

Information in (and not in) the term structure

*Gregory R. Duffee**

Haas School of Business
University of California – Berkeley

First draft: March 2008
This Draft: June 23, 2008

ABSTRACT

Casual intuition says that today's term structure reflects all information investors have about expected future yields. However, this is not required by finance theory, nor is it consistent with observed Treasury yield behavior. Kalman filter estimation uncovers a factor that has an almost imperceptible effect on yields, but has clear forecast power for future short-term interest rates and substantial forecast power for future excess bond returns. The factor appears to be related to short-run fluctuations in economic activity.

*Voice 510-642-1435, email duffee@haas.berkeley.edu. Address correspondence to 545 Student Services Building #1900, Berkeley, CA 94720. The first version of this paper was written while the author was visiting Wharton. Thanks to seminar participants at the Federal Reserve Board and Wharton.

1 Introduction

Standard intuition says that the term structure of zero-coupon bond yields contains all information relevant to predicting both future returns to bonds and future bond yields. The simplest explanation is that investors' beliefs about future prices determine what they are willing to pay for bonds today. Yet even if investors are poorly informed, an accounting identity equates a zero-coupon bond's yield with the average log return to the bond during its remaining life. Taking the expectation of the identity reveals that changes over time in the shape of the term structure are equivalent to changes in expected future log returns. Since log returns to zero-coupon bonds can be expressed in terms of yields, expectations of future log returns contain the same information as expectations of future yields.

Researchers commonly invoke this intuition when building and estimating term structure models. One important application is the choice of the dimension of the model. If yield dynamics are captured by the cross-section of yields, then the same factors that explain the cross-section explain dynamics. Hence factor analyses of the unconditional covariance matrix of yields (or differenced yields) are often used to pin down the number of factors. Another application is model estimation. Estimation often assumes there is a one-to-one mapping from the factors to an equal number of bond yields. This assumption makes explicit the notion that the cross-section of bond yields follows a Markov process. We know the assumption is not literally true, because yields on individual bonds appear to have idiosyncratic components associated with market imperfections. But this noise is too small to alter the core of the standard intuition: the important determinants of expected future yields are the important determinants of current yields.

However, two empirical observations by Cochrane and Piazzesi (2005) cast some doubt on this view. First, they find that the forward rate from year four to year five contains substantial information about future excess bond returns, even though the contribution of this forward rate to the overall volatility of cross-section of bond yields is very small. Second, they find that lagged bond yields contain information about future excess bond returns not found in current bond yields. Cochrane and Piazzesi suggest that the noise in bond yields may play a role in these results. I offer a different interpretation of this wedge between determinants of the cross-section and determinants of expectations (although noise plays a role in my interpretation as well).

I show that it is easy to build a multifactor model in which one of the factors plays an important role in determining investors' expectations of future yields, yet has zero effect on current yields. The factor must have opposite effects on expected future interest rates and bond risk premia. Consider, for example, economic news that raises risk premia and

simultaneously leads investors to believe the Fed will soon cut short-term interest rates. The increase in risk premia induces an immediate increase in long-term bond yields, while the expected drop in short rates induces an immediate decrease in these yields. In a Gaussian term structure model, a single parameter restriction equates these effects, leaving the current term structure—but not expected future term structures—unaffected by the news. More generally, factors that drive risk premia and expected short rates in opposite directions can have arbitrarily small effects on the cross-section of yields, yet large effects on yield dynamics.

In principle, this result complicates substantially our efforts to model the term structure. We cannot choose the model’s number of factors based on the number of factors that explain the cross-section of yields. It also prevents us from using estimation techniques that rely on the ability to infer time- t factors from time- t yields. Even if there is no factor that has an exactly zero effect on the time- t term structure, its effect on yields may be too small to readily distinguish from idiosyncratic noise. But these concerns will be more theoretical than practical if we have no reason to believe that such hidden factors exist.

I look for hidden factors by fitting a five-factor Gaussian term structure model to monthly Treasury yields over the period 1964 through 2007. The Kalman filter allows us to infer the presence of hidden factors from term structure dynamics. Estimation uncovers a term structure factor that has a trivial effect on the cross-section of Treasury yields but contains substantial information about both expected future short rates and—necessarily—expected excess bond returns. Based on the model’s point estimates, a one standard deviation change in the factor has an almost imperceptible effect on the term structure (on the order of a few basis points), lowers the expected one-year-ahead short rate by about 35 basis points, and raises the expected excess return to a five-year bond over the next year by about 1.3 percent. This “expectation” factor accounts for about 30 percent of the total variance in expected excess bond returns.

Not surprisingly, there is substantial uncertainty in these point estimates. If we relied only on the results of the estimation, a skeptic easily could argue that the model is overfitting observed data, and the expectation factor is spurious. However, evidence from the Survey of Professional Forecasters confirms that survey-based expectations of future short rates move contemporaneously with filtered estimates of the factor. Moreover, the factor is related to short-run fluctuations in economic activity. An increase in the factor corresponds to lower expected future short rates, higher risk premia, and lower growth in industrial production.

I also investigate properties of regressions that use the term structure to forecast future excess annual bond returns. Under the maintained hypothesis that the estimated five-factor model is correct, such regressions are incapable of capturing all of the true variation in expected excess returns because they cannot capture fully the expectations factor. The

Cochrane and Piazzesi (2005) regression that uses five forward rates as forecasting variables slightly outperforms a regression that uses measures of level, slope, and curvature. In population, the difference in R^2 is about one percentage point. It also slightly outperforms in long (44-year) finite samples. However, in these finite samples it is also easy to conclude, mistakenly, that the two regressions capture substantially different amounts of information. Although the mean difference in these finite-sample R^2 s is only one percentage point, there is a twelve percentage point range from one end of the 95th percentile bound to the other.

The term structure model is presented in the next section. Section 3 summarizes properties of the estimated model. Section 4 compares the expectation factor to survey evidence on expectations and links the factor to the macroeconomy. Finite-sample properties on forecasting regressions are in Section 5. Concluding comments are in Section 6.

2 The modeling framework

The objective of this section is to explain why the important determinants of the cross-section of bond yields need not correspond to the important determinants of yield dynamics. To make this point in the starkest terms, I build a model in which n factors are necessary to model term structure dynamics, but only $n - 1$ factors appear in the term structure.

I follow much of the modern term structure literature by abstracting from standard economic concepts such as utility functions and consumption dynamics. Instead, both the short rate and the nominal pricing kernel are functions of a latent state vector. The factors and their dynamics can be viewed as reduced-form representations of inflation, business cycles, and market clearing.

2.1 The standard Gaussian model

I use a standard discrete time Gaussian term structure framework. The use of discrete time is innocuous. The role played by the Gaussian assumption is discussed in Section 2.5. The one-period interest rate is r_t . This rate is continuously compounded and expressed per period. (For example, if a period is a month, $r_t = 0.01$ corresponds to twelve percent/year.) Interest rate dynamics are driven by a length- n state vector x_t . The relation between the short rate and the state vector is

$$r_t = \delta_0 + \delta_1' x_t. \tag{1}$$

The state vector has first-order Markov dynamics

$$x_{t+1} = \mu + Kx_t + \Sigma\epsilon_{t+1}, \quad E_t(\epsilon_{t+1}\epsilon_{t+1}') \sim N(0, I). \tag{2}$$

The period- t price of a zero-coupon bond that pays a dollar at $t + m$ is denoted $P_t^{(m)}$. The corresponding continuously-compounded yield is $y_t^{(m)}$. Bond prices satisfy the law of one price

$$P_t^{(m)} = E_t \left(M_{t+1} P_{t+1}^{(m-1)} \right) \quad (3)$$

where M_{t+1} is the pricing kernel. The pricing kernel has the log linear form

$$\log M_{t+1} = -r_t - \Lambda_t' \epsilon_{t+1} - \frac{1}{2} \Lambda_t' \Lambda_t. \quad (4)$$

The vector Λ_t is the compensation investors require to face shocks to state vector. The price of risk satisfies

$$\Sigma \Lambda_t = \lambda_0 + \Lambda_1 x_t, \quad (5)$$

which is the essentially affine form introduced in Duffee (2002). Bonds are priced using the equivalent-martingale dynamics

$$x_{t+1} = \mu^q + K^q x_t + \Sigma \epsilon_{t+1}^q, \quad (6)$$

where the equivalent-martingale parameters are

$$\mu^q = \mu - \lambda_0, \quad K^q = K - \lambda_1. \quad (7)$$

The discrete-time analogues of the restrictions in Duffee and Kan (1996) imply that zero-coupon bond yields can be written as

$$y_t^{(m)} = A_m + B_m' x_t, \quad (8)$$

where the scalar A_m and the n -vector B_m are functions of the parameters in (1) and (6). The focus of this paper is on yield factor loadings, which can be written as

$$\begin{aligned} B_m' &= \frac{1}{m} \delta_1' (I + K^q + (K^q)^2 + \dots + (K^q)^{m-1}) \\ &= \frac{1}{m} \delta_1' (I - K^q)^{-1} (I - (K^q)^m). \end{aligned} \quad (9)$$

2.2 The information in the term structure

In the absence of specific parameter restrictions, the period- t state vector can be inferred from a cross-section of period- t bond yields. Stack the yields on n zero-coupon bonds in the

vector y_t^a . We can write this vector as

$$y_t^a = A^a + B^a x_t \quad (10)$$

where A^a is a length- n vector containing A_m for each of the n bonds and B^a is a square matrix with rows B'_m for each bond. In general, B^a is invertible. Put differently, element i of the state vector affects the n bond yields in a way that cannot be duplicated by a combination of the other elements. With invertibility, the term structure contains the same information as x_t . We can write

$$x_t = (B^a)^{-1} (y_t^a - A^a). \quad (11)$$

Since x_t is Markov and the term structure of yields contains the same information as x_t , the term structure is also first-order Markov.

We now investigate special cases of this Gaussian framework where B^a has rank less than n , so that the state vector cannot be extracted from the term structure. An example illustrates the mathematics and the economic intuition.

