equals

$$\begin{aligned} w(t) = \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}(\mu - \iota r) + \left(1 - \frac{1}{\gamma} \right) (\sigma\sigma^\top)^{-1}(\sigma\sigma_\Pi^\top) \\ + (\sigma\sigma^\top)^{-1}(\sigma\sigma_X^\top) \frac{1}{F} (F_X)^\top. \quad (30) \end{aligned}$$

The remainder of the investor's wealth, $1-w(t)^\top \iota$, is invested in the nominal riskfree asset.

The proof is given in Appendix B. The minmax price of risk equals the price of risk spanned by the existing assets Λ^* , plus an adjustment term. The adjustment, ν , equals $1-\gamma$ times the unhedgeable part of inflation risk. ν is thus an investor-specific measure of market incompleteness.

Equation (30) shows that the investor can be viewed as investing in $m+2$ risky asset "funds". The first fund is the portfolio that is instantaneously mean-variance efficient. It is straightforward to check that this portfolio achieves the maximum Sharpe ratio $\sqrt{(\Lambda^*)^\top \Lambda^*}$. The second fund adjusts for the fact that the first fund is mean-variance efficient in nominal rather than real terms. Together, these portfolios constitute what is known as "myopic demand", namely the optimal allocation if the investor ignores the future investment opportunity set.

It is the last term in (30) that is the focus of this study. This term represents the sum of the m hedge portfolios:

$$(\sigma\sigma^\top)^{-1}(\sigma\sigma_X^\top) \frac{1}{F} (F_X)^\top = \frac{1}{F} \sum_{j=1}^M (\sigma\sigma^\top)^{-1}(\sigma\sigma_{X_j}^\top) F_{X_j}$$

Hedge portfolio j is formed by projecting state variable j onto the available assets. Scaling the portfolio is the sensitivity of wealth to state variables j , $\frac{1}{F} (F_{X_j})^\top$. If

⁹ $\text{tr}(\cdot)$ denotes the trace. F_{XX} is the $m \times m$ matrix of second derivatives.

increases in state variable j increase wealth in the future, then the investor allocates a positive amount to the hedge portfolio $(\sigma\sigma^\top)^{-1}(\sigma\sigma_{X_j}^\top)$, a negative amount if the effect on wealth is negative. Because we have assumed that there are as many non-redundant bonds as state variables, it is possible to completely hedge the state variables by trading in the underlying assets. Moreover, hedging demand for bonds will be nonzero. Because bonds are the discounted value of \$1, their prices covary with the variables that affect the investment opportunity set, namely $X(t)$.

Also of interest is the investor's indirect utility. Cox and Huang (1989) show that it is possible to derive indirect utility from the expression for wealth. Corollary 2 generalizes this result to the case where there is unexpected inflation (and specializes to the case of power utility):

Corollary 2 *Define the investor's indirect utility function as follows:*

$$J(W(t), \Pi(t), X(t), t, T) = E_t \left[\frac{1}{1-\gamma} \left(\frac{W(T)}{\Pi(T)} \right)^{1-\gamma} \right] \quad (31)$$

Then $J(W, \Pi, X, t, T)$ takes the form

$$J(W(t), \Pi(t), X(t), t, T) = \frac{1}{1-\gamma} \left(\frac{W(t)}{\Pi(t)} \right)^{1-\gamma} F(X(t), t, T)^\gamma$$

where $F(X(t), t, T)$ is defined in Theorem 1.

The proof of Corollary 2 can be found in Appendix B.

These results generalize to the case where the investor has utility over consumption between times 0 and T . At each time, besides allocating wealth among assets, the investor also decides what proportion of wealth to consume. The investor solves

$$\max E \left[\int_0^T e^{-\rho t} \frac{(c(t)/\Pi(t))^{1-\gamma}}{1-\gamma} dt \right] \quad (32)$$

$$\text{s.t. } dW(t) = \left(w(t)^\top (\mu(t) - r(t) \iota) + r(t) \right) W(t) dt + w(t)^\top \sigma W(t) dz - c(t) dt$$

$$W(T) \geq 0$$

As shown in Wachter (2002a), computing the solution to this case does not require solving a new partial differential equation.¹⁰ As in the case of terminal

¹⁰While the results in Wachter (2002a) assumed that markets were complete, the same reasoning can be applied here because the adjustment for incomplete markets in the minmax pricing kernel (28) takes a particularly simple form.

wealth, the dynamic problem can be recast as static problem for an endogenous pricing kernel. Using arguments similar to those in the proof of Theorem 1, it can be shown that, when the only market incompleteness comes from inflation, the investor-specific pricing kernel (ϕ_ν) for the case of intermediate consumption takes the same form as the investor-specific pricing kernel for terminal wealth. The static budget constraint is therefore equal to:

$$E \left[\int_0^T c(t) \phi_\nu(t) dt \right] = W(0) \quad (33)$$

The following corollary describes the form of the investor's consumption policy, optimal wealth, and portfolio allocation.

Corollary 3 *The optimal consumption policy $c(t)$ satisfies:*

$$\frac{c(t)}{\Pi(t)} = (l \phi_\nu(t) \Pi(t))^{-\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma} t}, \quad (34)$$

where l is the Lagrange multiplier that allows (33) to hold. Optimal wealth is given by

$$W(t) = Z(t)^{\frac{1}{\gamma}} \Pi(t) \int_t^T F(X(t), t, s) e^{-\frac{\rho}{\gamma}(s-t)} ds, \quad (35)$$

where $Z(t)$ is defined by (25), and F satisfies the partial differential equation (29). The optimal portfolio weights are given by (30) with F replaced by $\int_t^T F e^{-\frac{\rho}{\gamma}(s-t)}$.

Theorem 1 generalizes the well-known result that the price of risk and the riskfree rate are sufficient statistics for the investment opportunity set when markets are complete. Unless either the price of risk or the interest rate vary, the optimal portfolio rule is myopic. Moreover, if two economies have the same process for the price of risk and riskfree rate, the optimal consumption and wealth process of the agent will be the same, even though the weights will depend on the specific assets that trade.

Theorem 1 shows that in the setting of unhedgeable inflation risk, the minmax price of risk $\Lambda^* + \nu$ and the difference between the nominal interest rate and expected inflation $r - \pi$ are sufficient statistics for the investment opportunity set. Assuming that both of these are constants results in a function F that is identically 1, as the partial differential equation (29) shows. Not only does time-variation in the price of risk matter, but so does ν , the component of inflation risk that investors are not able to hedge. The other important quantity for investors is $r - \pi$. For convenience,

we will abuse terminology slightly and refer to this as the real riskfree rate, keeping in mind that no asset that is riskfree in real terms exists.¹¹

3.2 Portfolio allocation when the nominal term structure is affine

Theorem 1, Corollary 2, and Corollary 3 do not require that bond yields be affine. They hold generally, as long as the investor has power utility over terminal wealth. The following corollary explicitly solves for the portfolio weights, given the assumptions on Λ , r , and π .

Corollary 4 *Assume Λ and $r - \pi$ are linear in the state variables $X(t)$, and that inflation and asset prices are homoscedastic, and the investor has utility over terminal wealth given by (20). Then F takes the form:*

$$F(X(t), t, T) = \exp \left\{ \frac{1}{\gamma} \left(\frac{1}{2} X(t)^\top B_3(\tau) X(t) + B_2(\tau) X(t) + B_1(\tau) \right) \right\}, \quad (36)$$

where $\tau = T - t$ and the matrix B_3 , the vector B_2 , and the scalar B_1 satisfy a system of ordinary differential equations. The optimal portfolio rule equals:

$$w(t) = \frac{1}{\gamma} (\sigma \sigma^\top)^{-1} (\mu - \iota r) + \frac{\gamma - 1}{\gamma} (\sigma \sigma^\top)^{-1} (\sigma \sigma_\Pi^\top) + \frac{1}{\gamma} (\sigma \sigma^\top)^{-1} (\sigma \sigma_X^\top) \left(\frac{B_3(\tau) + B_3(\tau)^\top}{2} X(t) + B_2(\tau) \right). \quad (37)$$

The remainder of the investor's wealth, $1 - w(t)^\top \iota$, is invested in the nominal riskfree asset.

The proof of Corollary 4 and the differential equations for B_3 , B_2 , and B_1 can be found in Appendix B. A noteworthy special case arises when risk premia are constant. Then $B_3(\tau) = 0$, as can be checked by setting $\lambda_2^* = 0$ into the differential equation for B_3 . The optimal portfolio allocation is constant, and F is exponential-affine. A two-factor version of this case is considered by Brennan and Xia (2002).

Why do time-varying risk premia produce a functional form that is exponential-quadratic? As Campbell and Viceira (1999) discuss, the reason is that the investor can profit both when risk premia $\sigma \Lambda$ are especially high and positive, and when they are especially low and negative. A function for wealth that is quadratic in $X(t)$

¹¹Indeed, the results in Section 2 show that this will only be the real riskfree rate if markets are completed such that the price of inflation risk is zero.

captures this quality. Note that exponential-quadratic wealth implies a portfolio rule that is linear in the state variables.

Using Corollary 3, it is also possible to write down an explicit formula for the optimal portfolio for an investor with utility over consumption.

Corollary 5 *Assume Λ and $r - \pi$ are linear in the state variables $X(t)$, and that inflation and asset prices are homoscedastic. Suppose the investor has utility over consumption. The optimal portfolio weights equal:*

$$w(t) = \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}(\mu - \iota r) + \frac{\gamma - 1}{\gamma}(\sigma\sigma^\top)^{-1}(\sigma\sigma_{\Pi}^\top) + \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}(\sigma\sigma_X^\top) \left(\frac{\int_t^T F(t, t + \tau) \left(\frac{1}{2}(B_3(\tau) + B_3(\tau)^\top)X(t) + B_2(\tau)^\top \right) e^{-\frac{\rho}{\gamma}\tau} d\tau}{\int_t^T F(t, t + \tau) e^{-\frac{\rho}{\gamma}\tau} d\tau} \right)$$

The results above show that wealth, indirect utility, and the optimal allocation are available in closed form up to the solution of ordinary differential equations. In the following sections, we estimate the parameters of the model and evaluate the implications for portfolio choice.

4 Estimation

The previous sections described optimal portfolio choice when the nominal term structure is affine and the investor has access to stock as well as bonds. In this section we estimate a three-factor term structure model that has been shown to perform well in out-of-sample forecasting (Duffee (2002)), and in replicating the failure of the expectations hypothesis seen in the data (Dai and Singleton (2002a))¹². Our estimation differs from that in previous studies in that we incorporate data on equity returns, and most importantly, on inflation.

