Information in (and not in) the term structure

Gregory R. Duffee[*]

Johns Hopkins

First draft: March 2008
This Draft: August 7, 2009

ABSTRACT

Standard approaches to building and estimating dynamic term structure models rely on the assumption that yields can serve as the factors. However, the assumption is neither theoretically necessary nor empirically supported. This paper documents the presence of a factor that has an almost imperceptible effect on the cross section of yields, but has strong forecast power for future short-term interest rates and excess bond returns. The factor appears to be related to short-run fluctuations in economic activity.

[*]Voice 410-516-8828, email duffee@jhu.edu. Address correspondence to 440 Mergenthaler Hall, 3400 N. Charles St., Baltimore, MD 21218. The first version of this paper was written while the author was visiting Wharton. Thanks to Pierre Collin-Dufresne, Wayne Ferson, Scott Joslin, Sydney Ludvigson, Jacob Sagi, Dan Thornton, and seminar participants at the Federal Reserve Board, USC, Wharton, the 2008 NBER Summer Institute, and the 2009 WFA Annual Meetings for helpful conversations and comments.

1 Introduction

This paper advocates a significant change in the construction and estimation of multifactor term structure models. In a literature spanning more than two decades, researchers have almost universally assumed that the factors driving term structure dynamics can be represented as functions of yields. The assumption plays a critical role in all aspects of estimation. However, because it rules out a potentially important class of term structure dynamics, we need research methodologies that do not rely on the assumption.

The intuition behind the standard approach is so obvious that it is seldom mentioned. Investors' beliefs about future bond prices determine what investors are willing to pay for bonds today. This suggests that today's term structure contains all information relevant to predicting both future returns to bonds and future bond yields. Put somewhat differently, the term structure follows a Markov process.

Empirical work exploits this Markov structure in many ways. It helps researchers choose the dimension of a model, because the same factors that determine the cross section of yields also determine yield dynamics. Therefore factor analysis of unconditional covariances among yields (the cross section) pins down the length of the state vector. It also simplifies considerably the search for time-varying expected bond returns, because it implies that time-t conditional expectations of returns can be expressed entirely in terms of forward rates observed at t. Other data are unnecessary to model yield dynamics. In addition, the one-to-one mapping from factors to yields implied by the Markov structure leads to tractable estimation of very complicated term structure models.

Yet recent empirical evidence calls this assumption into question. Ludvigson and Ng (2009) and Cooper and Priestly (2008) conclude that various measures of macroeconomic activity contain information about future excess bond returns that is not in forward rates. Cochrane and Piazzesi (2005) find that lagged forward rates contain information about future excess bond returns that is not in current forward rates. One possible explanation, as noted by Cochrane and Piazzesi, is that measurement error in yields obscures the Markov structure.

In other words, these empirical results hinge on our inability to precisely observe yields. But plausible measurement error in Treasury yields is on the order of only a few basis points. Thus it is incumbent upon us to attempt to understand, from a formal perspective, why tiny measurement errors can cover up important information contained in the cross section of yields.

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. I refer to such a factor as a "hidden" factor, in the sense that a snapshot of the time-t yield curve conveys no information about its level. A hidden factor has 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.

This theoretical result, although not well-known, can be inferred from the existing term structure literature. Duffee (2002) contains an example in which the physical and equivalent-martingale dynamics are driven by state vectors with different dimensions. But its implications for empirical work have not been recognized until this paper, and contemporaneous and independent work by Joslin, Priebsch, and Singleton (2009). We take this idea in different directions. In a nutshell, I use a filtering approach to ask whether there are hidden factors. Their work assumes the existence of two hidden factors that are linear combinations of observed inflation and industrial production, and estimate the resulting model using both yield and macroeconomic data.

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 factor that has an almost imperceptible 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 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 factor accounts for about 30 percent of the total variance in expected excess bond returns, yet there is no linear combination of yields that captures most of the variation in the factor.

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 hidden factor is spurious. However, evidence from the Survey of Professional Forecasters confirms that survey-based expectations of future short rates move contemporaneously with 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.

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

2 The modeling framework

This section explains why some important determinants of the yield dynamics may be undetectable in the cross section. 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 yields.

I follow much of the modern term structure literature by abstracting from standard economic concepts such as utility functions and production technologies. 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 A 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 2.6. 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

$$x_{t+1} = \mu + K x_t + \Sigma \epsilon_{t+1}, \quad \epsilon_{t+1} | x_t \sim N(0, I). \tag{2}$$

The state vector is latent, hence identifying restrictions are typically imposed in estimation. A convenient normalization is described in the paper's empirical section.

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

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

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

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

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

$$\Sigma \Lambda_t = \lambda_0 + \lambda_1 x_t, \tag{5}$$

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

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

where the equivalent martingale parameters are

$$\mu^q = \mu - \lambda_0, \quad K^q = K - \lambda_1. \tag{7}$$

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

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

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

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

2.2 Information in the cross section

Absent 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

$$y_t^a = A^a + B^a x_t \tag{10}$$

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

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

Since x_t follows a first-order Markov process, the term structure of yields also follows a first-order Markov process.

Although this result is derived here in a Gaussian setting, it applies more generally to the class of affine term structure models. The entire empirical literature on dynamic term structure models (setting aside the current paper and Joslin et al. (2009)) takes it for granted. For example, the handbook treatment of Piazzesi (2009) does not mention that B^a may not invertible. The next subsection explains why, from an empirical perspective, invertibility is a very useful property.

2.3 The role of invertibility in empirical analysis

Invertibility allows us to infer the dimension of the state vector n from properties of the cross section of yields. One method, introduced by Stambaugh (1988), studies the predictability of excess bond returns. He infers n by using a condition equivalent to (11): conditional expectations of excess bond returns are functions of n forward rates. This methodology remains at the leading edge of the literature through Cochrane and Piazzesi (2005). Another method to infer n is factor analysis of the unconditional covariance matrix of yields or differenced yields. Litterman and Scheinkman (1991) conclude three factors explain, in a statistical sense, all but a negligible fraction of the variation in the term structure. Duffee (2002) and Brandt and Chapman (2003) use this result and (11) to justify the choice of $n = 3$.

Equation (11) implies that maximum likelihood estimation of affine term structure models requires only a panel of n yields and the density function of the state vector.[1] In fact, Piazzesi (2009) defines likelihood-based estimation of affine models in terms of (11). Pearson and Sun (1994) are the first to exploit this result. Chen and Scott (1993) expand the panel's cross section to d yields by assuming that n linear combinations of yields are observed without error and $d - n$ are observed with error. In the special case of Gaussian models, maximum likelihood estimation is also feasible when all d yields are observed with measurement error. Yet even with Gaussian models, estimation is simplified considerably when factors are treated as linear combinations of yields. Cochrane and Piazzesi (2008) and Joslin, Singleton, and Zhu (2009) are recent applications that use (11) to estimate Gaussian models.

