Term Structure of Expectations and Expectations of Term Structures*

Pooya Molavi† Alireza Tahbaz-Salehi‡ Andrea Vedolin§

January 2025

Abstract

According to the expectations hypothesis (EH) of the term structure of interest rates, yields on long-term bonds are determined by the bond market's forecast of the short-term interest rates. Common empirical tests of EH proxy the expectations of short rates by their realizations, thus effectively jointly testing expectations and rational expectations hypotheses. In this paper, we test EH under market participants' subjective expectations. To this end, we develop a method to extract the entire term structure of expectations of short rates from survey data. We find that the yields implied by EH under subjective expectations closely match yield data, indicating that time-varying risk premiums are quantitatively small.

Keywords: forecasts.

JEL Classification: F31, G15

*Preliminary and incomplete. Comments welcome.

†Northwestern University, pmolavi@kellogg.northwestern.edu.
‡Northwestern University and CEPR, alirezat@kellogg.northwestern.edu.
§Boston University, NBER, and CEPR, avedolin@bu.edu.

1 Introduction

Yields on long-term bonds consist of two components: the anticipated returns from shorter-term instruments over the same period, and an additional factor known as the term premium. The term premium is essentially the extra return that investors require as compensation for the risks associated with holding a long-term bond.

Canonical models in asset pricing attribute these risk premiums to "maturity habitats" (Vayanos and Vila, 2021), long-run macroeconomic risk (Bansal and Shaliastovich, 2013), among many others. In this paper, we argue that time-varying risk premiums on long-dated bonds are potentially small.

Term premiums are not observable, and there are two ways to estimate them. Assume interest rates are continuously compounded and expressed in percent per period and denote the nominal short-rate by r_t. No arbitrage leads to an expression for an n-period yield:

$$y_t^{(n)} = \frac{1}{n}\mathbb{E}_t\left(\sum_{k=0}^{n-1} r_{t+k}\right) + \underbrace{\frac{1}{n}\mathbb{E}_t\left(\sum_{k=0}^{n-2} rx_{t+k+1}^{(\tau-k)}\right)}_{\text{term premium}}$$

To calculate the RHS of this equation, we would need to know the future path of interest rates. Unfortunately, in general, we do not know the future path of the short-rate, r_t. The literature has proposed two ways to addressing this issue. First, the term structure literature assumes that yields are driven by a small number of factors.[1] Imposing no-arbitrage restrictions, allows them to estimate yields both under the physical and risk-neutral dynamics, therefore being able to identify risk compensation.

Another way to calculate term premia is to subtract a survey measure of the average expected short rate from the observed bond yield. However, survey data are not updated frequently and (typically) include only a limited set of forecast horizons. For example, the Survey of Primary Dealers provides long-term forecasts of the short-rate only since 2014. The Survey of Professional Forecasters reports long-term average short-term expectations only once per year, and Blue Chip Financial Forecasts provides monthly forecasts of the short-rate with horizons less than two years or semi-annual longer-term forecasts.

To address these data shortcomings, we begin by developing a theoretical framework to construct the entire term structure of expectations from the handful of forecasts observed in the data. In general and in the absence of any assumptions, subjective forecasts at various horizons can be arbitrary. Therefore, as our starting point, we impose two assumptions on how agents' form their forecasts and show that these two assumptions together pin down the entire term structure of subjective expectations.

Our first assumption maintains that agents' forecasts are affine functions of the past realizations of a state variable, which in our baseline specification we take to be the short-term interest rate itself.

[1] For example, Adrian, Crump, and Moench (2013) assume that yields are driven by the first four principal components of yields only, Kim and Wright (2005) add survey data on three-month interest rates, Christensen and Rudebusch (2012) model a three-factor, affine Gaussian term structure model centered around the affine arbitrage-free class of Nelson-Siegel term structure models.

As our second assumption, we assume that agents' forecasts at different horizons are internally consistent with one another, in the sense that agents' subjective expectations of the short-term interest rate satisfy the law of iterated expectations.

As our main theoretical result, we show that these two assumptions—linearity and internal consistency of expectations—imply a tight relationship between expectations at different horizons. More specifically, we show that the coefficients of the linear representation of forecasts at any horizon τ can be expressed recursively in terms of forecasts at horizons $\tau - 1$ and horizon 1 (corresponding to the shortest possible forecast horizon). This recursive relationship implies that one can construct the entire term structure of expectations from forecasts at the shortest horizon. This construction will be the basis of the empirical analysis that follows in subsequent sections.

Our empirical results can be summarized as follows. Using survey data on the Federal Funds Rate from Blue Chip Financial Forecasts running from Q1 1988 to Q4 2023, we first construct the term structure of expected short-rates for arbitrary horizons. As a first test of our model, we study whether the assumptions of our model are satisfied, that is, whether agents' subjective expectations at different horizons are internally consistent with one another. To this end, we study the relationship between the realized forecasts in the data and the implied forecasts. We find that the two are strongly aligned. We then use the path of the short-rate to construct implied yields. By construction, these yields satisfy the expectations hypothesis. Comparing our model-implied yields to those in the data, we find that (not surprisingly), differences at short maturities are insignificant while differences increase with the bond maturity. Interestingly, we find that the difference between the 10-year realized yield and our model-implied yield is almost perfectly correlated with measures of the term premium but significantly smaller in size. For example, at the ten-year maturity, we find that the average wedge is only around 30bps while the average term premium is over 120bps for the 1988 to 2023 period. We also find that our wedges are negative since 2015.

Literature Review. The paper closest to ours are Crump, Eusepi, and Moench (2024) and Farmer, Nakamura, and Steinsson (2023). Using a rich dataset of professional forecasts covering output growth, inflation, and interest rates, Crump et al. (2024) show that deviations from rationality and drifting beliefs cannot reconcile the disconnect between forward rates and short-term rate expectations. The study decisively rejects the EH, showing that the wedge between forward rates and expected short rates explains most of the variability in longer-maturity yields. Like us, the authors allow for model misspecification but impose specific models for both agents' subjective forecasts and the true process, with parameters estimated from survey data.

Similar to us, Farmer et al. (2023) study deviations from the expectations hypothesis without relying on time-varying risk premiums. The authors employ a Bayesian learning framework focused on slow parameter learning which induces deviations from rational expectations. They show that this approach accounts for observed facts both in survey and interest rate data. Different from us, the authors assume a parametric model for the process that governs economic fundamentals, say, short-term interest rates, assuming that the process determines both the data-generating process and agents' forecasts, hence, there is no model misspecification.