2.3 A two-factor example

Consider the two-factor Gaussian model. Because the latent factors in this model can be arbitrarily rotated, the state vector can be transformed into the short rate and some other factor, denoted f_t . For this rotation, the dynamics of the state vector are (explicitly indicating the elements of the feedback matrix)

$$\begin{pmatrix} r_{t+1} \\ f_{t+1} \end{pmatrix} = \mu + \begin{pmatrix} k_{11} & k_{12} \\ k_{21} & k_{22} \end{pmatrix} \begin{pmatrix} r_t \\ f_t \end{pmatrix} + \Sigma \epsilon_{t+1}. \quad (12)$$

When k_{12} is not restricted to zero, time- t expectations of future short rates depend on both r_t and f_t . Thus we can think of f_t as all information about future short rates that is not captured by the current short rate.

If investors were risk-neutral, the level of f_t would necessarily affect the term structure through expectations of future changes in the short rate. But if risk premia also vary with f_t , the net effect of f_t on yields is ambiguous. The restriction adopted in this example is that changes in risk premia exactly cancel expectations of future short rates, leaving yields unaffected by f_t . Formally, the requirement is $k_{12}^q = 0$, or $k_{12} = \lambda_{1(12)}$. Then the equivalent-

martingale dynamics of the state are

$$\begin{pmatrix} r_{t+1} \\ f_{t+1} \end{pmatrix} = \mu^q + \begin{pmatrix} k_{11}^q & 0 \\ k_{21}^q & k_{22}^q \end{pmatrix} \begin{pmatrix} r_t \\ f_t \end{pmatrix} + \Sigma \epsilon_{t+1}^q. \quad (13)$$

A glance at (13) reveals that under the equivalent-martingale measure, the short rate follows a (scalar) first-order Markov process. The loading of the m -period bond yield on the state vector is, from (9),

$$B_m = \begin{pmatrix} \frac{1}{m} (1 - k_{11}^q)^{-1} (1 - (k_{11}^q)^m) \\ 0 \end{pmatrix}. \quad (14)$$

Thus the matrix B^a in (10) cannot be inverted because it has a column of zeros. The factor f_t cannot be inferred from the period- t term structure.

Although the factor does not affect yields, investors observe it. They take it into account when setting bond prices and forming expectations of future yields (or equivalently, future returns to holding bonds). For concreteness, consider the case $k_{12} > 0$. Then for fixed r_t , an increase in f_t raises investors' expectations of future short rates. For example, consider macroeconomic news, such as unexpectedly high GDP growth, that raises the likelihood of future tightening by the Federal Reserve. If investors' willingness to bear interest risk did not change with f_t , this news would raise current long-maturity bond yields. But with the restriction $k_{12} = \lambda_{1(12)}$, investors accept lower expected excess bond returns. The change in willingness to bear risk offsets exactly the news about expected future short rates, leaving yields unaffected.

The functional relation between expected excess returns and f_t can be seen in the formula for the expected excess log return, from t to $t + 1$, on a bond with maturity m at period t . (Here, "excess" is in excess of the short rate.) The period- t expectation is

$$\begin{aligned} E_t \left(x r_{t,t+1}^{(m)} \right) &\equiv m y_t^{(m)} - (m - 1) E_t \left(y_{t+1}^{(m-1)} \right) - r_t \\ &= m A_m - (m - 1) A_{m-1} \\ &\quad + (1 - k_{11}^q)^{-1} \left[(1 - (k_{11}^q)^m) - (1 - (k_{11}^q)^{(m-1)}) k_{11} - 1 \right] r_t \\ &\quad - (1 - k_{11}^q)^{-1} (1 - (k_{11}^q)^{(m-1)}) k_{12} f_t. \end{aligned} \quad (15)$$

The final term in (15) captures the dependence of expected excess returns on f_t .

Even if an econometrician knows the parameters of the model, she cannot infer f_t from the cross-section of yields at t . Nor can f_t be backed out of the price of some other fixed-income instrument, such as bond options. The econometrician can, however, use a panel of data to form filtered estimates of f_t . The filtering approach is discussed again in Section 3.3.

The intuition behind filtering is easier to grasp if we call it learning by the econometrician. The period- t forecast error (the difference between realized yields and the econometrician's $t-1$ forecast) is produced by both true period- t shocks and the error in the econometrician's $t-1$ prediction of f_{t-1} . The cross-sectional pattern of the period- t forecast errors helps the econometrician revise her prediction of f_{t-1} and form her prediction of f_t .

In this example, the short rate follows a two-factor Markov process under the physical measure and a one-factor Markov process under the equivalent martingale measure. A single parameter restriction is required to generate this structure. Armed with the intuition of this example, it is straightforward to proceed to the more general case in which the short rate follows an n -factor Markov process under the physical measure and an $(n-1)$ -factor Markov process under the equivalent martingale measure. As in the two-factor case, a single parameter restriction is required.

2.4 The n -factor version

Latent state vectors in affine term structure models are inherently arbitrary. Dai and Singleton (2000) describe in detail how they can be translated and rotated without observable consequences. One particular rotation simplifies considerably the analysis here. Beginning with the standard n -factor Gaussian model of Section 2.1, diagonalize the equivalent-martingale feedback matrix K^q into

$$K^q = PVP^{-1} \quad (16)$$

where the columns of P are eigenvectors and V is a diagonal matrix of eigenvalues. Define a rotated state vector

$$x_t^* = Px_t. \quad (17)$$

The equivalent-martingale dynamics of the rotated state vector are

$$x_{t+1}^* = P\mu^q + Vx_t^* + P\Sigma\epsilon_{t+1}^q. \quad (18)$$

With this rotation, each individual factor follows its own univariate first-order Markov process because V is diagonal. Innovations among the factors can be correlated. The loading of the short rate on the rotated state vector is

$$(\delta_1^*)' = \delta_1' P^{-1} \quad (19)$$

Here, as in the two-factor case, a single parameter restriction produces a model where physical dynamics of the short rate follow an n -factor process and equivalent-martingale

dynamics follow an $(n - 1)$ -factor process. The restriction is that for some i ,

$$\delta_{1,i}^* = 0. \tag{20}$$

This restriction implies that element i of the state vector drops out of the equivalent-martingale dynamics of the short rate. It is immediate from (20) that the period- t values of the other $n - 1$ factors are sufficient to determine the period- t short rate. Similarly, the short rate at $t + \tau$ depends only on the period- $(t + \tau)$ values of $n - 1$ factors. Since each factor follows a univariate Markov process under the equivalent-martingale measure, the period- t equivalent-martingale expectation of the short rate at $t + \tau$ depends only on the period- t values of those same $n - 1$ factors. Therefore period- t yields depend only $n - 1$ factors.

As in the two-factor case, physical dynamics of the short rate depend on all n factors. The physical dynamics of the rotated state vector are

$$x_{t+1}^* = P\mu + PKP^{-1}x_t^* + P\Sigma\epsilon_{t+1}. \tag{21}$$

As long as risk premia vary with the state vector ($\lambda_1 \neq 0$), the matrix P that diagonalizes K^q will not diagonalize K . Then in general, each factor in the state vector contains information about the evolution of the short rate.

2.5 The role of the Gaussian setting

Section 2.4 shows that with an appropriate restriction on a term structure model, only $n - 1$ factors of an n -dimensional state vector affect bond yields. Models that exhibit unspanned stochastic volatility (USV), as described in Collin-Dufresne and Goldstein (2002), can be described similarly. Here I clarify the relation between the approach here and the USV approach.

Here, short rate dynamics are described by an n -factor Markov process under the physical measure and an $(n - 1)$ -factor Markov process under the equivalent martingale measure. All $n - 1$ factors that appear in the equivalent-martingale process affect bond yields. Thus we can say that under the equivalent-martingale measure, the term structure follows an $(n - 1)$ factor Markov process. By contrast, the USV framework is concerned only with the equivalent martingale measure. The physical measure is not specified. Under the equivalent martingale measure, the short rate follows an n -factor Markov process. Bond yields nonetheless do not depend on all n factors. (Prices of some other fixed-income instruments will depend on all n factors.) Thus under the equivalent-martingale measure, the term structure does not follow a Markov process.

The economic interpretations of the two sets of relevant parameter restrictions differ substantially. Here, variations in expected future short rates are offset by variations in risk premia. With USV, variations in equivalent-martingale expectations of future short rates are offset by variations in the Jensen’s inequality component of bond yields. Stochastic volatility is thus critical to USV models (hence the name of the model class), but does not appear here.

Although USV models appear to have little in common with the model here, they can provide an alternative mechanism driving a wedge between the factors driving dynamics of yields and the cross-section of yields. Set risk premia to zero so that physical and equivalent-martingale measures coincide. Then n factors are necessary to capture yield dynamics, while $n - 1$ factors affect bond yields. I do not pursue this approach because the parameter restrictions necessary in a USV model are very tight.

In fact, one reason I use the Gaussian framework is to avoid complications associated with stochastic volatility. Reconsider the two-factor example of Section 2.3. If the conditional covariance matrix of factor innovations is allowed to be linear in f_t (a discrete-time approximation to a square-root diffusion model), then the level of f_t affects bond yields even when $k_{12}^q = 0$. Variations in risk premia can offset variations in expected future short rates, but do not offset variations in the Jensen’s inequality component of yields. This problem does not arise in the two-factor example if conditional variances are allowed to depend on the short rate instead of f_t .

2.6 From theory to practice

This model illustrates that factors driving the dynamics of yields need not also drive the cross section. But the model offers no motivation for the necessary parameter restriction. In fact, it is probably unreasonable to assume that there is some factor for which variations in expected future short rates are exactly offset by variations in required expected returns. Thus there seems to be no theoretical justification to either *a priori* impose the constraint or to test statistically whether the constraint is consistent with observed yields.

The more important lesson to take from the model is that there can be a large wedge between the importance of a factor in the cross section and its importance in dynamics. It is easy to tell stories in which types of news have opposite effects on expected future short rates and investors’ required expected excess returns. For example, the Taylor (1993) rule and its variants (see, e.g., Clarida, Galí, and Gertler (2000)) suggest that good news about future output is also news that future short rates are likely to rise. If willingness to bear interest rate risk covaries positively with the business cycle, the immediate effect of such

news on bond yields will not accurately reflect the importance of the news in forecasting future short rates.