Invertibility implies that only yields are necessary to estimate affine models, but it does not rule out the use of other data. Ang and Piazzesi (2003) introduced macroeconomic variables into Gaussian term structure models, leading to an explosion of macro-finance research. This literature is not designed to produce more accurate term structure models, but rather to explicitly link the term structure to its fundamental determinants, such as inflation and monetary policy.

Although invertibility is widely assumed and useful, it need not hold. I now consider special cases of the 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

[1] The likelihood function for discretely observed observations may need to be evaluated numerically.

and the economic intuition.

2.4 A two-factor example

Consider a 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)

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

When k_{12} does not equal 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 are risk-neutral, the level of f_t necessarily affects 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 $k^q_{12} = 0$, or $k_{12} = \lambda_{1(12)}$. Then the equivalent martingale dynamics of the state are

$$\begin{pmatrix} r_{t+1} \\ f_{t+1} \end{pmatrix} = \mu^q + \begin{pmatrix} k^q_{11} & 0 \\ k^q_{21} & k^q_{22} \end{pmatrix} \begin{pmatrix} r_t \\ f_t \end{pmatrix} + \Sigma \epsilon^q_{t+1}. \tag{13}$$

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),

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

Thus the matrix B^a in (10) cannot be inverted because it has a column of zeros. The factor f_t is hidden, in the sense that it has no effect on the period-t term structure. 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.

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 does not change with f_t, this news raises 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

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

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

2.5 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

$$K^q = PVP^{-1} \tag{16}$$

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

$$x_t^* = Px_t. \tag{17}$$

The equivalent martingale dynamics of the rotated state vector are

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

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

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

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.

2.6 The role of the Gaussian setting

Section 2.5 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 exhibiting 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.

In this model, 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 of a USV model, the short rate is determined by a n-dimensional state vector that follows a 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 parameter restrictions differ substantially. In this model, 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 those driving 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.

One reason I use the Gaussian framework is to avoid complications associated with stochastic volatility. Reconsider the two-factor example of Section 2.4. 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.

2.7 From theory to practice

There is no measurement error in the yield equation (8). Therefore the parameter restriction (20) is knife-edge. If the parameter differs from zero by an arbitrarily small amount, then the exact mapping from factors to n yields in (10) implies that all factors can be inferred from the cross section using (11).

However, in real-world data the mapping is not exact. Equation (10) implies that the unconditional covariance matrix of $d > n$ bond yields has a rank no greater than n. (It equals n in the standard case and $n - 1$ in when (20) holds.) 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 $m_1, \ldots, m_d$ is expressed as

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

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 and is roughly uncorrelated across the widely-spaced maturities used in empirical analysis.

Measurement error eliminates the knife-edge nature of (20). A factor may be an important determinant of expected future yields, yet have an infinitesimally small effect on the current term structure—so small that the effect is lost in the noise η_t in (22). One way to measure the extent to which factor i is hidden is the regression

$$x_{i,t} = b_0 + b_1' y_t + \eta_{i,t}. \tag{23}$$

Absent measurement error, the R^2 of this regression is either one (the standard case) or zero (i is a hidden factor). An econometrician cannot estimate (23) because the state vector is unobserved, but population values coefficients can be computed from an estimated dynamic term structure model.

2.8 Implications for term structure estimation

If we take seriously the possibility that more factors affect yield dynamics than affect the cross section of yields, how should we estimate dynamic term structure models? One requirement is to build necessary flexibility into the model through the dimension of the state vector. Three state variables are needed to describe the cross section of Treasury yields. Thus if there is some variable hidden from the cross section, a model without at least four state variables is misspecified.

Given a sufficiently-flexible model, there are two broad paths to follow. They differ primarily in the data used in estimation. The direction taken in this paper is to infer the presence of hidden factors from the dynamics of yields, which we can call a "yields-only" approach. Alternatively, yield data can be augmented by other data that contain independent information about factors that drive yield dynamics, which I call a "yields-plus" approach.

Yields-only estimation of models with hidden factors can be done with filtering. Pennacchi (1991) introduces filtering into affine term structure estimation. The usual motivation, as noted in Piazzesi (2009), is to extract information about the period-t state vector from the entire period-t cross section, thus avoiding the ad hoc assumption that exactly n yields are observed without error. But filtering also uses dynamics to infer this vector. Intuitively, filtering is equivalent to learning by the econometrician. The period-t forecast error (the difference between realized yields and the econometrician's $t-1$ forecast of yields) is produced by both true period-t shocks and the error in the econometrician's $t-1$ prediction of the $t-1$ state vector. The cross sectional pattern of the period-t forecast errors helps the econometrician revise her prediction of the state vector at $t-1$ and form her prediction of the state vector at t.

In adition, it may be possible to infer the period-t state vector from period-t observations of non-yield data. Recall that hidden factors have equal and opposite effects on expected future short-term interest rates and risk premia. Data that depend separately on these two components (or weight them differently) will reveal such factors. Perhaps the most obvious choice is survey data on interest rate forecasts, such as that used by Kim and Orphanides

(2005).[2] In line with the literature's recent focus on macro-finance models, Joslin et al. (2009) use inflation and output growth. They assume two hidden factors are linear combinations of these variables.

The tradeoffs between a yields-only and a yields-plus approach are straightforward. Estimation using additional data is a more powerful approach, but also at greater risk of misspecification. Holding the sample length constant, and under the maintained hypotheses that (1) the additional data are functions only of state vector that drives bond yields, and (2) the data reveal otherwise hidden factors, the yields-plus approach will produce more precise estimates of the term structure model. It is much easier to infer period-t factors from direct observation of variables at t than from teasing them out of dynamics.

However, samples of survey data are shorter than samples of bond yields. Long time series of macro data are available, but the requirement that the macro variables are spanned by the variables that drive yields is often problematic. In particular, the relation between the macroeconomy and time-varying risk premia is an active area of research, but one with with few uncontroversial conclusions. The yields-only approach does not take a stand on the relation between the term structure and the macroeconomy, and thus avoids the possibility of misspecifying the relation.

The empirical analysis in this paper is a bit of a hybrid. The next section uses a long sample of yields to estimate a term structure model. Section 4 links the hidden factor uncovered through this estimation to both a shorter sample of surveys of interest rate forecasts and a long sample of industrial production data.

3 Empirical analysis

This section estimates a five-factor Gaussian term structure model using only yield data. The main question is whether there are factors that have little or no effect on the cross

[2]Kim and Orphanides use a three-factor model, which does not give them the flexibility to capture both the cross section and potential hidden factors.

section of yields yet are important for modeling dynamics. The conclusion is that a hidden factor drives a substantial fraction of the predictability of excess returns, although there is considerable statistical uncertainty in the estimates. Section 4 confirms the importance of the factor using non-yield data.