As Dai and Singleton (2000) discuss, the processes for X , Λ and r have too many degrees of freedom to be identified by the data. For example, it is not possible to simultaneously identify θ and δ_0 . For a given number of factors, Dai and Singleton (2000) specify a canonical form that can be identified. We follow their approach,

¹²In the notation of these papers, the model we estimate is known as $A_0(3)$, because it contains three factors and no square root processes.

modified as necessary for inflation and stock returns. We assume

$$\begin{aligned}\theta &= 0_{3 \times 1} \\ \sigma_X &= \begin{bmatrix} 0_{1 \times 3} & I_{3 \times 3} & 0_{1 \times 3} \end{bmatrix}.\end{aligned}$$

In addition, the canonical form requires that K be lower triangular. To reduce the number of parameters required for estimation, and to ameliorate concerns of overfitting, we estimate the preferred model of Duffee (2002), which involves setting some elements of K and of λ_2 to zero, as described further below.

We make the assumption that realized inflation is instantaneously uncorrelated with bond prices:

$$\sigma_{\Pi} = \begin{bmatrix} \sigma_{\Pi(1)} & 0_{1 \times 4} \end{bmatrix}. \quad (38)$$

In addition we assume that stock prices and inflation are instantaneously uncorrelated, but allow the covariance between stocks and bonds to be unrestricted:

$$\sigma_S = \begin{bmatrix} 0 & \sigma_{S(1)} & \sigma_{S(2)} & \sigma_{S(3)} & \sigma_{S(4)} \end{bmatrix}$$

Even though inflation is instantaneously uncorrelated with stock and bond prices, over finite intervals, it will in general be correlated with both. The implications of our continuous-time processes for data observed at finite intervals is discussed in Appendix C.

An advantage of assuming this form of σ and σ_{Π} is that it simplifies the estimation of λ_1^* and λ_2^* . We define

$$\lambda_1^* = \begin{bmatrix} 0 & \lambda_{1(1)}^* & \dots & \lambda_{1(4)}^* \end{bmatrix}^{\top}$$

and similarly

$$\lambda_2^* = \begin{pmatrix} 0 & 0 & 0 \\ \lambda_{2(1,1)}^* & \dots & \lambda_{2(1,3)}^* \\ \vdots & & \vdots \\ \lambda_{2(4,1)}^* & \dots & \lambda_{2(4,3)}^* \end{pmatrix}$$

We make this assumptions for two reasons. First, as discussed in Section 2, the price of inflation risk is indeterminate. Because σ_{Π} takes the form (38), this says that the first entry of Λ , and hence the first entry of λ_1 and the first row of λ_2 are indeterminate. Fixing these entries at a constant value allows Λ to be identified. Second, optimal portfolio choice requires an estimate of Λ^* , the unique price of risk spanned by asset returns. Setting the indeterminate entries to zero insures that Λ

equals Λ^* , rather than some other valid price of risk. In addition, we assume that excess stock returns are not predictable, though this is easy to relax. Bond data pin down three dimensions of Λ^* , and the equity premium determines the fourth.

Our bond data consist of monthly observations on zero-coupon yields for 3 month, 6 month, 1, 2, 5, and 10 year U.S. government bonds. The bond data is available from the website of Gregory Duffee. Monthly observations on the CPI and on returns on a broad stock index are available from CRSP. The sample begins in 1952 and ends in 1998. Following Duffee (2002), we assume that prices on the 3 month, 1 year, and 5-year bonds are measured with normally distributed errors. We then estimate the parameters of the model using maximum likelihood as described in Appendix C.

Table 1 describes the results from our estimation. Because the yields are in annual terms, time is in years. The parameters ζ_0 and δ_0 equal the sample means of inflation and the nominal interest rate. Both of these parameters equal their sample means from the data. While this may seem like a natural property, as Campbell and Viceira (2001) discuss, it is not guaranteed that the models fit the time series mean. In fact, the affine models investigated by Duffee (2002) all result in a sample mean for the nominal interest rate that is too low.¹³ Surprisingly, including inflation in the estimation helps to estimate this parameter. Table 1 also shows that the volatility of the inflation residual σ_{Π} is estimated to be 0.93%. This is close to, but smaller than the volatility of realized inflation in the data (1.17%). This makes sense; the state variables add information and reduce the variance of unexplained volatility.

Other than δ_0 described above, the parameters that we estimate for the term structure are very close to those found by Duffee (2002).¹⁴ The components of λ_1^* are significantly negative, but estimated with noise. As shown by (3), λ_1^* corresponds to the mean of the price of risk corresponding to each state variable. Because bond prices load negatively on the state variables, negative values of λ_1^* imply positive, but noisy, average risk premia on bonds. The estimates of λ_2^* imply that two factors determine time-varying risk premia on bonds. The first is given by the transitory factor X_2 , while the second is a linear combination of X_1 , X_2 and X_3 , and hence is more persistent. The final panel of Table 1 shows the estimated risk premium

¹³Duffee ends his sample in 1994. This does not account for the difference however. We estimate the $A_0(3)$ model without inflation, and find $\delta_0 = 4.4\%$, even when we include the last four years of the sample.

¹⁴The variance covariance matrix for the errors, which we do not report, is nearly identical to that found by Duffee (2002).

on equity to be 7.5%, with large standard errors. The second, third and fourth elements of σ_S are nonzero, and the second and fourth are significantly negative. This implies a nonzero correlation between stock and bond returns, and suggests (because bonds depend negatively on the state variables) that this correlation will be positive. Indeed, Table 2, which shows the correlations between the assets, shows that the correlation between bonds and stocks is about 0.2. The five and the ten-year bond are very highly correlated, as are the five and the one-year bond.

Figures 1-3 illustrate the implications of the model for average yield spreads, standard deviations of yield spreads, and Campbell-Shiller long-rate regressions. Each figure plots the values in the data (“sample”) and the values implied by the model and the parameters in Table 1 (“population”). Following Dai and Singleton (2002a), we construct 95% confidence bands by simulating 500 sample paths from our model with length equal to the sample path in the data. Figures 1 and 2 show that the model implies average yield spreads and standard deviations of yield spreads close to those found in the data. The confidence bands reflect the well-known result that means are estimated much more imprecisely than variances. In both cases, the data falls well within the error bands implied by the model. We conclude that the model does a reasonable job of fitting the cross-sectional moments of bond yields – not a guarantee as the model must fit cross-sectional and time-series moments together.

Because our aim is to study the implications of the expectations puzzle for investors, it is especially important to determine whether the model accounts for the expectations puzzle found in the data. To do so, we follow the approach of Dai and Singleton (2002a) and check whether the model replicates the empirical findings of Campbell and Shiller (1991). Dai and Singleton explain the connection between the Campbell-Shiller regressions and time-variation in risk premia in detail.

Figure 3 plots the slope coefficients from regressions of quarterly changes in yields on the scaled yield spread, as described in Campbell and Shiller (1991). If the expectations hypothesis held, the coefficients would be identically equal to 1. Instead, Campbell and Shiller find coefficients that are negative and decrease with maturity. Figure 3 replicates this result in our data, and shows that the model captures both the negative coefficients and the downward slope. Except for values at the very short end of the term structure, the data falls within the 95% confidence bands implied by the model. It is apparent from Figure 3 that the model captures the failure of the expectations hypothesis found in the data. To the extent that the

failure of the expectations hypothesis is a bit less extreme in the model than the data, we may understate the implications for long-run investors.¹⁵

Figure 4 plots the time series of monthly realized inflation, and our expected inflation series constructed from the state variables using the relationship

$$\pi(t) = \zeta_0 + \zeta X(t),$$

where values for ζ_0 and ζ come from the maximum likelihood estimation described above, and are given in Table 1. Our joint estimation procedure allows inflation to influence the dynamics of state variables. In practice, however, this effect is small, and except for the effect on δ_0 described above, our parameter values are close to what we would find by first estimating the term structure model, and then regressing realized inflation on the factors. This latter strategy would, of course, understate the standard errors on ζ .

Figure 4 shows that our expected inflation series does a surprisingly good job in accounting for changes in realized inflation. In fact, expected inflation accounts for 37% of the variance of realized inflation. It is worth emphasizing that these results come about even though the factors $X(t)$ are linear combinations of yields alone. Thus long-term bond yields contain substantial information about future inflation.

Figure 5 plots the time series for the nominal interest rate $r(t)$ implied by the parameters estimated in Table 1. While not shown in the graph, $r(t)$ is essentially equal to the three-month yield. Also shown in Figure 5 is expected inflation $\pi(t)$. Both series are highly persistent, and become larger and more volatile in the late 1970's to the early 1980's. The difference between the nominal interest rate $r(t)$ and $\pi(t)$, which we informally refer to as the real interest rate, is positive through nearly the entire sample. Thus the expected inflation and real riskfree rate implied by the model have reasonable time-series properties.

The results in Section 3 show that the real interest rate and the price of risk Λ are the important quantities for investors. Figure 6 plots the time series of risk

¹⁵The literature has identified a number of econometric difficulties with this regression. Non-exogenous regressors bias the coefficients upward, causing the hypothesis to be rejected less strongly than it should be (Bekaert, Hodrick, and Marshall (1997), Stambaugh (1999)), while Peso problems (Bekaert, Hodrick, and Marshall (2001)) result in increased dispersion of the estimates, leading the model to be rejected too strongly. Bekaert and Hodrick (2001) argue that standard tests tend to reject the null of the expectations hypothesis even when it is true. They find, however, that the data remain inconsistent with the expectations hypothesis, even after adjusting for small-sample properties. Accounting for these biases within the investment decision is beyond the scope of this manuscript, but will be pursued in future work.

premia (a linear transformation of Λ) for the one, five, and ten year bonds implied by the model. As Figure 6 shows, risk premia are highly volatile, especially in the latter half of the sample. Table 1 implies that there are two factors driving risk premia: the first is the highly transitory second state variable, the second is a linear combination of all three state variables that is much more persistent. Nonetheless, all three risk premia appear to move closely together. This is consistent with the findings of Cochrane and Piazzesi (2002), who show that a single factor can explain much of the time-variation in expected excess returns on bonds.

Taking the results in this section together, we conclude that our model succeeds in capturing important features of the term structure and of inflation. The next section considers the implications of our parameter estimates for portfolio choice.

5 Portfolio choice under the failure of the Expectations Hypothesis

This section combines the theoretical results from Section 3 with the parameter estimates from Section 4 to evaluate the implications of the failure of the expectations hypothesis for long-horizon investors. The failure of the expectations hypothesis can affect the optimal portfolio in two possible ways. First, the myopic (mean-variance efficient) component of the optimal portfolio, $\frac{1}{\gamma}(\sigma\sigma^\top)^{-1}(\mu - \iota r)$, depends directly on risk premia. If risk premia vary, so will myopic demand. Second, time-varying risk premia imply that investment opportunities vary over time (as long as changes in risk premia are not directly offset by changes in volatility). As Merton (1971) shows, the investor hedges these changes in the investment opportunity set, implying that the optimal allocation is not mean-variance efficient. Hedging demand causes the optimal portfolio for a long-horizon investor to differ from the optimal portfolio for an investor with a short horizon. Both effects are present in theory. The question is, are they economically significant?