If all n factors affecting dynamics also affect the cross section, the mapping from factors to n yields in (10) implies that all factors can be inferred from the cross section using (11). However, the exact mapping does not hold in practice. Equation (10) implies that the unconditional covariance matrix of d bond yields is singular for $d > n$. Yet in the data, sample covariance matrices of zero-coupon bond Treasury yields are nonsingular for even large d (say, greater than ten). One interpretation of this result is that n is large, perhaps even infinite, as in Collin-Dufresne and Goldstein (2003). But from a variety of perspectives, it is more appealing to view bond yields as contaminated by small, transitory, idiosyncratic noise.

This noise is generated from three sources. First, there are market imperfections that distort bond prices, such as bid/ask spreads. Second, there are market imperfections that distort payoffs to bonds (and thus distort what investors will pay for bonds), such as special RP rates. Third, there are distortions created by the mechanical interpolation of zero-coupon bond prices from coupon bond prices.

I model the noise as classic measurement error. A vector of d period- t yields on bonds with maturities m_1, \dots, m_d is expressed as

$$y_t = A + Bx_t + \eta_t, \quad \eta_t \sim \sigma_\eta^2 N(0, I) \quad (22)$$

where η_t is a vector of measurement errors. For simplicity, in (22) the measurement error for each yield has the same variance. Element i of the vector A contains A_{m_i} and row i of the matrix B contains B'_{m_i} .

Equation (22) cannot be pushed to its logical limits. Since the measurement error is uncorrelated across maturities and time, (22) suggests that using either more points on the term structure or higher frequency data eliminates the effects of noise. Instead, the specification should be viewed as an approximation to a world in which noise dies out quickly (within a month) and is roughly uncorrelated across the widely-spaced maturities used in empirical analysis.

The presence of noise weakens the knife-edge distinction between a model in which all factors affect the cross section and a model in which one or more factors affects only dynamics. One factor may be an important determinant of expected future yields, yet have a small effect on the current term structure—so small that it is difficult or impossible to distinguish the factor’s influence on the cross section from the noise η_t in (22).

The most obvious conclusion to draw from this discussion is that we cannot rule out a

priori the existence of factors that have a minimal effect on the cross-section of yields yet have important effects on yield dynamics. The next sections asks whether empirical analysis can confirm or deny the presence of such factors.

3 Empirical analysis

The core of this empirical analysis is the interpretation of an estimated five-factor Gaussian term structure model. The model is used in two ways. First, I ask whether there are factors that have little effect on the cross-section of yields yet are important for modeling dynamics. Second, assuming that the estimated model is correct, I study properties of regressions that forecast excess bond returns using the cross-section of bond yields. Before getting into details of model estimation, I discuss why a five-factor model is used.

3.1 How many factors?

The goal of the empirical work in this section is to explore the possibility that some factor(s) are important for term structure modeling but have little effect on the cross-section of yields. Therefore the number of factors must exceed the usual choice of three—a choice motivated by cross-sectional properties. Unfortunately, the model of Section 2 does not tell us how many factors should be used.

Two considerations motivate the choice of five factors. First, Cochrane and Piazzesi (2005) use information from five points on the yield curve to form forecasts of excess bond returns. If five points are needed, the underlying model should have at least five factors. Second, the number of free parameters is unmanageable for six or more factors. A five-factor Gaussian canonical term structure model has 52 free parameters. As we will see, extracting information about each of these parameters is close to (or beyond) the limits imposed by available data and estimation techniques. Convincing a skeptic that they have something to learn from a model with 52 parameters is also difficult. A six-factor canonical model has more than 70 free parameters. It is beyond the ability of the author to convince anyone to take seriously the parameter estimates of a 70+ parameter model.

3.2 Data

Treasury bond yields are from the Center for Research in Security Prices (CRSP). The yield on a three-month Treasury bill is from the Riskfree Rate file (bid/ask average). Artificially-constructed yields on zero-coupon bonds with maturities of one, two, three, four, and five years are from the Fama-Bliss file. Yields are observed at the end of each month from

January 1964 through December 2007. The first observation is chosen to align with the sample studied by Cochrane and Piazzesi (2005).

Panel A of Table 1 reports means and standard deviations of the yields. Panel B reports the magnitudes of the first five principal components of the six yields, monthly changes in the yields, and annual returns to bonds with initial maturities of one through five years. The characteristics of the principal components are well-known. The first three principal components of levels explain more than 99.9 percent of their total variation. The corresponding percentages for monthly changes and annual returns are 98.6 and 99.9 respectively. Not shown in Table 1 are the shapes of the principal components. It is well-known that the first three are the level, slope, and curvature of the term structure. We look at these shapes in detail below.

Similar results popularized by Litterman and Scheinkman (1991) are typically used to motivate the number of factors included in formal term structure models. For example, the choice of three factors in Duffee (2002) is explicitly justified by this result. One of the questions studied here is whether the relative importance of factors in the cross-section matches their relative importance in dynamics.

3.3 Model estimation

A factor that is hard or impossible to extract from the cross section should be inferred using filtering techniques instead of assuming that all period- t factors are linear combinations of period- t yields. The Kalman filter produces correct conditional means and covariance matrices when the underlying model fits into the Gaussian term structure framework. Thus I estimate a five-factor Gaussian model with maximum likelihood (ML) by applying the Kalman filter.

The model of Section 2 is written in terms of dynamics under physical and equivalent-martingale measures. That form allows us to understand the economics underlying factor models of the term structure. However, for the purpose of estimation, it is convenient to use a slightly different parameterization. Following the language of the Kalman filter, write the model in the form of a transition equation and a measurement equation. The state vector used in the estimated model is denoted x_t^\dagger . The transition equation is

$$x_{t+1}^\dagger = D^\dagger x_t^\dagger + \Sigma^\dagger \epsilon_{t+1}. \quad (23)$$

In (23), D^\dagger is a diagonal matrix and Σ^\dagger is lower triangular with ones along the diagonal. The forms of D^\dagger and Σ^\dagger are normalizations, as is the state vector's unconditional mean of zero. There are five latent factors in the state vector x_t^\dagger . Therefore there are a total of 15

free parameters in (23). The measurement equation is

$$y_t = A + B^\dagger x_t^\dagger + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2). \quad (24)$$

In (24), A is a 6×1 vector and B^\dagger is a 6×5 matrix. There is also a single standard deviation of measurement error, resulting in a total of 52 parameters.

Lurking behind the parameters of the measurement equation are the equivalent-martingale dynamics of x_t . Because there are five factors to explain six bond yields, A and B^\dagger exactly identify the unconstrained parameters of the no-arbitrage model δ_0 , δ_1 , μ^q , and K^q . As discussed in Duffee (2008), numerical optimization of the likelihood function is faster and more reliable when the estimated parameters are A and B^\dagger than when they are the parameters of the no-arbitrage model. Here I follow exactly the optimization procedure used in that paper.

3.4 A principal components factor rotation

The state-vector rotation implied by (23) and (24) is convenient for estimation. A rotation based on principal components is more useful for interpreting the results. Denote ‘uncontaminated’ yields—yields without measurement error—by \tilde{y}_t . Drop the three-year bond, denoting the vector of the remaining five yields by $\tilde{y}_{\setminus 3,t}$. The loadings of these yields on the factors are denoted $B_{\setminus 3}^\dagger$, a 5×5 matrix. Estimates of the parameters of (23) and (24) imply a population covariance matrix of $y_{\setminus 3,t}$. (As the model in Section 2.4 illustrates, there are parameterizations for which this covariance matrix is singular, but the parameter estimates do not happen to satisfy the restriction necessary for singularity.) Diagonalize this covariance matrix into

$$\text{Var}(\tilde{y}_{\setminus 3,t}) = C_0 \Omega C_0^{-1}. \quad (25)$$

Define the 5×5 matrix Γ as

$$\Gamma = C_0^{-1} B_{\setminus 3}^\dagger. \quad (26)$$

The state vector that is easy to interpret is

$$x_t = \Gamma x_t^\dagger. \quad (27)$$

The factors in this vector are all five principal components of the yields on bonds with maturities of three months, one, two, four, and five years. Their unconditional covariance matrix is the diagonal matrix of eigenvalues

$$\text{Var}(x_t) = \Omega. \quad (28)$$

The dynamics of the rotated state vector are

$$x_{t+1} = Kx_t + \Sigma\epsilon_{t+1}, \quad (29)$$

where the parameters are defined by

$$K = \Gamma D^\dagger \Gamma^{-1}, \quad \Sigma = \text{chol}(\Gamma \Sigma^\dagger \Sigma^\dagger \Gamma'). \quad (30)$$

The relation between bond yields and the rotated factors is

$$y_t = A + Bx_t + \eta_t, \quad (31)$$

where the new factor loadings are

$$B = B^\dagger \Gamma^{-1}. \quad (32)$$

These factor loadings (for all but the three-year bond yield) are the eigenvectors of the diagonalization (25).

Table 2 reports the point estimates of the model for this principal components rotation. There are 77 parameters in the table, although the model has only 52 free parameters. There are 15 restrictions built into these parameters that derive from the requirement that the factors are principal components of the yields. Standard errors are in parentheses. They are constructed from Monte Carlo simulations. Assuming that the estimated model is true, 528 months of yields are randomly generated for a given simulation. The model is estimated with maximum likelihood using these data and the parameter estimates are stored. This procedure is repeated 1000 times to construct the standard errors in Table 2. The covariance matrix of the 77 parameter estimates has rank 52.

3.5 Estimates of the factors' role in the cross section

The estimates in Table 2 are reported for only for completeness. There is not much to be learned from the individual parameters. Instead, I summarize the important properties of the estimated model. This subsection focuses on the cross-sectional properties. A quick summary is that only the first three factors play a noticeable role in the cross section. The remaining factors are hard to disentangle from noise in yields.

Table 3 describes the cross-sectional relation between the factors and bond yields. Since the factors are, by construction, principal components of yields, it is not surprising that the first few factors explain almost all of the variation in yields. We see in the first column that population standard deviations of these orthogonal factors range from 6.02 for the first factor

to 0.04 for the fifth. Standard errors of these population standard deviations, computed from Monte Carlo simulations, are in parentheses.