3.1 The choice of five factors

Section 2.8 argues that, since three factors are needed to explain the cross section of yields, at least four factors are necessary to uncover the presence of a hidden factor. However, Cochrane and Piazzesi (2005) find that information from five points on the yield curve to form forecasts of excess bond returns. If we take this result literally, a six-factor model is needed here. Unfortunately, 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. Moreover, it is difficult to convince a skeptic that they have something to learn from a model with 52 parameters. 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 An unrestricted model

The model is estimated without imposing any parameter restrictions that produce hidden factors. The main reason for estimating an unrestricted model is that it is hard to ask whether there are hidden factors using a model that *a priori* imposes their existence. In addition, such parameter restrictions are difficult to motivate on economic grounds. It seems quite unlikely that there truly is a factor for which variations in expected future short rates are *exactly* offset by variations in required expected returns.

The more important practical question is whether there are factors that have effects on yields that are indistinguishable from measurement error. It is easy to tell stories in which news has 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 is ambiguous and might be very close to zero.

3.3 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).

An alternative source of zero-coupon bond yields is a panel produced by the Federal Reserve Board. The advantage of the Fed data is that yields are available for maturities greater than five years. The crucial disadvantage, at least for the purposes of this paper, is that Gurkaynak, Sack, and Wright (2006) produce yields by fitting a smooth function to the term structure. Cochrane and Piazzesi (2008) find that this smoothing removes a significant component of the forecasting power of the term structure. In some preliminary work with only the Federal Reserve yields that match the maturities of the CRSP yields, I confirm their result and find that the smoothing also weakens the evidence for a hidden factor. (These results are not reported in detail in the paper.)

Panel A of Table 1 reports means and standard deviations of the CRSP 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. Note that in the table, yields are expressed in percent per year, while the model is written in terms of decimal points per month. This transformation is used in all of the tables.

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. The first three are the level, slope, and curvature of the term structure. Section 3.6 takes a detailed look at these shapes.

3.4 An identifying normalization

The factors can be arbitrarily rotated. To help interpret the extent to which any factors are hidden from the cross section, I rotate them to equal principal components of uncontaminated yields—yields without measurement error. Denote the vector of six observed yields by y_t and its uncontaminated counterpart by $\tilde{y}_t$. Drop the three-year bond from the latter vector, denoting the vector of remaining five uncontaminated yields by $\tilde{y}_{\backslash 3,t}$.

The estimated model will imply a population covariance matrix of the vector $\tilde{y}_{\backslash 3,t}$. Because the parameter restriction for a hidden factor is not imposed, this matrix will nonsingular. Diagonalize it into

$$\mathrm{Var}\left(\tilde{y}_{\backslash 3,t}\right) = C_0 \Omega C_0^{-1} \tag{24}$$

where Ω is the diagonal matrix of the covariance eigenvalues.

The state vector x_t is normalized to equal the vector of mean-zero principal components of $\tilde{y}_{\backslash 3,t}$, ordered from largest to smallest. Its dynamics are

$$x_{t+1} = K x_t + \Sigma \epsilon_{t+1}, \tag{25}$$

where

$$\mathrm{Var}(x_t) = \Omega. \tag{26}$$

The appendix explains how the normalization is imposed in estimation.

3.5 Model estimation

Estimation is with the Kalman filter, which produces correct conditional means and covariances in a Gaussian setting. The transition equation is (25) and the measurement equation is (22).

Table 2 reports the point estimates of these two equations. 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 equation (26), 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.

Note that Table 2 does not report estimates of the short-rate equation (1) or the equivalent-martingale parameters of (6). They are not of direct interest here, and are reported in the appendix.

3.6 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.[3] 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 annualized 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 less than half a basis point of monthly yields, or about about five and a half basis points of annualized yields. Although economically small, this measurement error is enough to obscure the effects of these factors on yields. One way to see this is through the R^2 of the population regression of factors on yields in equation (23). The point estimates of the model allow analytic calculation of the R^2s.

The second column of statistics in Table 3 reports the R^2s for each factor. The effects of the first three factors on yields are sufficiently large to dominate measurement error. The R^2s for these factors range from 1.0 to 0.95. However, the R^2s for the fourth and fifth factors

[3]Recall that the model is estimated using yields expressed in decimal form per month. The values here have been multiplied by 1200 to put them in terms of percent per year.

are only 0.62 and 0.43 respectively. Hence most of the variation in the fifth factor cannot be explained by contemporaneous yields; it is largely hidden from the cross section.

Because it uses information from dynamics, 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 three 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 smoothed estimates of the factors. These correlations are 0.87 and 0.86 for the fourth and fifth factors. Naturally, smoothed estimates of the factors are more closely related to observed yields than are true factors (since observed yields are used in the smoothing), 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.

3.7 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)

$$\begin{aligned} E_t\left(y_{t+j}\right) &= A + BE_t\left(x_{t+j}\right) \\ &= A + BK^j x_t. \end{aligned} \tag{27}$$

The unconditional covariance matrix of these forecasts is

$$\operatorname{Var}\left(E_t(y_{t+j})\right) = BK^j \Omega B'(K^j)'. \tag{28}$$

Because the unconditional covariance matrix of the factors Ω is diagonal, the variance in (28) 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. Because the factor has no immediate effect on the term structure, I refer to this fifth factor as the hidden 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, as noted above, the 'true' model implies a standard deviation of three-month yield forecasts of 2.28 percentage points. The mean standard deviation from the Monte Carlo simulations is only 1.86 percentage points, as displayed in parentheses. The 2.5 and 97.5 percentile values are 0.87 and 3.15 percentage points respectively, as displayed in brackets.

There is substantial statistical uncertainty about the role of the hidden 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 hidden 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 30 basis points. A 95 percent confidence interval ranges from 3 to 64 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 hidden factor. Below I also draw on evidence from the Survey of Professional Forecasters and the growth of industrial production.

Because the hidden 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 smoothed 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.

3.8 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. Expressed in terms of yields, the observed excess return is

$$xr_{t,t+j}^{(m)} \equiv my_t^{(m)} - (m-j)y_{t+j}^{(m-j)} - jy_t^{(j)}. \tag{29}$$

Using the model's description of yield dynamics, this return is

$$\begin{aligned} xr_{t,t+j}^{(m)} \quad = \quad & mA_m - (m-j)A_{m-j} - jA_j \\ & + \left(mB'_m - (m-j)B'_{m-j}K^j - jB'_j\right) x_t \\ & -(m-j)B'_{m-j}\left(\sum_{i=1}^{j} K^{j-i}\epsilon_{t+i}\right) \\ & +m\eta_t^{(m)} - j\eta_t^{(j)} - (m-j)\eta_{t+j}^{(m-j)}. \end{aligned} \tag{30}$$

The four lines on the right side of (30) 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 20 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 fifth, hidden factor explains up to 30 percent of the predictable variance. Again, a glance at the two figures explains why. The hidden factor lowers expected future short rates while leaving long-term yields unchanged.