5.1 Optimal allocation between a long-term bond and the nominal riskfree asset

To investigate the effect of time-varying risk premia on optimal portfolios, we first consider the case where the investor has access to a single long-term bond and a nominally riskfree asset. This case allows us to temporarily abstract from questions pertaining to the optimal composition of the bond portfolio, and focus on the horizon

properties taking the composition as given. The results in Theorem 1 apply only to the case where nominal markets are complete, namely when there are the same number of long-term bonds as state variables. However, they are easily modified for the case of incomplete nominal markets. Optimal wealth and allocation to long-term bonds still take the same form as in Corollary 4. Theorem 1 and Corollary 4 are extended to the incomplete-market case in Appendix C.¹⁶

Figure 7 plots the optimal allocation for the investor who allocates wealth between a five-year bond and the nominally riskfree asset. The investor is assumed to have utility over wealth at the end of the horizon, and risk aversion γ of 10. Both myopic demand and hedging demand depend on the current premia on bonds over the riskfree rate. Thus the optimal allocation is a function of the state as well as horizon. In order to understand how the optimal portfolio varies with the state, we plot the optimal allocation when the state variables are equal to their long-run mean of zero, and then we vary each state variable by two unconditional standard deviations. The results are similar in each case, so we discuss only the effects of varying X_1 .¹⁷

The parameter estimates in Table 1 imply that the price of risk Λ is increasing in X_1 : $\lambda_{2(3,1)} > 0$. Therefore bond premia are decreasing in X_1 because bond prices are negatively correlated with the state variables. The risk premium on the ten-year bond equals 2% per annum when the state variables are at their long-run mean, 11% when X_1 is two standard deviations below its long-run mean, and -7% when X_1 is two standard deviations above its long-run mean.

Not surprisingly, Figure 7 shows that the lower is X_1 (and the greater are risk premia), the greater is the myopic allocation to the five-year bond. The myopic (mean-variance efficient) allocation is equal to the y-intercept, because, under power utility, the myopic allocation is independent of horizon. In the case of a single risky asset, the myopic allocation takes a simple form: it is proportional to the risk premium divided by the variance. The higher the risk premium, the greater the myopic allocation.

¹⁶The results for utility over consumption (Corollary 5) have no straightforward extension.

¹⁷The unconditional variance-covariance matrix of the state variables can be calculated using the results of Appendix E. The unconditional standard deviation is 0.93 for X_1 , 0.39 for X_2 , and 3.0 for X_3 . Varying X_3 has smaller effects on myopic demand, which can be seen by comparing $0.93\lambda_{2(3,1)}$ to $3\lambda_{2(3,3)}$. Because X_3 is a more persistent variable, the effects on hedging demand are larger. By contrast, X_2 has a larger effect on risk premia, and thus on myopic demand. However, its effects on hedging demand are smaller because it is much less persistent.

There are also strong horizon effects for long-term bonds. For values of X_1 implying positive risk premia, the allocation to the five-year bond rises steadily with the horizon. When risk premia are negative (implying a negative allocation to the five-year bond), the optimal allocation initially falls, but then rises after a horizon of about one year. The difference between short-horizon and long-horizon investors is economically large; when X_1 is at its long-run mean, the myopic investor allocates 40% of his wealth to the long-term bond. An investor with a horizon of 20 years, by contrast, allocates over 100% of his wealth to the long-term bond.

Is it the failure of the expectations hypothesis that leads a long-horizon investor to allocate more wealth to the long-term bond? As discussed in Section 3, hedging demand arises from two sources. One is time-variation in risk premia (the failure of the expectations hypothesis). The other is time-variation in the real riskfree rate, $r - \pi$. To further understand hedging demand, we consider each of these effects in isolation.

First we consider the optimal allocation when the investor has the correct myopic demand, but sets hedging demand assuming that risk premia are constant. Thus the optimal allocation corresponds to (37), with λ_2 set equal to zero in the ordinary differential equations defining $B_3(\tau)$ and $B_2(\tau)$. As noted in Section 3, $B_3(\tau) \equiv 0$ when $\lambda_2 = 0$. When the investor only hedges changes in the real interest rate, hedging demand is constant over time.

In Figure 7, the allocation when the investor hedges only the real riskfree rate is marked with circles. The optimal allocation is still increasing in horizon, but by much less than the optimal allocation. If a ten-year investor hedged only time-variation in the riskfree rate, he would put about 50% of his wealth in the long-term bond, rather than 80%. Clearly time variation in risk premia has a large effect on the optimal portfolio.

What causes the upward slope when the investor hedges the real riskfree rate? Hedging demand represents the investor's desire to hedge changes in the investment opportunity set. A multiperiod investor chooses the optimal portfolio not only to maximize his Sharpe ratio, but also so that realizations in his wealth have the "right" correlation with the real interest rate. If $\gamma > 1$, the investor has lower marginal utility of wealth when the real interest rate is high; the income effect dominates (a higher real interest rate makes him richer, he can afford a lower payoff in those states). If $\gamma < 1$, the investor has lower marginal utility of wealth when the real interest rate is low; the substitution effect dominates (wealth is more valuable when

the interest rate is higher because it can be invested at a higher rate). Suppose for concreteness that $\gamma > 1$. Then the investor will over-weight (relative to the mean-variance efficient allocation) assets that have a negative covariance with changes in the interest rate. These assets pay off when the interest rate is low, thus giving the investor more wealth when marginal utility for wealth is highest.

A number of studies have argued (e.g. Brennan and Xia (2000), Sorensen (1999), Wachter (2002b)) that a time-varying riskfree rate leads investors with longer horizons to allocate a greater percentage of their portfolio to long-term bonds. According to this argument, long-term bonds are negatively correlated with the interest rate, and thus should be over-weighted in the portfolios of investors with risk aversion greater than one. The limitation with this argument is that it requires bonds to be real. Nominal bonds are negatively correlated with the *nominal* interest rate, but the investor desires to hedge the real interest rate $r - \pi$, and nominal bonds may not be negatively correlated with the real interest rate. For our calibration, long-term bonds are indeed negatively correlated with the real interest rate, though it is important to note that this is an empirical, not a theoretical result. Thus the investor with risk aversion greater than one chooses to increase her allocation to long-term bonds relative to the myopic portfolio. Because changes in the real riskfree rate are persistent, the longer the investor's horizon, the greater the effect of the riskfree rate on indirect utility, and the larger is hedging demand.

We now consider the optimal allocation when the investor hedges the risk premium, but not the riskfree rate. This is calculated by setting $\zeta = \delta$ in the equations for B_3 , B_2 , and B_1 . This allocation is shown in Figure 7 and marked with plus signs. Note that time-varying risk premia also cause hedging demand to increase when risk premia are positive. When risk premia are negative, hedging demand is negative at short horizons and positive at long horizons.

The intuition is similar to that for a time-varying interest rate. Consider the case of a single risky asset, and suppose that the risk premium on this asset is positive, so that the investor holds a positive amount in his portfolio. The income effect leads the investor to prefer assets that fall in price when the risk premium rises, because he can afford to have less wealth when there are greater investment opportunities. The substitution effect leads the investor to prefer assets that rise in price when the risk premium rises, because wealth can be invested at a higher rate. For $\gamma > 1$, the first effect dominates, for $\gamma < 1$, the second effect dominates. Supposing that $\gamma > 1$, the investor will over-weight an asset that has a negative correlation with

the risk premium. Because the return on a bond is negatively correlated with its risk premium, this effect leads the investor to allocate more to long-term bonds. Changes in risk premia are persistent, as are changes to the real interest rates. Thus the longer is the investor's horizon, the greater is hedging demand, and the larger is the total allocation to the long-term bond.

This reasoning also explains why hedging demand for the long-term bond can be negative. Figure 7 shows that when risk premia are negative, hedging demand arising from time-variation in risk premia causes the allocation to fall in the horizon before increasing again. Thus hedging demand is negative for some investors. When the investor is short the long-term bond, *decreases* in the risk premium represent improvements in the investment opportunity set. In order to hedge these changes, the investor has a more negative allocation to long-term bonds than the myopic investor. However, rather than steadily decreasing in the horizon, hedging demand begins to increase after a horizon of about two years, and eventually becomes positive.

Campbell and Viceira (1999) and Kim and Omberg (1996) noted the same effect for allocations to stocks. If the risk premium on stocks was negative and close to zero, hedging demand would still be positive. The precise value where hedging demand switched signs was horizon-dependent. This is because if risk premia are negative but close to zero, increases, rather than decreases represent improvements in investment opportunities. This counter-intuitive result arises because the average risk premium is positive. Because the risk premium reverts to its long-term average, if the risk premium is negative, it must pass through zero. A long-horizon investor cares not only about risk premia today, but risk premia at every point in the future as well. All else equal, a long-term investor would prefer positive risk premia because they are likely to stay positive, rather than going through zero, which is the least advantageous value for the investor. Figure 7 shows that this effect is operative in the case of bonds as well.

The solid line in Figure 7 plots the fully optimal allocation. This allocation is not simply a sum of the two effects mentioned above; it arises from a nonlinear interaction between them. Because the investor uses the long-term bond to hedge time-variation in the real riskfree rate, she has an additional reason to prefer positive risk premia in the long run. This effect, and the effect described in the paragraph above, imply that when the risk premium is small and negative, the investor with a sufficiently long horizon would prefer it to become positive, rather than more negative. To hedge the possibility that risk premia fall further, the investor allocates

more wealth to the long-term bond; hedging demands are positive. Mathematically, this can be seen from the equation for $B_2(\tau)$ in Appendix B. From (30), it follows that if $B_2(\tau)$ were zero, then the hedging demand as a function of X_1 would be symmetric around zero. This is not the case because $\lambda_1^* > 0$, namely because the mean risk premium is not zero but positive, and because $\delta - \zeta \neq 0$, namely that the real interest rate is time-varying.¹⁸ Thus because risk premia are positive on average, and because bond returns are negatively correlated with the real interest rate, hedging demand may be positive even when myopic demand is negative.

This section has shown that accounting for time-variation in the risk premia on long-term bonds has two effects on the investor’s optimal portfolio. First, it induces investors to time the bond market. A lower risk premium on a long-term bond leads the investor to allocate less wealth to the bond at all horizons. The second effect arises from the investor’s wish to hedge changes in the risk premium. This causes the optimal portfolio to increase with horizon. This effect is qualitatively large. Thus the failure of the expectations hypothesis “matters” for long-term investors, at least in the case where the investor has access to a single long-term bond. The following section generalizes these results to the case where the investor has access to multiple long-term bonds.

5.2 Optimal allocation to multiple long-term bonds

Figure 8 plots the optimal allocation when the investor has access to a three-year bond, a ten-year bond, and a nominally riskless asset. As in the previous section, we determine the optimal allocation for the long-run mean of the state variables, and for the state variables plus and minus one standard deviation. We report only the effects of varying X_1 .

