Expectations and the Term Structure of Interest Rates*

Pooya Molavi† Alireza Tahbaz-Salehi‡ Andrea Vedolin§

August 2026

Abstract

This paper studies the relationship between investors' subjective expectations and the term structure of interest rates. Departing from rational expectations, we allow investors to hold arbitrary—and potentially heterogeneous—beliefs about future interest rates, and derive the relationship that expected and realized interest rates must satisfy irrespective of how expectations are formed. We use this relationship to decompose the sensitivity of long-term yields to changes in the short rate into contributions from expectations and risk premia, without imposing any assumption on expectation formation. We find that, for short- and medium-maturity bonds, this sensitivity is driven primarily by movements in expectations rather than risk premia. We then study the sensitivity of expectations themselves and show that investors' subjective expectations across different maturities and forecast horizons are inconsistent with the structural relationships linking different interest rates. Specifically, we find evidence of overreaction: investors' short-horizon expectations of long-term yields overreact relative to their long-horizon expectations of short-term yields.

JEL Classification: E43, D84, G41
Keywords: subjective expectations, term structure of interest rates

*We are grateful to Camilla Silvestri and Samir Trabelsi for excellent research assistance. We thank Marios Angeletos, Leland Farmer, Christian Heyerdahl-Larsen, Matteo Leombroni, Yueran Ma, Roberto Marfè, Stefan Nagel, Hao Pang, Monika Piazzesi, Carolin Pflueger, Juliana Salomao, Martin Schneider, Yeji Sung, and Chen Wang for helpful comments and suggestions.

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

1 Introduction

The sensitivity of long-term interest rates to movements in the short rate plays a central role in the transmission of monetary policy. While central banks' primary policy instrument is the short-term interest rate, many economically relevant rates—such as mortgage rates and corporate borrowing costs—are determined by longer-term interest rates that are largely outside the central bank's direct control. Understanding how changes in the short rate pass through to long-term yields is therefore essential for assessing the effectiveness of monetary policy.

A natural framework for studying this sensitivity is the expectations hypothesis of the term structure of interest rates. In its canonical form, which assumes rational expectations, the yield $y_t^{(n)}$ on a bond that matures in n periods satisfies

$$y_t^{(n)} = \frac{1}{n}\sum_{s=0}^{n-1} \mathbb{E}_t^*[y_{t+s}^{(1)}] + \text{risk premium}, \tag{1}$$

where $\mathbb{E}_t^*[\cdot]$ denotes expectations under the true data-generating process. That is, the yield on a long-term bond is equal to the average of expected future short-term interest rates over the bond's lifetime plus a term representing the bond's risk premium. Seen through the lens of this relationship, the sensitivity of long-term yields to changes in the short rate operates through two distinct channels: revisions in investors' expectations of future short rates and adjustments in the risk premium (see, for instance, Hanson and Stein (2015) and Hanson, Lucca, and Wright (2021)).

Understanding the mechanism underlying the sensitivity of long-term yields to the short rate, however, requires taking a stand on how expectations are formed. If the assumed expectation formation process is misspecified, the decomposition into expectations and risk-premium components may misattribute the sensitivity between the two channels. For example, what appears to be a movement in the risk premium may in fact reflect a departure from rational expectations, or vice versa.

In this paper, we examine how investors' subjective expectations shape the term structure of interest rates. Rather than imposing rational expectations, we allow investors to hold arbitrary beliefs about future interest rates, and derive the relationship that expected and realized interest rates must satisfy irrespective of how expectations are formed. We use this relationship to decompose the sensitivity of long-term yields to changes in the short rate into contributions from expectations and risk premia, without imposing any assumptions on expectation formation. We find that, for short- and medium-maturity bonds, this sensitivity is driven primarily by movements in expectations rather than risk premia. We then study the sensitivity of expectations themselves and show that investors' subjective expectations across maturities and forecast horizons are inconsistent with the equilibrium relationships linking different interest rates.

We derive our results in the context of a simple discrete-time economy populated by long-lived investors who can hold positions in bonds of varying maturities. Unlike most of the prior literature, we do not impose rational expectations; instead, we allow investors to hold arbitrary and potentially heterogeneous beliefs about future bond prices. Specifically, we focus on what is often referred to as the economy's "temporary equilibrium," which imposes individual rationality and market clearing but treats expectations as model primitives (i.e., without modeling how they are formed). Accordingly,

heterogeneity in investors' expectations in our framework may stem from different information sets, different models for how various (endogenous or exogenous) variables evolve, different understandings of the economy's structural equations, or any combination thereof.

Using the model, we characterize how investors' subjective beliefs shape the term structure of interest rates to a second-order approximation. We show that, as long as disagreement among investors is of the same order of magnitude as their subjective uncertainty—a condition necessary for equilibrium existence—bond prices can be expressed in terms of the short-term interest rate, terms capturing risk premia, and investors' subjective expectations of bonds' resale values next period. These expectations enter prices through a specific aggregation: bonds are priced as if the economy were populated by a representative investor holding *consensus expectations*, defined as a weighted average of individual expectations with greater weights placed on investors with higher risk-bearing capacities.

This characterization leads to an intuitive but important implication: irrespective of how expectations are formed, realized and expected bond yields must satisfy

$$y_t^{(n)} = \frac{1}{n} y_t^{(1)} + \frac{n-1}{n} \overline{\mathbb{E}}_t[y_{t+1}^{(n-1)}] + \text{risk premium}, \tag{2}$$

where $\overline{\mathbb{E}}_t[\cdot]$ denotes consensus expectations. Crucially, this is the only relationship among interest rates and expectations that follows from individual portfolio optimization and market clearing. Any other relationship, including the standard formulation of the expectations hypothesis in (1) or its variants with subjective expectations, necessarily rests on additional assumptions on how expectations are formed.

Since equation (2) holds irrespective of how expectations are formed, it implies a decomposition of the sensitivity of long-term yields to the short rate into expectation and risk-premium components that is robust to misspecification of the expectation formation process. Implementing this decomposition using survey data from Blue Chip Financial Forecasts (BCFF) and Treasury yield data, we find that for short- and medium-maturity bonds, the risk-premium component accounts for only a small share of the overall sensitivity, with movements in expectations responsible for most of the response. This changes markedly at longer maturities, where movements in risk premia account for a nontrivial share of the total sensitivity to the short rate.

While our decomposition highlights the importance of movements in yield expectations, equation (2)—on which this decomposition rests—is silent on how expectations themselves respond to changes in the short rate. We therefore turn to studying the sensitivity of expectations directly. We show that, under the additional assumption that investors understand the equilibrium relationship linking interest rates, the sensitivity of expected yields of any maturity to the short rate must equal a specific weighted average of the sensitivities of expected yields of shorter maturities at staggered horizons. This consistency restriction places tight constraints on how investors' subjective expectations across different maturities and forecast horizons can respond to changes in the short rate.

Empirically, we document pervasive violations of the consistency restrictions discussed above across maturities and forecast horizons. Moreover, we find that these violations point in a systematic direction: investors' short-horizon expectations of long-term yields overreact relative to their own long-horizon expectations of short-term yields. This finding is distinct from the commonly documented evidence that investors' forecast errors are predictable: it concerns the internal consistency of investors' beliefs

across the yield curve, rather than their accuracy relative to the true data-generating process.[1]

We conclude by developing a model of expectation formation to illustrate how such cross-maturity internal inconsistencies may arise endogenously. We model investors as forming expectations using a statistical model of the yield curve estimated from historical data and show that when investors' models are misspecified—in the sense that the true data-generating process does not fall within the class of statistical models they consider—their expectations may be inconsistent with the equilibrium relationships linking different interest rates. Crucially, these violations arise even though expectations are endogenously determined and optimally estimated from realized yield data that satisfy those very relationships. We also show that while richer, but still misspecified, statistical models can improve forecast accuracy, this may come at the expense of consistency with the structural restrictions implied by the bond market. Thus, rather than reflecting suboptimal forecasting, these violations can emerge as a byproduct of forming optimal forecasts under model misspecification.

Related Literature. Our paper belongs to the literature that uses survey data to study the relationship between investors' subjective expectations and the term structure of interest rates. Early contributions include Friedman (1979) and Froot (1989), who argue that empirical rejections of the expectations hypothesis may reflect violations of rational expectations rather than of the hypothesis itself. More recently, Piazzesi, Salomao, and Schneider (2015) and Nagel and Xu (2023) document that subjective bond return expectations extracted from survey data are less cyclical than those implied by predictive regressions using realized returns, while Leombroni, Pflueger, and Sunderam (2026) document that subjective expectations of excess returns on long-term government bonds are tightly linked to the bonds' stock market betas.[2] We contribute to this literature by deriving the relationships between realized and expected interest rates that hold irrespective of how expectations are formed, and by using survey data to study how investors' expectations across different maturities and forecast horizons relate to one another.