Table 5 documents substantial statistical uncertainty about the contribution of the hidden factor to expected excess returns. This mirrors the results for yield dynamics in Table 4. For example, when the hidden 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 1.15 percentage points. A 95 percent confidence interval ranges from 24 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 hidden 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 2.5. 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 hidden 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 Additional evidence of the hidden factor

Is the estimated hidden 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 investors' actual 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 hidden term structure factor. Here, "contemporaneous" means the smoothed estimate for the end of the first month in the quarter.

If the smoothed factor is spurious, forecasters' contemporaneous expectations should be unrelated to it. For example, assume the quarter-t level of the smoothed hidden 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 hidden 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 4 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 smoothed estimates of the fifth factor are not.[4] The October 2001 estimate of this factor is about 0.8 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 1.7 percent for the five-year bond. In 2004, the expectations are negative, ranging from −0.4 percent to −1.2 percent. Differences in expected excess returns are 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 65 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 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

[4]In particular, the months were not chosen based on the contemporaneous survey forecasts.

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 particpants) 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 hidden 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 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's yield convention. Denote quarter-t smoothed estimates of the hidden factor as MODEL_HIDDEN$_t$. Following the timing convention of yields, I define the quarterly factor as the smoothed estimate 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

$$\text{MODEL\_HIDDEN}_t = b_0 + b_1 \text{SPF\_EXPECT}(t, j) + e_{j,t} \tag{31}$$

for forecast horizons of one through four quarters ($j = 1, \ldots,$). Under the null hypothesis that the smoothed estimate of the hidden 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 smoothed estimate of the hidden 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 hidden 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 (31), 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 -3.0 to -4.2. The point estimates are less than the model predicts, ranging from -0.5 to -1.3. 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 smoothed values of the level and hidden factors is very close to zero. But from 1981Q3 through 2007Q4, the sample correlation is about 0.27. As Figure 2 shows, both the level and hidden 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 (31) 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 (31) and add the estimated level factor as an additional explanatory variable. The regression is

$$\text{SPF\_EXPECT}(t, j) = b_0 + b_1 \text{MODEL\_LEVEL}_t + b_2 \text{MODEL\_HIDDEN}_t + e_{j,t}. \tag{32}$$

Both explanatory variables are generated regressors. Because the hidden 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 hidden factor and survey expectations does not disappear when the level factor is included. The asymptotic t statistics for the coefficients on the hidden factor range between -2.2 and -2.6.

This evidence supports the model's conclusion that the hidden 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 hidden factor can be linked to the business cycle.

4.3 The hidden factor and economic activity

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

$$100(\log(\text{IP}_t) - \log(\text{IP}_{t-1})) = b_{0,i} + b_{1,i} \text{MODEL\_HIDDEN}_{t-i} + e_{t,i}, \quad i = -6, \ldots, 6. \tag{33}$$

The change in IP lags the hidden 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 (33) 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 hidden factor is normalized by its population standard deviation.

Estimation results are in Table 7. There is an inverse relation between industrial production and the hidden 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 hidden factor, continuing for a couple of months after the increase in the hidden factor. If the smoothed hidden factor is a standard deviation above its mean in month t, monthly growth in industrial production in months $t-4$ through $t+2$ averages about 10 basis points per month below average. (To put the 10 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 Conclusion

In the context of a Gaussian dynamic term structure model, the evidence presented here points to the presence of a hidden factor. The factor contains substantial information about expected future yields but has a negligible immediate effect on the term structure. The factor is related to both real activity and survey reports of investors' interest rate expectations.

An important lesson to draw from this evidence is that an econometrician should not rely on estimation techniques that extract information exclusively from the cross section. Such techniques are standard in the literature on dynamic term structure models. Instead, she needs to build models that accommodate hidden factors, and use estimation techniques that are robust to the presence of these factors. The method adopted here is filtering, which uses information from yield dynamics to infer factor properties. Another potentially valuable approach is to use information from sources other than bond yields.

Appendix

This appendix presents some details of model estimation and factor rotation. For convenience, the model is estimated imposing the normalized dynamics

$$x^{\dagger}_{t+1} = D^{\dagger} x^{\dagger}_t + \Sigma^{\dagger} \epsilon^{\dagger}_{t+1}. \tag{34}$$

In (34), $D^{\dagger}$ is a diagonal matrix and $\Sigma^{\dagger}$ is lower triangular with ones along the diagonal. Rather than estimating directly both the physical and equivalent-martingale parameters, I estimate the physical parameters in (34) and the parameters of an unrestricted measurement equation

$$y_t = A + B^{\dagger} x^{\dagger}_t + \eta_t, \quad \eta_t \sim N(0, \sigma^2_{\eta}). \tag{35}$$

In (35), A is a 6×1 vector and $B^{\dagger}$ is a 6×5 matrix. Lurking behind the parameters of this measurement equation are the equivalent martingale dynamics of $x^{\dagger}_t$. Because there are five factors to explain six bond yields, A and $B^{\dagger}$ exactly identify the parameters of the no-arbitrage model δ_0, δ_1, μ^q, and K^q. As discussed in Duffee (2009), 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. Further details of the optimization procedure followed here are in Duffee (2009).

After estimation, the factors are rotated into principal components as described in Section 3.4. Define the 5×5 matrix Γ as

$$\Gamma = C_0^{-1} B^{\dagger}_{\setminus 3} \tag{36}$$

where $B^{\dagger}_{\setminus 3}$ is the matrix $B^{\dagger}$ excluding the row corresponding to the three-year bond. The rotated state vector is

$$x_t = \Gamma x_t^{\dagger}. \tag{37}$$

The parameters of its dynamics in equation (25) are defined by

$$K = \Gamma D^{\dagger} \Gamma^{-1}, \quad \Sigma = \text{chol}\left(\Gamma \Sigma^{\dagger} \Sigma^{\dagger\prime} \Gamma'\right). \tag{38}$$

The relation between bond yields and the rotated factors is equation (22) where the new factor loadings are

$$B = B^{\dagger} \Gamma^{-1}. \tag{39}$$

These factor loadings (for all but the three-year bond yield) are the eigenvectors of the diagonalization (24).

