

Information in (and not in) the term structure

*Gregory R. Duffee**

Johns Hopkins

First draft: March 2008

This Draft: September 17, 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 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, Scott Joslin, Sydney Ludvigson, Jacob Sagi, and seminar participants at the Federal Reserve Board, Wharton, and the 2008 NBER Summer Institute for helpful conversations and comments.

1 Introduction

Investors' beliefs about future bond prices determine what investors are willing to pay for bonds today. This truism suggests that the term structure of bond yields contains all information relevant to predicting both future returns to bonds and future bond yields. Put somewhat differently, any variable that helps to explain yield dynamics must also help explain the cross section of yields.

Researchers commonly invoke this standard intuition when building and estimating term structure models, such as when choosing the dimension of the model. Factor analysis of the unconditional covariance matrix of yields (or differenced yields) tells us how many factors drive the cross section. Such evidence is often used to specify the number of factors driving dynamics. Another application is model estimation, which often assumes a one-to-one mapping from the factors to an equal number of bond yields.

The one-to-one assumption makes explicit the notion that the cross section of bond yields follows a Markov process. 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 the 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.

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 have not been recognized. 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. However, these concerns are 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 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 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.

I also investigate properties of regressions that use the term structure to forecast future excess annual bond returns. Cochrane and Piazzesi (2005) find that at least 35 percent of the variance of annual excess bond returns is predicted with regressions using the cross section of yields. The model's factors estimated here contain even more information about future excess returns. However, the estimated model also tells us that the high R^2 s are just

a consequence of sampling error. Population R^2 s for the regressions are actually less than 20 percent.

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. Finite-sample properties on forecasting regressions are in Section 5. This section also considers whether the estimated factors capture information about future bond yields that is contained in macroeconomic data. Concluding comments are in Section 6.

2 The modeling framework

This section explains 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 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 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

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

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

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

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

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

We now turn to 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} 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 $k_{12}^q = 0$, 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 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

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

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 ,

$$\delta_{1,i}^* = 0. \quad (20)$$

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

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

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 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 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 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 2.3. 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 I cannot point to a reasonable theory of investor behavior that motivates the necessary parameter restriction. 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. Thus there seems to be no reason 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 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 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 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 the cross section is affected by all factors and a model in which it is affected by a subset of factors. 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 two sections ask whether empirical analysis can confirm or deny the presence of such factors.

3 Empirical analysis

The core of the empirical analysis in this paper is an estimated Gaussian term structure model. This section discusses the estimation procedure and describes features of the es-

estimated model. The main economic question is whether there are factors that have little effect on the cross section of yields yet are important for modeling dynamics. According to the estimated model, 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 both macroeconomic data and surveys of interest-rate forecasters.

3.1 The choice of five factors

We want to determine if a factor that affects yield dynamics is hidden from the cross section. Since we know three factors are necessary to capture the cross section, the estimated model must have more than three factors. 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. 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 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. Because the smoothing process can artificially erase the a factor from the

cross section of yields, I use the CRSP data.

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 3.5 takes a detailed look at these shapes.

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. 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, for 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 . These parameters are reported in the Appendix. 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 from this vector, denoting the vector of remaining five uncontaminated yields by $\tilde{y}_{\setminus 3,t}$. Take the corresponding five rows from B^\dagger and place them in $B_{\setminus 3}^\dagger$, a 5×5 matrix of factor loadings.

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

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

where Ω is the diagonal matrix of the covariance eigenvalues.

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

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

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.¹ 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 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 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 param-

¹Recall 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.

ters 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.

3.6 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. 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.

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

$$xr_{t,t+j}^{(m)} \equiv my_t^{(m)} - (m - j)y_{t+j}^{(m-j)} - jy_t^{(j)}. \quad (35)$$

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

$$\begin{aligned} xr_{t,t+j}^{(m)} &= mA_m - (m - j)A_{m-j} - jA_j \\ &\quad + (mB'_m - (m - j)B'_{m-j}K^j - jB'_j) x_t \\ &\quad - (m - j)B'_{m-j} \left(\sum_{i=1}^j K^{j-i} \epsilon_{t+i} \right) \\ &\quad + m\eta_t^{(m)} - j\eta_t^{(j)} - (m - j)\eta_{t+j}^{(m-j)}. \end{aligned} \quad (36)$$

The four lines on the right side of (36) 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 2.4. 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 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.² 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

²In particular, the months were not chosen based on the contemporaneous survey forecasts.

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 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 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 $\text{SPF_EXPECT}(t, j)$. 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'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

$$\text{MODEL_HIDDEN}_t = b_0 + b_1 \text{SPF_EXPECT}(t, j) + e_{j,t} \quad (37)$$

for forecast horizons of one through four quarters ($j = 1, \dots$). 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 (37), 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 (37) 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 (37) and add the estimated level factor as an additional explanatory variable. The regression is

$$\text{SPF_EXPECT}(t, j) = b_0 + b_1 \text{MODEL_LEVEL}_t + b_2 \text{MODEL_HIDDEN}_t + e_{j,t}. \quad (38)$$

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

$$100(\log(\text{IP}_t) - \log(\text{IP}_{t-1})) = b_{0,i} + b_{1,i}\text{MODEL_HIDDEN}_{t-i} + e_{t,i}, \quad i = -6, \dots, 6. \quad (39)$$

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 (39) 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 Properties of predictive regressions

Cochrane and Piazzesi (2005) find surprisingly strong predictability of excess returns to Treasury bonds. Using five forward rates to forecast annual excess returns, they (hereafter CP) uncover R^2 s around 35 percent. This section uses the estimated model to evaluate the accuracy of predictive regressions, and in particular to study the plausibility of this high R^2 .