A large literature documents violations of the rational-expectations hypothesis in forecast data and offers alternative models of expectation formation. A central question in this literature is whether households, firms, and professional forecasters under- or over-extrapolate from new information. Coibion and Gorodnichenko (2015) provide evidence of under-extrapolation in consensus forecasts for professional forecasters. Bordalo, Gennaioli, Ma, and Shleifer (2020) show that individual forecasts of professional forecasters over-extrapolate from recent news. Angeletos, Huo, and Sastry (2021) find evidence of under-extrapolation at short horizons and over-extrapolation at longer horizons, whereas Afrouzi, Kwon, Landier, Ma, and Thesmar (2023) find evidence of over-extrapolation in a lab setting, and Broer and Kohlhas (2024) find evidence for both under- and over-extrapolation depending on the aggregate variable being studied. In parallel with this empirical literature, many papers have proposed theoretical models of under- and over-extrapolation. Natural expectations, e.g., Fuster, Laibson, and Mendel (2010) and Fuster, Hebert, and Laibson (2012), and diagnostic expectations, e.g., Bordalo, Gennaioli, and Shleifer (2018), are examples of models where agents over-extrapolate from the recent past. Cognitive discounting of Gabaix (2020) and level-k thinking, e.g., García-Schmidt and Woodford (2019) and Farhi and Werning (2019), are examples of models that feature under-extrapolation. Our methodology allows us to non-parametrically identify the forecasters' subjective model of the short rate without taking a stand on whether forecasters exhibit over- or under- extrapolation (or other behavioral biases). Furthermore, we can use the forecasters' estimated subjective model to rule in and rule out specific models of expectation formation (such as adaptive learning or diagnostic expectations).

Outline. The rest of the paper is organized as follows. Section 2 presents our theoretical framework for constructing the term structure of expectations of the short-term interest rate from a handful of observed forecasts. Section 3 describes the datasets we use in our analysis. Our main empirical results in presented in Section 4. Section 5 concludes.

2 Theoretical Framework

Notation. Time is discrete and infinite. At each time t, there are zero-coupon, default-free bonds with face value one maturing in $n \geq 1$ periods. Let $p_t^{(n)} = \log P_t^{(n)}$ denote the logarithm price of the n-period bond at time t and define the bond's yield as

$$y_t^{(n)} = -p_t^{(n)}/n.$$

We refer to the yield of the 1-period bond as the short rate and denote it by $r_t = y_t^{(1)}$. The one-period holding return on the n-period bond for $n \geq 2$ is defined as $R_{t+1}^{(n)} = P_{t+1}^{(n-1)}/P_t^{(n)} - 1$. Similarly, the return of the short maturity bond is given by $R_{t+1}^{(1)} = 1/P_t^{(1)} - 1 = R_t$. Finally, define the *forward rate* as the current yield of an n-period bond maturing in $n+m$ periods:

$$f_t^{(n,m)} = \frac{1}{n}\left[(n+m)y_t^{(n+m)} - my_t^{(m)}\right] = \frac{1}{n}(p_t^{(m)} - p_t^{(n+m)}). \tag{1}$$

and the *excess return* of the n-period bond as its return above and beyond the short rate:

$$\mathrm{rx}_{t+1}^{(n)} = r_{t+1}^{(n)} - r_t = p_{t+1}^{(n-1)} - p_t^{(n)} - r_t.$$

The Expectations Hypothesis. According to the expectation hypothesis of interest rates, the yield on an n-period bond is equal the average expected value of the yield on a 1-period bond over the bond's lifetime:

$$y_t^{(n)} = \frac{1}{n}\sum_{\tau=0}^{n-1} \mathbb{E}_t[r_{t+\tau}]. \tag{2}$$

Since expected future interest rates are not observable, common empirical tests of EH (such as Fama (1984) and Campbell and Shiller (1991)) proxy the expectations of short rates, $\mathbb{E}_t[r_{t+\tau}]$, by their realizations, $r_{t+\tau}$, thus effectively jointly testing expectations and rational expectations (RE) hypotheses. So, even though these tests consistently reject the relationship in (2)—indicating the presence of a systematic wedge between the observed yields and those implied by the average expected value of the short rate—such rejections may either be because of a violation of EH, RE, or both (Froot, 1989). Put differently, a direct test of EH—without also imposing RE—requires having access to the entire term structure of subjective expectations.

In what follows, we develop a theoretical framework to extract the entire term structure of expectations of the short rate from survey data without imposing rational expectations or any form of consistency between subjective and objective expectations.

2.1 The Term Structure of Expectations

In general—and absent imposing any assumptions on how agents form their forecasts—the term structure of expectations can be arbitrary. Therefore, our starting point is to impose two assumptions on how agents' subjective forecasts are formed.

Our first assumption states that agents' form their subjective forecasts at all horizons using a linear model. To formally state this assumption, let $x_t \in \mathbb{R}^n$ denote the vector of observables.

Assumption 1. Subjective expectations at any horizon $\tau \geq 1$ are given by

$$\mathbb{E}_t[x_{t+\tau}] = \phi_0^{(\tau)} + \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+1}^{(\tau)} x_{t-s}, \tag{3}$$

where $\mathbf{\Phi}_s^{(\tau)} \in \mathbb{R}^{n\times n}$ are matrices that do not depend on the realizations of x or their expectations and $\phi_0^{(\tau)} \in \mathbb{R}^n$ is a vector.

As already mentioned, according to this assumption, agents' forecasts at any given horizon τ are affine functions of the past realizations of vector of observables x_t. Importantly, this assumption does not impose any restrictions on the matrix of coefficients $\mathbf{\Phi}_s^{(\tau)}$ nor on the relationship between agents' subjective forecasts and the true data-generating process. Nonetheless, Assumption 1 still imposes two important restrictions on how subjective expectations are formed. First, it assumes

that the elements of vector x_t are not misspecified. in the sense that we are not leaving out any variable that the agent is using for forecasting. Second, it assumes that the agent has a linear (and hence smooth) forecasting model that does not experience regime switches.

Our second assumption requires that the agent's expectations at different horizons are mutually consistent, in the sense that they satisfy the law of iterated expectations, as formalized below:

Assumption 2. $\mathbb{E}_t\mathbb{E}_{t+k}[x_{t+k+\tau}] = \mathbb{E}_t[x_{t+k+\tau}]$ for all $k, \tau \geq 0$.

As in Assumption 1, the above assumption does not impose any relationship between the agents' subjective model and the true data-generating process. Instead, what it does is to require that the agents' subjective expectations at different horizons are internally consistent with one another: the agents' time-t expectations of their own future forecasts are equal to their current forecasts.

Our main result in this section shows that, together, Assumptions 1 and 2 impose a tight restriction on the term structure of agents' expectations.

Proposition 1. *Suppose Assumptions 1 and 2 are satisfied. Then,*

$$\phi_0^{(\tau+k)} = \phi_0^{(\tau)} + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \phi_0^{(k-r)} \tag{4}$$

$$\mathbf{\Phi}_s^{(\tau+k)} = \mathbf{\Phi}_{s+k}^{(\tau)} + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \mathbf{\Phi}_s^{(k-r)} \tag{5}$$

for all $s, \tau, k \geq 1$.