The precise mapping from factors to yields is displayed in Figure 1, which plots the matrix of estimated factor loadings B scaled by the factor standard deviations. The first panel plots loadings on the first three factors. They are the usual level, slope, and curvature factors. For example, a one standard deviation increase in the first factor raises all yields by about 2.5 percentage points. The second panel plots loadings on the fourth and fifth factors. There is no obvious cross-sectional interpretation for these two factors, which appear to be economically tiny. Note the difference in scale between the two panels. A one standard deviation in the fifth factor does not change any yield by more than four basis points.

Because of measurement error, it is difficult to extract the final two factors from the cross-section of the term structure, even if we know the model’s parameters. Table 2 reports the estimated standard deviation of measurement error is about five and a half basis points (annualized yields). Although economically small, it is enough to obscure the effects of these factors on yields. One way to see this is to imagine a regression, using an infinite time series, of a factor on contemporaneous yields. (An econometrician cannot estimate this regression because she does not directly observe the factors.) The point estimates of the model allow analytic calculation of the R^2 for the regression.

The second column of statistics in Table 3 reports the R^2 s for each factor regressed on bond yields with maturities of three months and one through five years. The effects of the first three factors on yields are sufficiently large to dominate measurement error. The R^2 s for these factors range from 1.0 to 0.95. However, the R^2 s for the fourth and fifth factors are only 0.62 and 0.43 respectively. Put differently, the correlations between the true and OLS-fitted estimates of the factors are 0.79 and 0.66.

Kalman filtering produces more accurate estimates of the factors. Population properties of the Kalman filter are proxied by simulating one million months of bond yields (the maturities are six months and one through five years), where the “true” model is the model estimated with ML. The Kalman filter is then applied to these data, using the true parameters in the filter. The final column of Table 3 reports correlations between true and Kalman filtered estimates of the factors. These correlations are 0.84 and 0.80 for the fourth and fifth factors. Naturally, filtered estimates of the factors are more closely related to observed yields than are true factors (since observed yields are used in the filtering), as documented in the third column of statistics in Table 3.

Since only the first three factors make noticeable contributions to the cross-section of yields, why should we care about our ability to infer the other factors from the data? The reason is that according to the model’s point estimates, the fifth factor plays an important

role in yield dynamics and expected excess bond returns.

3.6 Estimates of the factors' role in yield dynamics

Consider investors' j -month-ahead forecast of the yields used in estimation of the model. The vector of forecasts is (recall that investors know the true state vector)

$$\begin{aligned} E_t(y_{t+j}) &= A + BE_t(x_{t+j}) \\ &= A + BK^j x_t. \end{aligned} \tag{33}$$

The unconditional covariance matrix of these forecasts is

$$\text{Var}(E_t(y_{t+j})) = BK^j \Omega B' (K^j)'. \tag{34}$$

Because the unconditional covariance matrix of the factors Ω is diagonal, the variance in (34) can be unambiguously expressed as the sum of components attributable to each of the five factors.

Table 4 reports information about this decomposition. To simplify interpretation, the table reports standard deviations rather than variances. To illustrate the results, consider the first row. The table reports that twelve-month-ahead forecasts of the three-month annualized bill yield have a standard deviation of 2.28 percentage points. More than 95 percent of the variance is due to the first, "level" factor. The standard deviation of twelve-month-ahead forecasts attributable to this factor is 2.23 percentage points. Standard deviations attributable to all other factors are much smaller.

The surprising result in this first row is that much of the remaining variance in twelve-month-ahead forecasts is captured by the fifth factor. The standard deviation of the forecast attributable to this factor is 36 basis points, which is larger than the amount attributable to any other non-level factor. This pattern holds for all maturities included in the table. The vast majority of the variation in twelve-month-ahead forecasts is driven by the level factor, while the fifth factor picks up most of the remainder.

Visual evidence of the contributions of the factors to short-rate forecasts is in Figure 2. The figure displays impulse responses of the three-month bill yield to one standard deviation changes in each factor. For example, in the first panel the month-zero yield is 2.73 percentage points above its mean. Two years later, the yield remains 1.72 percentage points above its mean. The second (slope) factor corresponds to an immediate drop in the short rate of about 60 basis points, half of which has disappeared after a year. The third and fourth factors contribute little to current or future short rates. The effect of the fifth factor is qualitatively

different from all of the other factors. It has no effect on the short rate at month zero. One year later, the short rate has dropped 35 basis points, where it remains for the next year. Accordingly, I label this fifth factor the “expectation” factor.

I use Monte Carlo simulations to calculate the bias and uncertainty in Table 4’s point estimates. An individual simulation begins by assuming the model estimated here is correct. Then a panel of 528 months of yields is simulated. Using these simulated data, the model is estimated with the Kalman filter. The simulations reveal that the total standard deviations of twelve-month-ahead forecasts are downward biased. For example, the ‘true’ model implies a standard deviation of three-month yield forecasts of 2.28 percentage points (from the top row of the table). The mean standard deviation from the Monte Carlo simulations is only 1.98 percentage points, as displayed in parentheses. The 2.5 and 97.5 percentile values are 0.86 and 3.40 percentage points respectively, as displayed in brackets.

There is substantial statistical uncertainty about the role of the expectations factor in yield dynamics. Under the null that the estimated model is true, point estimates of the contribution of the factor to twelve-month-ahead forecasts are downward biased. Their confidence intervals are also very wide. For example, when the expectations factor truly accounts for 36 basis points of standard deviation in twelve-month-ahead forecasts of the short rate, ML estimation using 528 months of data produces a mean point estimate of 31 basis points. A 95 percent confidence interval ranges from 3 to 63 basis points. Thus if we restrict ourselves to using only bond yields, it is probably impossible to make even qualitative statements about the role of the expectations factor. Below I also draw on evidence from the Survey of Professional Forecasters and the growth of industrial production.

Because the expectation factor plays the central role in the remainder of the paper, it is useful to take a quick look at its time-series behavior. Figure 3 plots filtered estimates of this factor over the sample period 1964 through 2007. The factor is normalized by its model-implied population standard deviation. Its persistence is fairly low. The model’s parameter estimates imply that a shock to the factor (holding all other factors constant) has a half life of five months. Any relation between the factor and economic fluctuations is not obvious from this figure, which also displays NBER turning points. Section 4.3 uncovers a relatively high-frequency relation between the factor and economic activity.

3.7 Estimates of the factors’ role in excess return dynamics

Although the level factor is the dominant driver of yields, it plays a much less important role in expected excess returns. In this section I focus on the behavior of the log return from t to $t + j$ on a bond with period- t maturity m , in excess of the log return on a j -period bond.

The observed excess return, expressed in terms of factors and measurement error, is

$$\begin{aligned}
xr_{t,t+j}^{(m)} &= mA_m - (m-j)A_{m-j} - jA_j \\
&+ (mB'_m - (m-j)B'_{m-j}K^j - jB'_j) x_t \\
&- (m-j)B'_{m-j} \left(\sum_{i=1}^j K^{j-i} \epsilon_{t+i} \right) \\
&+ m\eta_t^{(m)} - j\eta_t^{(m)} - (m-j)\eta_{t+j}^{(m-j)}. \tag{35}
\end{aligned}$$

The four lines on the right side of (35) are, respectively, the unconditional mean, the variation in the conditional mean owing to the period- t state vector, the return innovation owing to shocks to the state vector, and the measurement error component.

The estimates of A and B allow to study directly the population properties of this excess return for a one-year horizon ($j = 12$) and for bonds with maturities of two, three, four, and five years. Panel A of Table 5 reports unconditional means and standard deviations of these returns. Standard deviations are calculated for both true returns (i.e., excluding measurement error) and observed returns. The panel also reports the fraction of the total variance attributable to factor-driven variations in the conditional mean.

Unconditional mean excess annual returns are less than one percent for all of these bonds. Population standard deviations of the returns range from 1.8 percent for the two-year bond to 5.6 percent for the five-year bond. We see in the panel that measurement error contributes very little to the volatility of observed returns; differences in standard deviations between true and observed returns are at most a basis point.

Panel A also reports that predictable variations in returns account for about 21 percent of total return variance. Panel B decomposes this predictable variance into components attributable to each factor. The structure of Panel B mirrors that of Table 4. Consider, for example, the month- t expectation of the annual excess log return to a five-year bond. The estimated unconditional standard deviation of this expectation is 2.53 percent. Most of this variation is due to “slope” factor. The standard deviation attributable to this factor is 1.97 percent.

Given the well-known relation between the slope of the term structure and expected excess bond returns, it is not surprising that for each bond, the slope factor accounts for over half of the predictable variance. A glance at Figures 1 and 2 explains why. The slope factor simultaneously raises long-term bond yields and lowers expected future short rates. The more interesting result in Panel B is that the expectation factor explains up to 30 percent of the predictable variance. Again, a glance at the two figures explains why. The expectation factor lowers expected future short rates while leaving long-term yields unchanged.

Figure 4 displays the sensitivity of expected excess annual log returns to the level, slope, and expectation factors. For example, a one-standard-deviation increase in the slope factor raises the expected excess return to a two-year bond by about 60 basis points. The corresponding change in expected return to a five-year bond is about 200 basis points. These values are plotted with the dashed line in the figure. The solid line plots changes in expected excess returns for a one-standard deviation increase in the level factor and the dashed line plots changes for the expectation factor.

Table 5 documents substantial statistical uncertainty about the contribution of the expectations factor to expected excess returns. This mirrors the results for yield dynamics in Table 4. For example, when the expectations factor truly accounts for 1.35 percentage points of standard deviation in annual excess returns to a five-year bond, ML estimation using 528 months of data produces a mean point estimate of 1.16 percentage points. A 95 percent confidence interval ranges from 20 basis points to 2.12 percentage points.

These results, along with the results in the previous subsections, lead to two main conclusions. First, the point estimates imply an economically important role for the expectations factor. It drives both expectations of future yields and excess returns, although its role in the cross section is negligible. Put differently, factors that are most important for determining the shape of the term structure are not the most important in determining expected excess bond returns. This conclusion is consistent with the theory of Section 2.4. Second, the uncertainty in these point estimates is very large. Based only on this evidence, we cannot be confident that the results are not spurious.

From a statistical perspective, the main problem is that the expectations factor is difficult to infer from a panel of yields. We need to look at other sources of information to learn more about this factor.