A summary of the results helps to put the details into context. The five factors extracted from the data by the Kalman filter contain more information than the forward rates about future excess returns. For example, if smoothed factors replace the forward rates in CP regressions, the R^2 s are nearly 50 percent. But the dynamic structure of the model allows us to conclude that these high R^2 s are largely driven by sampling error. The model's population properties imply that no more than 20 percent of annual excess returns are predictable. Under the null hypothesis that the model is correct, ML estimation of the complete model (and its 50+ parameters) produces more accurate finite-sample estimates of expected excess returns than does the CP regression. Finally, I conclude that the filtering procedure used here does not come close to capturing all information investors have about expected future yields. There is independent information in macroeconomic variables.

5.1 Two sets of conflicting parameter estimates

The explanatory variables used by CP are the month- t forward rates for borrowing from year i to year $i + 1$, $i = 0, \dots, 4$. Denote these annualized forward rates by $F(0, 1)_t, \dots, F(4, 5)_t$. Stack them in the vector F_t . The predictive regression is

$$xr_{t,t+12}^{(m)} = b_{m,0} + b'_m F_t + e_{t,t+12}^{(m)} \quad (40)$$

where the excess return is defined as in equation (35). Panel A of Table 8 reports results of estimating the regression with ordinary least-squares (OLS) over the sample period used in this paper.³ To conserve space, results are displayed only for a five-year bond ($m=60$). The R^2 of 33 percent is slightly lower than that reported by CP, who use data through 2003. The difference in sample period does not affect the tent shape in the coefficient estimates that is highlighted by CP. The point estimates are significantly different from zero, at least as measured by asymptotic Newey-West standard errors with 18 lags. Unreported results for other maturities are similar.

An alternative approach to estimating the parameters of (40) is to construct population

³The index t spans 516 observations from 1964:1 through 2006:12. The final year of data is reserved for leads.

properties of (40) from the term structure model estimated in Section 3. I refer to the first set of estimates as OLS estimates and the alternative estimates as population estimates.⁴ Population estimates are calculated analytically by expressing excess returns and the forward rate vector as functions of the model's parameters, factors, and innovations. The expression for the annual excess log return to a five-year bond is

$$xr_{t,t+12}^{(60)} = c_0 + c'_1 x_t + c'_2 \eta_t + c'_3 \eta_{t+12} + \sum_{i=1}^{12} c'_{3+i} \epsilon_{t+i}, \quad (41)$$

$$\begin{aligned} c_0 &= 12(5A_{60} - 4A_{48} - A_{12}), \quad c'_1 = 12(5B'_{60} - 4B'_{48}K^{12} - B'_{12}), \quad c'_{3+i} = -48B'_{48}K^{12-i}, \\ c'_2 &= 12 \begin{pmatrix} 0 & -1 & 0 & 0 & 0 & 5 \end{pmatrix}, \quad c'_3 = 12 \begin{pmatrix} 0 & 0 & 0 & 0 & -4 & 0 \end{pmatrix}. \end{aligned} \quad (42)$$

The vectors c_2 and c_3 have six elements because the model was estimated using a cross section of six yields with maturities ranging from three months to five years. The forward rate vector is

$$F_t = D_0 + D_1 x_t + D_2 \eta_t, \quad (43)$$

$$\begin{aligned} D_0 &= 12 \begin{pmatrix} A_{12} \\ 2A_{24} - A_{12} \\ 3A_{36} - 2A_{24} \\ 4A_{48} - 3A_{36} \\ 5A_{60} - 4A_{48} \end{pmatrix}, \quad D_1 = 12 \begin{pmatrix} B'_{12} \\ 2B'_{24} - B'_{12} \\ 3B'_{36} - 2B'_{24} \\ 4B'_{48} - 3B'_{36} \\ 5B'_{60} - 4B'_{48} \end{pmatrix}, \\ D_2 &= 12 \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 2 & 0 & 0 & 0 \\ 0 & 0 & -2 & 3 & 0 & 0 \\ 0 & 0 & 0 & -3 & 4 & 0 \\ 0 & 0 & 0 & 0 & -4 & 5 \end{pmatrix}. \end{aligned} \quad (44)$$

These expressions for excess returns and forward rates, combined with the unconditional covariance matrices of x_t , ϵ_t , and η_t , can be plugged into

$$\beta_{60} = \text{Var}(F_t)^{-1} \text{Cov}(F_t, xr_{t,t+12}^{(60)}), \quad (45)$$

which is the population counterpart of b_m for $m = 60$ in (40).