The above result is a simple consistency equation that follows from the law of iterated expectations. It establishes that when agents' expectations are internally consistent, there is a tight relationship between the matrices of coefficient $\mathbf{\Phi}_s^{(\tau)}$ in equation (3). As we will show in the subsequent sections, the tight relationship implied by Proposition 1 allows us to construct the entire term structure of expectations even if only a handful of forecast horizons are available in the data.

This can be seen most transparently by imposing $k = 1$ in Proposition 1, in which case equations (4) and (5) take the following simple recursive forms:

$$\begin{aligned} \phi_0^{(\tau+1)} &= \phi_0^{(\tau)} + \mathbf{\Phi}_1^{(\tau)} \phi_0^{(1)} \\ \mathbf{\Phi}_s^{(\tau+1)} &= \mathbf{\Phi}_{s+1}^{(\tau)} + \mathbf{\Phi}_1^{(\tau)} \mathbf{\Phi}_s^{(1)} \qquad \text{for all } s \geq 1. \end{aligned} \tag{6}$$

According to equations, which form a recursion in τ, one can construct the entire term structure of the forecasts for all horizons τ from the forecasts at the shortest horizon ($\tau = 1$).

3 Data

We use data from various sources.

Interest Rates. The monthly effective FFR is from the FRED Economic Database available through the Federal Reserve Bank of St. Louis. For longer maturity yields, we use zero-coupon nominal Treasury yields from Gürkaynak, Sack, and Wright (2007) available on the Federal Reserve Board's (FRB) Web site. The FFR and interest rates are expressed in percent per year.

Survey data. We use forecasts from Blue Chip Financial Forecasts (BCFF). On average around 40 financial professionals provide forecasts about interest rates as well as other macroeconomic variables on around the third week of each month. For our analysis, we use forecasts about the level of the Federal Funds Rate (FFR). Forecasters are asked to provide their estimate of the average FFR for the current quarter (nowcast) up to five quarters ahead. Given this, forecast horizons shrink from one month to the next. For example, in January 2024, forecasters estimate the average FFR for the quarter running from January to March 2024, in February 2024, forecasters still have to estimate the average but now the January 2024 value is known. Given the structure of the forecasts, we therefore have forecasts with horizons ranging between one to 15 months. From these, we then calculate a consensus forecast as the median of all forecasts made each month. We use data starting in January 1988 to December 2023.

Estimates of Term Premia. To get an estimate of the term premia in the data, we use two different estimates from Adrian, Crump, and Moench (2013) available through the webpage of the Federal Reserve Bank of New York and Christensen and Rudebusch (2012) available through the webpage of the Federal Reserve Bank of San Francisco. Both consider affine-term structure models in which yields are driven by a small number of pricing factors and which evolve over time according to a vector autoregression (VAR). In these models, the dynamics of the pricing kernel and the market prices of risk are pinned down. In particular, it is assumed that the pricing kernel is exponentially affine in the shocks that drive the economy and that market prices of risk are affine in the state variables. Parameters that determine the relationship between the pricing factors and the yields are restricted by no-arbitrage constraints.

From the estimated parameters, term premia can then be calculated as the difference between the model-implied fitted yield and the model-implied average expected future short-term interest rate over the corresponding horizon. Adrian et al. (2013) estimate an affine-term structure model using five principal components from yields as pricing factors. Christensen and Rudebusch (2012) estimate an affine Nelson-Siegel term structure model using three principal components from yields as pricing factors. [2]

4 Empirical Analysis

According to Proposition 1, as long as subjective expectations satisfy the linearity and internal consistency assumptions (Assumptions 1 and 2, respectively), forecasts at various horizons are interrelated to one another via equations (4) and (5). In fact, as laid out by equations (6), one can construct the entire term structure of expectations solely from forecasts at the shortest horizon.

[2] Notice that we cannot use the Kim and Wright (2005) term premia estimates since they are estimated using BCFF survey data on the three-month T-Bill. Since we want to compare our nonparametric estimates to those from parametric term structure models, we omit models where yields are functions of survey data since it may induce a mechanical link.

4.1 Estimation

We estimate $\phi_0^{(1)}$ and $\{\mathbf{\Phi}_{s+1}^{(1)}\}_s$ using the generalized method of moments (GMM). Once these parameters are estimated, the implied forecasts for any horizon can be constructed using Proposition 1.

Let $\mathbf{\Phi}^{(\tau)} = (\phi_0^{(\tau)}, \{\mathbf{\Phi}_{s+1}^{(\tau)}\}_{s=0}^{\infty})$ and let $\mathbf{x}_t = (x_t, x_{t-1}, \dots)$. By Assumption 1, for any horizon $\tau \geq 1$ and any $\mathbf{\Phi}^{(\tau)}$, the time-t model-implied forecasts at horizon τ can be written as a linear function of $\mathbf{\Phi}^{(\tau)}$ and $\mathbf{x}_t$. On the other hand, by Proposition 1, $\mathbf{\Phi}^{(\tau)}$ can be written as a function of $\mathbf{\Phi}^{(1)}$. Therefore, we can write the time-t model-implied forecasts at horizon τ as a (nonlinear) function of $\mathbf{\Phi}^{(1)}$ and $\mathbf{x}_t$:

$$\mathbb{E}_t[x_{t+\tau}] = F^{(\tau)}(\mathbf{\Phi}^{(1)}, \mathbf{x}_t).$$

We use the above display to estimate the $\mathbf{\Phi}^{(\tau)}$. However, to take this expression to data, we need to make two adjustments that deal with the finiteness of our sample. First, $\mathbf{\Phi}^{(\tau)}$ and $\mathbf{x}_t$ are infinite dimensional objects. We replace them with their finite dimensional truncations up to lag L. Second, forecast are observed in data with noise. We deal with this problem by including an additive measurement noise term, thus ending up with the following specification:

$$f_t[x_{t+\tau}] = F^{(\tau,L)}\left(\mathbf{\Phi}^{(1,L)}, \mathbf{x}_t^{(L)}\right) + \epsilon_t^{(\tau)}, \tag{7}$$

where $f_t[x_{t+\tau}]$ is the realized time-t forecast at horizon τ, $\mathbf{\Phi}^{(1,L)} = (\phi_0^{(\tau)}, \{\mathbf{\Phi}_{s+1}^{(\tau)}\}_{s=0}^{L-1})$, $\mathbf{x}_t^{(L)} = (x_t, x_{t-1}, \dots, x_{t-L+1})$, $\epsilon_t^{(\tau)}$ is a mean zero measurement error that is orthogonal to $\mathbf{x}_t^{(L)}$, and $F^{(\tau,L)}$ is the L-lag truncation of $F^{(\tau)}$.