The parameters of the short-rate equation and the equivalent-martingale dynamics of the state vector are not directly used in the paper. They are reported here for convenience, using the principal components factor rotation. Thus they correspond to the physical measure parameters reported in Table 2.

$$\delta_0 = 0.0043$$

$$\delta_1 = \begin{pmatrix} 0.527 & \text{-0.662} & \text{-0.182} & 2.495 & \text{-3.898} \end{pmatrix}'$$

$$\mu^q = 10^{-4} \begin{pmatrix} 0.520 & 0.094 & 0.452 & -0.204 & 0.066 \end{pmatrix}'$$

$$K^q = \begin{pmatrix} 0.973 & 0.096 & -0.029 & -0.962 & 1.388 \\ 0.028 & 0.990 & 0.052 & 1.103 & -1.870 \\ 0.032 & 0.042 & 1.239 & 1.231 & -1.764 \\ -0.012 & -0.013 & -0.149 & 0.597 & 0.816 \\ 0.002 & 0.003 & 0.016 & 0.035 & 0.994 \end{pmatrix}$$

The implied term structure dynamics are well-behaved over the maturity range used in estimation (out to five years). Beyond that range–say, at ten years–the estimated term structure dynamics are unrealistic. This is a common problem with estimated Gaussian term structure models. Empirical estimates of K^q matrices typically have the largest eigenvalue of K^q approximately equal to one. If the value exceeds one, factor dynamics are explosive as maturity increases. Here, the largest eigenvalue is 1.046.

References

Ang, Andrew, and Monika Piazzesi, 2003, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics* 50, 745-787.

Brandt, Michael W., and David A. Chapman, 2003, Comparing multifactor models of the term structure, Working paper, Duke University.

Chen, Ren-Raw, and Louis Scott, 1993, Maximum likelihood estimation for a multifactor equilibrium model of the term structure of interest rates, *Journal of Fixed Income* 3, 14-31.

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.

Cochrane, John H., and Monika Piazzesi, 2008, Decomposing the yield curve, Working paper, Chicago Booth.

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.

Cooper, Ilan, and Richard Priestly, 2008, Time-varying risk premiums and the output gap, *Review of Financial Studies*, forthcoming.

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., 2009, Forecasting with the term structure: The role of no-arbitrage restrictions, Working paper, Johns Hopkins.

Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379-406.

Gurkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2006, The U.S. Treasury yield curve: 1961 to the present, working paper, Federal Reserve Board.

Joslin, Scott, Marcel Priebsch, and Kenneth J. Singleton, 2009, Risk premium accounting in macro-dynamic term structure models, Working paper, Stanford GSB.

Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu, 2009, Characterizing Gaussian DTSMs: Why seemingly different models are (nearly) identical, Working paper, Stanford GSB.

Kim, Don H., and Athanasios Orphpanides, 2005, Term structure estimation with survey data on interest rate forecasts, FEDS working paper 2005-48, Federal Reserve Board.

Litterman, Robert, and Jose Scheinkman, 1991, Common factors affecting bond returns, *Journal of Fixed Income* 1, 54-61.

Ludvigson, Sydney C., and Serene Ng, 2009, Macro factors in bond risk premia, *Review of Financial Studies*, forthcoming.

Pearson, Neil D., and Tong-sheng Sun, 1994, Exploiting the conditional density in estimating the term structure: An application to the Cox, Ingersoll, and Ross model, *Journal of Finance* 49, 1279-1304.

Pennacchi, George G., 1991, Identifying the dynamics of real interest rates and inflation: evidence using survey data, *Review of Financial Studies* 4, 53-86.

Piazzesi, Monika, 2009, Affine term structure models, *Handbook of Financial Econometrics* 1, Yacine Ait-Sahalia and Lars Hansen, eds., North-Holland, forthcoming.

Stambaugh, Robert F., 1988, The information in forward rates, *Journal of Financial Economics* 21, 41-70.

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 log 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. Variances of principal components

	Index of component				
	1	2	3	4	5
Yield levels	40.405	0.930	0.068	0.008	0.005
Monthly changes	1.119	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

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

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

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

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	−0.018	−0.172	0.987	−3.355	
	(0.009)	(0.049)	(0.204)	(0.872)	(1.446)	
	0.003	0.936	−0.301	0.213	−0.033	
	(0.002)	(0.020)	(0.067)	(0.266)	(0.486)	
	−0.001	−0.003	0.820	0.506	0.065	
	(0.001)	(0.005)	(0.033)	(0.121)	(0.242)	
	0.000	−0.002	0.026	0.692	−0.024	
	(0.000)	(0.002)	(0.010)	(0.049)	(0.093)	
	0.000	−0.001	0.001	0.018	0.869	
	(0.000)	(0.001)	(0.006)	(0.035)	(0.046)	
$\Sigma \times 10^4$	7.913	0	0	0	0	
	(0.256)					
	−0.713	2.538	0	0	0	
	(0.304)	(0.091)				
	0.379	0.590	0.928	0	0	
	(0.084)	(0.109)	(0.044)			
	−0.027	0.025	−0.224	0.328	0	
	(0.030)	(0.031)	(0.035)	(0.033)		
	0.007	0.033	−0.039	−0.034	0.165	
	(0.022)	(0.024)	(0.029)	(0.031)	(0.028)	
	A×10^3	B(:,1)	B(:,2)	B(:,3)	B(:,4)	B(:,5)
3 mon	4.258	0.459	−0.650	−0.598	0.102	0.003
	(1.036)	(0.024)	(0.028)	(0.029)	(0.018)	(0.016)
1 year	4.564	0.464	−0.314	0.602	−0.567	0.055
	(1.061)	(0.015)	(0.037)	(0.023)	(0.022)	(0.071)
2 year	4.713	0.457	0.065	0.405	0.742	−0.268
	(1.085)	(0.004)	(0.032)	(0.025)	(0.038)	(0.101)
3 year	4.851	0.443	0.287	0.167	0.433	0.618
	(1.077)	(0.011)	(0.023)	(0.030)	(0.089)	(0.115)
4 year	4.960	0.432	0.437	−0.132	0.047	0.776
	(1.075)	(0.018)	(0.018)	(0.031)	(0.102)	(0.012)
5 year	5.027	0.422	0.533	−0.316	−0.340	−0.568
	(1.071)	(0.024)	(0.021)	(0.028)	(0.075)	(0.048)
$\sigma_\eta \times 10^5$	4.612					
	(0.114)					

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 smoothed estimates of the factors in place of the true factors. The fourth column reports population correlations between true factors and smoothed estimates of the factors.

		R^2's of factors on yields		Correl of true,
Factor	Std dev	True factors	Smoothed factors	smoothed factor
1	6.017	1.000	1.000	1.000
	(1.182)			
2	0.925	0.997	1.000	0.998
	(0.110)			
3	0.251	0.954	0.993	0.982
	(0.020)			
4	0.067	0.623	0.821	0.871
	(0.005)			
5	0.043	0.433	0.591	0.858
	(0.004)			

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.