4 The expectations factor: additional evidence

Is the estimated expectations factor truly capturing investors' expectations, or is it simply the consequence of overfitting a particular sample? A natural way to answer this question is to compare the factor to independent observations of investors' forecasts. At the end of the first month of every quarter since 1981Q3, participants in the Survey of Professional Forecasters are asked for their forecasts of the average level of the three-month Treasury bill during each of the next four quarters. This section examines the relation between mean forecasts (where the mean is taken across the participants) and contemporaneous values of the expectation term structure factor. Here, "contemporaneous" means the filtered estimate for the end of the first month in the quarter.

If the expectation factor is spurious, forecasters' contemporaneous expectations should be unrelated to it. For example, assume the quarter- t level of the filtered expectation factor predicts that the short rate will decline over the next few quarters. If this prediction is simply an ex-post interpretation of the data by the maximum likelihood estimation, then the survey responses in quarter t will not anticipate a decline in rates. Thus we can test the null hypothesis that the expectation factor is entirely spurious by examining its covariation with survey forecasts of changes in rates.

Before presenting the regression results, it is instructive to study in detail two particular observations.

4.1 A tale of two Octobers

Panel A of Figure 5 displays term structures for the month-ends of October 2001 and October 2004. (The plotted points are yields for maturities of three months and one through five years.) The shapes of the term structures are similar. The three-month bill yields are both around two percent. The largest difference between the term structures is at the long end, where the October 2001 observation is 37 basis points above the October 2004 observation. The dates were chosen both because the term structures are similar and the filtered estimates of the fifth factor are not.¹ The October 2001 estimate of this factor is about 0.7 standard deviations, while the October 2004 estimate is about -1.1 standard deviations.

This large difference in estimates of the fifth factor corresponds to a large difference in expected excess bond returns. Panel B of the figure displays model-implied expectations, as of October 2001 and October 2004, of one-year log returns to bonds in excess of the yield on a one-year bond. In 2001, the expectations are positive for all of the plotted maturities (two through five years), from 0.4 percent for the two-year bond to 1.6 percent for the five-year bond. In 2004, the expectations are negative, ranging from -0.4 percent to -1.2 percent. The difference in expected excess returns is largely accounted for by the difference in the expected time path of the three-month bill rate. Panel C reports that for 2001, the bill rate is expected to decline slightly for a few months, then rise modestly. By contrast, in 2004 the bill rate is expected to rise substantially over the next year. The average difference between the two sets of forecasts over the upcoming year (November through December of the next year) is about 60 basis points.

Are these model-implied expectations reasonably consistent with investors' expectations at the time? According to the Survey of Professional Forecasters, they are. For the surveys returned in early November 2001, the mean forecasts of the three-month bill rate for the next

¹In particular, the months were not chosen based on the contemporaneous survey forecasts.

four quarters (2002Q1 through 2002Q4) are 1.9, 2.0, 2.4 percent, and 2.8 percent respectively. Three years later, mean forecasts are about 50 basis points higher. The forecasts for 2005Q1 through 2005Q4 are 2.3, 2.6, 2.9, and 3.2 percent. Investors (or at least those investors with beliefs similar to those embodied by the mean forecasts of the survey participants) anticipated lower expected excess returns in October 2004 than in October 2001.

Differences in expected excess returns across these two months may be related to anticipated macroeconomic activity. Forecasters responding to the 2001Q4 survey were much more pessimistic about near-term economic growth than were those responding to the 2004Q4 survey. The 2001Q1 mean forecast of real GDP growth in 2002 was 0.8 percent. By contrast, the 2004Q4 forecast of real GDP growth in 2005 was 3.5 percent. The link between the expectations factor and expected future economic growth is pursued in Section 4.3.

A single comparison of two months is illuminating, but not statistically compelling. The next subsection contains some regression evidence.

4.2 Regression results

Denote the quarter- t mean survey forecast of the three-month bill in quarter $t + j$ less the quarter- t bill yield as $\text{SPF_EXPECT}(t, j)$. To align the bill yield with the survey timing, the quarter- t yield is defined as the three-month yield as of the end of the first month in the quarter. The continuously compounded yield from CRSP is converted to a discount basis to match the survey. Denote quarter- t filtered estimates of the expectation factor as MODEL_EXPECT_t . These are estimates for the end of the first month in the quarter. To simplify interpretation of the estimated regression coefficients, this factor is normalized by its population standard deviation. The sample period is 1981Q3 through 2007Q4.

I first estimate the regression

$$\text{MODEL_EXPECT}_t = b_0 + b_1 \text{SPF_EXPECT}(t, j) + e_{j,t} \quad (36)$$

for forecast horizons of one through four quarters ($j = 1, \dots$). Under the null hypothesis that the filtered estimate of the expectation factor is spurious, the coefficient b_1 should be zero. Because quarterly survey forecasts are serially correlated, standard errors use the Newey-West adjustment for four lags of moving average residuals. Although the regression is probably more intuitive if the regressor and regressand are switched, there is a generated regressor problem when using the filtered estimate of the expectation factor as the explanatory variable.

The coefficient should be negative if the model's factor is not spurious. As shown in Figure 2, the model implies that a one standard deviation increase in the expectation factor

corresponds to an expected drop in the three-month bill rate of 35 basis points over the subsequent year. Reversing the order of this comparison for the purposes of (36), an expected increase in the bill rate of one percentage point corresponds to -2.9 standard deviations of the factor.

Coefficient estimates for each forecast horizon are displayed in Panel A of Table 6. The null hypothesis is overwhelmingly rejected. The point estimates are reliably negative, with asymptotic t statistics ranging from -2.9 to -5.6 . The point estimates are less than the model predicts, ranging from -0.5 to -1.2 . In other words, the estimated factors respond less to true variations in expected changes in short rates than the model implies.

These regressions are estimated over the entire sample for which forecasts are available from the Survey of Professional Forecasters. From a statistical perspective, one unfortunate feature of this sample is that the estimated term structure factors are not uncorrelated. Over the entire 1964 through 2007 sample, the sample correlation between filtered values of the level and expectation factors is very close to zero. But from 1981Q3 through 2007Q4, the sample correlation is about 0.22. As Figure 2 shows, both the level and expectation factors have the same qualitative effect on expected future short rates. When the factors are high, short rates are expected to decline. Hence it is possible that the negative point estimates for (36) are proxying for the relation between the level of rates and expected future changes in rates. (Note, though, that this proxy story does not explain the tale of two Octobers.)

To control for the level of the term structure, I reverse (36) and add the estimated level factor as an additional explanatory variable. The regression is

$$\text{SPF_EXPECT}(t, j) = b_0 + b_1 \text{MODEL_LEVEL}_t + b_2 \text{MODEL_EXPECT}_t + e_{j,t}. \quad (37)$$

Both explanatory variables are generated regressors. Because the expectation factor is harder to extract from the yield curve than is the level factor, there is likely to be more noise in the model's estimate of the former factor than the latter.

Coefficient estimates for each forecast horizon are displayed in Panel B of Table 6. Both factors are negatively associated with survey expectations of future changes in the bill yield. More importantly, the statistical significance of the relation between the expectation factor and survey expectations does not disappear when the level factor is included. The asymptotic t statistics for the coefficients on the expectation factor are about -3.1 for one-quarter-ahead and two-quarter-ahead forecasts, -2.4 for three-quarter-ahead forecasts, and -2.0 for four-quarter-ahead forecasts.

This evidence supports the model's conclusion that the expectations factor is known by investors. In order for this factor to not affect the term structure, its predictive power for

future short rates must be offset by variations in risk premia. Such a story is more plausible if the expectations factor can be linked to the business cycle.

4.3 The expectations factor and economic activity

I examine the lead/lag relation between filtered estimates of the expectation factor and monthly changes in log industrial production. The estimated regression is

$$100(\log(\text{IP}_t) - \log(\text{IP}_{t-1})) = b_{0,i} + b_{1,i}\text{MODEL_EXPECT}_{t-i} + e_{t,i}, \quad i = -6, \dots, 6. \quad (38)$$

The change in IP lags the expectation factor for $i < 0$ and leads it for $i > 0$. Log changes in IP are serially correlated. A typical serial correlation of fitted residuals for (38) is about 0.3. I therefore report Newey-West standard errors adjusted for two lags of moving average residuals. As in Section 4.2, the expectations factor is normalized by its population standard deviation.

Estimation results are in Table 7. There is a strong, statistically significant inverse relation between industrial production and the expectations factor. In other words, low growth in industrial production corresponds to high risk premia accompanied by expected future declines in short-term rates. Growth in industrial production begins to drop a few months prior to the increase in the expectations factor, continuing for a couple of months after the increase in the expectations factor. If the filtered expectations factor is a standard deviation above its mean in month t , monthly growth in industrial production in months $t - 5$ through $t + 2$ averages about 11 basis points per month below average. (To put the 11 basis points in perspective, the standard deviation of monthly IP growth is about 70 basis points.)

These results are comforting because they are qualitatively consistent with a simple story. Investors believe that the Fed will attempt to offset some types of short-lived macroeconomic shocks with monetary policy actions. The Fed action is not anticipated to be immediate; short rates may not change for a number of months. The same macroeconomic shocks change investors' willingness to bear risk. Thus the net effect of the macro shocks on current yields is muted because the expected change in short rates and the change in risk premia work in opposite directions.

5 Forecasting with the cross section

The empirical evidence in the previous section tells us there is a term structure factor that can predict future yields and excess bond returns, yet is difficult to detect in the cross section.

This conclusion leads to a natural question. If we follow standard practice by forecasting excess returns using only information in the cross section of yields, how accurate will we be?

According to the estimated model (and as summarized in Table 5), about 21 percent of the total variance in annual excess log returns to bonds corresponds to variation in conditional mean returns. An econometrician cannot reproduce a time series of conditional expected returns, even given an infinite time series, because she must filter the factors instead of observe them. This section considers how accurately she estimates expected excess using standard predictive regressions instead of filtering. The forecasting variables are taken from cross-section of the term structure.