The assumption that the measurement error is orthogonal to the $\mathbf{x}_t^{(L)}$ leads to the following moment condition:

$$\mathbb{E}^*\left[\left(f_t[x_{t+\tau}] - F^{(\tau,L)}\left(\mathbf{\Phi}^{(1,L)}, \mathbf{x}_t^{(L)}\right)\right)\mathbf{x}_t^{(L)\prime}\right] = 0,$$

where $\mathbb{E}^*$ denotes expectation with respect to the true data-generating process. We can replace the above population moment condition with its sample counterpart:

$$\frac{1}{T}\sum_{t=1}^{T}\left[\left(f_t[x_{t+\tau}] - F^{(\tau,L)}\left(\mathbf{\Phi}^{(1,L)}, \mathbf{x}_t^{(L)}\right)\right)\mathbf{x}_t^{(L)\prime}\right] = 0. \tag{8}$$

We use versions of this equation to estimate $\mathbf{\Phi}^{(1,L)}$.

This empirical framework allows us to combine data on forecasts at different horizons and test Assumptions 1 and 2 using the resulting over-identifying restrictions. Observation of forecasts at each horizon gives us $n + Ln^2$ moment conditions for $n + Ln^2$ parameters—n parameters for the intercept, and n^2 parameters for every lag included. Moreover, the parameters are identified when $\tau = 1$. Therefore, if we use data on forecasts for H horizons—one of them being the $\tau = 1$ horizon—we end up with $H(n + Ln^2)$ moment conditions for $n + Ln^2$ parameters and $(H-1)(n + Ln^2)$ over-identifying restrictions. We use the over-identifying restrictions to test Assumptions 1 and 2 using Sargan-Hansen's J test.

4.2 Term-Structure of Short-Rate Expectations

Table 1 reports summary statistics (mean and autocorrelation coefficient) for a selection of forecasts for horizons 1, 4, 8, 20, and 40 quarters. We find that on average the term structure of expectations is upward-sloping. This is intuitive since the inception of our data in 1988, the Fed Funds Rate is on a downward trajectory. Long-term expectations, therefore, are for the Fed Funds Rate to increase over the longer horizon. We also find the forecasts to be highly persistent as indicated by the AR(1) coefficient. Figure 1 plots the realized Fed Funds Rate together with forecasts at the 1, 4, 8, 20, and 40 quarter horizon. Recall that forecasts for horizons one to six quarters are available from Blue Chip.[3] The picture echoes the findings from Table 1. When the Fed Funds Rate increases, expectations about the future level of the Fed Funds Rate decline, when the Fed Funds Rate decreases, expectations fall.

Finally, we plot in Figure 2 the slope of the term structure of expectations. We notice that the slope, similar to the term structure of interest rates, drops before recessions (as indicated by the gray shaded bars). This implies that short-term expectations of the Fed Funds Rate are higher than long-term expectations shortly before recessions and then revert back to being lower.

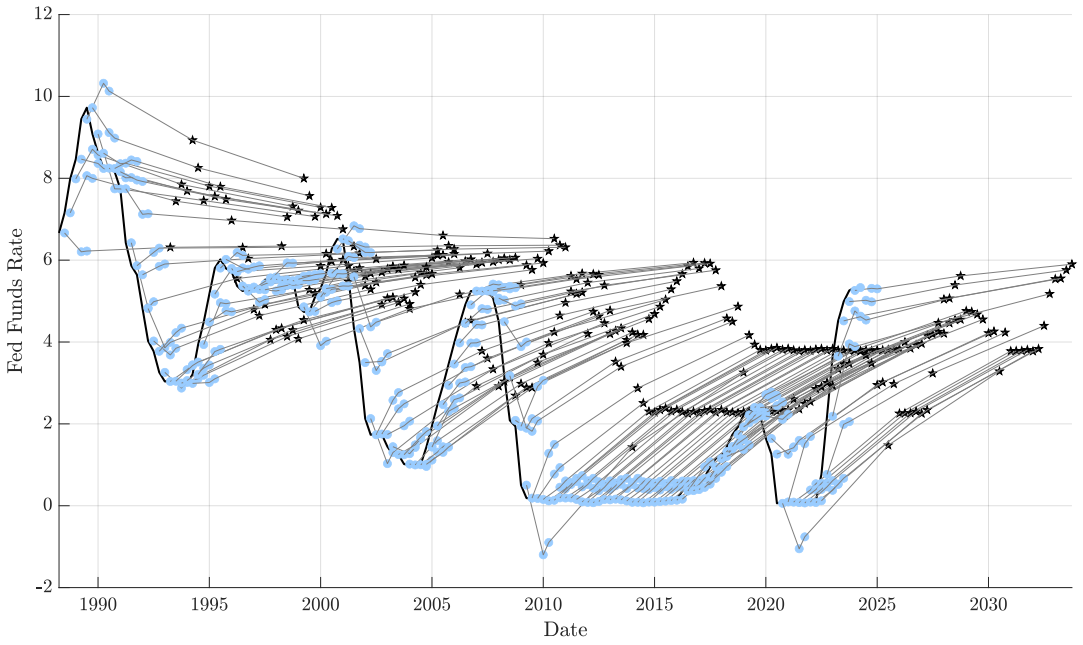


Figure 1. Fed Funds Rate and Forecasts

Notes: This figure plots the realized Fed Funds Rate with forecasts made for 1, 4, 8, 20, and 40 quarters. Data is quarterly and runs from Q1 1988 to Q4 2023.

[3] Our forecasts do not coincide with the actual Blue Chip Forecasts themselves, something that we test in a later section.

Table 1. Summary Statistics FFR Forecasts

	1Q	4Q	8Q	20Q	40Q
mean	3.12	3.34	3.62	4.29	5.06
AR(1)	0.98	0.98	0.98	0.99	0.99

Notes: This table reports mean and AR(1) coefficient for the estimated forecasts for the Fed Funds Rate at different horizons: 1, 4, 8, 20, and 40 quarters. Data is quarterly and runs from Q1 1988 to Q4 2023.