⁴This language is not perfect because the population estimates are population values of OLS coefficients, thus 'OLS' could be applied to both sets.

These population estimates, also reported in Panel A of Table 8, are at odds with the OLS estimates. The model generates the tent shape in coefficients, but the individual parameters are closer to zero. Put differently, the model says that the OLS estimates are picking the right linear combination of forecasting variables, but overestimating the total amount of predictability. The population R^2 is only 16 percent. This value might appear to contradict results in Table 5, which reports that 21 percent of the five-year bond's excess return population variance is predictable. But according to the model, the five forward rates do not capture all of the information in the state vector because a critical factor is nearly hidden by measurement error in yields.

5.2 The information in the state vector

The model says that the five forward rates contain less information about yield dynamics than is in the state vector, but the model might be wrong. It is conceivable that owing to model misspecification, the factors extracted with the Kalman filter are less useful in forecasting than are the forward rates used by CP. An informal test of this possibility replaces the forward rates on the right side of (40) with either filtered or smoothed estimates of the model's five state variables. OLS estimates are reported in Panel B of Table 8.

A comparison of R^2 s across the three estimated OLS regressions indicates that the model's factors contain more information about future excess returns than do the forward rates. According to the OLS results, filtered factors explain 38 percent of the variance of annual excess returns, while smoothed factors explain almost half of the return variance.

Naturally, these new OLS estimates are also at odds with the population properties of the estimated model. Population estimates of these regressions are unavailable analytically, thus I use a single simulation to compute them. One million months of yields are randomly generated using the estimated dynamic model. Filtered and smoothed factors are extracted from these yields using the Kalman filter, treating the model's parameters as known. Using OLS, the predictive regressions are then estimated on these million observations. Panel B of Table 8 reports that the population estimated coefficients are smaller than the OLS point estimates, and the population R^2 s are about half of the OLS estimated R^2 s.⁵

⁵The population R^2 of 23 percent using smoothed factors exceeds the predictability of 21 percent reported in Table 5. The reason is that even in an infinite sample, smoothed factors have a look-ahead feature that exaggerates predictability.

5.3 Dynamics embedded in the regressions

Forecasting annual log returns to five-year bonds is equivalent to forecasting the four-year yield twelve months ahead. The regressions focus precisely on this forecast. By contrast, the model can be used to forecast yields on many different bonds at many different horizons. Internal consistency is forced on these various forecasts. The requirement of consistency is the reason for the large wedge between OLS and model-based forecasts.

To emphasize the role of the forecast horizon, rewrite the annual log return to a five-year bond as

$$xr_{t,t+12}^{(60)} = 60y_t^{(60)} - 48y_t^{(48)} - 12y_t^{(12)} - 48 \sum_{i=1}^{12} \left(y_{t+i}^{(48)} - y_{t+i-1}^{(48)} \right). \quad (46)$$

Twelve separate monthly changes in the four-year yield appear in (46). Consider forecasting each by regressing individual monthly changes on the month- t forward rates used by CP. The 12 coefficients for the sample period 1964 through 2007 are displayed with circles in Fig. 5. Also displayed are plus and minus two standard error bounds on the coefficients. They are adjusted for generalized heteroskedasticity. Ignore the plotted diamonds for the moment. (To help interpret the coefficients, the regressions use annualized four-year yields.)

The contrast between the OLS regression and population properties of the model is illustrated by the coefficients on the forward rate $F(2,3)$. Consider an increase in this forward rate by one percentage point, holding the other four forward rates constant. By summing the twelve OLS point estimates in Panel C of Fig. 5, we find that the annualized four-year bond yield is expected to drop by almost one percentage point (98.75 basis points) during the next twelve months. This decline corresponds to an higher expected return to a five-year bond of $4(0.9875)$ or 3.95 percent. This is the OLS estimated coefficient on $F(2,3)$ that appears in the Cochrane-Piazzesi regression in Panel A of Table 8.

But the figure reveals an erratic pattern in the OLS-implied expected decline in the four-year yield. During the first three months, the yield is expected to decline a total of 42 basis points. The decline then levels off, with an additional total decline of less than ten basis points in months four through seven. In months eight through twelve, the decline accelerates. The yield is expected to drop an additional 47 basis points during these months.

The model cannot reproduce this pattern of OLS point estimates because of the smoothness built into the model's dynamics. Population regression coefficients implied by the model's parameters are displayed with diamonds. According to the estimated model, the four-year yield is expected to decline by 30 basis points over the first three months, an additional 23 basis points over months four through seven, and by 13 basis points over months eight through twelve. If we take the model at face value, the large estimated regression

coefficients for these final months are largely attributable to sampling error.

Since we cannot be sure that the model specification used here is correct, we cannot be sure that the OLS regressions overstate predictability. However, it is likely that reproducing the point estimates of Fig. 5 using a dynamic term structure model will require a substantially more flexible specification than the 52-parameter, five-factor version used here.