For all three values of the state variable, the myopic portfolio consists of a short position in at least one of the bonds. These leveraged positions arise because of the correlation structure of bond returns implied by the model (and found in the data). Table 2 shows the implied correlations in bond returns (Panel A), and correlations of monthly log bond returns from the data (Panel B).¹⁹ As Table 2 shows, bonds at all maturities are highly correlated. Thus any estimated difference in the risk-return

¹⁸Note that $\sigma_{\Pi} \lambda_2^* = 0$.

¹⁹Because yield data is unavailable for all maturities, the correlations in Panel B rely on approximating the yield on the 9 year, 11-month bond with the yield on the ten-year bond. Thus the correlations in Panel B are essentially correlations between changes in yields.

trade-off between the three and ten-year bond leads the investor to leverage the bonds off one another. In the context of a time-varying real interest rate, Brennan and Xia (2002) and Campbell and Viceira (2001) also find that the investor takes highly levered positions in long-term bonds.

When risk premia are high and positive, the investor takes a leveraged position in the ten-year bond, financed by a short position in the three-year bond and the riskfree asset. In this case, hedging demand makes the myopic allocation more extreme. Because the investor has a long position in the ten-year bond, decreases in the risk premium on the ten-year bond reflect deteriorations in the investment opportunity set. The investor hedges these changes in risk premia by allocating more to the ten-year bond. Because the investor has a short position in the three-year bond, increases in the risk premium reflect deteriorations in the investment opportunity set. Thus the investor allocates less to the three-year bond.

When risk premia are positive but closer to zero, the optimal allocation changes. Now the risk-return trade-offs are such that the myopic portfolio consists of a positive fraction of wealth in the three-year bond and a negative fraction in the ten-year bond. Hedging demands also reverse in sign. For short horizons, hedging demand is positive for the three-year bond and negative for the ten-year bond. At long horizons, however, hedging demand is positive for both the ten and the three year bonds. As we show below, the hedging demand is non-monotonic in horizon because the investor also hedges time-variation in the real interest rate.

Finally, when risk premia are negative, the investor holds a positive position in the three-year bond and a negative position in the ten-year bond. Hedging demands cause these positions to become more extreme. Investment opportunities deteriorate when the risk premium on the ten-year bond rises or the risk premium on the three-year bond falls. The investor chooses the optimal portfolio so that wealth is higher when this occurs. Note that the optimal allocation levels off and slightly decreases in magnitude as a function of horizon.

As discussed above, the non-monotonicity in hedging demands occur because the investor hedges both risk premia and the real interest rate. Figure 9 separates out these effects. The left panel shows the allocation when the investor hedges risk premia but not the real riskfree rate. At short horizons, the allocation is similar to the optimal allocation shown in Figure 8. Note however that the allocation is now monotonic in horizon. Time variation in risk premia unambiguously cause the optimal allocation to be more extreme than the myopic allocation.

The right panel in Figure 9 shows the allocation when the investor hedges only time-variation in the real riskfree rate, but not time-variation in risk premia. This allocation was discussed in more detail in Section 5.1 in the context of allocation to a single long-term bond. When the investor hedges only the real riskfree rate, hedging demands are much smaller in magnitude than when the investor hedges time-variation in risk premia. Hedging demand does not depend on the value of the state variables, and is positive for both the three and the five-year bond. Moreover, hedging demand increases monotonically with horizon. The optimal allocation, shown in Figure 8, clearly results from the interaction between time-variation in the riskfree rate and time-variation in risk premia. It does not arise from simply adding one to the other.

Figure 10 examines the case where the investor has access to three long-term bonds. Because the nominal market is complete in this last case, it does not matter for the investor's utility or wealth which three bonds are chosen. Thus without loss of generality, we assume that the investor has access to a one, five, and ten-year bond, as well as the nominally riskless asset. The caveat stated above for the case where the investor has access to two bonds applies to an even greater extent in this case. Because the three bonds are so highly correlated, the investor can achieve (perceived) high Sharpe ratios while taking on less risk than when he had access to fewer bonds. This leads to a highly leveraged myopic portfolio.

The results in this case have much in common with the result from the two-bond case. In general, hedging demand causes the optimal portfolio to be more extreme than the myopic portfolio. When risk premia are positive, the myopic allocation consists of a positive position in the ten-year bond and a negative position in the five-year bond. Hedging demand increases in horizon for the ten-year bond and decreases for the five-year bond. When risk premia are negative, the myopic allocation for the ten-year bond is below that for the five-year bond.²⁰ Hedging demand takes the opposite sign as when risk premia are positive: it is negative for the ten-year bond and positive for the five-year bond. As in the case of two long-term bonds, hedging demand is non-monotonic at long horizons. This is again the result of time-variation in the real riskfree rate. When this effect is taken away, as in the two-bond case, the optimal allocation flattens as the horizon lengthens. In one sense, the three-bond case is more complicated. The investor always takes a long

²⁰For the parameter values we consider, they are both negative. When risk premia become more negative, however, the allocation to the ten-year bond is negative and the five-year bond is positive

position in the one-year bond, regardless of the risk premium. Hedging demand for the one-year bond is generally opposite in sign to that of the five-year bond; it does not depend on whether the investor is long or short the one-year bond. The one-year bond may help the investor isolate the component of the ten and five-year bonds that are correlated with future expected returns, rather than hedge time-variation in the one-year bond itself.

In this section, we have assumed that equities are unavailable. We have repeated the analysis assuming the investor can also hold stock, with the parameters estimated in Table 1. In the two and three-bond cases, stock has a negligible effect on the optimal portfolio, because the opportunities in stock are small compared to those available from going long and short bonds. In the one-bond case, the availability of stock causes the optimal allocation to the bond to fall when the risk premium on the bond is positive and to rise when it is negative. However, the mean allocation to the long-term bond is still positive, and the results we report are qualitatively unchanged.

5.3 Utility costs of sub-optimal strategies

In order to assess the economic importance of the failure of the expectations hypothesis, we calculate utility costs under strategies that fail to take it into account. Two sub-optimal strategies are considered. In the first, the investor chooses the myopic (mean-variance efficient) portfolio. This is the strategy the investor would choose if he were solving a static problem; thus the utility cost of following this strategy represents the cost of treating the multi-period problem as static. It would also be possible to consider an even more sub-optimal strategy: one where the investor fails to recognize that the optimal myopic portfolio varies over time. It turns out that the indirect utility of following this strategy approaches negative infinity for many parameter values. Campbell and Viceira (1999) also find that in the case of predictability of stock returns, the unconditional myopic strategy results in infinite utility costs. Because we find high utility costs to following even the conditional myopic strategy, and because the costs of the unconditional myopic strategy are higher still, we do not report the costs associated with the unconditional myopic strategy.

The second strategy we evaluate is closer to the optimum than the myopic strategy. For this strategy, the investor chooses the myopic portfolio correctly, but hedges only time-variation in the real interest rate, not time-variation in risk premia on

bonds. The optimal portfolio rule when the investor follows this strategy was discussed in Section 5.1. For both the first strategy and the second strategy, the optimal portfolio rule takes the form

$$\hat{w}(t) = \alpha_0 + \alpha_1 X(t) \quad (39)$$

For the first strategy,

$$\begin{aligned} \alpha_0 &= \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}\sigma\lambda_1 + \frac{\gamma-1}{\gamma}(\sigma\sigma^\top)^{-1}(\sigma\sigma_\Pi^\top) \\ \alpha_1 &= \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}\sigma\lambda_2 \end{aligned}$$

For the second strategy,

$$\begin{aligned} \alpha_0 &= \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}\sigma\lambda_1 + \frac{\gamma-1}{\gamma}(\sigma\sigma^\top)^{-1}(\sigma\sigma_\Pi^\top) + \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}(\sigma\sigma_X^\top)B_2^*(\tau)^\top \\ \alpha_1 &= \frac{1}{\gamma}(\sigma\sigma^\top)^{-1}\sigma\lambda_2 \end{aligned}$$

where $B_2^*(\tau)$ is given by (53) in the case of complete nominal markets and (57) in the case of incomplete incomplete markets, with λ_2^* set equal to zero. Note that $B_3(\tau) = 0$ if $\lambda_2^* = 0$.

To calculate utility costs, we solve for indirect utility (31) when the investor follows a strategy of the form (39). Because indirect utility is an expectation of future direct utility it is a martingale and thus has zero drift. From the Markov property it is a function of wealth, the price level, $X(t)$, and the horizon. Thus indirect utility corresponding to the strategy $\hat{w}(t)$ must satisfy the partial differential equation:

$$J_t + \mathcal{L}J = 0 \quad (40)$$

where \mathcal{L} is the infinitesimal generator of J given by

$$\begin{aligned} \mathcal{L}J &= J_W W(\hat{w}^\top(\mu - r\iota) + r) + J_X \mu_X + J_\Pi \Pi \pi + \\ &\quad J_{WX} W \hat{w}(t)^\top \sigma \sigma_X^\top + J_{W\Pi} W \Pi \sigma_W \sigma_\Pi^\top + J_{X\Pi} \Pi \sigma_X \sigma_\Pi^\top \\ &\quad + \frac{1}{2} J_{WW} W^2 \hat{w}^\top \sigma \sigma^\top \hat{w} + \frac{1}{2} J_{\Pi\Pi} \Pi^2 \sigma_\Pi \sigma_\Pi^\top + \frac{1}{2} \text{tr}(J_{XX} \sigma_X \sigma_X^\top) \end{aligned} \quad (41)$$

For the cases where the allocation is linear in $X(t)$, the solution of (40) takes the same form as indirect utility when an investor follows an optimal strategy. Namely:

$$\hat{J}(W(t), \Pi(t), X(t), t, T) = \frac{1}{1-\gamma} \left(\frac{W(t)}{\Pi(t)} \right)^{1-\gamma} \hat{H}(X(t), t, T),$$

where $\hat{H}(X(t), t, T)$ is exponential quadratic. The coefficients solve ordinary differential equations given in Appendix D.

We are interested in the amount by which we would have to increase the wealth of an investor following a sub-optimal strategy so that he has expected utility equal to that of an investor who follows an optimal strategy. Let

$$\hat{J}(W(0), \Pi(0), X(0), 0, T) = \frac{1}{1-\gamma} \left(\frac{W(0)}{\Pi(0)} \right)^{1-\gamma} \hat{H}(X(0), 0, T).$$

equal the indirect utility from following a sub-optimal strategy. The certainty-equivalent gain from following the optimal strategy equals

$$\text{CER}(X(0), 0, T) = \left(\frac{H(X(0), 0, T)}{\hat{H}(X(0), 0, T)} \right)^{\frac{1}{1-\gamma}}.$$

When $\gamma > 1$, $H < \hat{H}$, implying that the investor requires more wealth to be as well off following the sub-optimal strategy.

Figure 11 plots the certainty-equivalent gain from following the optimal strategy when the investor allocates wealth between the riskfree asset and a five-year bond for various levels of risk aversion. Lines without circles represent the certainty equivalent gain relative to the myopic strategy; lines with circles represent the certainty equivalent gain relative to the strategy where the investor hedges only the real risk-free rate. For a given value of risk aversion γ , the line with circles lies below the line without, because the strategy of hedging only the interest rate is less sub-optimal than the myopic strategy.