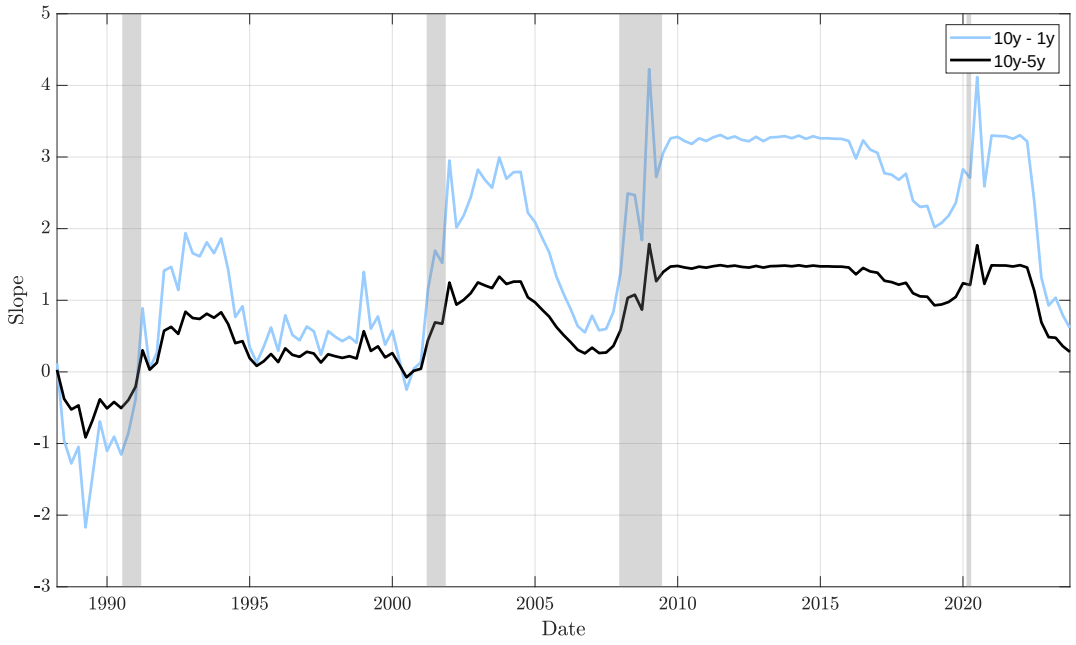


Figure 2. Fed Funds Rate and Forecasts

Notes: This figure plots the difference between the 40 quarter and (20) 4 quarter forecast of the Fed Funds Rate. Data is quarterly and runs from Q1 1988 to Q4 2023. Gray shaded areas represent recessions as defined by the NBER.

4.3 Testing Model's Underlying Assumptions

Before proceeding further, we test the empirical validity of Assumptions 1 and 2 that underpin the term structure of expectations constructed in the previous subsection in three ways.

To this end, we reconstruct the term structure of expectations, but this time only from the survey of short-term forecasts, namely, the recursive system of equations (6). The estimation procedure is identical to the one outlined in Subsection 4.1, except for the fact that we only use the moment condition corresponding to the shortest forecast horizon.

Having constructed the term structure of expectations from short-horizon survey data, we then first study the relationship between the constructed forecasts and their counterparts in the survey data for all observable horizons (except obviously for those corresponding to the shortest forecast horizon). For any forecast horizon $n \in \{2,3,4,5,6\}$ quarters available in the survey data, we then run the following regression:

$$\mathbb{E}_t^{\text{imp}}[r_{t+n}] - r_t = \alpha + \beta(\mathbb{E}_t^{\text{sur}}[r_{t+n}] - r_t) + \epsilon, \tag{9}$$

where $\mathbb{E}_t^{\text{imp}}[r_{t+n}]$ denotes the expected future short-term interest rate implied by Assumptions 1 and 2 (and hence constructed from equation (6)) while $\mathbb{E}_t^{\text{sur}}[r_{t+n}]$ is the corresponding forecast in survey data.

Figure 3 plots the relationship between the left- and the right-hand sides of regression (9), pooled across all forecast horizons $n \in \{2,3,4,5,6\}$. As the figure illustrates, there is a positive and significant relationship between the two variables, indicating that whenever expected future interest rates are higher than the current interest rate in survey data, so is the implied expected interest rate constructed under Assumptions 1 and 2.

Another test of the model's assumption is to regress the forecast errors, that is, the difference between actual and implied forecasts on the level of the short-rate:

$$\mathbb{E}_t^{\text{imp}}[r_{t+n}] - \mathbb{E}_t^{\text{sur}}[r_{t+n}] = \alpha + \beta r_t + \epsilon_t. \tag{10}$$

If Assumption 2 is satisfied, then the coefficient, β, of this regression should be equal to zero. Figure 4 depicts a pooled (across all horizons) scatter of the forecast error against the level of the short-rate. The estimated coefficient is -0.01 and the associated t-statistic is -0.98, indicating that the slope is not statistically different from zero.

A more stringent test of the model's assumptions is the Sargan–Hansen J test of the over-identifying restrictions. We use the observed forecasts for the current quarter, the next quarter, two quarters ahead, and three quarters ahead, thus ending up with $4(L+1)$ moments for $L+1$ parameters, where L denotes the number of lags used. Table 2 shows the test statistic and the corresponding p-value as we vary the number of lags. The results show that Assumptions 1 and 2 cannot be rejected using conventional significance thresholds.