Our paper is also related to the literature that studies the sensitivity of long-term interest rates to monetary policy shocks by decomposing yield movements into expectations and risk-premium components, including Gertler and Karadi (2015), Gilchrist, López-Salido, and Zakrajšek (2015), Hanson and Stein (2015), and Abrahams et al. (2016). Despite differences in methodology, all of these papers construct the expectations component by taking a stand on how investors form their beliefs about future interest rates, either explicitly assuming rational expectations or using model-implied forecasts as stand-ins for investors' subjective beliefs.[3] We contribute to this literature by dispensing with any

[1] This form of internal inconsistency also has direct implications for the decomposition of the sensitivity of long-term yields to the short rate into expectations and risk-premium components. Decompositions based on (1)—even those that replace rational expectations with survey-based forecasts—implicitly assume that investors' expectations across maturities and forecast horizons are consistent with the equilibrium restrictions implied by the bond market. Our finding that this assumption may be violated underscores the importance of the decomposition based on (2), which requires no such assumption.

[2] Also see Piatti and Whelan (2026), who find that subjective excess returns are on average negative and vary countercyclically with subjective expected consumption growth, and Bauer, Pflueger, and Sunderam (2024), who estimate the perceived monetary policy rule from survey forecasts and show that time variation in these perceptions shapes the sensitivity of long-term rates to macroeconomic news.

[3] Using a regression-based approach, Hanson and Stein (2015) attribute the predictability of distant forward rate movements to risk premia, an interpretation that hinges on rational expectations. Gertler and Karadi (2015) and Abrahams et al. (2016) instead construct the expectations component using model-implied forecasts of future interest rates: VAR-implied forecasts of future short rates in the former, and forecasts implied by an estimated affine term structure model in the latter. In both cases, these forecasts coincide with investors' subjective expectations only under the assumption that investors form their beliefs in

such assumption, decomposing the sensitivity of long-term yields to the short rate into expectations and risk-premium components without imposing restrictions on how expectations are formed. We also examine whether investors' expectations across maturities and forecast horizons are consistent with the equilibrium relationships linking different interest rates, and find that they are not.

Our paper also contributes to the body of work on over- and under-reaction of expectations in macroeconomics and finance (Coibion and Gorodnichenko, 2015; Bordalo, Gennaioli, Ma, and Shleifer, 2020). In the context of the bond market, and thus more directly related to our work, Cieslak (2018), d'Arienzo (2020), and Wang (2021) document systematic errors in market participants' forecasts of future interest rates. The common feature of these studies is their focus on how subjective forecasts relate to the true data-generating process. In contrast, we document a form of overreaction that is internal to investors' mental models of the yield curve: investors' short-horizon expectations of long-term yields overreact relative to their own long-horizon expectations of short-term yields, a finding that concerns the consistency of investors' beliefs across the yield curve rather than how those beliefs relate to realized outcomes.

A more recent strand of the literature examines whether the failure of the expectations hypothesis stems from misspecification in investors' beliefs about the process driving the short-term interest rate (Crump, Eusepi, and Moench, 2026), or from investors' inability to learn the parameters of that process within a reasonable timeframe (Farmer, Nakamura, and Steinsson, 2024). Like these papers, our work allows for departures from full-information rational expectations. However, whereas these studies focus on expectational errors about the short rate and assume that short- and long-term rates are linked through equation (1) with subjective expectations replacing rational ones, we focus on the relationship between market participants' expectations across different maturities and forecast horizons and find that these expectations are not necessarily consistent with (1).

Also related is the recent work by Shue, Townsend, and Wang (2025), who document that movements in realized long-term yields and professional forecasters' yield expectations are inconsistent with the expectations hypothesis. We depart from their work along two dimensions. First, our approach nests the restriction tested by Shue, Townsend, and Wang (2025) as a special case, while covering a broader set of consistency conditions across maturities and forecast horizons. Second, whereas they attribute the documented pattern to forecasters' tendency to lump different interest rates into a single, coarse category of "interest rates," we show that violations of this kind can also arise as a byproduct of optimally estimating a misspecified, but otherwise unbiased, statistical model of the yield curve.

Another related work is Giglio and Kelly (2018), who document that long-maturity prices across a range of asset classes are more volatile than are warranted by the behavior of short-maturity prices under an affine-Q specification. Since their test is based on prices rather than direct measures of expectations, the affine structure serves as a bridge that translates short-end price movements into a model-implied benchmark against which long-maturity prices are evaluated. By using survey-based forecasts directly, we instead test whether expectations across maturities and forecast horizons are consistent with the equilibrium relationships linking different interest rates, without specifying the dynamics of the pricing measure.

We also contribute to the theoretical literature on heterogeneous beliefs in bond markets, such

accordance with that same model.

as Xiong and Yan (2010), Ehling et al. (2018), and Buraschi and Whelan (2022). These papers model belief heterogeneity by assuming that investors "agree to disagree" about economic fundamentals, such as future interest rates or inflation. In addition, for analytical tractability, they restrict heterogeneity to two types (e.g., optimists and pessimists) and impose parametric assumptions on preferences and the underlying data-generating process. The tractability of our framework allows us to relax these assumptions. Using a second-order approximation, we characterize how investors' expectations, risk-bearing capacities, and subjective uncertainty jointly shape the term structure of interest rates—without imposing parametric assumptions, specifying the source of belief heterogeneity, or limiting the number of investor types.

Finally, our paper is related to the literature that studies the implications of constraints on agents' cognitive and computational abilities for asset prices. Using a simple behavioral framework, we argue that if investors form their expectations using misspecified statistical models, they may fail to uncover the structural relationships linking yields of different maturities, even when their models are estimated optimally from data that satisfy those very relationships. In this sense, our paper is related to Fuster, Hébert, and Laibson (2012), Molavi, Tahbaz-Salehi, and Vedolin (2024), and Molavi (2026), who study how model misspecification shapes equilibrium outcomes; to Eyster and Piccione (2013), who consider an asset pricing environment in which investors lack structural knowledge about the market; and to Gabaix (2014), who emphasizes the role of sparsity in agents' models in generating departures from full rationality. Unlike these papers, however, our focus is squarely on how such constraints on investors' statistical models shape the relationship between their expectations across different maturities and forecast horizons.

Outline. The rest of the paper is organized as follows. We present our conceptual framework in Section 2 and develop our theoretical results in Section 3. Section 4 contains our main empirical findings, while Section 5 presents the behavioral framework for yield curve expectation formation. The proofs and some additional theoretical and empirical results are presented in the Appendix.

2 Model

Consider a discrete-time, infinite-horizon economy, with time indexed by $t \in \{0, 1, \dots\}$. In every period, there are zero-coupon, default-free bonds with face value one that mature in $n \geq 1$ periods. Let $s_t^{(n)}$ denote the net supply of the n-period bond at time t, which is assumed to be inelastic for all $n \geq 2$. The short-term bond is supplied elastically at price $p_t^{(1)}$, which is set by the central bank. The short-term interest rate is thus given by $y_t^{(1)} = -\log p_t^{(1)}$.

The economy is populated by a unit mass of long-lived, risk-averse bond investors, indexed by $i \in [0, 1]$, who choose their bond holdings and can readjust their portfolios each period. Investor i enters the market at $t = 0$ with initial wealth $\widetilde{w}_i$. At any future point in time, the investor exits the market with probability $1 - \rho$ and is replaced by a new investor—also indexed i—who enters with the same initial wealth $\widetilde{w}_i$. As in Vayanos and Vila (2021) and Kekre, Lenel, and Mainardi (2026), we assume investors derive utility from their terminal nominal wealth at the time of exit. Unlike these papers, however, we allow investors to have arbitrary and potentially different expectations about future bond prices.

These differences in expectations may reflect investors' different information sets, different models for how various (endogenous or exogenous) variables evolve, different understandings of the economy's structural equations, or some combination thereof.