	Std dev of forecast	Std dev attributable to factor				
Maturity	(%/year)	1	2	3	4	5
3 mon	2.28	2.23	0.28	0.07	0.06	0.36
	(1.86)	(1.79)	(0.27)	(0.14)	(0.11)	(0.30)
	[0.87 3.15]	[0.77 3.10]	[0.02 0.58]	[0.00 0.39]	[0.00 0.31]	[0.03 0.64]
1 yr	2.32	2.28	0.15	0.04	0.07	0.37
	(1.88)	(1.82)	(0.18)	(0.12)	(0.11)	(0.31)
	[0.88 3.18]	[0.80 3.13]	[0.01 0.44]	[0.00 0.35]	[0.00 0.31]	[0.04 0.61]
2 yr	2.34	2.31	0.02	0.07	0.04	0.38
	(1.88)	(1.83)	(0.11)	(0.11)	(0.10)	(0.32)
	[0.87 3.18]	[0.82 3.13]	[0.00 0.31]	[0.01 0.32]	[0.00 0.28]	[0.06 0.60]
3 yr	2.33	2.29	0.12	0.14	0.01	0.37
	(1.86)	(1.81)	(0.11)	(0.13)	(0.08)	(0.32)
	[0.86 3.12]	[0.81 3.09]	[0.00 0.35]	[0.00 0.35]	[0.00 0.25]	[0.08 0.58]
4 yr	2.32	2.27	0.20	0.19	0.01	0.37
	(1.85)	(1.79)	(0.15)	(0.17)	(0.08)	(0.32)
	[0.86 3.11]	[0.79 3.07]	[0.01 0.41]	[0.01 0.39]	[0.00 0.23]	[0.09 0.56]
5 yr	2.30	2.24	0.25	0.24	0.03	0.37
	(1.83)	(1.77)	(0.19)	(0.20)	(0.08)	(0.32)
	[0.85 3.07]	[0.77 3.04]	[0.01 0.47]	[0.02 0.42]	[0.00 0.23]	[0.10 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 4.

Panel A. Univariate statistics

		True returns		Observed returns	
Maturity	Mean	Std dev	Predictable frac of var	Std dev	Predictable frac of var
2 yr	0.36	1.78	0.20	1.78	0.19
3 yr	0.68	3.24	0.20	3.24	0.20
4 yr	0.87	4.50	0.22	4.51	0.22
5 yr	0.88	5.58	0.21	5.59	0.21

Panel B. Decomposition of volatility of expected excess returns

	Std dev of conditional mean	Std dev attributable to factor				
Maturity	(%/year)	1	2	3	4	5
2 yr	0.79	0.43	0.56	0.02	0.07	0.35
	(0.83)	(0.47)	(0.55)	(0.11)	(0.11)	(0.29)
	[0.54 1.14]	[0.15 0.75]	[0.20 0.87]	[0.00 0.32]	[0.01 0.29]	[0.03 0.59]
3 yr	1.46	0.57	1.04	0.11	0.05	0.84
	(1.54)	(0.69)	(1.04)	(0.21)	(0.19)	(0.72)
	[1.03 2.13]	[0.18 1.15]	[0.40 1.62]	[0.01 0.62]	[0.01 0.54]	[0.21 1.28]
4 yr	2.12	0.74	1.54	0.13	0.02	1.24
	(2.21)	(0.90)	(1.54)	(0.30)	(0.26)	(1.07)
	[1.53 3.02]	[0.25 1.52]	[0.62 2.37]	[0.01 0.84]	[0.01 0.74]	[0.36 1.85]
5 yr	2.53	0.82	1.97	0.23	0.02	1.35
	(2.68)	(1.03)	(1.99)	(0.37)	(0.31)	(1.15)
	[1.84 3.64]	[0.22 1.79]	[0.85 2.99]	[0.01 1.04]	[0.01 0.90]	[0.24 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 $t + j, j = 1, \ldots 4$. The contemporaneous three-month yield is subtracted from the forecasts to produce forecasted changes in the yield. Contemporaneous filtered estimates of the "level" and "hidden" factors are taken from a five-factor term structure model. The factors are normalized to have population 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 hidden factor on the survey-based expected change

	Quarters ahead (j)			
	1	2	3	4
Coef	−1.332	−0.965	−0.711	−0.544
	(0.315)	(0.300)	(0.237)	(0.174)
AR(1) of residual	0.71	0.73	0.72	0.71

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

	Quarters ahead (j)			
	1	2	3	4
Coef on level	−0.101	−0.140	−0.198	−0.256
	(0.033)	(0.043)	(0.056)	(0.072)
Coef on hidde	−0.132	−0.147	−0.142	−0.149
	(0.052)	(0.057)	(0.058)	(0.069)
AR(1) of residual	0.27	0.51	0.56	0.58

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

The log change industrial production from month $t-1$ to month t is regressed on the month $t-i$ smoothed estimate of the hidden factor, for i $= -6, \ldots, 6$. 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	0.002	0.044	0.54
−5	−0.032	0.045	−0.71
−4	−0.056	0.047	−1.20
−3	−0.085	0.049	−1.76
−2	−0.102	0.049	−2.08
−1	−0.102	0.050	−2.05
0	−0.117	0.053	−2.20
1	−0.118	0.056	−2.10
2	−0.105	0.056	−1.88
3	−0.091	0.053	−1.73
4	−0.088	0.050	−1.76
5	−0.064	0.050	−1.29
6	−0.074	0.050	−1.49