5.4 Finite-sample evidence

Under the maintained hypothesis that the estimated model is true, which method produces more accurate finite-sample estimates of expected excess returns: OLS regressions or Kalman filter estimation of the complete model? For concreteness, consider in-sample forecasts of annual excess log returns to a five-year bonds, where the sample consists of 44 years of monthly data. One procedure is to construct fitted values from the CP regression using five forward rates. Another is to estimate the model with maximum likelihood, then use the Kalman filter to construct the fitted values. This approach can use either filtered or smoothed factors.

The tradeoffs are straightforward. The CP approach requires estimation of only six parameters. Full model estimation requires estimation of 52 parameters, but this estimation uses information from many yields and horizons. It can also infer the presence of hidden or nearly-hidden factors.

The accuracy of these approaches is evaluated with Monte Carlo simulations. Treating the model estimated in Section 3 as truth, 528 months of state vectors and bond yields are randomly generated. Using the month t state vector, the month- t true expectation of $xr_{t,t+12}^{(60)}$ is computed. A hypothetical econometrician does not observe the state vector. Instead, the econometrician estimates the CP regression using five forward rates to construct in-sample fitted values of this expected annual log return. Denote this fitted expectation by

$$E_t^{CP} \left(xr_{t,t+12}^{(60)} \right) = \hat{b}_{60,0} + \hat{b}'_{60} F_t. \quad (47)$$

Hats denote OLS parameter estimates.

The econometrician then estimates with maximum likelihood the five-factor term structure model. The corresponding forecasts of excess annual log returns using filtered states are

$$\hat{E}_t^{FIL} \left(xr_{t,t+12}^{(60)} \right) = 60y_t^{(60)} - 48 \left(\hat{A}_{48} + \hat{B}'_{48} \hat{K}^{12} \hat{x}_t \right) - 12y_t^{(12)} \quad (48)$$

where hats denote maximum likelihood parameter estimates. Filtered values of the fitted

values are also denoted with hats. The fitted expectation using smoothed factors is

$$\hat{E}_t^{SMO} \left(xr_{t,t+12}^{(60)} \right) = 60y_t^{(60)} - 48 \left(\hat{A}_{48} + \hat{B}'_{48} \hat{K}^{12} \tilde{x}_t \right) - 12y_t^{(12)} \quad (49)$$

where smoothed factors are denoted with tildes.

The performance of these fitted forecasts is evaluated with their root mean squared error (RMSE), defined as the difference between the true expectation and the fitted expectation. (Note that this is not the difference between the realization and the fitted expectation, which is a noisier measure of forecast accuracy.) For example, the RMSE of the CP regression forecast is

$$RMSE^{CP} = \sqrt{\frac{1}{528} \sum_{t=1}^{528} \left[(E_t - E_t^{CP}) xr_{t,t+12}^{(60)} \right]^2}. \quad (50)$$

For each simulated sample of 528 months, RMSEs are constructed for all three of these forecasting methods. For purposes of comparison, RMSEs are also constructed for three other methods. One assumes expected excess returns are constant over time and equal to the sample's mean excess return. The other two are forecasting regressions where the predictive variables are either filtered or smoothed factors from Kalman filter estimation.

Table 9 reports the mean, across 1000 Monte Carlo simulations, of these RMSEs. Because true expected excess returns vary substantially over time, the largest errors are associated with the assumption of constant expected excess returns. When sample means are used to estimate expected excess returns, the mean RMSE is 2.52 percent. The forward-rate regression of CP lowers this mean RMSE by more than 60 basis points, to 1.88 percent.

Although the CP regression captures some of the true variability of expected excess returns, its mean RMSE is disappointingly large. For example, imagine that for some month in the sample, the true expected excess return is one standard deviation above its unconditional mean. There is almost a nine percent probability that the CP regression will produce a fitted expected excess return *below* the unconditional mean.⁶ Forecasts produced with ML estimation and filtering are significantly more accurate. With smoothed factors, the mean RMSE is 1.43 percent. For the same hypothetical example, the probability that the fitted expected excess return is negative is less than four percent.

The improvement in fit has little to do with the ability of the Kalman filter to extract nearly hidden factors. Table 9 reports that forecasts produced using filtered or smoothed factors in predictive regressions have mean RMSEs similar to those of the CP regression. The greater accuracy of the Kalman filter procedure is thus presumably driven by its use of information in other yields and other return horizons.

⁶The cdf of the normal distribution for 2.53/1.88 is 0.9108.

5.5 The information in the state vector, revisited

When the term structure follows a first-order Markov process, all information about future yields and returns is contained in the current term structure. If, however, some factor is hidden (either exactly or imperfectly because of measurement error), information other than that in the cross section of yields will help form forecasts. Ludvigson and Ng (2007) find that factors constructed from measures of economic activity and inflation help predict excess bond returns, even when controlling for the forward rates used by Cochrane and Piazzesi. Their empirical evidence thus points to the presence of hidden factors.

The Kalman filter procedure extracts information from term structure dynamics, thus filtered and smoothed factors also contain more information than does the cross section of yields. The question I examine in this subsection is whether the factors produced by the Kalman filter subsume the information in Ludvigson and Ng’s economic factors.⁷

Their factors are the first eight principal components of 132 financial and nonfinancial time series that span 1964 through 2003. Ludvigson and Ng (hereafter LN) construct a nonlinear function of these principal components that predicts annual excess log returns to bonds. Because their function is chosen to optimally fit in-sample excess returns, a forecasting horse race between their function and filtered factors might appear to be biased in favor of their function. I therefore simply use all eight of their principal components in their raw (linear) form.