Following Cochrane and Piazzesi (2005) (hereafter CP), the regressions predict, as of month t , the excess log return to a bond from month t to month $t + 12$. One regression is inspired by standard three-factor term structure models. It uses month- t values of the level, slope, and curvature of the term structure. They are respectively defined as the five-year yield, five-year yield less three-month yield, and two-year yield less the average of the three-month and five-year yields. The other regression uses the five forward rates that CP found to contain substantial information about future excess returns. They are the month- t forward rates from year m to year $m + 1$ for $m = 0, \dots, 4$.

Population and finite-sample properties of the regressions are calculated by assuming that the term structure model estimated in Section 3 is correct. Population properties are determined analytically, while finite-sample properties are produced with Monte Carlo simulations. The length of each simulated sample is 528 months (44 years), which is the length of the sample used to estimate the model in Section 3. The results here are based on 10,000 simulations.

We first take a close look at predictions of annual excess returns to the five-year bond. The relevant information is in Table 8. Panel A contains results for the level, slope, and curvature regression, while Panel B contains results for the forward-rate regression. The population R^2 s of the two regressions are similar; 14 percent using level, slope, and curvature, and 16 percent using five forward rates. For comparison, Table 5 reports that 21 percent of the excess return variance is truly predictable. Because factors cannot be precisely inferred from the term structure, the regressions do not capture everything that investors know, even with an infinite time series. The population coefficients on the five forward rates display the tent shape uncovered by CP. However, these coefficients are closer to zero than are those estimated by CP. Similarly, the population R^2 here is less than the R^2 of 36 percent estimated by CP.

These differences in magnitudes are probably a consequence of model misspecification. The sample period used to estimate the model in Section 3 is almost identical to that used

by CP. Thus one way to interpret these differences is that the model does not reproduce the strong annual-horizon predictability that is a feature of the data used to estimate the model. However, there is relatively little information in this estimated predictability. The finite-sample evidence in Table 8 shows that both parameter estimates and R^2 s have very large sampling uncertainty. For example, the 95 percent confidence range for the forward-rate regression's R^2 is from 8 percent to 39 percent. The confidence range for the other regression's R^2 is similar. In a given sample, the two regressions produce substantially different R^2 s. The evidence is in Panel C of Table 7. The mean and median difference is only one percentage point, but a 95 percent confidence interval for the difference ranges from -0.041 (i.e., the three-factor R^2 is 4.1 percentage points *greater* than the forward-rate R^2) to 8.0%.

Moreover, although the tent shape is a population property of the model, it is not a result that leaps out of a random sample of data generated by the model. In more than 40 percent of the simulations, the forward-rate regression parameter estimates do not satisfy the tent shape.² By contrast, the coefficients on level and slope for the other regression are almost always positive, as indicated by their 95 percent confidence bounds in the table.

In population, the forward-rate regression captures slightly more of the true predictability of excess returns than does the three-factor regression. Is this also true for the sample size studied here? More precisely, is the average difference between the fitted prediction and the true expectation smaller for the forward-rate regression than for the three-factor regression? To answer this question, I use a root mean squared error metric with the following notation. Individual simulations are indexed by $i = 1, \dots, 10,000$. Annual log excess returns are predicted with OLS regressions for maturities of one through five years. Maturities are measured in months by m . For simulation i , define the root mean squared error for an m -maturity bond as

$$RMSE(i, m, p) = \sqrt{\frac{1}{528 - 12} \sum_{t=1}^{528-12} \left(E_t(xr_{t,t+12}^m) - E_t^{[p]}(xr_{t,t+12}^m) \right)^2}. \quad (39)$$

The first term on the right of (39) is the true conditional expected excess return to an m -month bond from t to $t + 12$. The second term is the fitted in-sample expectation from a predictive regression. In-sample forecasts for the three-factor regression are indicated with $p = 1$ and the corresponding forecasts for the forward-rate regression are indicated with $p = 2$. Note that this forecast error measures the distance between two forecasts, not the

²This result is not found in any table. The estimates are defined to satisfy the tent shape if the five estimated coefficients satisfy $1^{st} < 2^{nd} < 3^{rd} > 4^{th} > 5^{th}$.

difference between forecasts and realizations.

Table 9 contains the relevant results. The first column of statistics contains the population standard deviations of conditional expected annual returns. These standard deviations are useful benchmarks for the reported RMSEs because they can also be interpreted as RMSEs. Consider constant forecasts of annual excess returns equal to the unconditional mean excess returns. The RMSE of this forecast (again, defining RMSE relative to the true conditional forecast, not relative to realized returns) equals the standard deviation of the conditional expected excess return.

According to the RMSE metric, both regressions outperform the constant-expectation benchmark, and the forward-rate regression outperforms the three-factor regression. Consider, for example, forecasts of annual excess returns to a five-year bond. The population standard deviation of expected excess returns is 2.41 percent. The mean RMSEs for the three-factor and forward-rate regressions are 1.93 percent and 1.83 percent respectively. The three-factor RMSE exceeds the forward-rate RMSE in 84 percent of the individual simulations. The straightforward conclusion to draw from these results is that the additional information in two additional points on the term structure outweighs the problem of overfitting, at least for samples of the size studied here.

6 Conclusion

This paper shows that an econometrician cannot extract from bond yields all information investors have about expected future yields. An “expectation” factor contains information about expected future yields but is hidden, in the sense that it has a negligible effect on the term structure. Estimation procedures that explicitly look for a hidden factor, such as filtering, are helpful, but are no substitute for direct observation. One lesson to draw from these results is that information from sources other than bond yields can be valuable in uncovering hidden factors. The evidence here shows that the expectation factor is related to both real activity and investors’ reported expectations from surveys.

Moreover, nothing in the theory implies that hidden factor for one type of asset must also be hidden from the perspective of other types of assets. An important question—but one that is outside the scope of this paper—is whether information from stock and foreign exchange markets can be used to build more accurate term structure models.

References

- Clarida, Richard, Jordi Galí, and Mark Gertler, 2000, Monetary policy rules and macroeconomic stability: evidence and some theory, *Quarterly Journal of Economics* 115, 147-180.
- Cochrane, John H., and Monika Piazzesi, 2005, Bond risk premia, *American Economic Review* 95, 138-160.
- Collin-Dufresne, Pierre, and Robert S. Goldstein, 2002, Do bonds span the fixed income markets? Theory and evidence for unspanned stochastic volatility, *Journal of Finance* 57, 1685-1730.
- Collin-Dufresne, Pierre, and Robert S. Goldstein, 2003, Generalizing the affine framework to HJM and random field models, Working paper, Haas School of Business, University of California at Berkeley.
- Dai, Qiang, and Kenneth J. Singleton, 2000, Specification analysis of affine term structure models, *Journal of Finance* 55, 1943-1978.
- Duffee, Gregory R., 2002, Term premia and interest rate forecasts in affine models, *Journal of Finance* 57, 405-443.
- Duffee, Gregory R., 2008, Forecasting with the term structure: The role of no-arbitrage, Working paper, Haas School of Business, University of California at Berkeley.
- Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379-406.
- Litterman, Robert, and Jose Scheinkman, 1991, Common factors affecting bond returns, *Journal of Fixed Income* 1, 54-61.
- Taylor, John B., 1993, Discretion versus policy rules in practice, *Carnegie-Rochester Conference Series on Public Policy* 39, 195-214.

Table 1. Summary statistics for Treasury yields.

Month-end yields on six zero-coupon Treasury bonds are from CRSP. The sample is 528 observations from January 1964 through December 2007. Yields are continuously compounded and expressed in percent per year. Panel B reports five eigenvalues of covariance matrices. For “Yield levels,” the data are the six yields. For “Monthly changes,” the data are monthly changes in the six yields. For “Annual returns,” the data are overlapping observations of annual returns to the five bonds with initial maturities of one through five years.

Panel A. Univariate statistics

	Maturity					
	3 mon	1 yr	2 yr	3 yr	4 yr	5 yr
Mean	5.87	6.26	6.47	6.64	6.77	6.85
Std dev	2.77	2.74	2.66	2.58	2.53	2.49

Panel B. Principal components

	Index of component				
	1	2	3	4	5
Yield levels	40.405	0.930	0.068	0.008	0.005
Monthly changes	1.120	0.128	0.021	0.008	0.005
Annual returns	108.048	8.072	0.314	0.079	0.047

Table 2. An estimated dynamic term structure model

A length-five state vector x_t has dynamics

$$x_{t+1} = Kx_t + \Sigma\epsilon_{t+1}, \quad \epsilon_{t+1} \sim N(0, I).$$

Yields on bonds with maturities of three months and one through five years are stacked in the vector y_t . The measurement equation is

$$y_t = A + Bx_t + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2 I).$$

The model is estimated with maximum likelihood and the Kalman filter using month-end yields from 1964 through 2007. The factors are normalized to equal the five principal components of yields on bonds with maturities of three months and one, two, four, and five years. The table reports parameter estimates and standard errors. The standard errors are computed from Monte Carlo simulations under the null hypothesis that the estimated model is true.