Even for the one-bond case, the gains from following the optimal strategy are economically large. Relative to the myopic allocation, the gains for an investor with a horizon of twenty years are 8% for a risk aversion of 4, 13% for risk aversion of 10, and 24% of wealth for a risk aversion of 25. Relative to the allocation where the investor hedges only the riskfree rate, the gains are smaller but still significant. For risk aversions of 4 and 10, the gains are 5% of wealth. For a risk aversion of 25, the gain is 7% of wealth.

Figures 12 and 13 show the gain from following the optimal strategy when the investor has access to two and three long-term bonds respectively. In these cases, the gains from hedging time-variation in risk premia are larger than when the investor has access to only one long-term bond. For example, an investor with relative risk aversion of 10 and a horizon of twenty years who has access to two long-term bonds would require a 30% increase in wealth to be as well off following the myopic strategy,

and a 20% increase in wealth to be as well off following the strategy where she hedges only the real interest rate. The gains for an investor who has access to three bonds are even larger.

This section has shown that following a strategy that is optimal for a one-period investor carries high utility costs if the true problem is multi-period. The multi-period problem differs from the single-period problem for two reasons. First, the optimal portfolio hedges time-variation in the real interest rate $r - \pi$, second, the optimal portfolio hedges time-variation in risk premia. We find that both are important in the sense that the costs associated with only hedging time-variation in the real interest rate are very high. This effect does not rely on the investor taking large offsetting positions in bonds of different maturities; it is present even when the investor allocates wealth between the nominally riskfree asset and the long-term bond. Thus the failure of the expectations hypothesis is important for long-term investors; treating risk premia as if they will be constant over the life of the investment results in economically significant costs.

6 Conclusion

We have shown that the failure of the expectations hypothesis has potentially important consequences for the portfolios of long-term investors. For an investor who allocates wealth between a long and a short-term bond, time-variation in risk premia induces hedging demand that is large and positive. We find that long horizon investors should hold a greater fraction of their portfolio in the long-term bond; an effect that persists beyond a horizon of twenty years. When the investor has access to multiple long-term bonds, hedging demands make the optimal allocation more extreme. We find that failing to hedge time-variation in return predictability carries large certainty equivalent costs for the long-term investor.

We establish these results by extending the affine term structure literature to account for expected inflation. Jointly estimating a process for inflation and bond prices produces a series for expected inflation that can account for a large portion of the variance of realized inflation, even though it is constructed from bond yields alone. Including inflation in the estimation actually allows the term structure model to be estimated more accurately.

Our framework is rich enough to include time-variation in the real interest rate, in risk premia on stock returns, and in expected inflation, but at the same time

admits explicit solutions in near-to-closed form. Multiple extensions of our model are possible. For example, we have assumed for simplicity that the equity premium is constant. We could easily extend our results to the case where stock returns are predictable by the yield spread, as well as the dividend-yield. Ait-Sahalia and Brandt (2001) show that investors are unable to hedge changes in the yield spread using stocks alone, due to the low contemporaneous correlation between stocks and bond yields. Our results suggest that bonds could play an important role in hedging changes in risk premia on stocks. We could also modify our model to allow for parameter uncertainty, as in Barberis (2000), or learning, as in Xia (2002). Clearly there are important aspects of the portfolio choice problem that we do not address. Transaction costs, parameter uncertainty, and non-expected utility preferences have all been fruitfully explored in the context of stock-return predictability. Bonds present a similar, yet richer framework to explore these same issues.

Appendix

A Bond Prices

Following Cox, Ingersoll, and Ross (1985), we assume that bond prices are smooth functions of the state variables $X(t)$ and of time. That is, $P(X(t), t, T) \in C^{2,1}(R^M \times [0, \infty])$. No-arbitrage implies that P satisfies

$$P_X K(\theta - X(t)) + \frac{1}{2} \text{tr} \left(P_{XX} \sigma_X \sigma_X^\top \right) + P_t - r(t)P = P_X \sigma_X \Lambda(t) \quad (42)$$

with boundary condition $P(X(t), t, t) = 0$. Equation (42) follows from equating the instantaneous expected excess return to the volatility multiplied by the price of risk.

Conjecture that

$$P(X(t), t, T) = \exp \{ A_2(\tau) X(t) + A_1(\tau) \}, \quad (43)$$

where $\tau = T - t$. Substituting back into (42) and matching coefficients on $X(t)$ and the constants, produces the following system of ordinary differential equations for the row vector $A_2(\tau)$ and the scalar $A_1(\tau)$:

$$A_2'(\tau) = -A_2(\tau) (K + \sigma_X \lambda_2) - \delta \quad (44)$$

$$A_1'(\tau) = A_2(\tau) (K\theta - \sigma_X \lambda_1) + \frac{1}{2} A_2(\tau) \sigma_X \sigma_X^\top A_2(\tau)^\top - \delta_0 \quad (45)$$

The boundary conditions are $A_2(0) = 0_{1 \times m}$ and $A_1(0) = 0$.

B Optimal portfolio allocation

Proof of Theorem 1:

It follows from the Markov property of (Π, Z, X) that wealth may be written as

$$\begin{aligned} G(\Pi(t), Z(t), X(t), t, T) &= W(t) \\ &= \Pi(t) Z(t)^{\frac{1}{\gamma}} F(X(t), t, T) \end{aligned}$$

Because wealth is an asset, it satisfies a no-arbitrage differential equation analogous to that of bonds. Applying Ito's lemma to G and matching the instantaneous expected excess return on wealth to its volatility times the price of risk produces:²¹

$$\mathcal{L}G + G_t - rG = \left(G_Z Z((\Lambda^* + \nu)^\top - \sigma_\Pi) + G_\Pi \Pi \sigma_\Pi + G_X \sigma_X \right) (\Lambda^* + \nu), \quad (46)$$

²¹From Ito's lemma we can write

$$dZ(t) = \mu_Z dt + \sigma_Z dz$$

$$\begin{aligned} \mathcal{L}G = & ZG_Z\mu_Z + G_\Pi\Pi\pi + G_XK(\theta - X) + ZG_{ZX}\sigma_X\sigma_Z^\top + \Pi G_{\Pi X}\sigma_X\sigma_\Pi^\top + \\ & \frac{1}{2} \left(Z^2G_{ZZ}\sigma_Z\sigma_Z^\top + \Pi^2G_{\Pi\Pi}\sigma_\Pi\sigma_\Pi^\top + \text{tr} \left(G_{XX}\sigma_X\sigma_X^\top \right) \right), \end{aligned}$$

with boundary condition

$$G(\Pi(T), Z(T), X(T), T, T) = \Pi(T)Z(T)^{\frac{1}{\gamma}}.$$

Note that the no-arbitrage relationship for G only holds for the min-max pricing kernel ϕ_ν , while, by the bond pricing equation holds for any pricing kernel. Substituting (27) into (46) results in the partial differential equation for F given in the text.

In order that optimal wealth satisfy the dynamic budget constraint (21), the diffusion terms from the two processes must match. Therefore the price of risk and the optimal portfolio must jointly satisfy:

$$\frac{1}{\gamma}(\Lambda^* + \nu)^\top + \frac{\gamma - 1}{\gamma}\sigma_\Pi + \frac{F_X}{F}\sigma_X = \alpha^\top\sigma, \quad (47)$$

where α is the $N \times 1$ vector of portfolio weights. The left-hand side follows from Ito's lemma applied to G . Inflation risk σ_Π is not spanned by the row vectors of σ , thus for general ν , this equation will not have a solution.

We need to find ν so that the unhedgeable part of σ_Π drops out.²² Rewrite σ_Π as

$$\sigma_\Pi = (\sigma_\Pi\sigma^\top)(\sigma\sigma^\top)^{-1}\sigma + \left(\sigma_\Pi - (\sigma_\Pi\sigma^\top)(\sigma\sigma^\top)^{-1}\sigma \right). \quad (48)$$

The first term is the projection of σ_Π onto the traded assets. The second term is orthogonal to the traded assets. In order for (55) to have a solution, ν must satisfy

$$\frac{1}{\gamma}\nu^\top = \frac{1 - \gamma}{\gamma} \left(\sigma_\Pi - (\sigma_\Pi\sigma^\top)(\sigma\sigma^\top)^{-1}\sigma \right)$$

Therefore,

$$\nu = (1 - \gamma) \left(\sigma_\Pi^\top - \sigma^\top(\sigma\sigma^\top)^{-1}\sigma\sigma_\Pi^\top \right)^\top. \quad (49)$$

with

$$\begin{aligned} \mu_Z &= \left(r(t) - \pi(t) + (\Lambda^* + \nu)^\top(\Lambda^* + \nu) + \sigma_\Pi\sigma_\Pi^\top + \sigma_\Pi\Lambda \right) Z(t) \\ \sigma_Z &= ((\Lambda^* + \nu)^\top - \sigma_\Pi)Z(t) \end{aligned}$$

²² ν does not have to cancel out the unhedgeable parts of Λ^* , because the columns of Λ^* are spanned by the rows of σ . In fact, this is the reason for defining Λ^* as a projection of Λ onto the available assets.

Because ν is orthogonal to the basis assets, $\Lambda^* + \nu$, where Λ^* is given by (17), is indeed a valid price of risk.

Substituting (49) back into (55) produces²³

$$\frac{1}{\gamma}(\mu - \iota r)^\top (\sigma\sigma^\top)^{-1}\sigma + \frac{\gamma - 1}{\gamma}(\sigma_\Pi\sigma^\top)(\sigma\sigma^\top)^{-1}\sigma + \frac{1}{F}F_X(\sigma_X\sigma^\top)(\sigma\sigma^\top)^{-1}\sigma = \alpha^\top\sigma.$$

The equation for the optimal allocation (30) follows from multiplying both sides of the equation by $\sigma^\top(\sigma\sigma^\top)^{-1}$ and taking the transpose. This completes the proof of Theorem 1. \square

Proof of Corollary 2:

The argument follows that of Cox and Huang (1989), generalized to the case of unexpected inflation. The investor's problem at time 0 can equivalently be written as

$$\max_{W(t)>0} E_0 [J(W(t), \Pi(t), X(t), t, T)]$$

subject to the static budget constraint. The first order condition is given by

$$J_W(t) = l\phi_{\hat{\nu}}(t)^{-1}$$

where $\phi_{\hat{\nu}}(t)$ is the min-max pricing kernel. We do not know *a priori* that $\phi_{\hat{\nu}} = \phi_\nu$. As is well-known, the solution to (31) takes the form:

$$J(W(t), \Pi(t), X(t), t, T) = \frac{1}{1 - \gamma} \left(\frac{W(t)}{\Pi(t)} \right)^{1-\gamma} H(X(t), t, T). \quad (50)$$

Our goal is to prove the relationship between the functions H and F .