I also diverge from LN, as well as CP, by forecasting monthly excess returns instead of annual returns. I therefore avoid the use of overlapping observations. The resulting test statistics should have finite-sample properties close to their asymptotic properties. The return data, from CRSP, are returns to maturity-sorted portfolios of Treasury bonds. The month- t simple return to the portfolio with maturities less than a year is subtracted from simple returns to longer-maturity portfolios to form excess returns.

Denote the excess return to portfolio p in month $t + 1$ by $R_{t,t+1}^p$. The regression is

$$R_{t,t+1}^p = b_{p,0} + b'_{p,1}\tilde{x}_t + b'_{p,2}MF_t + e_{p,t,t+1}, \quad (51)$$

where \tilde{x}_t is the vector of smoothed states and MF_t is the vector of eight LN macro-financial factors. The sample is 1964 through 2003, although the smoothed factors are from ML estimation of the term structure model using data through 2007.

Table 10 displays the results. Across all five maturity buckets, the LN factors are better at predicting variations in excess returns than are the term structure factors. When only the smoothed factors are included in the regression, the R^2 s range from 6.4 percent (for

⁷Thanks very much to the authors for making their data available.

maturities between one and two years) to 8.9 percent (for maturities between five and ten years). When only the LN factors are included, the same range is from 12.6 percent to 10.0 percent. Naturally, we should treat these precise R^2 s with caution. As seen in Section 5.4, OLS regression estimates of predictability at a single horizon are imprecise.

The more important question here is whether the smoothed factors and the LN factors are picking up the same predictability. The evidence in Table 10 is that they are not. When both sets of factors are included in the regression, the R^2 s are much larger than they are when only one set of factors is included. For example, for the portfolio containing bonds with maturities between five and ten years, using both sets of factors results in an R^2 of 16.1 percent. The sum of the R^2 s of the two separate regressions is only slightly larger, at 18.9 percent. Standard χ^2 tests overwhelmingly reject the hypotheses that each set of factors can be excluded from the joint regression.

This strong evidence tells us that researchers should look beyond bond yields when estimating dynamic term structure models. Ang and Piazzesi (2003) develop a framework for incorporating macroeconomic variables into such models. Kim and Orphanides (2005) and Chernov and Mueller (2008) extend this framework to include survey forecasts. The results of Section 4 suggest that this approach can go a long way towards capturing a forecasting factor that is hidden from the cross section of yields.

6 Conclusion

This paper shows that an econometrician cannot extract from the cross section of bond yields all information investors have about expected future yields. There is a factor that 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 hidden 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.

Appendix

The calculations in this paper do not directly use the equivalent-martingale parameters of the estimated term structure model. They are reported here for completeness. These parameters are based on the principal components factor rotation, and thus correspond to the physical measure parameters reported in Table 2.

$$\delta_0 = 0.0043$$

$$\delta_1 = \left(0.527 \quad -0.662 \quad -0.182 \quad 2.495 \quad -3.898 \right)'$$

$$\mu^q = 10^{-4} \left(0.520 \quad 0.094 \quad 0.452 \quad -0.204 \quad 0.066 \right)'$$

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

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.
- Chernov, Mikhail, and Philippe Mueller, 2008, The term structure of inflation expectations, working paper, LBS.
- 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.
- 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.
- Kim, Don H., and Athanasios Orphanides, 2005, Term structure estimation with survey data on interest rate forecasts, working paper, 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, 2007, Macro factors in bond risk premia, Working paper, NYU.

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

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

Yields on bonds with maturities of three months and one through five years are stacked in the vector y_t . Yields are expressed in decimal form per month. 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	-0.018	-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)	
	$A \times 10^3$	$B(:,1)$	$B(:,2)$	$B(:,3)$	$B(:,4)$	$B(:,5)$
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.

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

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
	(1.86)	(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

Maturity	Mean	True returns		Observed returns	
		Std dev	Predictable frac of var	Std dev	Predictable frac of var
2 yr	0.36	1.78	0.20	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

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
	(0.83)	(0.47)	(0.55)	(0.11)	(0.11)	(0.29)
3 yr	[0.54 1.14]	[0.15 0.75]	[0.20 0.87]	[0.00 0.32]	[0.01 0.29]	[0.03 0.59]
	1.46	0.57	1.04	0.11	0.05	0.84
4 yr	(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]
5 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)
5 yr	[1.53 3.02]	[0.25 1.52]	[0.62 2.37]	[0.01 0.84]	[0.01 0.74]	[0.36 1.85]
	2.53	0.82	1.97	0.23	0.02	1.35
5 yr	(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 $t + j, j = 1, \dots, 4$. The contemporaneous three-month yield is subtracted from the forecasts to produce forecasted changes in the yield. Contemporaneous filtered estimates of the “level” and “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 (0.315)	-0.965 (0.300)	-0.711 (0.237)	-0.544 (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 (0.033)	-0.140 (0.043)	-0.198 (0.056)	-0.256 (0.072)
Coef on hidde	-0.132 (0.052)	-0.147 (0.057)	-0.142 (0.058)	-0.149 (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 $i = -6, \dots, 6$. 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	0.002	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