The investor indexed i thus faces the following optimization problem at time t:

$$\begin{aligned} \max_{x_{i\tau}^{(n)}} \quad & (1-\rho)\sum_{\tau=t}^{\infty}\rho^{\tau-t}\mathbb{E}_{it}[u_i(w_{i\tau+1})] \\ \text{s.t.} \quad & w_{i\tau+1} = w_{i\tau}/p_\tau^{(1)} + \sum_{n=2}^{\infty} x_{i\tau}^{(n)}(p_{\tau+1}^{(n-1)} - p_\tau^{(n)}/p_\tau^{(1)}), \end{aligned} \tag{3}$$

where $p_\tau^{(n)}$ is the time-τ price of the bond that matures in n periods, $x_{i\tau}^{(n)}$ denotes i's position in the bond, $w_{i\tau}$ is i's wealth at time τ, and $u_i(\cdot)$ denotes the investor's strictly concave and thrice-differentiable utility function. Operator $\mathbb{E}_{it}[\cdot]$ in (3) represents investor i's subjective expectation, which—unless otherwise noted—we treat as a model primitive that need not coincide with rational expectations. In particular, $\mathbb{E}_{it}[\cdot]$ may be inconsistent with the economy's structural equations (such as market clearing) as well as with the processes driving economic fundamentals (such as shocks to the short-term interest rate or to the supply of bonds). While we do not take a stance on how expectations are formed, we assume that $\mathbb{E}_{it}[\cdot]$ represents a valid expectation operator: it is linear, satisfies $\mathbb{E}_{it}[1] = 1$, and $\mathbb{E}_{it}[z] \geq 0$ whenever $z \geq 0$. We also assume that investors' beliefs are internally consistent, in the sense that each investor's expectations satisfy the law of iterated expectations, that is, $\mathbb{E}_{it}\mathbb{E}_{it+h}[\cdot] = \mathbb{E}_{it}[\cdot]$ for all $h \geq 0$ and all t.[4]

To analyze how investors' expectations shape the term structure of interest rates, we approximate the economy around a deterministic steady state in which (i) the short-term interest rate is constant and equal to zero in all periods, and (ii) investors face no uncertainty about future bond prices. Formally, we parameterize investors' *subjective* uncertainty by assuming

$$\operatorname{cov}_{it}(\log p_{t+h}^{(n)}, \log p_{t+s}^{(m)}) = \sigma^2 \Sigma_{t,h,s}^{(n,m)} \tag{4}$$

and approximate the economy as $\sigma \to 0$.[5] As σ approaches zero, uncertainty about future prices vanishes at a common rate σ^2. Throughout, we assume that the subjective conditional covariance matrix of future log prices is invertible when $\sigma > 0$.[6]

We also impose a restriction on the extent of disagreement among investors: if investors disagree with one another in the limit as uncertainty vanishes, they would want to take infinitely sized bets against one another, thereby ruling out equilibrium existence. We thus restrict the extent of disagreement by assuming that for any random variable z that decays to zero at rate σ^k, the disagreement $\mathbb{E}_{it}[z] - \mathbb{E}_{jt}[z]$

[4] The law of iterated expectations is a consistency restriction on expectations of the *same* variable, i.e., $\mathbb{E}_{it}\mathbb{E}_{it+h}[z] = \mathbb{E}_{it}[z]$ for any z. It places no requirement on investor i's expectations of distinct variables—say, $p_{t+s}^{(n)}$ and $p_{t+\ell}^{(m)}$—being consistent or even related to one another. We impose this assumption to ensure that, given an investor's expectations, the corresponding portfolio optimization problem is dynamically consistent.

[5] The assumption that the short-term interest rate is equal to zero in steady state corresponds to $\lim_{\sigma\to 0} \log p_t^{(1)} = 0$. We also assume that the subjective expectation of any random variable z whose realizations are of order σ^k also decays to zero at rate σ^k. This technical condition guarantees that the log-quadratic approximation of equilibrium prices depends only on the first two moments of investors' beliefs.

[6] Since the right-hand side of (4) is independent of i, this specification also implies that all investors face the same subjective uncertainty at each point in time. We make this assumption for simplicity. As we discuss below, the characterization result in Proposition 1 generalizes to an economy with heterogeneous subjective uncertainty.

between any two investors $i \neq j$ about that variable decays to zero at rate σ^{k+1} or faster. Specializing to $z = \log p_{t+h}^{(n)}$, which decays at rate σ, this implies

$$\lim_{\sigma \to 0} \frac{1}{\sigma^2} (\mathbb{E}_{it}[\log p_{t+h}^{(n)}] - \mathbb{E}_{jt}[\log p_{t+h}^{(n)}]) < \infty \tag{5}$$

for all $i \neq j$, so that investors' disagreement about future bond prices is at most of the same order of magnitude as their subjective uncertainty. This condition ensures that investors' bond-holding positions remain finite even as uncertainty vanishes.

We now characterize equilibrium bond prices in terms of model primitives. Let

$$\theta_{it} = -u_i''(w_{it})/u_i'(w_{it}) \qquad \text{and} \qquad \gamma_{it} = -w_{it} u_i''(w_{it})/u_i'(w_{it}) \tag{6}$$

denote investor i's time-t coefficients of absolute and relative risk aversion, respectively. Also, let $\bar{\theta}_t$ denote the harmonic average of investors' coefficients of absolute risk aversion at time t. We have the following result.

Proposition 1. *The log price of the n-period bond at time t is given by*

$$\begin{aligned} \log p_t^{(n)} = \log p_t^{(1)} + \overline{\mathbb{E}}_t[\log p_{t+1}^{(n-1)}] + \frac{1}{2}\mathrm{var}_t(\log p_{t+1}^{(n-1)}) - \bar{\theta}_t \sum_{m=2}^{\infty} s_t^{(m)} \mathrm{cov}_t(\log p_{t+1}^{(n-1)}, \log p_{t+1}^{(m-1)}) \\ - \left(\int_0^1 (\bar{\theta}_t/\theta_{it})(1-\gamma_{it}) \mathrm{d}i \right) \sum_{h=1}^{\infty} \rho^h \mathrm{cov}_t(\log p_{t+1}^{(n-1)}, \log p_{t+h}^{(1)}) \end{aligned} \tag{7}$$

to a second-order approximation as $\sigma \to 0$, where

$$\overline{\mathbb{E}}_t[\cdot] = \int_0^1 (\bar{\theta}_t/\theta_{it}) \mathbb{E}_{it}[\cdot] \mathrm{d}i \tag{8}$$

denotes investors' consensus expectations.

This result characterizes equilibrium bond prices in terms of investors' subjective expectations, their risk-bearing capacities, the short-term interest rate, and the supply of bonds. Crucially for our purposes, Proposition 1 holds regardless of how investors form their expectations, as it follows directly from individual portfolio optimization and market clearing. In particular, the expression in (7) remains valid whether investors hold identical or heterogeneous beliefs, regardless of how their expectations relate (if at all) to the process governing the short-term interest rate, and irrespective of their understanding of the relationship among different bond prices—including equation (7) itself. As a result, this proposition allows us to study the implications of investors' subjective beliefs for the term structure of interest rates without taking an explicit or implicit stance on how those beliefs are formed. In the language of Adam and Marcet (2011), Proposition 1 relies on investors' *internal rationality*—that is, they understand their own objective and optimize their portfolios accordingly, given their subjective beliefs—but not on their *external rationality* (i.e., whether they understand the market structure, other investors' behavior, etc.).

We now examine each term on the right-hand side of (7) in turn. The first two terms show that the price of the n-period bond depends on the short-term interest rate and on a weighted average of investors' expectations. Notably, despite the fact that investors are long-lived, only their one-period-ahead forecasts of the bond's resale price, $\mathbb{E}_{it}[p_{t+1}^{(n-1)}]$, enter the expression; expectations over longer

horizons or those of other bond prices do not. This reflects the assumption that investors can rebalance their portfolios each period.

The second term on the right-hand side of (7) indicates that bonds are priced as if the economy were populated by a representative investor holding the *consensus expectations* defined in (8). These expectations are weighted averages of investors' subjective forecasts, with greater weights assigned to investors with lower coefficients of absolute risk aversion, θ_{it}.[7] This reflects the intuitive idea that beliefs of investors with higher risk-bearing capacity exert greater influence on asset prices.[8] It is worth noting that, despite (5), the cross-sectional distribution of expectations shapes bond prices in a nontrivial way: while (5) rules out disagreements of order σ, it still allows for disagreements of order σ^2, which affect consensus expectations in (8) to a second-order approximation. In other words, investor disagreement, though bounded, plays a meaningful role in determining equilibrium prices.

The third term on the right-hand side of (7) is a standard convexity adjustment that arises from expressing the relationship between different bond prices in logs.

The remaining two terms on the right-hand side of (7) are also standard and capture the bond's risk premia. The fourth term represents the "myopic component" of the compensation investors demand for bearing the bond's price risk. This compensation increases with the harmonic average of investors' absolute risk aversion, the subjective covariance between the bond's resale price, $\log p_{t+1}^{(n-1)}$, and the prices of other bonds, and investors' positions in those bonds—which, by market clearing, equal the bonds' net supply, $s_t^{(m)}$. The fifth term is a reflection of investors' intertemporal hedging demand due to their uncertainty about future interest rates: conservative investors (with $\gamma_{it} > 1$) demand more of the n-period bond when they believe its resale price tends to be high in states where future short-term interest rates are expected to be low (Campbell and Viceira, 2001). This term vanishes when all investors are either short-lived or have logarithmic preferences.

An immediate implication of Proposition 1 is that, as long as the convexity adjustment term in (7) is either small or constant over time,

$$\Delta y_t^{(n)} = \frac{1}{n}\Delta y_t^{(1)} + \frac{n-1}{n}\Delta\overline{\mathbb{E}}_t[y_{t+1}^{(n-1)}] + \Delta \mathrm{rp}_t^{(n)}, \tag{9}$$

where $y_t^{(n)} = -\frac{1}{n}\log p_t^{(n)}$ is the yield on the n-period bond, $\Delta x_t = x_t - x_{t-1}$, and $\mathrm{rp}_t^{(n)}$ denotes the bond's risk premium, corresponding to the last two terms on the right-hand side of (7). Since (9) follows directly from individual portfolio optimization and market clearing, it is the only relationship among interest rates, expectations, and risk premia that holds for an arbitrary specification of investors' subjective expectations. Three aspects of this relationship are worth highlighting.