Define \hat{Z} analogously to (25) as:

$$\hat{Z}(t) = (l\phi_{\hat{\nu}}(t)\Pi(t))^{-1}.$$

Then the investor's first-order condition can be re-written as

$$J_W(t) = \hat{Z}(t)^{-1}\Pi(t)^{-1}$$

²³Because σ_X is spanned by the rows of σ ,

$$\sigma_X = (\sigma_X\sigma^\top)(\sigma\sigma^\top)^{-1}\sigma$$

This is the only place in the argument where we need that the rows of σ span σ_X . If we drop this assumption, then we have two layers of market incompleteness. We could use similar arguments, except that η would be endogenous.

Substituting in from (50) implies that

$$W(t) = \hat{Z}(t)^{\frac{1}{\gamma}} \Pi(t) H(X(t), t, T)^{\frac{1}{\gamma}}. \quad (51)$$

Because $W(t)$ is an asset, it must satisfy partial differential equation (46). Comparing (51) with (27), it follows that $H^{\frac{1}{\gamma}}$ and $\hat{\nu}$ must jointly satisfy the partial differential equations (29) and (55). Therefore, $\hat{\nu}$ must equal ν and $H^{\frac{1}{\gamma}}$ must equal F . \square

Proof of Corollary 4:

To solve for F , we conjecture the form of it and then we verify. Our conjecture is that

$$F(X(t), t, T) = \exp \left\{ \frac{1}{\gamma} \left(\frac{1}{2} X_t^\top B_3(\tau) X_t + B_2(\tau) X_t + B_1(\tau) \right) \right\}$$

where $\tau = T - t$, $B_1(\tau)$ is a matrix, $B_2(\tau)$ is a row vector, and $B_3(\tau)$ is a scalar. Plugging the hypothesized solution back into the PDE (29) and matching coefficients on $X_t^\top [\cdot] X_t$, X_t , and the constant terms, leads to a system of ordinary differential equations:

$$\begin{aligned} B_3'(\tau) &= (B_3(\tau) + B_3(\tau)^\top) \left[\left(\frac{1}{\gamma} - 1 \right) \sigma_X \lambda_2^* - K \right] \\ &+ \left(\frac{1}{4\gamma} (B_3(\tau) + B_3(\tau)^\top) \sigma_X \sigma_X^\top (B_3(\tau) + B_3(\tau)^\top) + \left(\frac{1}{\gamma} - 1 \right) \lambda_2^{*\top} \lambda_2^* \right) \end{aligned} \quad (52)$$

$$\begin{aligned} B_2'(\tau) &= B_2(\tau) \left[\left(\frac{1}{\gamma} - 1 \right) \sigma_X \lambda_2^* - K + \frac{1}{2\gamma} \sigma_X \sigma_X^\top (B_3(\tau) + B_3(\tau)^\top) \right] + \\ &\frac{1}{2} \left[\theta^\top K^\top + \left(\frac{1}{\gamma} - 1 \right) \lambda_1^{*\top} \sigma_X^\top + \left(1 - \frac{1}{\gamma} \right) \sigma_\Pi \sigma_X^\top \right] (B_3(\tau) + B_3(\tau)^\top) \\ &+ (1 - \gamma)(\delta - \zeta) + \left(\frac{1}{\gamma} - 1 \right) \lambda_1^{*\top} \lambda_2^* + (\gamma - 1) \sigma_\Pi \lambda_2^* \end{aligned} \quad (53)$$

$$\begin{aligned} B_1'(\tau) &= B_2(\tau) \left[K\theta + \left(\frac{1}{\gamma} - 1 \right) \sigma_X \lambda_1^* + \left(1 - \frac{1}{\gamma} \right) \sigma_X \sigma_\Pi^\top \right] \\ &+ \frac{1}{2\gamma} B_2(\tau) \sigma_X \sigma_X^\top B_2(\tau)^\top + \frac{1}{4} \text{tr} \left((B_3(\tau) + B_3(\tau)^\top) \sigma_X \sigma_X^\top \right) \\ &+ \frac{1}{2} \left(\frac{1}{\gamma} - 1 \right) (\lambda_1^{*\top} \lambda_1^* + \nu^{*\top} \nu^*) + \frac{\gamma}{2} \sigma_\Pi \sigma_\Pi^\top \\ &+ (1 - \gamma) \sigma_\Pi \lambda_1^* + (1 - \gamma)(\delta_0 - \zeta_0) \end{aligned} \quad (54)$$

C Optimal portfolio allocation under incomplete nominal markets

This Appendix modifies the results above to the case where the investor has fewer bonds than state variables. In this case, nominal markets are incomplete. To determine the minmax price of risk in this case we start from the equation:

$$\frac{1}{\gamma}(\Lambda^* + \nu)^\top + \frac{\gamma - 1}{\gamma}\sigma_\Pi + \frac{F_X}{F}\sigma_X = w^\top \sigma, \quad (55)$$

We then project σ_Π and σ_X on the available assets:

$$\begin{aligned} \sigma_\Pi &= (\sigma_\Pi \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma + (\sigma_\Pi - (\sigma_\Pi \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma) \\ \sigma_X &= (\sigma_X \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma + (\sigma_X - (\sigma_X \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma). \end{aligned}$$

It is useful to define the residual of the projections as

$$\begin{aligned} (\sigma_\Pi^\perp) &= \sigma_\Pi - (\sigma_\Pi \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma \\ (\sigma_X^\perp) &= \sigma_X - (\sigma_X \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma \end{aligned}$$

Following the same reasoning as before we find that ν takes the form

$$\nu = (1 - \gamma) \left(\sigma_\Pi - (\sigma_\Pi \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma \right)^\top - \gamma \left(\sigma_X - (\sigma_X \sigma^\top)(\sigma \sigma^\top)^{-1} \sigma \right)^\top \frac{F_X^\top}{F}.$$

Substituting into the PDE for F in Theorem 1 we find the following ODE's:

$$\begin{aligned} B'_3(\tau) &= \{ \quad \} + \left(\frac{1}{\gamma} - 1 \right) \left(\frac{1}{4} (B_3(\tau) + B_3(\tau)^\top)(\sigma_X^\perp)(\sigma_X^\perp)^\top (B_3(\tau) + B_3(\tau)^\top) \right) \\ &\quad - \left(\frac{1}{\gamma} - 1 \right) \left(\frac{1}{2} (B_3(\tau) + B_3(\tau)^\top) \sigma_X (\sigma_X^\perp)^\top (B_3(\tau) + B_3(\tau)^\top) \right) \quad (56) \end{aligned}$$

$$\begin{aligned} B'_2(\tau) &= \{ \quad \} + \\ &\left(\frac{1}{\gamma} - 1 \right) \left(\frac{1 - \gamma}{2} (\sigma_\Pi^\perp) \sigma_X^\top (B_3(\tau) + B_3(\tau)^\top) - B_2(\tau) \sigma_X (\sigma_X^\perp)^\top (B_3(\tau) + B_3(\tau)^\top) \right) \\ &\quad + \left(\frac{1}{\gamma} - 1 \right) \left(\frac{1}{2} (B_2(\tau) (\sigma_X^\perp) - (1 - \gamma) (\sigma_\Pi^\perp)) (\sigma_X^\perp)^\top (B_3(\tau) + B_3(\tau)^\top) \right) \\ &\quad + \frac{\gamma - 1}{2} \sigma_\Pi (\sigma_X^\perp) (B_3(\tau) + B_3(\tau)^\top) \quad (57) \end{aligned}$$

$$\begin{aligned}
B_1'(\tau) = \{ \quad \} + \left(\frac{1}{\gamma} - 1\right) B_2(\tau) \sigma_X \left((1 - \gamma)(\sigma_{\Pi^\perp}) - (\sigma_X^\perp) B_2(\tau) \right)^\top + \\
\frac{1}{2} \frac{1 - \gamma}{\gamma} B_2(\tau) (\sigma_X^\perp) (\sigma_X^\perp)^\top B_2(\tau)^\top - \left(\frac{1}{\gamma} - 1\right) (1 - \gamma) B_2(\tau) (\sigma_X^\perp) (\sigma_{\Pi^\perp})^\top + \\
(\gamma - 1) \sigma_{\Pi} (\sigma_X^\perp)^\top B_2(\tau)^\top \quad (58)
\end{aligned}$$

The terms $\{ \quad \}$ represents the quantity on the right hand side of equations (52), (53), and (54) respectively.

Note that when markets are complete, the new terms on the right hand side of (56), (57), and (58) reduce to zero. In particular, $(\sigma_{\Pi^\perp}) \sigma_X^\top = 0$ because σ_X^\top is now within the span of σ .

D Indirect utility for sub-optimal strategies.

It follows from the partial differential equation (40) that indirect utility takes the form:

$$J(W(t), \Pi(t), X(t), T) = \frac{1}{1 - \gamma} \left(\frac{W(t)}{\Pi(t)} \right)^{1 - \gamma} H(X(t), t, T).$$

where $H(X(t), t, T)$ satisfies the partial differential equation

$$\begin{aligned}
H_t + (1 - \gamma) H \left(w(t)^\top (\mu(t) - \iota r(t)) + r(t) - \pi(t) \right) \\
- \gamma H w(t)^\top \sigma \sigma^\top w(t) - (1 - \gamma) H w(t)^\top \sigma \sigma_{\Pi}^\top - \frac{\gamma - 2}{2} H \sigma_{\Pi} \sigma_{\Pi}^\top \\
+ H_X \left(K(\theta - X(t)) + (1 - \gamma) \sigma_X \sigma^\top w(t) - (1 - \gamma) \sigma_X \sigma_{\Pi}^\top \right) + \frac{1}{2} \text{tr}(H_{XX} \sigma_X \sigma_X^\top) = 0. \quad (59)
\end{aligned}$$