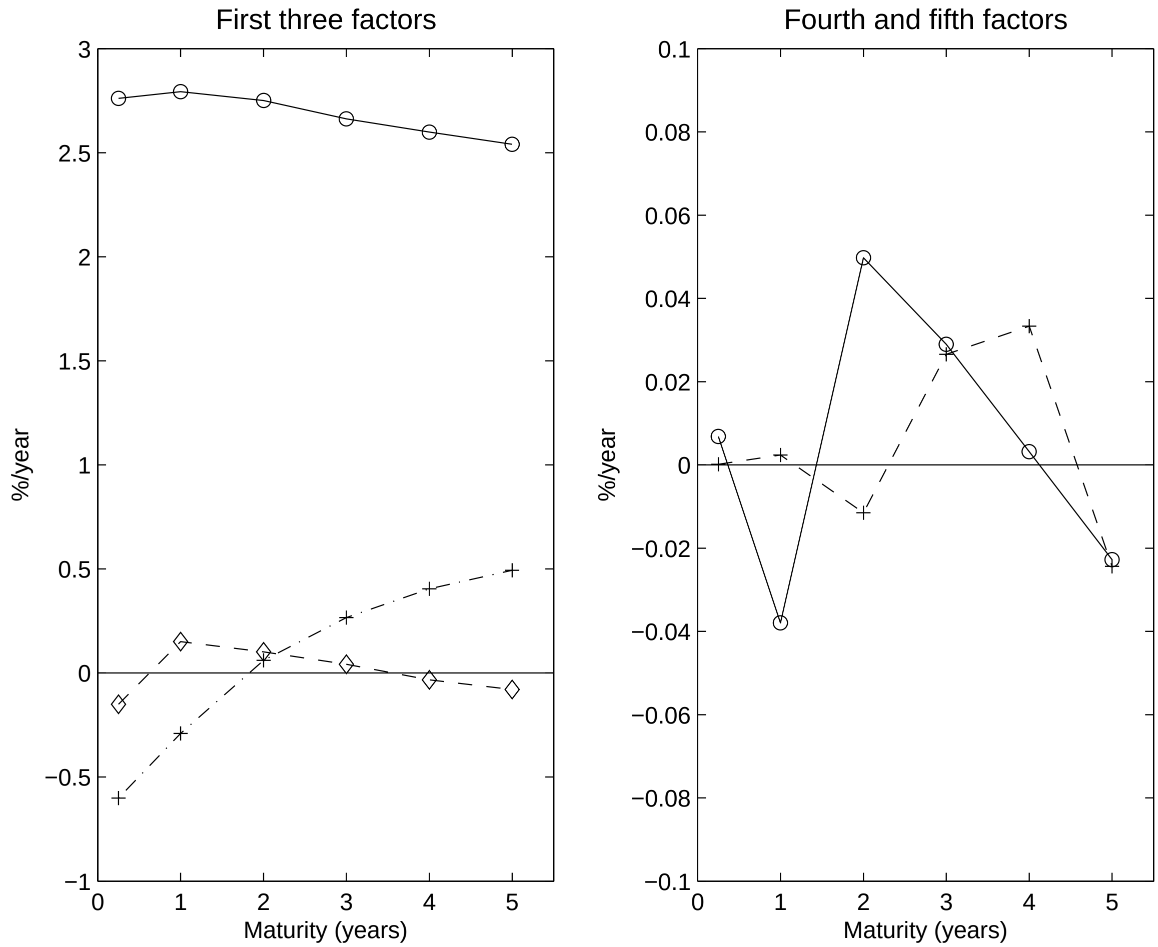

Fig. 1. Estimated loadings of annualized yields on the five factors of a term structure model. Each line represents the response of the term structure to a one standard deviation change in the given factor.

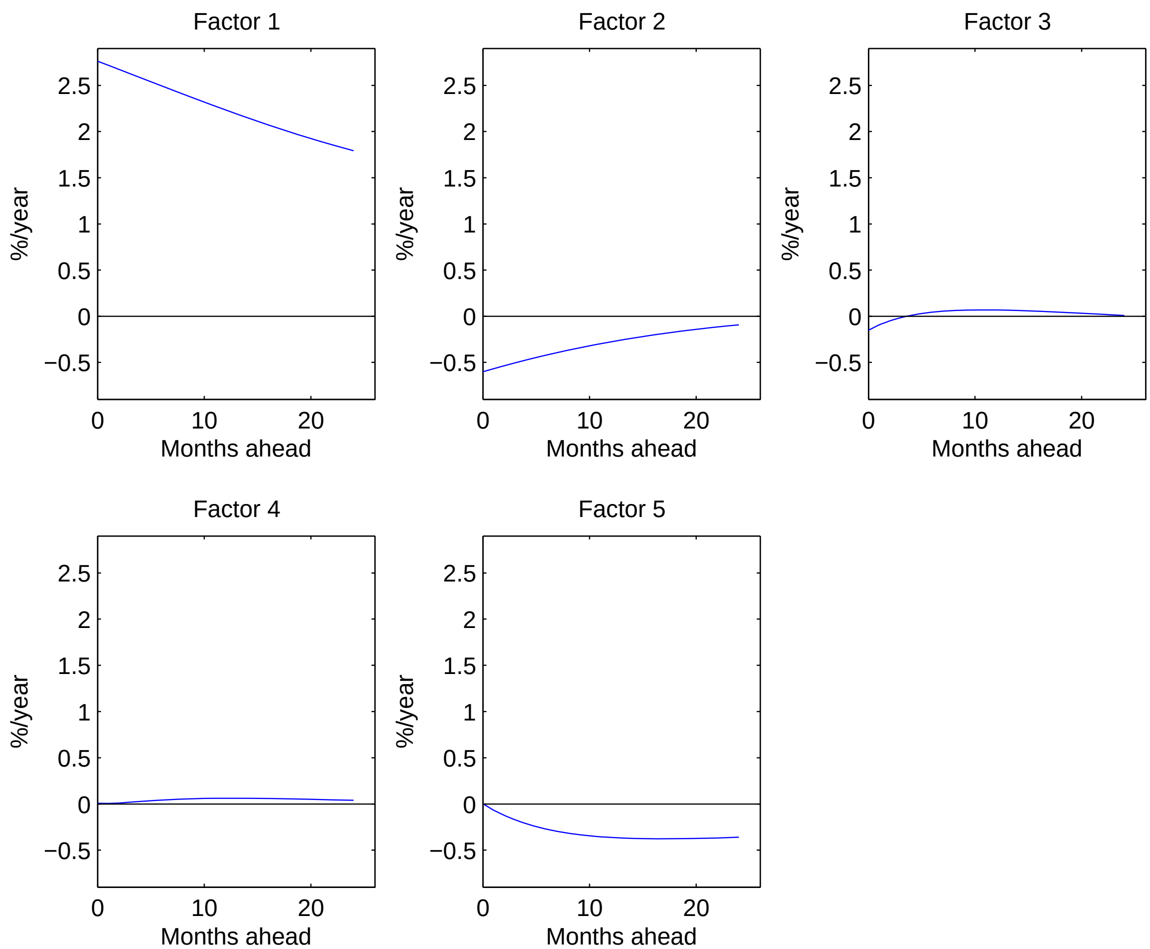

Fig. 2. Responses of the three-month bill rate to term structure factors. Each panel plots the expected time path of the three-month bill yield, assuming that at month zero the specified factor is one standard deviation above its mean. All other factors are set to their unconditional means.