First, equation (9) implies that, if yield expectations are observable, one can decompose movements in the yield curve into those driven by changes in expectations and those driven by changes in risk premia, without imposing any assumptions beyond those required for Proposition 1. We exploit this observation in the next section to assess the relative contribution of movements in expectations and risk premia to the transmission of short-rate changes to long-term yields.

[7] If investors were also allowed to have heterogeneous subjective uncertainties, these weights would reflect not only their degree of risk aversion but also scale inversely with their subjective uncertainty.

[8] For example, in the special case of logarithmic utility, the weighted average on the right-hand side of (7) reduces to the wealth-weighted average of investors' expectations, in line with Xiong and Yan (2010). Also see Barillas and Nimark (2017), who derive an equation similar to (7) under the additional assumptions that all investors (i) are short-lived, (ii) have logarithmic preferences, and (iii) form rational expectations under incomplete information about the short-term interest rate.

Second, while (9) pins down the relationship between realized yields and one-period-ahead forecasts of adjacent maturities, it places no restrictions on how expectations across other maturities and forecast horizons relate to one another. In particular, absent further restrictions, investors' subjective expectations of short- and long-term interest rates may be entirely disconnected.

Third, since (9) is the only relationship among bond yields and their expectations that holds for an arbitrary specification of subjective expectations, any other such relationship—including the standard formulation of the expectations hypothesis that links long-term yields to expected future short rates—necessarily rests on additional assumptions about how expectations are formed. Put differently, whereas $\mathrm{rp}_t^{(n)}$ in (9), which coincides with the consensus subjective expectation of excess returns of the n-period bond, represents the actual compensation investors demand for bearing interest rate risk, the "term premium" typically defined as the residual in decompositions such as (1) need not reflect compensation for risk except under additional assumptions on investors' expectations.

3 Expectations and the Term Structure of Interest Rates

The sensitivity of long-term yields to movements in the short rate is one of the key determinants of the transmission of monetary policy to the broader economy. In this section, we examine this sensitivity through the lens of the term structure of investors' expectations. Using equation (9) as our starting point, we first decompose the sensitivity of long-term yields to changes in the short rate into an expectations component and a risk-premium component. We then study the expectations component directly by examining how investors' subjective forecasts, across both maturities and forecast horizons, respond to movements in the short rate. All theoretical results are presented to the same order of approximation as in Proposition 1. For notational simplicity, we suppress the approximation error terms throughout the exposition, while accounting for them explicitly in the proofs.

3.1 Sensitivity of Long-Term Interest Rates

We measure the response of long-term interest rates to changes in the short rate using the regression proposed by Hanson, Lucca, and Wright (2021),

$$\Delta y_t^{(n)} = \alpha + \beta_{\mathrm{HLW}}^{(n)} \Delta y_t^{(1)} + \varepsilon_t, \tag{10}$$

where $\beta_{\mathrm{HLW}}^{(n)}$ captures the contemporaneous sensitivity of the n-period yield to movements in the short rate. Equation (9) immediately implies a decomposition of this sensitivity into three distinct components,

$$\beta_{\mathrm{HLW}}^{(n)} = \beta_{\mathrm{short}}^{(n)} + \beta_{\mathrm{exp}}^{(n)} + \beta_{\mathrm{rp}}^{(n)}, \tag{11}$$

where $\beta_{\mathrm{short}}^{(n)} = 1/n$ is the effect of the short rate on the n-period yield holding expectations and risk premia fixed, $\beta_{\mathrm{exp}}^{(n)}$ and $\beta_{\mathrm{rp}}^{(n)}$ are the slope coefficients of regressions

$$\frac{n-1}{n} \Delta \overline{\mathbb{E}}_t[y_{t+1}^{(n-1)}] = \alpha + \beta_{\mathrm{exp}}^{(n)} \Delta y_t^{(1)} + \varepsilon_t, \tag{12}$$

$$\Delta \mathrm{rp}_t^{(n)} = \alpha + \beta_{\mathrm{rp}}^{(n)} \Delta y_t^{(1)} + \varepsilon_t, \tag{13}$$

respectively, and $\Delta \mathrm{rp}_t^{(n)}$ is recovered from realized yields and consensus expectations via equation (9). The components in (11) capture three distinct sources of the sensitivity of long-term yields to changes in the short rate: a mechanical effect that operates even when expectations and risk premia are unchanged, a revision in investors' expectations of future yields, and an adjustment in the bond's risk premium.

In principle, the sensitivity of long-term yields to changes in the short rate can always be decomposed into an expectations component and a residual, but such a residual need not have a structural interpretation as compensation for risk. The key point is that, by Proposition 1, the residual $\Delta \mathrm{rp}_t^{(n)} = \Delta y_t^{(n)} - \frac{1}{n}\Delta y_t^{(1)} - \frac{n-1}{n}\Delta \overline{\mathbb{E}}_t[y_{t+1}^{(n-1)}]$ is exactly the change in the bond's risk premium, so that $\beta_{\mathrm{rp}}^{(n)}$ captures the sensitivity of risk premia to changes in the short rate, irrespective of how expectations are formed. In particular, this decomposition requires no assumption that investors' expectations are consistent with the true data-generating process, nor that their expectations across different maturities and forecast horizons are mutually consistent with one another.

By contrast, any decomposition relying on a relationship between yields and expectations other than (9) necessarily imposes additional assumptions on how expectations are formed, implying that the resulting sensitivity of the residual to changes in the short rate corresponds to the sensitivity of the bond's risk premium only under specific models of expectation formation. As an illustrative example, consider the approach of Hanson and Stein (2015), who attribute the predictability of distant forward rate movements to time variation in risk premia. Adapted to our setting, their test takes the form

$$f_{t+h}^{(1,m)} - f_t^{(1,m+h)} = \alpha + \beta_{\mathrm{HS}}^{(1,m)} \Delta y_t^{(1)} + \varepsilon_{t,t+h},$$

where $f_t^{(n,m)} = (1+m/n)y_t^{(n+m)} - (m/n)y_t^{(m)}$ is the forward yield of an n-period bond maturing in $n+m$ periods. Hanson and Stein (2015) argue that $\beta_{\mathrm{HS}}^{(1,m)}$ measures the contribution of risk premia to forward rate movements, as this coefficient must be zero if risk premia are constant. However, this interpretation relies on two distinct assumptions on how expectations are formed. First, since the left-hand side uses realized forward rates $f_{t+h}^{(1,m)}$ rather than expected future forward rates, the coefficient is informative about risk premia only if investors' forecast errors are unpredictable given time-t information, as would be the case under rational expectations. Second, even if the realized forward rate $f_{t+h}^{(1,m)}$ is replaced with subjective expectations $\overline{\mathbb{E}}_t[f_{t+h}^{(1,m)}]$ on the left-hand side, the corresponding regression coefficient does not in general reflect movements in risk premia: since (9) is silent on how expectations at horizon $h > 1$ relate to one another and to realized yields, the quantity $\overline{\mathbb{E}}_t[f_{t+h}^{(1,m)}] - f_t^{(1,m+h)}$ may be predictable by time-t variables even when risk premia are constant, simply due to the way investors' expectations across different maturities and forecast horizons are formed.[9]

3.2 Sensitivity of Yield Expectations

As already discussed, equation (11) decomposes the sensitivity of long-term yields to the short rate into three distinct components, highlighting the potential importance of movements in yield expectations. At the same time, equation (9)—on which this decomposition rests—is silent on how expectations themselves respond to changes in the short rate. In this subsection, we turn to this question directly.

[9] To be precise, Hanson and Stein (2015) use the change in the forward rate in the prior period, $f_t^{(1,m)} - f_{t-1}^{(1,m)}$, as their right-hand side variable. The argument above holds irrespective of this choice.

To this end, we extend the regression of Hanson, Lucca, and Wright (2021) to the family of regressions

$$\Delta\overline{\mathbb{E}}_t[y_{t+h}^{(n)}] = \alpha + \gamma_h^{(n)}\Delta y_t^{(1)} + \varepsilon_t, \tag{14}$$

where coefficients $\gamma_h^{(n)}$ capture how investors' subjective expectations of n-period yields at horizon $h \geq 1$ respond to changes in the short rate. These coefficients, indexed across maturities and forecast horizons, trace out the full contemporaneous response of the term structure of expectations to the short-term interest rate. Comparing (14) to (12), it is immediate that $\gamma_1^{(n)} = (1 + 1/n)\beta_{\text{exp}}^{(n+1)}$.

Absent further assumptions on how expectations are formed, nothing constrains how the sensitivity coefficients $\gamma_h^{(n)}$ relate to one another across maturities and forecast horizons. For example, there is no guarantee that the sensitivity of investors' expectations of the n-period yield to the short rate bears any relationship to the sensitivity of their expectations of future short rates over the bond's lifetime, as would be required by the expectations hypothesis. Characterizing any relationship among these coefficients requires taking a stand on investors' mental models of yield curve dynamics. Our next result does so under a natural benchmark: that investors understand the bond market structure and believe risk premia to be constant.