Of interest is indirect utility when the investor follows a myopic strategy, and a strategy that optimally hedges time-variation in the real interest rate, but not time-variation in risk premia. Both strategies can be expressed as

$$w(t) = \alpha_0 + \alpha_1 X(t), \quad (60)$$

for some constant scalar α_0 and vector α_1 . When the trading strategy can be expressed as (60), it follows from (59) that $H(X(t), t, T)$ is exponential quadratic:

$$H(X(t), t, T) = \exp \left\{ X(t)^\top \Gamma_3 X(t) + \Gamma_2 X(t) + \Gamma_1 \right\}.$$

where Γ_3 , Γ_2 , and Γ_1 satisfy the following system of ordinary differential equations:

$$\begin{aligned}
\Gamma_3' = (\Gamma_3 + \Gamma_3^\top) \left[(1 - \gamma) \sigma_X \sigma^\top \alpha_1 - K \right] + \frac{\Gamma_3 + \Gamma_3^\top}{2} \sigma_X \sigma_X^\top \frac{\Gamma_3 + \Gamma_3^\top}{2} \\
+ 2(1 - \gamma) \alpha_1^\top \sigma \lambda_2 - \gamma(1 - \gamma) \alpha_1^\top \sigma \sigma^\top \alpha_1 \quad (61)
\end{aligned}$$

$$\begin{aligned}
\Gamma'_2 = & \Gamma_2 \left[(1-\gamma)\sigma_X\sigma^\top\alpha_1 + \sigma_X\sigma_X^\top \frac{\Gamma_3 + \Gamma_3^\top}{2} - K \right] \\
& + \left[\theta^\top K^\top + (1-\gamma)\alpha_0^\top\sigma\sigma_X^\top - (1-\gamma)\sigma_\Pi\sigma_X^\top \right] \frac{\Gamma_3 + \Gamma_3^\top}{2} \\
& + (1-\gamma) \left[\alpha_0^\top\sigma\lambda_2 + \delta - \zeta + \lambda_1^\top\sigma^\top\alpha_1 \right] - (1-\gamma)^2\sigma_\Pi\sigma^\top\alpha_1 - \gamma(1-\gamma)\alpha_0^\top\sigma\sigma^\top\alpha_0
\end{aligned} \tag{62}$$

$$\begin{aligned}
\Gamma'_1 = & \Gamma_2 \left[K\theta + (1-\gamma)\sigma_X\sigma^\top\alpha_0 - (1-\gamma)\sigma_X\sigma_\Pi^\top \right] + \frac{1}{2}\Gamma_2\sigma_X\sigma_X^\top\Gamma_2^\top \\
& + (1-\gamma)(\alpha_0^\top\sigma\lambda_1 + \delta_0 - \zeta_0) - (1-\gamma)^2\sigma_\Pi\sigma^\top\alpha_0 - \frac{\gamma(1-\gamma)}{2}\alpha_0^\top\sigma\sigma^\top\alpha_0 \\
& - \frac{(1-\gamma)(\gamma-2)}{2}\sigma_\Pi\sigma_\Pi^\top + \frac{1}{2}\text{tr} \left(\frac{\Gamma_3 + \Gamma_3^\top}{2}\sigma_X\sigma_X^\top \right)
\end{aligned} \tag{63}$$

E Estimation

This section extends the results of Duffee (2002) to include inflation and stock return data in the estimation of bond yields. For convenience, it is assumed that the state variables are Gaussian (as in the body of the paper). Duffee's quasi-maximum likelihood results for square-root models can be extended in a similar fashion. In what follows, let e^Q denote the matrix exponential of Q , let $(x_i)_i$ denote the diagonal matrix with diagonal elements x_i , and let $(x_{i,j})_{i,j}$ denote the matrix with the (i, j) element equal to $x_{i,j}$. It is assumed that K is diagonalizable.

Let $Y(t)$ denote the vector of perfectly observed yields at time t . Namely

$$Y(t) = \begin{pmatrix} y(X(t), t, \tau_1) \\ \vdots \\ y(X(t), t, \tau_m) \end{pmatrix}$$

for maturities (τ_1, \dots, τ_m) , where y is defined in (6). Let \tilde{Y} denote the vector of yields which are observed imperfectly. From (5), it follows that the perfectly observed yields can be inverted to find the state variables:

$$X(t) = L_1^{-1}(Y(t) - L_0).$$

where L_1 is an $m \times m$ matrix with row i given by $-A_2(\tau_i)/\tau_i$, and L_0 is a vector with elements $-A_1(\tau_i)/\tau_i$. Let $f(\cdot | \cdot)$ denote (with slight abuse of notation), the

conditional likelihood function. Then the likelihood function for yields can be related to the likelihood function for the state variables by

$$f(Y(t+1), \Pi(t+1), S(t+1) | Y(t), \Pi(t), S(t)) = \frac{1}{\det [L_1]} f(X(t+1), \Pi(t+1), S(t+1) | X(t), \Pi(t), S(t)). \quad (64)$$

Let $\epsilon(t)$ denote the observation errors on the yields that are not perfectly observed. We assume that $\epsilon(t)$ is independent of innovations to the state variables or to inflation. Under this assumption, the full likelihood is given by:

$$\begin{aligned} l_t(\Theta) &= \log f(Y(t), \Pi(t), S(t) | Y(t-1), \Pi(t-1), S(t-1)) + \\ &\quad \log f(\tilde{Y}(t) | Y(t), \Pi(t), S(t)) \\ &= \log f(Y(t), \Pi(t), S(t) | Y(t-1), \Pi(t-1), S(t-1)) + \log f(\epsilon(t) | Y(t)) \end{aligned}$$

It therefore suffices to specify $f(X(t+1), \Pi(t+1) | X(t), \Pi(t))$.

We show that $f(\log \Pi(t+1), X(t+1), \log S(t+1) | \log \Pi(t), X(t), \log S(t))$ is multivariate normal, and calculate the mean and variance. Consider the augmented state vector

$$\hat{X}(t) = \begin{bmatrix} \log \Pi(t) \\ X(t) \\ \log S(t) \end{bmatrix}$$

Then we can write the continuous time dynamics of this vector as

$$d\hat{X}(t) = (\kappa_1 \hat{X} + \kappa_2)dt + \sigma_{\hat{X}} dz, \quad (65)$$

where

$$\kappa_1 = \begin{bmatrix} 0 & \zeta & 0 \\ \mathbf{0} & -K & \mathbf{0} \\ 0 & \eta + \delta & 0 \end{bmatrix}, \quad \kappa_2 = \begin{bmatrix} \zeta_0 - \frac{1}{2}\sigma_{\Pi}\sigma_{\Pi}^{\top} \\ K\theta \\ \eta_0 + \delta_0 - \frac{1}{2}\sigma_S\sigma_S^{\top} \end{bmatrix}, \quad \sigma_{\hat{X}} = \begin{bmatrix} \sigma_{\Pi} \\ \sigma_X \\ \sigma_S \end{bmatrix}$$

Applying Ito's lemma to the process $e^{-\kappa_1 t} \hat{X}_t$, it follows that:

$$\hat{X}(T) = e^{\kappa_1(T-t)} \hat{X}_t + \int_t^T e^{\kappa_1(T-s)} \kappa_2 ds + \int_t^T e^{\kappa_1(T-s)} \sigma_{\hat{X}} dw(s). \quad (66)$$

Which shows that \hat{X}_T is normally distributed conditional on \hat{X}_t .

If K can be diagonalizable, κ_1 can also be diagonalizable, with the first and last eigenvalues equal to 0, and the remaining eigenvalues equal to that of K . Let U be such that

$$\kappa_1 = UDU^{-1}, \quad D \text{ diagonal.}$$

From the definition of the matrix exponential and (66), it follows that

$$E_t(\hat{X}(T)) = e^{\kappa_1(T-t)} \hat{X}(t) + \left(\int_t^T U e^{D(T-s)} U^{-1} ds \right) \kappa_2.$$

Note that $e^{D(T-t)} = (e^{d_i(T-t)})_i$. Performing the integration element-by-element produces:

$$E_t(\hat{X}(T)) = e^{\kappa_1(T-t)} \hat{X}_t + U (f(d_i, T-t))_i U^{-1} \kappa_2.$$

where

$$f(d_i, T-t) = \begin{cases} -\frac{1}{d_i}(1 - e^{d_i(T-t)}) & d_i \neq 0 \\ T-t & d_i = 0 \end{cases}$$

This completes the derivation of the conditional mean.

From (66), the conditional variance satisfies:

$$\begin{aligned} \text{Var}_t(\hat{X}(T)) &= E_t \left[\left(\int_t^T e^{\kappa_1(T-u)} \sigma_{\hat{X}} dw_u \right) \left(\int_t^T e^{\kappa_1(T-u)} \sigma_{\hat{X}} dw_u \right)^\top \right] = \\ &= E_t \left[\int_t^T e^{\kappa_1(T-u)} \sigma_{\hat{X}} \sigma_{\hat{X}}^\top e^{\kappa_1(T-u)\top} du \right] = \\ &= \int_t^T e^{\kappa_1(T-u)} \sigma_{\hat{X}} \sigma_{\hat{X}}^\top \left(e^{\kappa_1(T-u)} \right)^\top du. \end{aligned}$$

Let $\Omega = U^{-1} \sigma_{\hat{X}} \sigma_{\hat{X}}^\top (U^{-1})^\top$. Integrating the above equation element-by-element produces:

$$\begin{aligned} \text{Var}_t(\hat{X}(T)) &= \int_t^T U e^{D(T-u)} \Omega e^{D(T-u)} U^\top du \\ &= U [g(d_i, d_j, T-u) \Omega_{i,j}]_{i,j} U^\top, \end{aligned}$$

where

$$g(d_i, d_j, T-t) = \begin{cases} -\frac{1}{d_i+d_j}(1 - e^{(d_i+d_j)(T-t)}) & d_i \neq 0 \text{ or } d_j \neq 0 \\ T-t & d_i = d_j = 0 \end{cases}$$

This completes the derivation of the conditional variance-covariance matrix.

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.
- Ait-Sahalia, Yacine, and Michael Brandt, 2001, Variable Selection for Portfolio Choice, *Journal of Finance* 56, 1297–1351.
- Balduzzi, Pierluigi, and Anthony W. Lynch, 1999, Transaction Costs and Predictability: Some Utility Cost Calculations, *Journal of Financial Economics* 52, 47–78.
- Bansal, Ravi, and Hao Zhou, 2002, Term Structure of Interest Rates with Regime Shifts, *Journal of Finance* 57, 1997–2044.
- Barberis, Nicholas, 2000, Investing for the Long Run when Returns are Predictable, *Journal of Finance* 55, 225–264.
- Bekaert, Geert, and Robert J. Hodrick, 2001, Expectations hypothesis tests, *Journal of Finance* 56, 1357–1394.
- Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, On biases in tests of the expectations hypothesis of the term structure of interest rates, *Journal of Financial Economics* 44, 309–348.
- Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 2001, Peso problem explanations for term structure anomalies, *Journal of Monetary Economics* 48, 241–270.
- Brandt, Michael W., 1999, Estimating portfolio and consumption choice: A conditional Euler equations approach, *Journal of Finance* 54, 1609–1645.
- Brandt, Michael W., and David A. Chapman, 2002, Comparing Multifactor Models of the Term Structure, Manuscript, University of Texas and the Wharton School.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic Asset Allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Brennan, Michael J., Ashley W. Wang, and Yihong Xia, 2002, A Simple Model of Intertemporal Capital Asset Pricing and its Implications for the Fama-French Three-Factor Model, Working Paper, UCLA and the Wharton School.