Table 8. Regression forecasts of excess returns to five-year bonds

Excess log returns to a five-year bond from month t to month $t + 12$ are predicted with the month- t term structure using regressions. One regression uses five forward rates as in Cochrane and Piazzesi (2005). The other two use filtered and smoothed state vectors extracted from bond yields by ML estimation of a five-factor term structure model. The sample period is January 1964 through December 2007. Newey-West standard errors, adjusted for 18 lags of moving-average residuals, are in parentheses. The estimated term structure model is used to calculate population properties of each regression. The notation $F(m, n)$ denotes the forward rate from year m to year n .

Panel A. Forward rates

	Coefficient on:					R^2
	$F(0, 1)$	$F(1, 2)$	$F(2, 3)$	$F(3, 4)$	$F(4, 5)$	
OLS estimate	-2.914 (0.504)	0.760 (0.861)	3.950 (0.811)	1.329 (0.657)	-2.603 (0.602)	0.328
Population value	-1.861	0.465	2.658	0.289	-1.096	0.157

Panel B. Factors from a model estimated with the Kalman filter

	Coefficient times 100 on factor:					R^2
	1	2	3	4	5	
Filtered factors	0.932 (0.538)	2.786 (0.484)	-0.265 (0.411)	0.346 (0.327)	2.541 (0.529)	0.377
Population value	0.812	1.973	0.219	-0.038	1.280	0.182
Smoothed factors	0.919 (0.487)	2.781 (0.431)	-0.244 (0.405)	0.411 (0.320)	3.029 (0.471)	0.459
Population value	0.810	1.972	0.229	-0.130	1.830	0.227

Table 9. Finite-sample accuracy of excess return forecasts

Excess log returns to a five-year bond from month t to month $t + 12$ are predicted with four OLS regressions and Kalman filter estimation of a term structure model. One regression uses only a constant term. Another uses five forward rates as in Cochrane and Piazzesi (2005). The final two use filtered and smoothed state vectors from the Kalman filter. 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 in-sample 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 reports the mean RMSE across 1000 simulations. All values are in percent per year.

Forecast method	Mean RMSE
OLS w/constant	2.523
OLS with five forward rates	1.883
Kalman filter w/filtered states	1.488
Kalman filter w/smoothed states	1.429
OLS w/filtered states	1.796
OLS w/smoothed states	1.833

Table 10. Forecasting excess returns with term structure and macro factors

Excess returns from month t to month $t+1$ to five maturity-sorted portfolios of Treasury bonds are regressed on month- t term structure and macro factors. Five of the factors are smoothed state vectors from Kalman filter estimation of a term structure model. The others are from Ludvigson and Ng (2007), who use dynamic factor analysis to extract eight common factors from 132 measures of economic activity. The table reports R^2 s and tests that the either set of factors can be excluded from regressions that include both. Asymptotic test statistics are adjusted for generalized heteroskedasticity. p -values are in brackets. The sample period is January 1964 through December 2003. Each portfolio is identified by the shortest and longest maturities included in it.

	Portfolio maturities (years)				
	1-2	2-3	3-4	4-5	5-10
R^2 for five smoothed factors	0.064	0.067	0.082	0.077	0.089
R^2 for eight macro factors	0.126	0.106	0.108	0.104	0.100
R^2 for all factors	0.167	0.149	0.161	0.153	0.161
$\chi^2(5)$ test excluding smoothed factors	13.44 [.020]	12.93 [.024]	16.14 [.006]	17.45 [.004]	20.87 [.001]
$\chi^2(8)$ test excluding macro factors	36.31 [.000]	29.34 [.000]	33.24 [.000]	33.37 [.000]	32.58 [.000]