Proposition 2. *Suppose investors (i) understand market clearing and other investors' ability to solve their optimization problems, and (ii) believe that risk premia are constant. Then,*

$$\delta_h^{(n)}\langle n_1, \ldots, n_\ell\rangle := \gamma_h^{(n)} - \frac{1}{n}\sum_{s=1}^{\ell} n_s \gamma_{h+n_1+\cdots+n_{s-1}}^{(n_s)} = 0 \tag{15}$$

for any $n, h, \ell \geq 1$ *and positive integers* $n_1, \ldots, n_\ell$ *such that* $n = n_1 + \cdots + n_\ell$.

This result rests on the observation that when investors understand the bond market structure, they also understand that yields satisfy the equilibrium condition in (9), and hence their subjective expectations of its left- and right-hand sides must coincide. This, in turn, disciplines investors' mental models of the yield curve, constraining how their subjective expectations across different maturities and forecast horizons relate to one another. Specifically, it implies that $\delta_h^{(n)}\langle n_1, \ldots, n_\ell\rangle$ defined in (15) equals the sensitivity to the short rate of the perceived cumulative risk-premium differential between holding the n-period bond from date $t + h$ to maturity and implementing the corresponding sequence of investments in bonds of maturities $n_1, \ldots, n_\ell$.[10] Together with the assumption that investors believe risk premia to be constant, this forces the sensitivity coefficients $\gamma_h^{(n)}$ to satisfy (15). We refer to this family of restrictions as *structural consistency* of expectations, since they require investors' expectations across different maturities and forecast horizons to conform to the equilibrium relationships linking different interest rates, given their belief that risk premia are constant. Importantly, this notion makes no reference to the true data-generating process or the accuracy of investors' forecasts, as reflected in the fact that all terms on either side of (15) are sensitivities of subjective yield expectations to the short rate, rather than of realized yields.

The restrictions in Proposition 2 have a straightforward interpretation: if investors understand the bond market structure, their expectations must be consistent with equilibrium pricing across different

[10] Iterating (9) forward under the assumption that investors understand the market structure implies $n\,\overline{\mathbb{E}}_t[y_{t+h}^{(n)}] = \sum_{s=0}^{n-1}\overline{\mathbb{E}}_t[y_{t+h+s}^{(1)}] + \sum_{m=2}^{n} m\,\overline{\mathbb{E}}_t[\text{rp}_{t+h+n-m}^{(m)}]$.

trading strategies. In particular, their expected return from buying an $(n_1 + \cdots + n_\ell)$-period bond at date $t + h$ and holding it to maturity must equal their expected return from a sequence of rolling trades in bonds of maturities $n_1, \ldots, n_\ell$ over the same window, up to a constant reflecting investors' belief that risk premia are time-invariant. Equation (15) expresses this requirement in terms of sensitivity coefficients. For instance, setting $\ell = n$ and $n_1 = \cdots = n_\ell = 1$, it implies that

$$\gamma_h^{(n)} = \frac{1}{n} \sum_{s=0}^{n-1} \gamma_{h+s}^{(1)}, \tag{16}$$

that is, the sensitivity of the h-period-ahead expected n-period yield to the short rate equals the average of the sensitivities of expected future short rates $\overline{\mathbb{E}}_t[y_{t+h+s}^{(1)}]$ over the bond's lifetime—precisely the restriction implied by the expectations hypothesis under subjective beliefs. Importantly, since (16) follows only after imposing the assumptions of Proposition 2, using it to describe the relationship between expectation sensitivity coefficients across maturities and forecast horizons may not be valid absent those assumptions.

We conclude this discussion with a few remarks. First, the notion of structural consistency of expectations is distinct from forecast accuracy: whereas the former concerns the internal coherence of subjective expectations across the yield curve, the latter concerns how those expectations relate to the true data-generating process. As such, investors' forecast errors may be predictable even when their expectations satisfy (15).[11]

Second, it is worth highlighting that although (9) holds in equilibrium irrespective of investors' beliefs, the linearity of the subjective expectation operator $\mathbb{E}_{it}[\cdot]$ alone does not guarantee that investor i's time-t subjective expectations of the left- and right-hand sides of (9) at some future date $t + h$ coincide. Linearity only ensures that the operator respects linear combinations of quantities the investor recognizes as related; it places no restriction on how the investor's expectations of the distinct quantities appearing on the two sides of (9) relate to one another unless the investor understands that the two sides are equal at all times. It is precisely the assumption that investors understand market clearing and other investors' ability to solve their optimization problems that guarantees this and forces the two sides to agree under $\mathbb{E}_{it}[\cdot]$.

Third, since under the assumptions of Proposition 2 each investor i's subjective expectations of both sides of (9) coincide, aggregating across investors implies that the consensus expectations of both sides are equal as well. Applying consensus expectations to both sides of (9) at a future date $t + h$, however, results in an iteration of the consensus expectation operator on the right-hand side in the form of $\overline{\mathbb{E}}_t\overline{\mathbb{E}}_{t+h}[\cdot]$. While consensus expectations do not satisfy the law of iterated expectations in general, our assumption that investor disagreement decays to zero at a faster rate than the variables they disagree about ensures that any violation is of a higher order and does not affect our results to the relevant order of approximation. As a result, the restrictions in (15) are expressed entirely in terms of the sensitivities of investors' first-order consensus expectations, rather than their higher-order (that is, iterated) expectations.

[11] Farmer, Nakamura, and Steinsson (2024), for instance, construct long-term yields as the average of expected future short rates over the bond's lifetime plus a constant risk premium, with subjective rather than full-information rational expectations—thereby imposing structural consistency across the yield curve while allowing for the possibility of predictable forecast errors.

Finally, since each coefficient in (15) can be estimated via regression (14), these restrictions provide an implementable test of the joint assumptions of Proposition 2: a violation of (15) indicates that investors' expectations across different maturities and forecast horizons cannot be jointly reconciled with an understanding of the equilibrium relationship in (9) and the belief that risk premia are constant. We exploit this observation to assess the structural consistency of investors' subjective expectations in the next section.

4 Empirical Analysis

In this section, we use survey and interest rate data to study the sensitivity of long-term yields and investors' subjective expectations to changes in the short rate. We proceed in two steps. First, we exploit the decomposition in (11) to disentangle the contributions of expectations and risk premia to the sensitivity of long-term yields to the short rate. Second, we estimate the sensitivity coefficients $\gamma_h^{(n)}$ in (14) and test whether they satisfy the structural consistency restrictions in (15), thereby assessing the assumptions underpinning Proposition 2.

4.1 Data

Estimating regressions (10) and (12)–(14) requires data on realized yields and consensus expectations of future interest rates. To this end, we combine data on zero-coupon Treasury yields with survey-based expectations, and rely on investor wealth data to construct consensus expectations as wealth-weighted averages of individual forecasts. We describe each data source and the construction of consensus expectations below.

Interest rates. We use zero-coupon yield curve estimates from Liu and Wu (2021), who construct daily yields for maturities from 1 month to 30 years starting in 1961, converted to quarterly frequency by averaging within each quarter. To align with the survey and wealth data described below, we restrict the sample to 2001:Q1–2024:Q4 and focus on maturities from 1 to 41 quarters. Throughout, we use the 1-quarter yield as our measure of the short rate, consistent with the quarterly frequency at which we sample the data.

Survey data. We use forecasts from the Blue Chip Financial Forecasts (BCFF) survey, in which approximately 60 financial professionals submit projections for interest rates and other macroeconomic variables toward the end of each month. More than half of all respondents are banks or broker-dealers; the remainder are research institutions, rating agencies, or mortgage lenders. Respondents provide estimates for the current quarter and up to five quarters ahead, with forecasts released on the first day of the following month. Although the survey dates back to the 1980s, we begin our sample in 2001:Q1 to ensure a consistent panel of interest rate forecasts across maturities and horizons and to match forecasters with balance-sheet information. We sample the data at a quarterly frequency through 2024:Q4. Our data cover forecasts of 1, 2, 4, 8, 20, and 40-quarter yields at horizons $h = 1, \ldots, 6$ quarters.[12]

[12] Although the survey is conducted monthly, respondents forecast calendar-quarter outcomes. Using every monthly release would mean that forecasts of the same target quarter cover different horizons depending on when they are made. It would

Consensus expectations. Implementing our regression specifications requires a measure of consensus expectations $\overline{\mathbb{E}}_t[y_{t+h}^{(n)}]$. By Proposition 1, these consensus expectations are weighted averages of individual forecasts, with greater weight placed on investors with lower coefficients of absolute risk aversion. Since risk aversion is in general not observable, we construct consensus expectations from survey data as a wealth-weighted average of individual forecasts, which corresponds to assuming that investors have CRRA preferences with a common coefficient of relative risk aversion.