K	0.987 (0.009)	0.018 (0.053)	0.172 (0.188)	0.987 (0.836)	3.355 (1.442)	
	-0.003 (0.002)	0.936 (0.021)	-0.301 (0.066)	-0.213 (0.250)	-0.033 (0.511)	
	0.001 (0.001)	-0.003 (0.005)	0.820 (0.032)	-0.506 (0.125)	0.065 (0.258)	
	0.000 (0.000)	0.002 (0.002)	-0.026 (0.010)	0.692 (0.053)	0.024 (0.100)	
	0.000 (0.000)	-0.001 (0.001)	0.001 (0.006)	-0.018 (0.037)	0.869 (0.049)	
$\Sigma \times 10$	9.820 (0.297)	0	0	0	0	
	-0.581 (0.363)	3.000 (0.112)	0	0	0	
	0.540 (0.096)	0.656 (0.131)	1.113 (0.054)	0	0	
	-0.005 (0.035)	0.077 (0.037)	-0.258 (0.043)	0.368 (0.040)	0	
	0.029 (0.028)	-0.001 (0.028)	0.044 (0.035)	-0.090 (0.040)	0.228 (0.036)	
	A	B(:,1)	B(:,2)	B(:,3)	B(:,4)	B(:,5)
3 mon	5.740 (1.219)	0.459 (0.024)	-0.655 (0.029)	-0.598 (0.030)	0.102 (0.018)	0.003 (0.016)
1 year	6.164 (1.252)	0.464 (0.015)	-0.314 (0.038)	0.602 (0.023)	-0.567 (0.022)	0.055 (0.071)
2 year	6.446 (1.279)	0.457 (0.004)	0.065 (0.031)	0.405 (0.024)	0.742 (0.037)	-0.268 (0.101)
3 year	6.644 (1.269)	0.443 (0.011)	0.287 (0.023)	0.167 (0.029)	0.433 (0.089)	0.618 (0.108)
4 year	6.808 (1.267)	0.432 (0.018)	0.437 (0.018)	-0.132 (0.032)	0.047 (0.102)	0.776 (0.012)
5 year	6.926 (1.262)	0.422 (0.024)	0.533 (0.022)	-0.136 (0.030)	-0.340 (0.075)	-0.568 (0.048)
$100\sigma_\eta$	5.559 (0.138)					

Table 3. Model-implied population properties of term structure factors

A five-factor Gaussian term structure model is estimated with the Kalman filter. True yields are affine functions of the unobserved factors. Observed yields are contaminated with iid measurement error. The data are month-end yields, from January 1964 through December 2007, on zero-coupon bonds with maturities of three months and one through five years. The factors are rotated to represent, in order, the first five principal components of the bond yields (expressed in percent per year). The first column of the table reports the population standard deviations of the factors. Standard errors, computed from Monte Carlo simulations, are in parentheses. The second column reports the population R^2 of a regression of the true, unobserved factors on contemporaneous values of all six observed bond yields. The third column reports the population R^2 of similar regressions using filtered estimates of the factors in place of the true factors. The fourth column reports population correlations between true factors and filtered estimates of the factors.

Factor	Std dev	R^2 's of factors on yields		Correl of true, filtered factor
		True factors	Filtered factors	
1	6.017 (1.287)	1.000	1.000	1.000
2	0.925 (0.114)	0.997	1.000	0.998
3	0.251 (0.019)	0.954	0.993	0.980
4	0.067 (0.005)	0.623	0.877	0.843
5	0.043 (0.004)	0.433	0.683	0.797

Table 4. Decomposition of volatility of 12-month-ahead yield forecasts

A five-factor Gaussian term structure model is estimated with the Kalman filter. The factors represent, in order, the first five principal components of the bond yields and are unconditionally uncorrelated. Parameter estimates are used to construct estimates of unconditional variances of 12-month-ahead expectations of bond yields, expressed in percent per year. These variances are the sums of estimated variances attributable to each of the five factors. The table reports the square roots of these estimated variances. Monte Carlo simulations are used to compute biases and uncertainty in these estimates. Using the null hypothesis that the estimated model is correct, the term structure model is estimated using simulated yields. Means and ninety-five percentile bounds on the estimated standard deviations are reported in parentheses and brackets respectively.

Maturity	Std dev of	Std dev attributable to factor				
	forecast (%/year)	1	2	3	4	5
3 mon	2.28	2.23	0.28	0.07	0.06	0.36
	(1.98)	(1.91)	(0.28)	(0.13)	(0.10)	(0.31)
	[0.86 3.40]	[0.80 3.36]	[0.01 0.65]	[0.00 0.40]	[0.00 0.32]	[0.03 0.63]
1 yr	2.32	2.28	0.15	0.04	0.07	0.37
	(2.00)	(1.94)	(0.19)	(0.12)	(0.10)	(0.32)
	[0.90 3.52]	[0.81 3.48]	[0.00 0.46]	[0.00 0.38]	[0.00 0.33]	[0.03 0.64]
2 yr	2.34	2.31	0.02	0.07	0.04	0.38
	(2.01)	(1.96)	(0.13)	(0.11)	(0.09)	(0.33)
	[0.89 3.48]	[0.82 3.46]	[0.01 0.34]	[0.00 0.34]	[0.00 0.30]	[0.06 0.61]
3 yr	2.33	2.28	0.12	0.14	0.01	0.37
	(1.98)	(1.93)	(0.12)	(0.14)	(0.08)	(0.32)
	[0.94 3.40]	[0.86 3.37]	[0.00 0.36]	[0.00 0.36]	[0.00 0.28]	[0.06 0.58]
4 yr	2.32	2.27	0.20	0.19	0.01	0.37
	(1.97)	(1.91)	(0.15)	(0.18)	(0.08)	(0.32)
	[0.91 3.35]	[0.82 3.32]	[0.00 0.43]	[0.01 0.40]	[0.00 0.26]	[0.07 0.56]
5 yr	2.30	2.24	0.25	0.24	0.03	0.37
	(1.95)	(1.88)	(0.19)	(0.21)	(0.08)	(0.33)
	[0.86 3.30]	[0.79 3.26]	[0.01 0.49]	[0.01 0.43]	[0.01 0.25]	[0.08 0.57]

Table 5. Model-implied properties of annual excess bond returns

A five-factor Gaussian term structure model is estimated with the Kalman filter. The factors represent, in order, the first five principal components of the bond yields and are unconditionally uncorrelated. Parameter estimates are used to calculate population properties of annual log returns to bonds in excess of the log return to a one-year bond. In Panel A, return variances are calculated for both true excess returns and observed excess returns. The latter are contaminated by measurement error. The columns labeled “Predictable frac of var” report the fraction of the variance attributable to time-variation in conditional means of true returns. Panel B decomposes the volatility of true conditional expected excess returns into components attributable to each factor. Its structure follows Table 3.

Panel A. Univariate statistics

Maturity	Mean	True returns		Observed returns	
		Std dev	Predictable frac of var	Std dev	Predictable frac of var
2 yr	0.36	1.78	0.196	1.78	0.195
3 yr	0.68	3.24	0.204	3.24	0.203
4 yr	0.87	4.50	0.221	4.51	0.220
5 yr	0.88	5.58	0.206	5.59	0.205

Panel B. Decomposition of volatility of expected excess returns

Maturity	Std dev of conditional mean (%/year)		Std dev attributable to factor				
	1	2	3	4	5		
2 yr	0.79	0.43	0.56	0.02	0.07	0.35	
	(0.82)	(0.44)	(0.55)	(0.11)	(0.11)	(0.30)	
3 yr	[0.50 1.17]	[0.12 0.72]	[0.18 0.94]	[0.00 0.33]	[0.01 0.31]	[0.02 0.61]	
	1.46	0.57	1.04	0.11	0.05	0.84	
4 yr	(1.54)	(0.64)	(1.05)	(0.22)	(0.19)	(0.73)	
	[0.98 2.17]	[0.14 1.13]	[0.34 1.69]	[0.00 0.64]	[0.01 0.58]	[0.18 1.28]	
5 yr	2.12	0.74	1.54	0.13	0.02	1.24	
	(2.22)	(0.84)	(1.55)	(0.30)	(0.26)	(1.09)	
5 yr	[1.47 3.02]	[0.19 1.46]	[0.59 2.44]	[0.00 0.90]	[0.01 0.84]	[0.31 1.85]	
	2.53	0.82	1.97	0.23	0.03	1.35	
5 yr	(2.68)	(0.96)	(2.00)	(0.38)	(0.31)	(1.16)	
	[1.75 3.69]	[0.14 1.71]	[0.79 3.10]	[0.01 1.10]	[0.01 0.97]	[0.20 2.12]	

Table 6. Model-implied expectations compared to survey forecasts

Quarterly observations of expectations of future Treasury bill yields are from the Survey of Professional Forecasters. The data used are quarter- t mean survey forecasts of the three-month T-bill yield during quarters $t+j$, $j = 1, \dots, 4$. The contemporaneous three-month yield is subtracted from the forecasts to produce forecasted changes in the yield. Contemporaneous filtered estimates of the “level” and “expectation” factors are taken from a five-factor term structure model. The factors are normalized to have standard deviations of one. All regressions are estimated from 1981Q3 through 2007Q4 (106 quarterly observations). Newey-West standard errors are in parentheses, adjusted for four lags of moving average residuals.

Panel A. Regressions of the expectation factor on the survey-based expected change

	Quarters ahead (j)			
	1	2	3	4
Coef	-1.154 (0.206)	-0.848 (0.207)	-0.612 (0.189)	-0.466 (0.161)
AR(1) of residual	0.55	0.58	0.58	0.59

Panel B. Regressions of the survey-based expected change on the level and expectation factors

	Quarters ahead (j)			
	1	2	3	4
Coef on level	-0.108 (0.036)	-0.147 (0.047)	-0.205 (0.059)	-0.263 (0.075)
Coef on expectation	-0.135 (0.043)	-0.155 (0.050)	-0.149 (0.062)	-0.158 (0.080)
AR(1) of residual	0.22	0.47	0.54	0.57

Table 7. The relation between industrial production and the expectations factor

The log change industrial production from month $t - 1$ to month t is regressed on the month $t - i$ filtered estimate of the expectations factor, for $i = -6, \dots, 6$. The log change is expressed in percent and the factor is normalized to have a standard deviation of one. Newey-West standard errors are calculated using two lags of moving average residuals. The sample period is 1964 through 2007.

Lead of $\Delta \log(\text{IP})$	Coef	Std error	t -statistic
-6	-0.025	0.048	-0.52
-5	-0.096	0.050	-1.92
-4	-0.098	0.051	-1.92
-3	-0.129	0.050	-2.56
-2	-0.140	0.049	-2.85
-1	-0.118	0.050	-2.36
0	-0.135	0.054	-2.48
1	-0.112	0.057	-1.97
2	-0.089	0.054	-1.66
3	-0.060	0.047	-1.26
4	-0.070	0.049	-1.43
5	-0.038	0.052	-0.74
6	-0.078	0.051	-1.52

Table 8. Population and finite-sample properties of predictive regressions

Excess log returns to a five-year bond from month t to month $t + 12$ are predicted with the month- t shape of the term structure using two OLS regressions. The first regression uses the level, slope, and curvature of the term structure, as defined in the text. The second regression uses five forward rates. The true data-generating process is this paper's estimated term structure model. Population values of the coefficients and R^2 s are calculated analytically. Finite-sample properties use simulations of 528 months of bond yields. The table summarizes results from 10,000 simulations. The notation $F(m, n)$ denotes the forward rate from year m to year n .