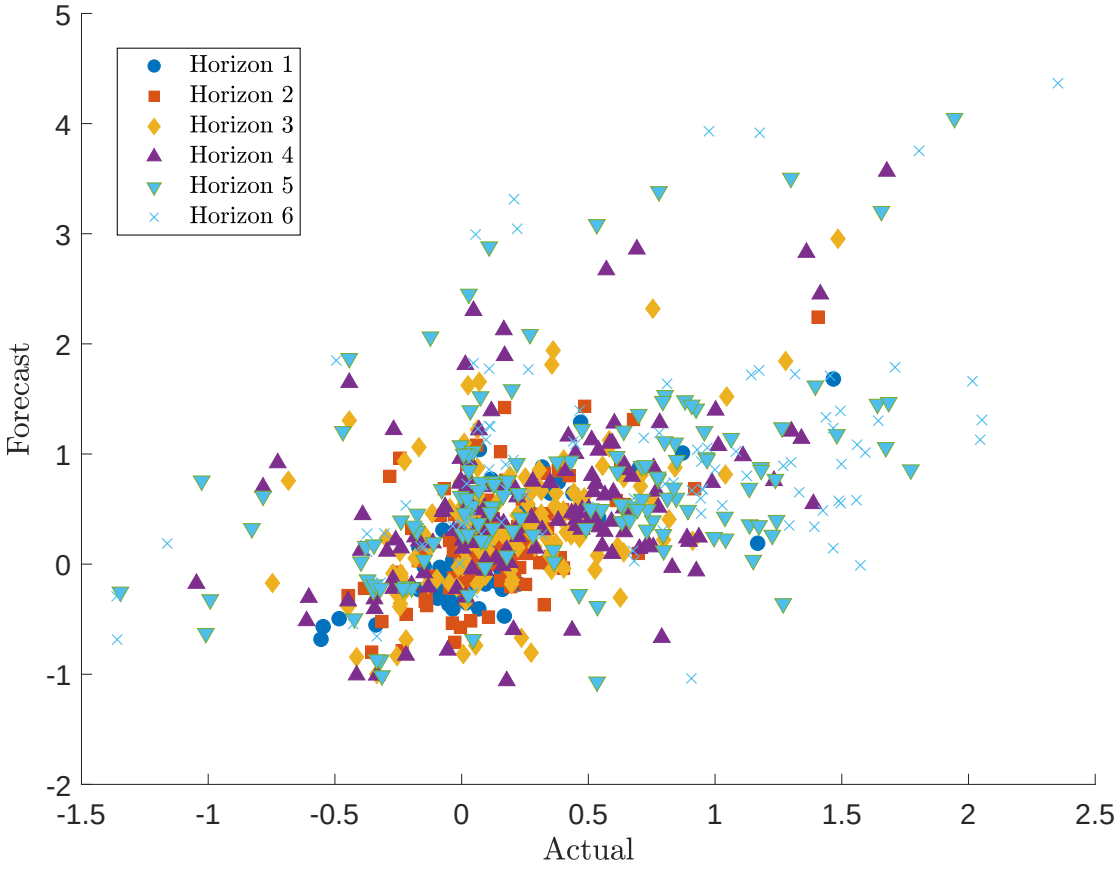


Figure 3. Actual vs Implied Forecasts

Notes: This figure plots actual forecasts for horizons 1 to 6 quarters from Blue Chip Financial Forecasts against implied forecasts for horizons 1 to 6 quarters. Data is quarterly and runs from Q1 1988 to Q4 2023.

Table 2. Sargan–Hansen J test

number of lags	2	3	4	5	6
test statistic	16.48	19.30	23.59	26.68	29.48
p value	0.057	0.081	0.072	0.085	0.103

Notes: This table reports the test statistic and p value for the Sargan-Hansen J test of model's over-identifying restrictions. Iterated GMM.

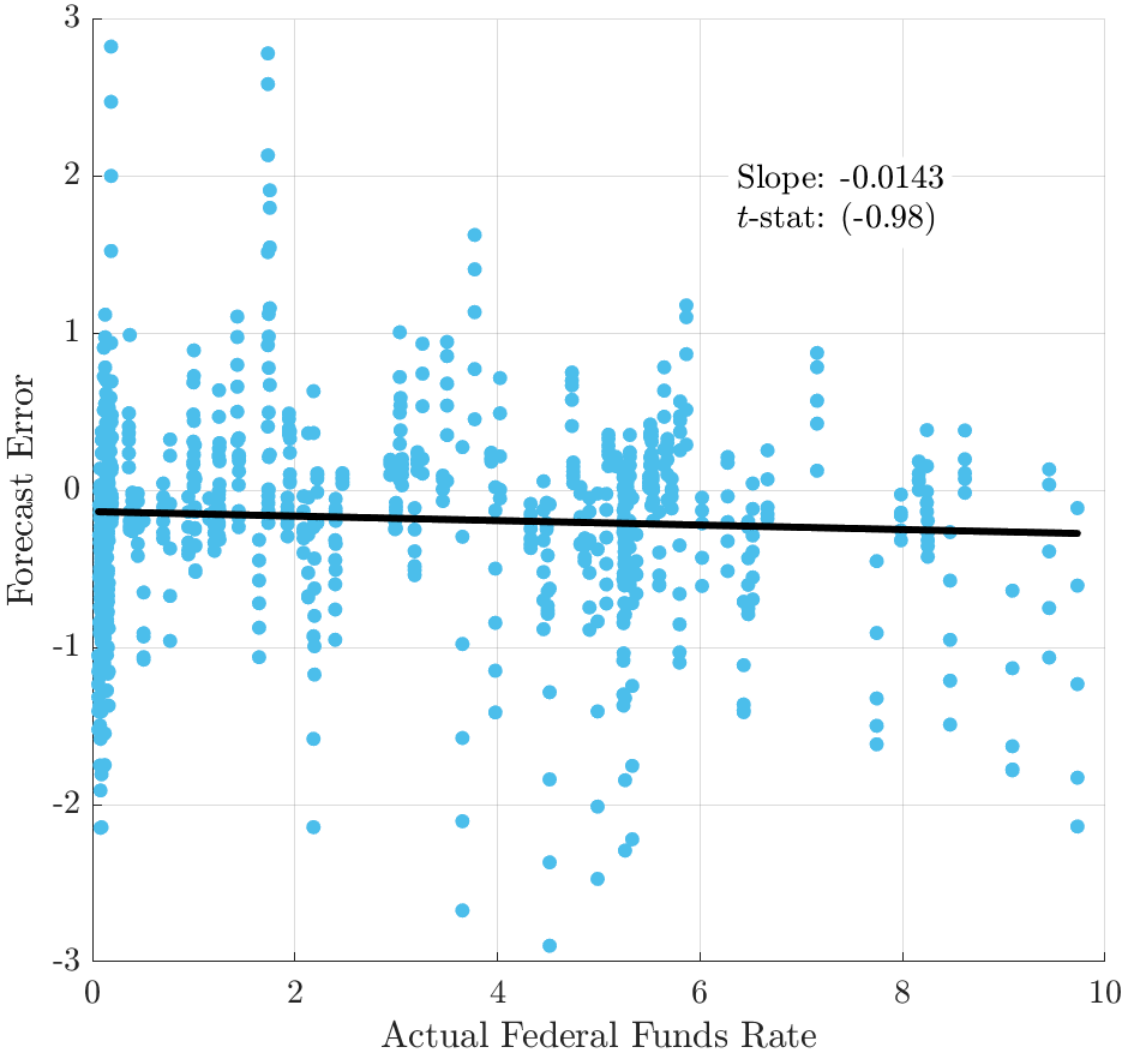


Figure 4. Forecast Error vs Realized Short-Rate

Notes: This figure plots forecast errors defined as the difference between the model-implied and forecasted FFR for different horizons, 1 to 6. Forecasted short-rates are from Blue Chip Financial Forecasts. Data is quarterly and runs from Q1 1988 to Q4 2023.

4.4 Model-Implied Yields

We now turn to studying yields and forward rates as implied by the expectations hypothesis. To this end, we calculate implied yields from our constructed path of future Fed Funds Rate:

$$y_t^{EH,(n)} = \frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}(r_{t+k}).$$

Similarly, we construct forward rates implied by the expectations hypothesis such that the forward rate equals the average expected nominal short-term interest rate over n months starting m months:

$$f_t^{EH,(n,m)} = \frac{1}{n}\sum_{k=m+1}^{n+m}\mathbb{E}(r_{t+k-1}). \tag{11}$$

Figure 5 plots the model-implied yields for 1y, 2y, 5y and 10y maturity. Not surprisingly, we find that the match at the 1 and 2y maturity is almost perfect, while the difference increases with maturity.