Wealth. To proxy for investor wealth, we focus on forecasters affiliated with primary dealers in the U.S. Treasury market, for which detailed balance-sheet information is available. Specifically, we use each dealer's holdings of U.S. Treasury securities as a market-specific proxy for wealth. Over our 2001–2024 sample period, we identify 24 institutions that served as primary dealers and have sufficiently broad coverage across time and forecast horizons: ABN AMRO, Amherst Pierpont, Banc One, Bank of America, Barclays, Bear Stearns, BMO, BNP Paribas, Credit Suisse, Daiwa, Deutsche Bank, Goldman Sachs, J.P. Morgan, Lehman Brothers, Merrill Lynch, Mizuho, Nomura, RBC, RBS Greenwich, Salomon Smith Barney, Scotiabank, Société Générale, UBS, and Wells Fargo. These institutions did not all serve as primary dealers simultaneously; on average, approximately 12 primary dealers are represented in our forecast sample each year. We obtain annual balance-sheet information from FOCUS reports (Form X-17A-5) filed by U.S. broker-dealers with the SEC and available through the EDGAR system. Unlike bank call reports, these filings report the financial condition of the broker-dealer entity itself rather than that of the consolidated bank holding company, facilitating comparisons across forecasters with different ownership structures. We extract the fair value of each dealer's holdings of U.S. Treasury securities and define its wealth weight as its share of the aggregate Treasury holdings of all primary dealers represented in our sample in a given year.

4.2 Sensitivity of Interest Rates

As our first empirical exercise, we measure the contributions of expectations and risk premia to the sensitivity of long-term yields to the short rate.

We begin by following Hanson, Lucca, and Wright (2021) and estimating regression (10) to measure the overall sensitivity of yields of different maturities to changes in the short rate. The estimated slope coefficients $\hat{\beta}_{\text{HLW}}^{(n)}$, together with 95% confidence intervals, are reported in Figure 1. As the figure indicates, movements in the short rate are associated with meaningful movements in yields across all maturities. This sensitivity, however, declines with maturity. For example, a 100 basis point (bps) increase in the short rate is associated with a 46 bps increase in the 5-year yield ($n = 20$) and only a 26 basis point increase in the 10-year yield ($n = 40$).

We next use the decomposition in (11) to determine the contributions of movements in expectations and risk premia to the overall sensitivity of long-term yields to the short rate. According to this decomposition, $\beta_{\text{HLW}}^{(n)} = \beta_{\text{short}}^{(n)} + \beta_{\text{exp}}^{(n)} + \beta_{\text{rp}}^{(n)}$, where $\beta_{\text{short}}^{(n)} = 1/n$ is the mechanical effect of the short rate on

also include mid-quarter projections of the current quarter. Sampling at the start of each quarter ensures that each forecast horizon remains the same across observations. See Appendix B for a more detailed discussion of the survey data, including the list of forecasters included in our sample and a comparison of primary dealer forecasts with those of other forecaster groups and surveys.

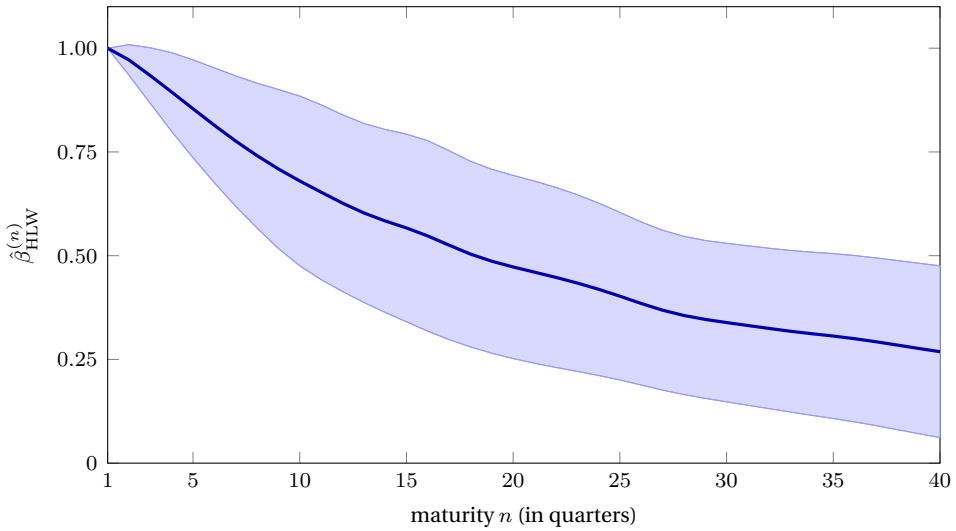


Figure 1. Sensitivity of n-Period Yields to Changes in the Short Rate

Notes: This figure plots the estimated slope coefficients $\beta^{(n)}_{\text{HLW}}$ from regression (10), together with 95 percent confidence intervals, for maturities $n \in \{1, \ldots, 40\}$ quarters. The dependent variable is the quarterly change in the n-period yield $\Delta y^{(n)}_t$ and the independent variable is the quarterly change in the 1-quarter yield $\Delta y^{(1)}_t$. Confidence intervals are based on EWC-HAC standard errors computed following Lazarus et al. (2018), with $\nu = 8$ cosine basis functions.

the n-period yield holding expectations and risk premia fixed, and $\beta^{(n)}_{\text{exp}}$ and $\beta^{(n)}_{\text{rp}}$ are the slope coefficients of regressions (12) and (13), capturing the contributions of revisions in investors' expectations and adjustments in the bond's risk premium, respectively. To implement this decomposition, we rely on the fact that, while risk premia are not directly observable, Proposition 1 implies that $\Delta \text{rp}^{(n)}_t$ can be recovered from realized yields and consensus expectations as a residual using equation (9). Because the BCFF survey contains forecasts for yields of maturities 1, 2, 4, 8, 20, and 40 quarters only, we implement this exercise for $n = 2, 3, 5, 9, 21$, and 41 quarters.

Table 1 and Figure 2 report the results. The first column of Table 1 shows the estimated sensitivity coefficients $\hat{\beta}^{(n)}_{\text{HLW}}$ from regression (10), also plotted in Figure 1, while the remaining columns decompose this sensitivity into its three components. With the exception of the longest-maturity bond ($n = 41$), the risk-premium component $\hat{\beta}^{(n)}_{\text{rp}}$ is statistically indistinguishable from zero across all maturities. For short-maturity bonds, it is also small in magnitude, with movements in expectations explaining nearly all of the response of long-term yields to changes in the short rate. In contrast, for the longest maturity bond, the risk-premium component is negative, statistically significant, and has a magnitude roughly half as large as the expectations component in absolute terms.

4.3 Sensitivity of Yield Expectations

The results in Table 1 establish that movements in yield expectations are the dominant driver of the sensitivity of yields to the short rate, with risk premia playing a more limited role, particularly at shorter maturities. The decomposition in (11), however, is silent on how expectations themselves are formed: it takes one-period-ahead forecasts of adjacent yields as given without placing any restrictions on how they move across maturities and forecast horizons. In this subsection, we turn to this question directly,

Table 1. Sensitivity Decomposition

n	$\hat{\beta}^{(n)}_{\text{HLW}}$	$\beta^{(n)}_{\text{short}}$	$\hat{\beta}^{(n)}_{\text{exp}}$	$\hat{\beta}^{(n)}_{\text{rp}}$
2	0.97^{***}	0.50	0.46^{***}	0.01
	(0.02)		(0.04)	(0.05)
3	0.93^{***}	0.33	0.70^{***}	-0.10
	(0.03)		(0.05)	(0.07)
5	0.85^{***}	0.20	0.79^{***}	-0.14
	(0.05)		(0.07)	(0.10)
9	0.71^{***}	0.11	0.75^{***}	-0.15
	(0.08)		(0.07)	(0.12)
21	0.46^{***}	0.05	0.59^{***}	-0.18
	(0.10)		(0.06)	(0.11)
41	0.26^{**}	0.02	0.42^{***}	-0.19^{**}
	(0.09)		(0.07)	(0.07)

Notes: This table reports the estimated slope coefficients from regressions (10), (12), and (13), with $\Delta y_t^{(1)}$ as the regressor, where by construction $\beta^{(n)}_{\text{short}} + \hat{\beta}^{(n)}_{\text{exp}} + \hat{\beta}^{(n)}_{\text{rp}} = \hat{\beta}^{(n)}_{\text{HLW}}$. The mechanical component $\beta^{(n)}_{\text{short}} = 1/n$ is known analytically. Consensus forecasts $\overline{\mathbb{E}}_t[\cdot]$ are constructed as the wealth-weighted average of individual forecasts. The numbers in parentheses report EWC-HAC standard errors according to Lazarus et al. (2018), with $\nu = 8$ cosine basis functions. Stars indicate statistical significance for the null hypothesis that the coefficient equals zero, with *, **, and *** denoting significance at the 10%, 5%, and 1% levels, respectively.