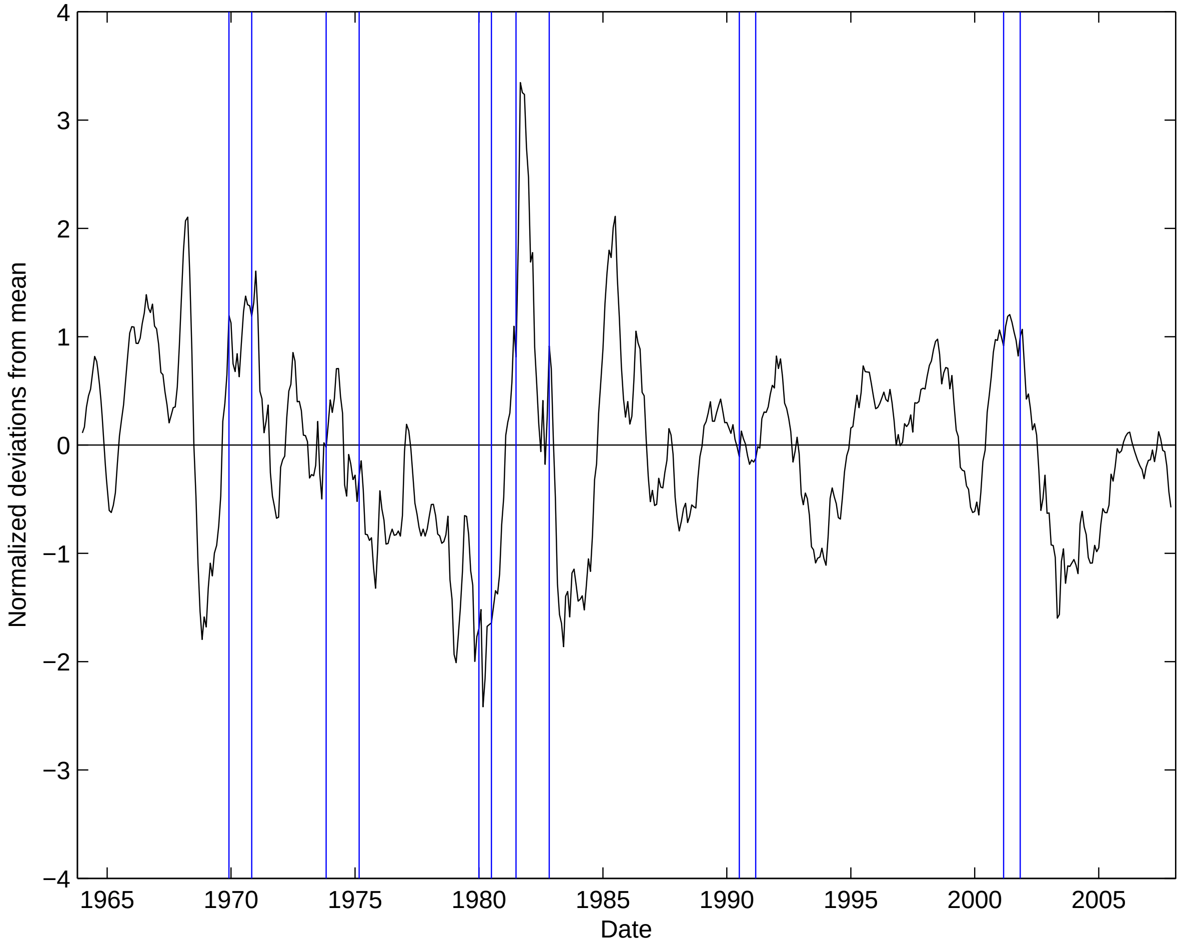

Fig. 3. Smoothed estimates of the "hidden" factor. The vertical lines are NBER business cycle break points.

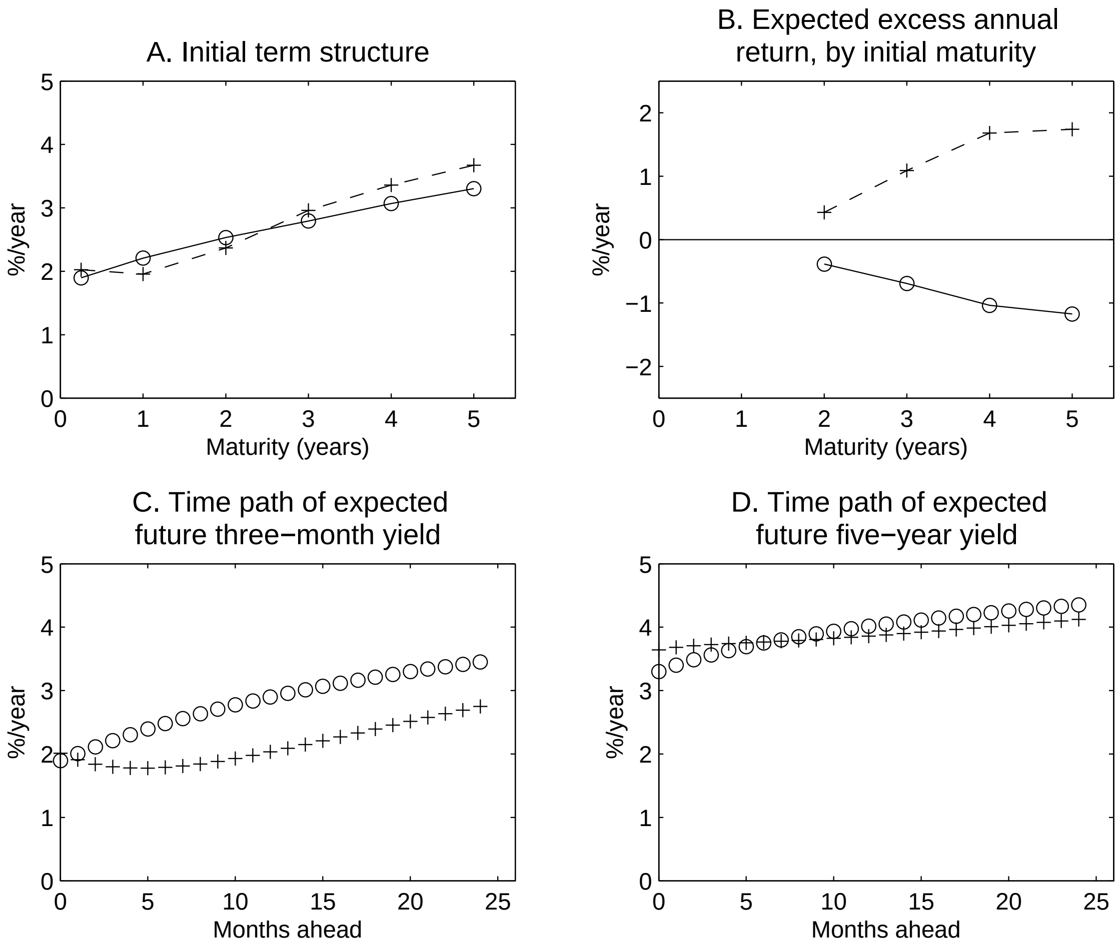

Fig. 4. A comparison of October 2001 and October 2004. Values for the two months are plotted with '+' and 'o' respectively. Panel A displays the month-end term structures. Panel B displays model-implied expected excess log returns (over the one-year yield) for bonds with maturities of two through five years. Panel C displays expected future three-month yields over the next 24 months, where month zero is October of either 2001 or 2004. Panel D displays expected future five-year yields.