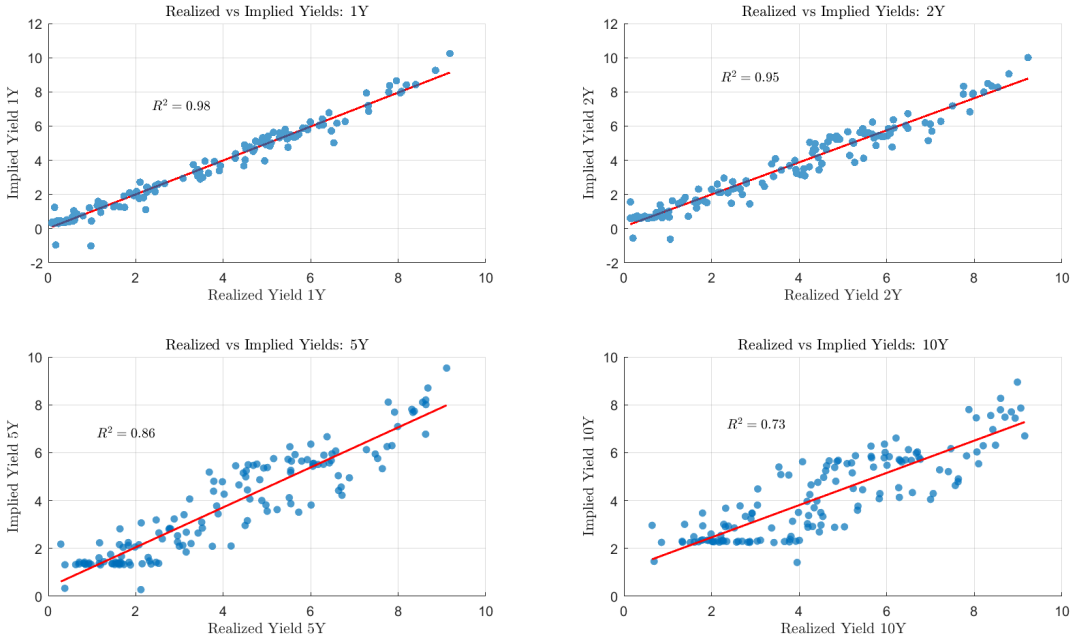


Figure 5. Expectation Hypothesis Implied Yields vs Realized Yields

Notes: This figure plots model-implied yields vis-à-vis realized yields from the data. Data is quarterly and runs from Q1 1988 to Q4 2023.

4.5 Wedges

We now turn to our main result. To this end, we calculate wedges as the difference between the realized yield with maturity n and the model-implied yield:

$$\text{wedge}_t = y_t^{(n)} - \frac{1}{n}\mathbb{E}_t\left(\sum_{k=0}^{n-1} r_{t+k}\right)$$

Figure 6 plots wedges for the 10y maturity together with the term premium estimate from Adrian et al. (2013). There are two noteworthy observations. First, the wedge is highly correlated with parametric estimates of the term premium.[4] Indeed, the correlation is over 90% over the full sample. Second, we also notice that the wedge is always smaller than the term premium estimate itself. In fact, it consistently turns negative before recessions and has remained negative since 2015.

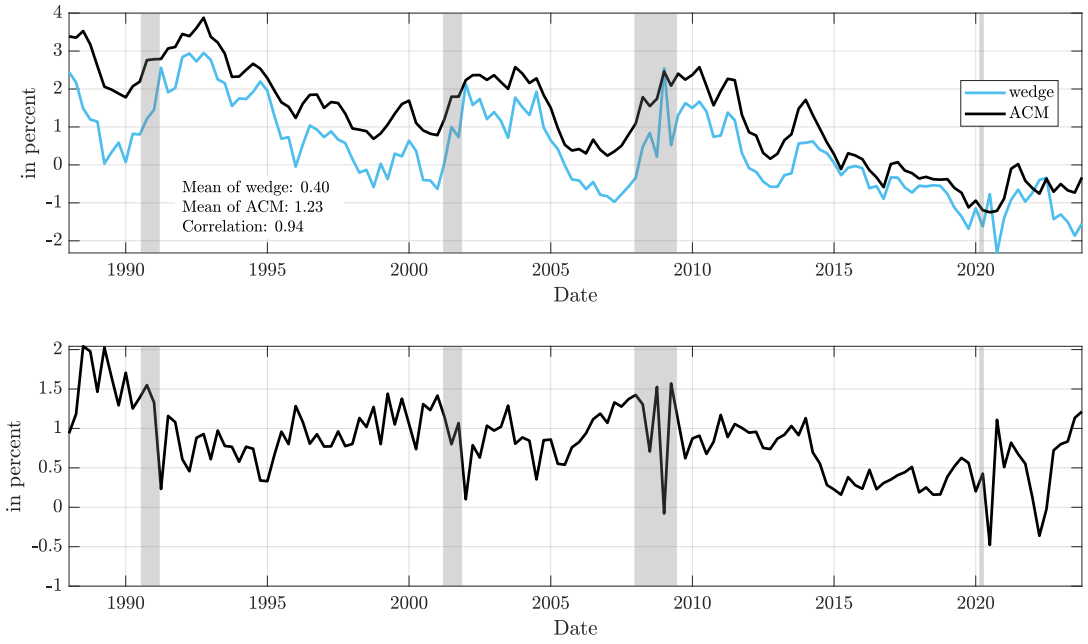


Figure 6. Expectation Hypothesis Implied Wedges vs Term Premia

Notes: The upper panel plots the 10y wedge, defined as the difference between the realized yield with maturity 10y and the implied yields together with the term premium from Adrian et al. (2013). The lower panel plots the difference between the term premium and the wedge. Data is quarterly and runs from Q1 1988 to Q4 2023. Gray shaded areas represent recessions as defined by the NBER.

4.6 Testing the Expectations Hypothesis

As a first test of the expectations hypothesis, we regress model-implied forward rates from equation (11) on realized forward rates. Since forward rates potentially contain a unit-root, we de-trend our series by subtracting the realized Fed Funds Rate at time t. More specifically, we run the following regression:

$$f_t^{EH,(n,m)} - r_t = \alpha_{n,m} + \beta_{n,m}\left(f_t^{(n,m)} - r_t\right) + \epsilon_t^{(n,m)},$$

where r_t is the realized Fed Funds Rate at time t. In our analysis, we focus on one-year rates in n years in the future. That is, $m = 4$, and $n = 2, \ldots, 9$. The regression coefficients $\beta_{n,m}$ are reported in Table 3.

[4]The results are quantitatively almost the same independent of the particular term premium estimate we use. For the Christensen and Rudebusch (2012) term premium, we find that the correlation between the 10y term premium and the wedge is over 80%.

Table 3. Forward-Rate Regressions

1Y	2Y1Y	3Y1Y	4Y1Y	5Y1Y	6Y1Y	7Y1Y	8Y1Y	9Y1Y
0.291	0.410	0.502	0.601	0.651	0.657	0.700	0.742	0.710

Notes: This table reports estimated slope coefficients from a regression that regresses model-implied forward rates (minus the realized Fed Funds Rate) on realized forward rates (minus the realized Fed Funds Rate). Data is quarterly and runs from Q1 1988 to Q4 2023.