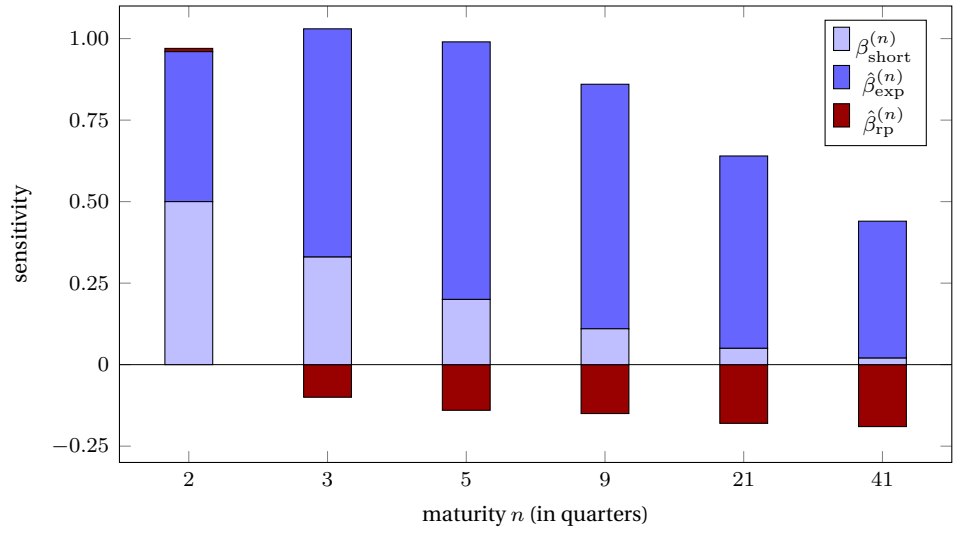


Figure 2. Decomposition of Sensitivity of n-Period Yields to Changes in the Short Rate

Notes: This figure plots the decomposition $\hat{\beta}^{(n)}_{\text{HLW}} = \beta^{(n)}_{\text{short}} + \hat{\beta}^{(n)}_{\text{exp}} + \hat{\beta}^{(n)}_{\text{rp}}$ for maturities $n \in \{2, 3, 5, 9, 21, 41\}$ quarters based on the point estimates reported in Table 1. Consensus forecasts $\overline{\mathbb{E}}_t[\cdot]$ are constructed as the wealth-weighted average of individual forecasts.

Table 2. Sensitivity of Yield Expectations to Changes in the Short Rate

maturity n (in quarters)	forecast horizon h (in quarters)					
	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
1	0.93^{***}	0.92^{***}	0.87^{***}	0.81^{***}	0.73^{***}	0.64^{***}
	(0.08)	(0.10)	(0.10)	(0.09)	(0.09)	(0.07)
2	1.05^{***}	1.04^{***}	0.97^{***}	0.85^{***}	0.72^{***}	0.56^{***}
	(0.08)	(0.09)	(0.08)	(0.07)	(0.07)	(0.07)
4	0.99^{***}	0.96^{***}	0.88^{***}	0.77^{***}	0.64^{***}	0.50^{***}
	(0.08)	(0.09)	(0.08)	(0.08)	(0.09)	(0.07)
8	0.85^{***}	0.83^{***}	0.76^{***}	0.65^{***}	0.58^{***}	0.49^{***}
	(0.07)	(0.08)	(0.08)	(0.08)	(0.07)	(0.06)
20	0.62^{***}	0.60^{***}	0.54^{***}	0.45^{***}	0.40^{***}	0.34^{***}
	(0.06)	(0.07)	(0.07)	(0.07)	(0.06)	(0.06)
40	0.44^{***}	0.42^{***}	0.37^{***}	0.32^{***}	0.28^{***}	0.22^{**}
	(0.07)	(0.08)	(0.08)	(0.08)	(0.08)	(0.10)

Notes: This table reports the estimated slope coefficients $\hat{\gamma}_h^{(n)}$ from regression (14) for maturities $n \in \{1, 2, 4, 8, 20, 40\}$ quarters and forecast horizons $h \in \{1, \ldots, 6\}$ quarters. Consensus forecasts $\overline{\mathbb{E}}_t[\cdot]$ are constructed as the wealth-weighted average of individual forecasts. The numbers in parentheses report EWC-HAC standard errors according to Lazarus et al. (2018), with $\nu = 8$ cosine basis functions. Stars indicate statistical significance for the null hypothesis that the coefficient equals zero, with *, **, and *** denoting significance at the 10%, 5%, and 1% levels, respectively.

studying the sensitivity of yield expectations to changes in the short rate and assessing whether these movements are consistent with the equilibrium relationships that link different interest rates to one another.

To this end, we estimate the family of regressions in (14), with results reported in Table 2. Since the BCFF survey covers forecasts of 1, 2, 4, 8, 20, and 40-quarter yields at horizons $h = 1, \ldots, 6$ quarters only, our estimates are limited to these maturities and forecast horizons. Movements in the short rate are associated with sizable and statistically significant movements in yield expectations across all maturities and horizons. This sensitivity, however, generally declines with both maturity and forecast horizon. For example, a 100 basis point increase in the short rate is associated with a 93 basis point revision in expectations of the 1-quarter yield one quarter ahead, but only a 44 basis point revision in expectations of the 10-year yield over the same horizon. Similarly, fixing maturity at $n = 1$ quarter, the sensitivity declines from 93 basis points for one-quarter-ahead forecasts to 64 basis points for six-quarter-ahead forecasts, suggesting that investors expect the yield to partially mean-revert over time.

With the estimated sensitivities of yield expectations to the short rate in hand, we next assess whether these sensitivities satisfy the structural consistency restrictions in (15). We start with the observation that the expression on the left-hand side of (15) can be estimated directly as the slope coefficient of the regression

$$\Delta\overline{\mathbb{E}}_t[y_{t+h}^{(n)}] - \frac{1}{n}\sum_{s=1}^{\ell} n_s \Delta\overline{\mathbb{E}}_t\left[y_{t+h+n_1+\cdots+n_{s-1}}^{(n_s)}\right] = \alpha + \delta_h\langle n_1, \ldots, n_\ell\rangle\, \Delta y_t^{(1)} + \varepsilon_t. \tag{17}$$

Table 3. Tests of Structural Consistency Restrictions

maturity n (in quarters)	partition $\langle n_1, \ldots, n_\ell \rangle$	forecast horizon h (in quarters)				
		1	2	3	4	5
2	$\langle 1,1 \rangle$	0.13	0.14	0.13	0.08	0.03
		(0.09)	(0.10)	(0.08)	(0.08)	(0.04)
4	$\langle 2,2 \rangle$	−0.02	0.01	0.03	0.06	
		(0.02)	(0.02)	(0.02)	(0.04)	
4	$\langle 1,1,2 \rangle$	0.04	0.08	0.10	0.11**	
		(0.06)	(0.07)	(0.06)	(0.04)	
4	$\langle 1,2,1 \rangle$	0.04	0.06	0.07		
		(0.05)	(0.06)	(0.05)		
4	$\langle 2,1,1 \rangle$	0.04	0.05	0.05		
		(0.06)	(0.05)	(0.04)		
4	$\langle 1,1,1,1 \rangle$	0.11	0.12	0.11		
		(0.10)	(0.11)	(0.08)		
8	$\langle 4,4 \rangle$	0.03	0.10			
		(0.06)	(0.07)			
8	$\langle 2,2,4 \rangle$	0.02	0.11			
		(0.06)	(0.07)			
8	$\langle 1,1,2,4 \rangle$	0.05	0.14**			
		(0.04)	(0.05)			
8	$\langle 1,2,1,4 \rangle$	0.05	0.13**			
		(0.04)	(0.05)			
8	$\langle 2,1,1,4 \rangle$	0.05	0.13**			
		(0.04)	(0.05)			
8	$\langle 1,1,1,1,4 \rangle$	0.08***	0.16***			
		(0.02)	(0.03)			

Notes: This table reports the estimated slope coefficients $\hat{\delta}_h\langle n_1, \ldots, n_\ell \rangle$ from regression (17). Consensus forecasts $\overline{\mathbb{E}}_t[\cdot]$ are constructed as the wealth-weighted average of individual forecasts. The numbers in parentheses report EWC-HAC standard errors according to Lazarus et al. (2018), with $\nu = 8$ cosine basis functions. Stars indicate statistical significance for the null hypothesis that the coefficient equals zero, with *, **, and *** denoting significance at the 10%, 5%, and 1% levels, respectively.