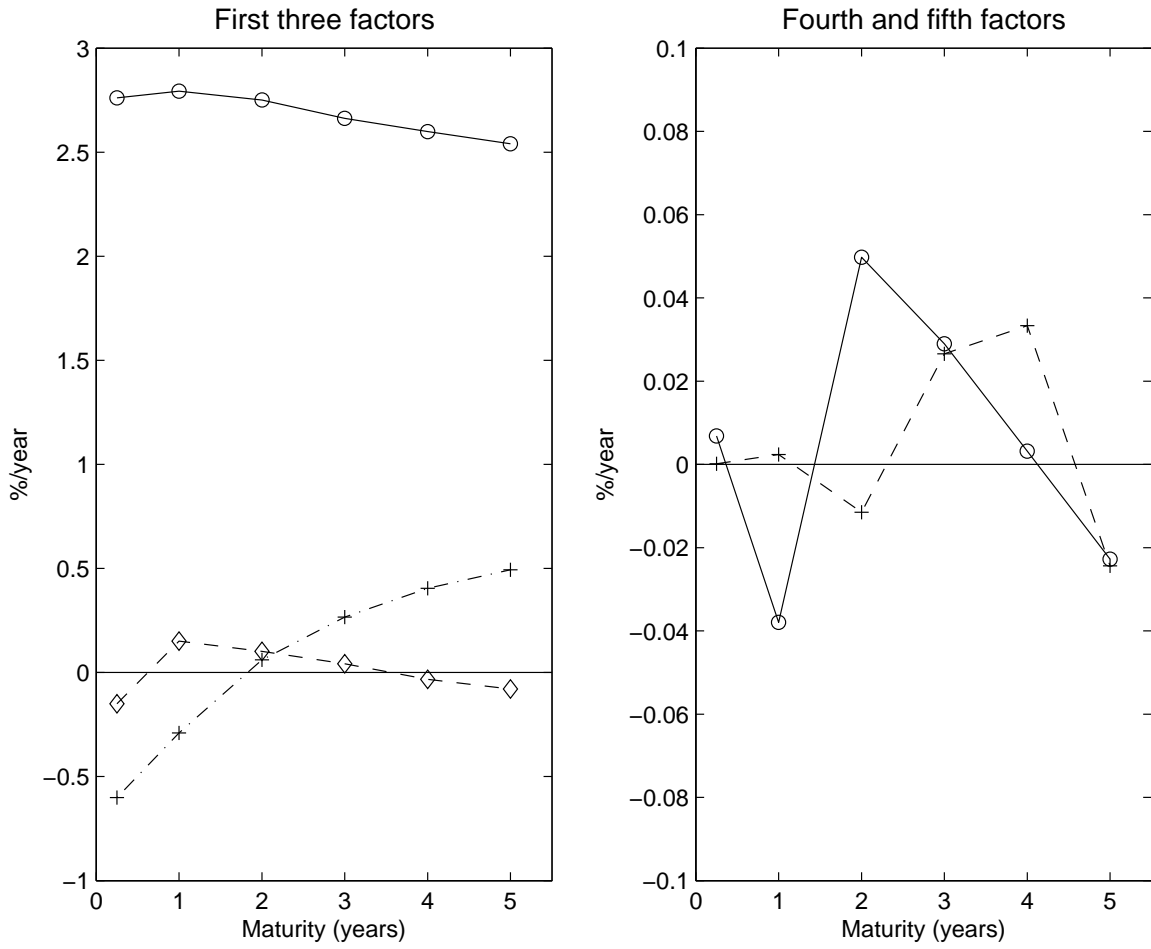


Fig. 1. Estimated loadings of annualized yields on the five factors of a term structure model. Each line represents the response of the term structure to a one standard deviation change in the given factor.




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.