5 Conclusion

Standard tests of expectations hypothesis of interest rates—such as the well-known regressions of Fama (1984) and Campbell and Shiller (1991)—are joint tests of expectations and rational expectations hypotheses. As such, a rejection of EH under such tests may indicate a violation of EH, RE, or both.

In this paper, we study deviations from the expectations hypothesis of the term structure of interest rates in an environment where agents may have specified beliefs about the process that generates the short-term interest rate. To this end, we first develop a theoretical framework to estimate agents' subjective expectations of the term structure of interest rates from a handful of observed forecasts from survey data. We show that under two parsimonious assumptions, namely linearity and internal consistency of subjective expectations, there is a tight relationship between agents' forecasts at different horizons. The characterization of these relationships is what allows us to construct the entire term structure of short-term interest rates from survey data, and as a consequence, the yields implied by agents' subjective expectations for all potential maturities.

Using a rich set of surveys of professional forecasters about the short-rate, we document the following facts. First, we show that model-implied subjective expectations of shot-term interest rates coincide with those observed in survey data. Specifically, we find that whenever expected future interest rates are higher than the current interest rate in survey data, so is the corresponding implied expected interest rate constructed under our two assumptions of linearity and internal consistency. We interpret this finding as evidence supporting the assumptions underlying the construction of forecasts. Second, comparing our model-implied yields to those in the data, we find no significant differences at shorter maturities. However, the wedge between two yields increases with the bond's maturity. In addition, we find that the difference between the 10-year realized yield and our model-implied yield is almost perfectly correlated with measures of the term premium but significantly smaller in size. Finally, we turn to testing the expectations hypothesis. We find that the yields implied by EH under subjective expectations closely match yield data, indicating that time-varying risk premiums are quantitatively small.

Appendix

A Proof of Proposition 1

Start with the observation that, by Assumption 1, time-t forecasts at horizon $\tau + k$ can be expressed as follows:

$$\mathbb{E}_t[x_{t+\tau+k}] = \phi_0^{(\tau+k)} + \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+1}^{(\tau+k)} x_{t-s}. \tag{12}$$

By the same logic,

$$\begin{aligned}
\mathbb{E}_t\mathbb{E}_{t+k}[x_{t+k+\tau}] &= \phi_0^{(\tau)} + \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+1}^{(\tau)} \mathbb{E}_t[x_{t+k-s}] \\
&= \phi_0^{(\tau)} + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \mathbb{E}_t\left[x_{t+k-r}\right] + \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+k+1}^{(\tau)} x_{t-s} \\
&= \phi_0^{(\tau)} + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \left(\phi_0^{(k-r)} + \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+1}^{(k-r)} x_{t-s} \right) + \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+k+1}^{(\tau)} x_{t-s}
\end{aligned}$$

and as a result,

$$\mathbb{E}_t\mathbb{E}_{t+k}[x_{t+k+\tau}] = \left[\phi_0^{(\tau)} + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \phi_0^{(k-r)} \right] + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+1}^{(k-r)} x_{t-s} + \sum_{s=0}^{\infty} \mathbf{\Phi}_{s+k+1}^{(\tau)} x_{t-s} \tag{13}$$

Next, note that, by Assumption 2, the expressions in equations (12) and (13) must be identical for all possible realizations of x_t. Equating the coefficients on both sides of the equations implies that

$$\begin{aligned}
\phi_0^{(\tau+k)} &= \phi_0^{(\tau)} + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \phi_0^{(k-r)} \\
\mathbf{\Phi}_s^{(\tau+k)} &= \mathbf{\Phi}_{s+k}^{(\tau)} + \sum_{r=0}^{k-1} \mathbf{\Phi}_{r+1}^{(\tau)} \mathbf{\Phi}_s^{(k-r)}
\end{aligned}$$

for all $s \geq 1$. □

References

Adrian, Tobias, Richard Crump, and Emanuel Moench (2013), "Pricing the term structure with linear regressions." *Journal of Financial Economics*, 110(1), 110–138.

Afrouzi, Hassan, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar (2023), "Overreaction in expectations: Evidence and theory." *The Quarterly Journal of Economics*, 138(3), 1713–1764.

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry (2021), "Imperfect Macroeconomic Expectations: Evidence and Theory." *NBER Macroeconomics Annual*, 35(1), 1–86.

Bansal, Ravi and Ivan Shaliastovich (2013), "A long-run risks explanation of predictability puzzles in bond and currency markets." *The Review of Financial Studies*, 26(1), 1–33.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020), "Overreaction in Macroeconomic Expectations." *American Economic Review*, 110(9), 2748–82.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018), "Diagnostic Expectations and Credit Cycles." *Journal of Finance*, 73(1), 199–227.

Broer, Tobias and Alexandre N. Kohlhas (2024), "Forecaster (Mis-)Behavior." *The Review of Economics and Statistics*, 106(5), 1334–1351.

Campbell, J. Y. and R. J. Shiller (1991), "Yield spreads and interest rate movements: A bird's eye view." *Review of Economic Studies*, 58, 495–514.

Christensen, Jens H.E. and Glenn D. Rudebusch (2012), "The response of interest rates to u.s. and u.k. quantitative easing." *Economic Journal*, 122, 385–414.

Coibion, Olivier and Yuriy Gorodnichenko (2015), "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8), 2644–78.

Crump, Richard, Stefano Eusepi, and Emanuel Moench (2024), "Is there hope for the expectations hypothesis?" *New York Federal Reserve Working Paper No. 1098*.

Fama, E. (1984), "The information in the term structure." *Journal of Financial Economics*, 13, 509–528.

Farhi, Emmanuel and Iván Werning (2019), "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review*, 109(11), 3887–3928.

Farmer, Leland, Emi Nakamura, and Jón Steinsson (2023), "Learning about the long run." *Journal of Political Economy*, 132, 3334–3377.

Froot, Kenneth A. (1989), "New hope for the expectations hypothesis of the term structure of interest rates." *The Journal of Finance*, 44(2), 283–305.

Fuster, Andreas, Benjamin Hebert, and David Laibson (2012), "Investment Dynamics with Natural Expectations." *International Journal of Central Banking*, 8(SUPPL. 1), 243–265.

Fuster, Andreas, David Laibson, and Brock Mendel (2010), "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives*, 24(4), 67–84.

Gabaix, Xavier (2020), "A Behavioral New Keynesian Model." *American Economic Review*, 110, 2271–2327. Working paper.

García-Schmidt, Mariana and Michael Woodford (2019), "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." *American Economic Review*, 109(1), 86–120.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright (2007), "The u.s. treasury yield curve: 1961 to the present." *Journal of Monetary Economics*, 54(8), 2291–2304.

Kim, D and Jonathan Wright (2005), "An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates." *Working Paper No. 2005-33*.

Vayanos, Dimitri and Jean-Luc Vila (2021), "A preferred-habitat model of the term structure of interest rates." *Econometrica*, 89(1), 77–112.