Panel A. Predicting excess returns with level, slope, and curvature

	Coefficient on:			R^2
	Level	Slope	Curvature	
Population value	0.429	1.520	0.592	0.14
Mean across sims	0.847	1.631	0.171	0.21
Std dev across sims	0.476	0.783	2.805	0.08
95 percent interval	[0.12 1.96]	[0.06 3.14]	[-5.45 5.66]	[0.06 0.38]

Panel B. Predicting excess returns with five forward rates

	Coefficient on:					R^2
	$F(0, 1)$	$F(1, 2)$	$F(2, 3)$	$F(3, 4)$	$F(4, 5)$	
Population value	-1.861	0.465	2.658	0.289	-1.096	0.16
Mean across sims	-1.786	0.450	2.627	0.446	-0.869	0.22
Std dev across sims	0.952	1.661	1.053	0.850	0.876	0.08
95 percent interval	[-3.63 0.12]	[-2.85 3.58]	[0.63 4.75]	[-1.20 2.15]	[-2.62 0.85]	[0.08 0.39]

Panel C. Probability distribution of the difference in finite-sample R^2 s

	Percentile				
	2.5%	25%	50%	75%	97.5%
Forward-rate R^2 less 3-factor R^2	-0.041	-0.002	0.014	0.033	0.080

Table 9. Finite-sample forecast accuracy of predictive regressions

Excess log returns to m -year bonds from month t to month $t+12$, $m = 1, \dots, 5$, are predicted with the month- t shape of the term structure using two OLS regressions. Regression [1] uses level, slope, and curvature. Regression [2] uses five forward rates. The true data-generating process is this paper's estimated term structure model. Finite-sample properties use simulations of 528 months of bond yields. The table summarizes monthly differences between fitted month- t forecasts and true month- t expectations of expected excess returns. For each simulation, the square root of the mean squared difference, denoted RMSE, is calculated for each regression. The table summarizes results from 10,000 simulations. All values are in percent per year.

Bond maturity	Std dev of true expectation	Mean RMSE across simulations		Fraction of sims with RMSE[1]>RMSE[2]
		Reg [1]	Reg [2]	
2 years	0.722	0.569	0.535	0.83
3 years	1.385	1.163	1.069	0.92
4 years	2.016	1.660	1.511	0.95
5 years	2.412	1.931	1.825	0.84




Fig. 1. Estimated loadings of yields on the five factors of a term structure model. Each line represents the response of the term structure to a one standard deviation variation in the given factor.

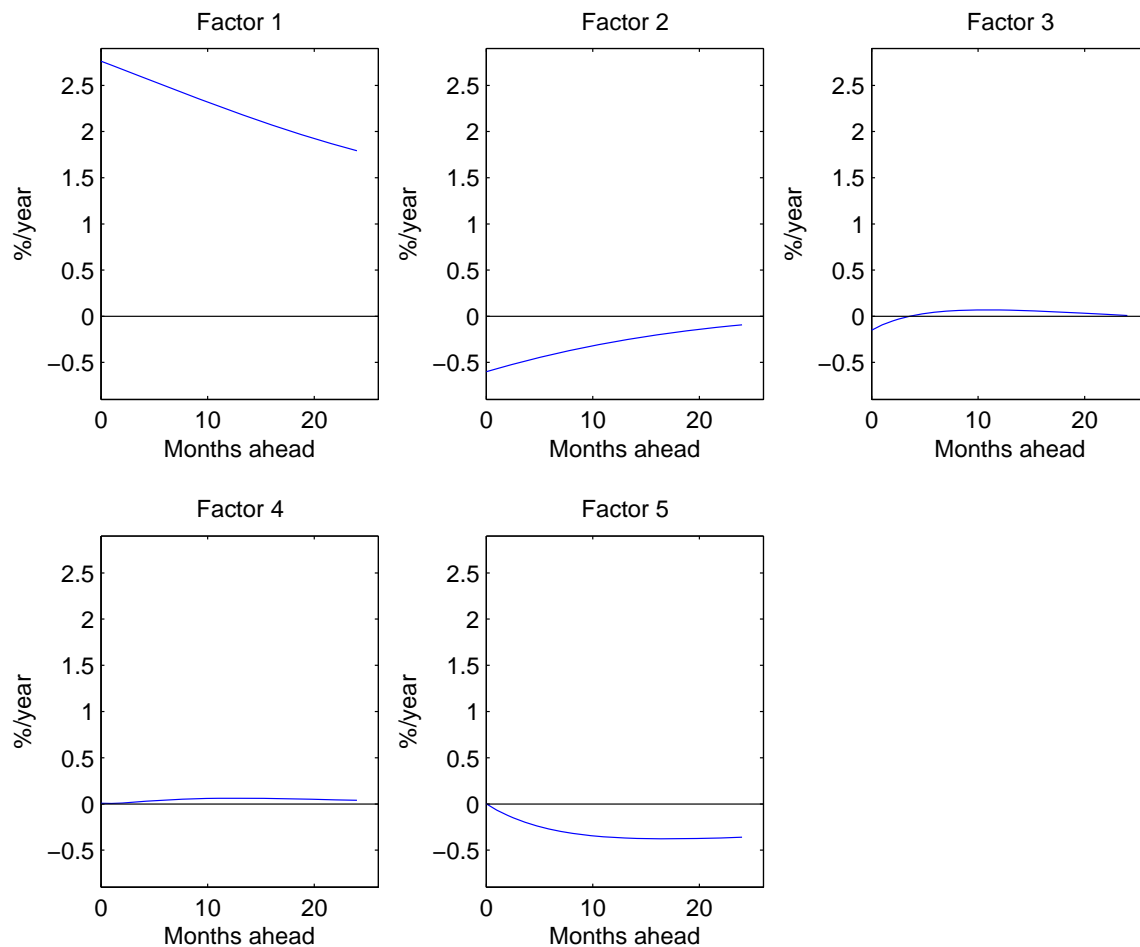


Fig. 2. Responses of the three-month bill rate to term structure factors. Each panel plots the expected time path of the three-month bill yield, assuming that at month zero the specified factor is one standard deviation above its mean. All other factors are set to their unconditional means.




Fig. 3. Filtered estimates of the “expectation” factor. The vertical lines are NBER business cycle break points.




Fig. 4. Sensitivity of expected excess bond returns to term structure factors. The month- t expected annual log return to a m -year bond less the log return to a one-year bond depends on the month- t values of the term structure factors. The figure plots, for $m = 2$ through $m = 5$, the sensitivity of the expected excess return to one-standard-deviation changes in the “level” factor (solid line), the “slope” factor (dashed line), and “expectation” factor (dotted line).

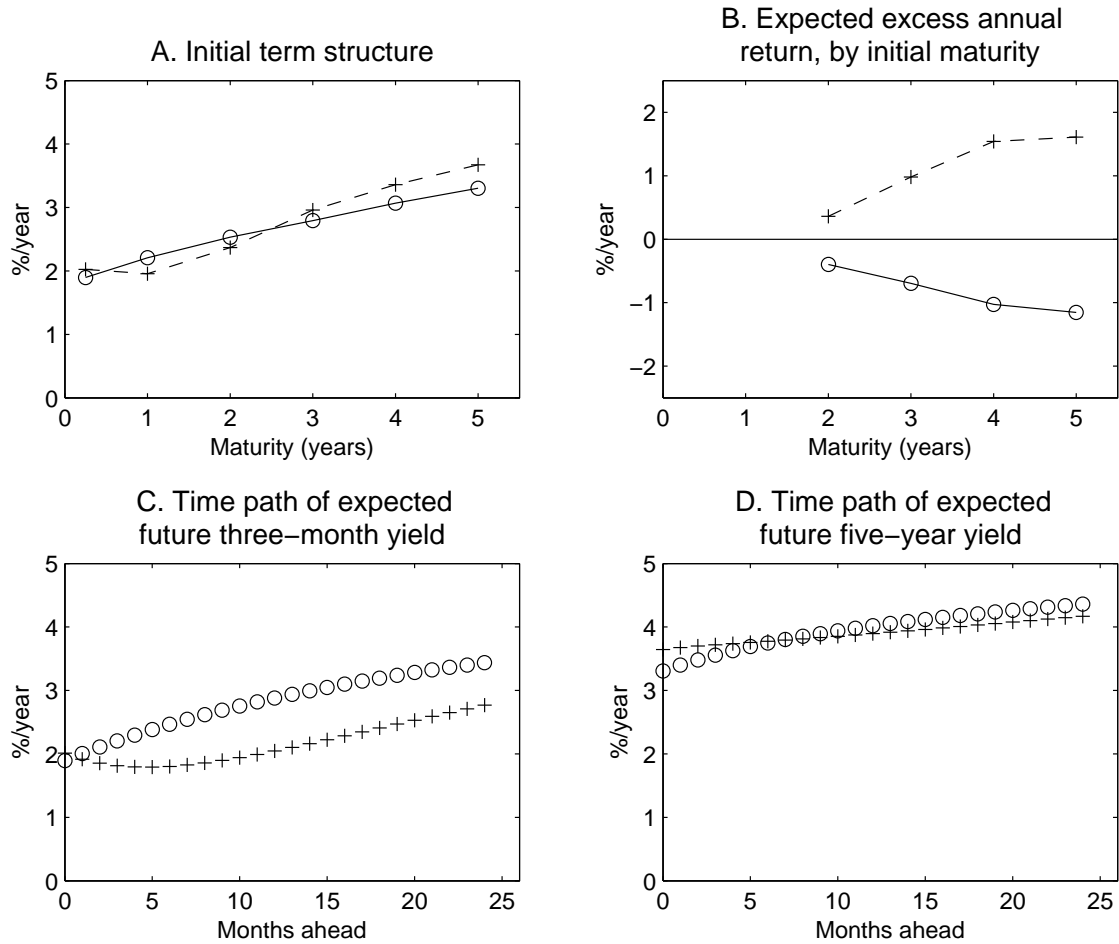


Fig. 5. A comparison of October 2001 and October 2004. Values for the two months are plotted with '+' and 'o' respectively. Panel A displays the month-end term structures. Panel B displays model-implied expected excess log returns (over the one-year yield) for bonds with maturities of two through five years. Panel C displays expected future three-month yields over the next 24 months, where month zero is October of 2001 and 2004 respectively. Panel D displays expected future five-year yields.