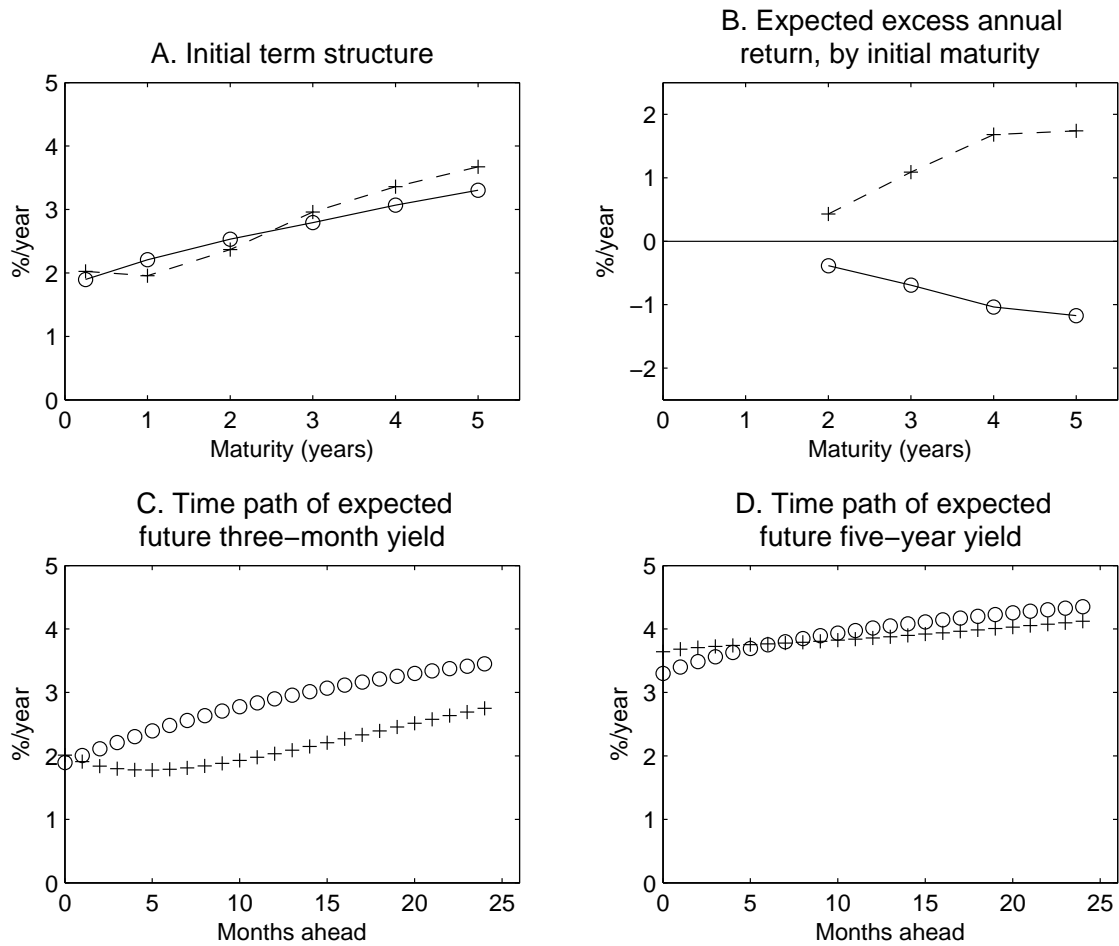


Fig. 4. A comparison of October 2001 and October 2004. Values for the two months are plotted with '+' and 'o' respectively. Panel A displays the month-end term structures. Panel B displays model-implied expected excess log returns (over the one-year yield) for bonds with maturities of two through five 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.

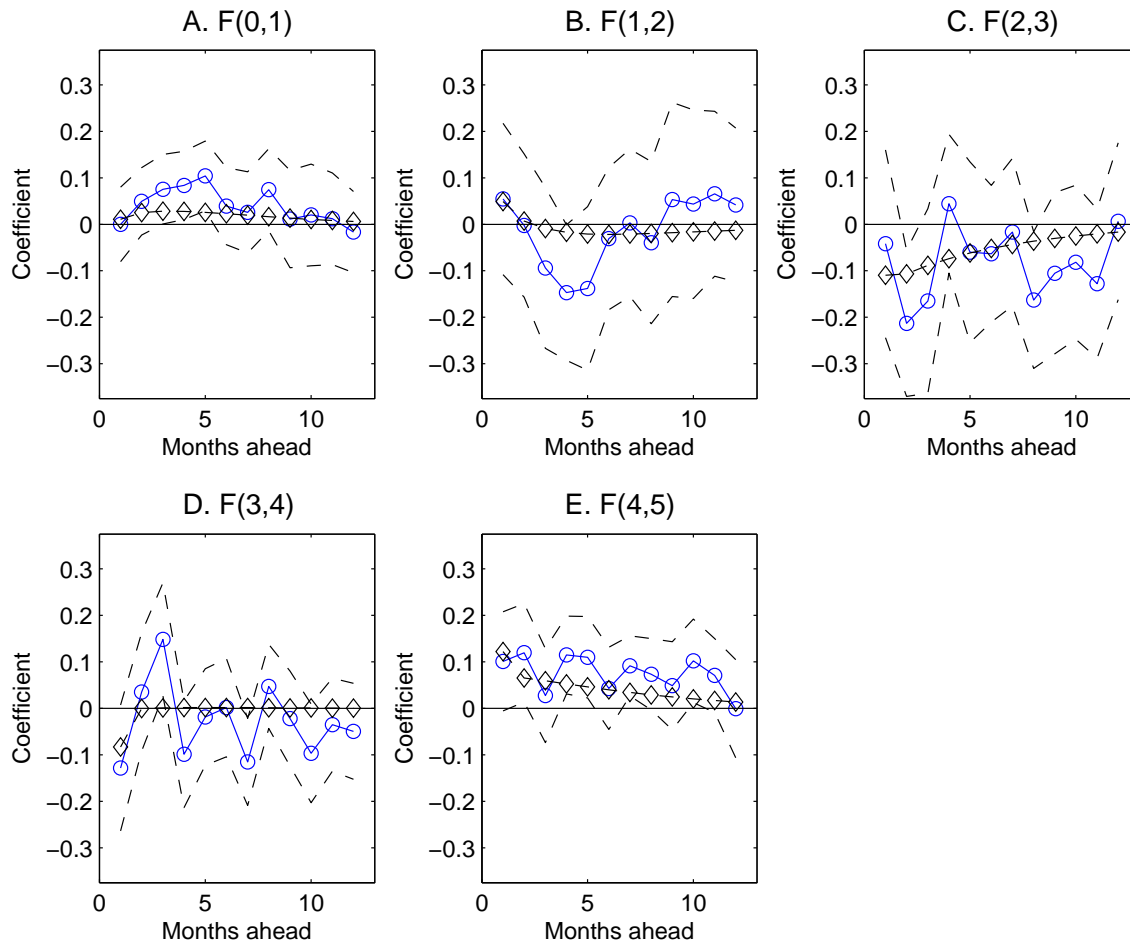


Fig. 5. Empirical and theoretical regression coefficients from predictive regressions. The change from month $t + i - 1$ to month $t + i$ in the four-year bond yield is regressed on five month- t forward rates. The headings on the panels are the forward rates, where $F(m, n)$ is the forward rate from year m to year n . The circles are estimated coefficients for $i = 1, \dots, 12$ using the sample January 1964 through December 2007. The dashed lines are \pm two standard error bounds on the coefficients, using heteroskedasticity-consistent standard errors. The diamonds are theoretical coefficients of population regressions implied by a term structure model.