- Brennan, Michael J., and Yihong Xia, 2000, Stochastic interest rates and the bond-stock mix, *European Finance Review* 4, 197–210.
- Brennan, Michael J., and Yihong Xia, 2002, Dynamic Asset Allocation under Inflation, *Journal of Finance* 57, 1201–1238.
- Campbell, John Y., Yeung Lewis Chan, and Luis M. Viceira, 2002, A multivariate model of strategic asset allocation, forthcoming, *Journal of Financial Economics*.
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 58, 495–514.
- 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., and Luis Viceira, 2001, Who should buy long-term bonds?, *American Economic Review* 91, 99–127.
- Campbell, John Y., and Luis M. Viceira, 1999, Consumption and Portfolio Decisions when Expected Returns are Time-varying, *Quarterly Journal of Economics* 114, 433–495.
- Cochrane, John, and Monika Piazzesi, 2002, Bond Risk Premia, NBER Working Paper # 9178.
- Cox, John C., and Chi-Fu Huang, 1989, Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process, *Journal of Economic Theory* 39, 33–83.
- Cox, John C., Jonathan C. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–408.
- Dai, Qiang, and Kenneth Singleton, 2000, Specification analysis of affine term structure models, *Journal of Finance* 55, 1943–1978.
- Dai, Qiang, and Kenneth Singleton, 2002a, Expectations puzzles, time-varying risk premia, and affine models of the term structure, *Journal of Financial Economics* 63, 415–442.
- Dai, Qiang, and Kenneth Singleton, 2002b, Term Structure Dynamics in Theory and Reality, forthcoming, *Review of Financial Studies*.

- Duffee, Gregory R., 2002, Term premia and interest rate forecasts in affine models, *Journal of Finance* 57, 369–443.
- Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379–406.
- Dybvig, Philip H., and Chi-Fu Huang, 1988, Nonnegative wealth, absence of arbitrage, and feasible consumption plans, *Review of Financial Studies* 1, 377–401.
- 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* 29, 23–49.
- Fisher, Mark, 1998, A simple model of the failure of the expectations hypothesis, Working paper, Federal Reserve Bank of Atlanta.
- He, Hua, and Neil D. Pearson, 1991, Consumption and Portfolio Policies with Incomplete Markets and Short-sale Constraints: The Infinite Dimensional Case, *Journal of Economic Theory* 54, 259–304.
- Karatzas, Ioannis, John P. Lehoczky, and Steven E. Shreve, 1987, Optimal Portfolio and Consumption Decisions for a Small Investor on a Finite Horizon, *SIAM Journal of Control and Optimization* 25, 1557–1586.
- Kim, Tong Suk, and Edward Omberg, 1996, Dynamic Nonmyopic Portfolio Behavior, *Review of Financial Studies* 9, 141–161.
- Liu, Jun, 1999, Portfolio Selection in Stochastic Environments, unpublished paper, Stanford University.
- Liu, Jun, and Jun Pan, 2002, Dynamic Derivative Strategies, Forthcoming, *Journal of Financial Economics*.
- Merton, Robert C., 1971, Optimum Consumption and Portfolio Rules in a Continuous-time Model, *Journal of Economic Theory* 3, 373–413.
- Piazzesi, Monika, 2002, Affine Term Structure Models, forthcoming *Handbook of Financial Econometrics*.

- Pliska, Stanley R., 1986, A Stochastic Calculus Model of Continuous Trading: Optimal Portfolios, *Mathematics of Operations Research* 11, 239–246.
- Schroder, Mark, and Costis Skiadas, 1999, Optimal Consumption and Portfolio Selection with Stochastic Differential Utility, *Journal of Economic Theory* 89, 68–126.
- Schroder, Mark, and Costis Skiadas, 2002, Optimal Lifetime Consumption-Portfolio Strategies under Trading Constraints and Generalized Recursive Preferences, Kellogg School of Management, Working Paper No. 285.
- Sorensen, Carsten, 1999, Dynamic asset allocation and fixed income management, *Journal of Financial and Quantitative Analysis* 34, 513–532.
- Stambaugh, Robert, 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.
- Vasicek, Oldrich, 1977, An equilibrium characterization of the term structure, *Journal of Financial Economics* 5, 177–188.
- Wachter, Jessica A., 2002a, Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets, *Journal of Financial and Quantitative Analysis* 37, 63–91.
- Wachter, Jessica A., 2002b, Risk aversion and allocation to long-term bonds, forthcoming, *Journal of Economic Theory*.
- Xia, Yihong, 2002, Long-Term Bond Markets and Investor Welfare, Working Paper, The Wharton School.

Table 1: Parameter estimates

Parameter values for the three-factor model described in Section 3 are estimated using monthly data on bond yields, inflation, and stock returns from 1952-1998. The risk premium on the stock index is assumed to be constant. Outer product standard errors are given in parentheses. Parameter values are in natural units.

Inflation							
Parameters	1		2		3		
σ_{Π}	0.0093	(0.0002)					
ζ_0	0.042	(0.026)					
ζ_i	0.018	(0.002)	0.018	(0.004)	0.0074	(0.0006)	
Bond							
Parameters	1		2		3		
δ_0	0.058	(0.034)					
δ_i	0.0182	(0.0003)	0.0074	(0.001)	0.0098	(0.0003)	
$K_{1,i}$	0.576	(0.026)	0		0		
$K_{2,i}$	0		3.308	(0.371)	0		
$K_{3,i}$	-0.375	(0.167)	0		0.076	(0.053)	
$\lambda_{1,i}^*$	-0.553	(0.208)	-0.243	(0.048)	-0.209	(0.048)	
$\lambda_{2(1,i)}^*$	0		1.752	(0.068)	0		
$\lambda_{2(2,i)}^*$	0		-1.790	(0.372)	0		
$\lambda_{2(3,i)}^*$	0.485	(0.172)	0.347	(0.094)	-0.075	(0.053)	
Stock							
Parameters	1		2		3		4
$\sigma_S \Lambda$	0.075	(0.025)					
σ_S	-0.013	(0.006)	0.006	(0.006)	-0.030	(0.006)	0.143 (0.003)

Table 2: Asset Correlations

Conditional correlations of asset prices implies by the parameter values in Table 1. Correlations are constructed using the instantaneous variance-covariance matrix $\sigma\sigma^\top$, where σ is defined as in (10).

Panel A: Model			
1-year Bond	5-year Bond	10-year Bond	Stock
1.000	0.880	0.744	0.194
	1.000	0.950	0.210
		1.000	0.215
			1.000

Panel B: Data			
1-year Bond	5-year Bond	10-year Bond	Stock
1.000	0.853	0.734	0.190
	1.000	0.932	0.192
		1.000	0.214
			1.000




Figure 1: Model-implied mean yield spreads, calculated using the parameters in Table 1. Yields are in annual terms, and defined as in (6). The short-term yield has maturity of 3 months. “Sample” refers to yield spreads calculated using data from 1953-1998 on bonds of selected maturities.




Figure 2: Yield spread standard deviations implied by the model and the parameters in Table 1. Yields are in annual terms, and defined as in (6). The short-term yield has maturity of 3 months. “Sample” refers to yield spreads calculated using data from 1953-1998 on bonds of selected maturities.




Figure 3: Model-implied coefficients on Campbell-Shiller (1991) long-rate regressions. Quarterly changes in yields $y(t, s) - y(t + \frac{1}{4}, s)$ are regressed on the spread between the $(s - t)$ -year bond, and the 3-month bond, scaled by $1/(4(s - t) - 1)$. “Sample” refers to yield spreads calculated using data from 1953-1998 on bonds of selected maturities.




Figure 4: Realized log inflation (from CRSP) and expected inflation implied by the parameters in Table 1.




Figure 5: Nominal interest rates and expected inflation implied by the parameters in Table 1.




Figure 6: Risk premia (in annual percentages) on long-term bonds implied by the parameters in Table 1.




Figure 7: Optimal allocation between the five-year bond and the nominal riskfree asset. Shown on the graph is the optimal allocation to the five-year bond; allocation to the riskfree asset is one minus this quantity. Allocation is plotted as a function of horizon for the investor with utility over terminal wealth. Lines without circles plot the optimal allocation, lines with circles plot the allocation when the investor hedges only the riskfree rate. Lines with plus signs plot the allocation when the investor hedges only the risk premium. X_2 and X_3 are set equal to zero while X_1 is varied by plus and minus two unconditional standard deviations. Risk premia are positive for $X_1 = -1.9$ and 0, and negative for $X_1 = 1.9$.






Figure 8: Optimal allocation between the ten-year and three-year bond and the nominal riskfree asset. Allocation to the three and ten-year bonds is plotted as a function of horizon for the investor with utility over terminal wealth. X_2 and X_3 are set equal to zero while X_1 is varied by plus and minus two unconditional standard deviation. “Premium” refers to the risk premium on the 10-year bond.

$X_1 = -1.9$, Premium = 11%



$X_1 = 0$, Premium = 2%



$X_1 = 1.9$, Premium = -7%






Figure 9: Allocation between the ten-year and three-year bond and the nominal riskfree asset when the investor hedges only risk premia (left panel) and when the investor hedges only the real riskfree rate (right panel). Allocation is plotted as a function of horizon for the investor with utility over terminal wealth. X_2 and X_3 are set equal to zero while X_1 is varied by plus and minus two unconditional standard deviation. “Premium” refers to the risk premium on the 10-year bond.

$X_1 = -1.9$, Premium = 11%



$X_1 = 0$, Premium = 2%



$X_1 = 1.9$, Premium = -7%




Figure 10: Optimal allocation between three bonds and the nominal riskfree asset. Allocation to the one, three, and ten-year bonds is plotted as a function of horizon for the investor with utility over terminal wealth. X_2 and X_3 are set equal to zero while X_1 is varied by plus and minus two unconditional standard deviation. “Premium” refers to the risk premium on the 10-year bond.




Figure 11: Certainty equivalents for sub-optimal strategies when the investor allocates wealth between a single long-term bond and the riskfree asset. Lines without circles represent the percent increase in wealth needed to make a myopic investor as well off as an investor who follows the optimal strategy. Lines with circles represent the percent increase in wealth needed to make an investor who hedges only the time-variation in the real interest rate as well off as in investor who follows an optimal strategy. γ refers to relative risk aversion. Certainty equivalents for $\gamma = 1$ are identically equal to 1. Certainty equivalents for the investor who hedges only time-variation in the riskfree rate are nearly equal for $\gamma = 4$ and $\gamma = 10$.




Figure 12: Certainty equivalents for sub-optimal strategies when the investor allocates wealth between a two long-term bonds and the riskfree asset. Lines without circles represent the percent increase in wealth needed to make a myopic investor as well off as an investor who follows the optimal strategy. Lines with circles represent the percent increase in wealth needed to make an investor who hedges only the time-variation in the real interest rate as well off as in investor who follows an optimal strategy. γ refers to relative risk aversion. Certainty equivalents for $\gamma = 1$ are identically equal to 1.




Figure 13: Certainty equivalents for sub-optimal strategies when the investor allocates wealth between a three long-term bonds and the riskfree asset. Lines without circles represent the percent increase in wealth needed to make a myopic investor as well off as an investor who follows the optimal strategy. Lines with circles represent the percent increase in wealth needed to make an investor who hedges only the time-variation in the real interest rate as well off as in investor who follows an optimal strategy. γ refers to relative risk aversion. Certainty equivalents for $\gamma = 1$ are identically equal to 1.