Estimating (17) requires survey forecasts of the n-period yield at horizon h and of each n_s-period yield at horizon $h + n_1 + \cdots + n_{s-1}$. Given our survey data, the testable restrictions are limited to those for which n and each n_s belong to the BCFF maturity grid $\{1, 2, 4, 8, 20, 40\}$ and all required forecast horizons fall within six quarters. These requirements lead to 34 restrictions, covering maturities of $n \in \{2, 4, 8\}$ quarters.[13]

Table 3 reports the estimated coefficients. The estimates are overwhelmingly positive across maturities, forecast horizons, and partitions, with several that are individually significant. By Proposition 2, if investors understand the market structure and believe risk premia to be constant, all

[13] No restriction is available at longer maturities: any partition of a 20- or 40-quarter bond would push at least one forecast beyond the six-quarter horizon available in the BCFF surveys.

of these coefficients must equal zero simultaneously. As a result, a single violation is enough to reject the assumptions underpinning Proposition 2. Individual p-values, however, do not account for the multiplicity of tests. To address this, we consider two complementary tests. First, if all 34 restrictions hold, the unweighted average of the 34 deviations must be zero. We estimate this average deviation by regressing the unweighted average of the left-hand sides of (17) on $\Delta y_t^{(1)}$, obtaining 0.077 (s.e. $= 0.024$; $p = 0.013$). Second, we apply the Bonferroni correction to the 34 individual p-values, which tests the joint null without relying on a particular aggregation of the restrictions. The smallest unadjusted p-value is 0.0012, corresponding to a Bonferroni-adjusted p-value of 0.042.[14] Both tests therefore reject the joint null at the 5 percent level: investors' expectations across maturities and forecast horizons cannot be reconciled with both an understanding of the equilibrium relationship in (9) and the belief that risk premia are constant.

Although the joint rejection does not by itself establish the direction of the violations, 33 of the 34 estimates in Table 3 are positive, including all those that are individually significant. So is the unweighted average of the deviations, reported above. These observations suggest that, when benchmarked against (15), investors' short-horizon expectations of longer-maturity yields overreact relative to their own longer-horizon expectations of shorter-maturity yields. On average across the 34 restrictions, a 100-basis-point increase in the short rate raises expected longer-maturity yields by 7.7 basis points more than the corresponding weighted average of expected shorter-maturity yields.

This violation of structural consistency documented above has implications for the decomposition of the sensitivity of yields to the short rate into expectations and risk-premium components. Specifically, our finding that (15) is violated in the data implies that any decomposition based on the expectations hypothesis (1) or its variants that link realized yields to multi-period forecasts of future yields—regardless of how those forecasts are formed—may misattribute movements in expectations to risk premia, or vice versa. This underscores the importance of the decomposition based on (11), which requires no assumption on how expectations are formed.

It is also worth emphasizing that the violations of (15) we document are statements about the investors' mental models, rather than about the accuracy of their forecasts. In particular, they are distinct from the evidence documented by Cieslak (2018), d'Arienzo (2020), and Wang (2021), who, in the spirit of Coibion and Gorodnichenko (2015), show that forecast errors are predictable, indicating that investors' subjective expectations are inconsistent with the true data-generating process. In contrast, as already discussed, (15) is a restriction on the consistency of investors' subjective expectations with equilibrium conditions that link different interest rates, without making a reference to the true data-generating process. Accordingly, it is possible for forecast errors to be predictable—say, because investors hold a misspecified model of the short rate—while investors' expectations remain mutually consistent across maturities and forecast horizons. What we document is that this mutual consistency fails: investors' expectations of yields across different maturities and horizons are inconsistent with the structural relationships linking different interest rates, regardless of how those expectations relate to the true data-generating process.

[14] The Bonferroni correction is conservative because it does not exploit the linear relationships among the 34 restrictions.

4.4 Discussion

We conclude this section with a brief discussion of the assumptions underlying our empirical analysis and the interpretation of our findings.

First, as in the vast majority of empirical studies on subjective expectations, we use survey data from professional forecasters as a proxy for investors' beliefs about future interest rates. At the same time, our theoretical results in Section 2 show that bond prices are determined by the consensus expectations of market participants actively trading in bond markets. Accordingly, our implementation of regressions (12)–(14) and (17) rests on the assumption that survey-based forecasts are informative proxies for the expectations of active market participants, with any potential measurement error orthogonal to the explanatory variables.

Second, according to Proposition 1, consensus expectations are weighted averages of investors' subjective forecasts, with greater weights assigned to investors with lower coefficients of absolute risk aversion. Constructing these consensus expectations therefore requires data on investors' risk-bearing capacities, which are in general not observable. In our implementation, we use relative wealth as a proxy for risk-bearing capacity, weighting each investor's forecasts by their share of U.S. Treasury holdings. This corresponds to assuming that all investors have CRRA preferences with a common coefficient of relative risk aversion.[15]

Third, as already mentioned, the BCFF survey provides forecasts for only a handful of maturities and horizons, which limits the number of regressions we were able to run. A richer dataset with forecasts covering a broader span of maturities and horizons would enable a more complete assessment of the relationship between subjective expectations and the term structure of interest rates.

Finally, note that the right-hand side of regressions (10) and (14) is the change in the short rate rather than an identified monetary policy shock. Accordingly, as in Hanson, Lucca, and Wright (2021), the estimated coefficients in Tables 1–3 should be interpreted as sensitivities rather than as the causal passthrough of monetary policy shocks to long-term yields and expectations.[16] That said, the structural consistency restrictions in (15) must hold whenever the assumptions of Proposition 2 are satisfied, regardless of whether changes in the short rate reflect monetary policy shocks or other sources of variation. The tests based on regression (17) are therefore valid even in the absence of monetary policy identification.

5 A Model of Yield Curve Expectations

The empirical evidence in Section 4 documents that investors' subjective expectations across maturities and forecast horizons are inconsistent with the cross-maturity structural relationships that would hold if

[15] In Appendix C, we consider an alternative specification as a robustness check by constructing an unweighted average of all respondents in the BCFF survey.

[16] Our focus on movements in realized and expected yields in response to the short rate, rather than to identified monetary policy shocks, is driven by a frequency mismatch: while identification of monetary policy shocks relies on high-frequency intraday interest-rate movements in narrow windows around policy announcements, survey data are only available at monthly or lower frequencies. Aggregating shocks over the horizon for which survey forecasts are available replaces the standard high-frequency identification assumption with a weaker one, requiring orthogonality not only to news within narrow announcement windows but also to broader macroeconomic forces unfolding over the entire horizon. See Hubert de Fraisse, Maxted, Sangani, and Sraer (2026) for a detailed discussion.

investors fully understood how bond prices are determined in equilibrium. Our theoretical framework in Section 2 allows for this possibility by treating expectations as model primitives, placing no restrictions on how they are formed. While this agnostic approach is well-suited for empirical analysis, it leaves open the question of what expectation formation process could give rise to such violations.

In this section, we develop a model of expectation formation to illustrate how such cross-maturity inconsistencies may arise endogenously. Rather than assuming inconsistent expectations directly, we model investors as forming expectations using a statistical model of the yield curve. We show that when investors' models do not explicitly encode the equilibrium relationships linking yields of different maturities, their expectations may fail to satisfy (15), even if their models are estimated optimally using realized yield curve data.

We work within the framework of Section 2, in which investors trade in the bond market by solving the portfolio optimization problem (3) given their expectations, while imposing two additional assumptions. First, we assume that risk premia are constant. As established in Proposition 1, this—together with individual portfolio optimization and market clearing—implies that realized and expected yields must satisfy

$$y_t^{(n)} = \frac{1}{n} y_t^{(1)} + \frac{n-1}{n} \overline{\mathbb{E}}_t[y_{t+1}^{(n-1)}] + \text{constant} \tag{18}$$

for all $n \geq 2$. Second, rather than treating expectations as model primitives, we augment the framework with an explicit model of expectation formation. Specifically, we assume that all investors share a common statistical model of the yield curve, forming expectations using a $\mathrm{VAR}(k)$ model with $k \geq 2$,

$$\overline{\mathbb{E}}_t[y_{t+1}] = \sum_{s=0}^{k-1} \mathbf{A}_s y_{t-s} + \text{constant}, \tag{19}$$

where $y_t = [y_t^{(1)}, \ldots, y_t^{(N)}]'$ denotes the vector of yields of all maturities and $\mathbf{A}_s \in \mathbb{R}^{N \times N}$ are coefficient matrices.[17] We assume that investors estimate these coefficients by regressing realized yields on their past k lags using ordinary least squares (OLS) and an infinitely large sample.[18] This closes the model by endogenizing expectations and ensuring that investors' forecasts are disciplined by the data, giving rise to a fixed point problem: the coefficient matrices $\mathbf{A}_s$ determine investors' expectations, which in turn determine equilibrium yields through (18), which are then used to estimate $\mathbf{A}_s$ via OLS. An equilibrium of this model is therefore a set of coefficient matrices $\mathbf{A}_s$ that are consistent with the yield dynamics they themselves generate.

The only remaining modeling choice is to specify which elements of $\mathbf{A}_s$ are assumed to be nonzero, corresponding to which covariates investors include in their forecasting models. Throughout, we focus on coefficient matrices $\mathbf{A}_s$ that are lower triangular with a nonzero first column: investors forecast each yield using the history of yields of the same or shorter maturities, with the short rate always included

[17] For horizons $h > 1$, subjective forecasts are obtained by iterating the one-step forecasting rule in (19).

[18] Throughout, we assume the short rate follows a stationary ARMA process, so that the OLS regressions are well defined. Our choice of OLS and the assumption of an infinitely large estimation sample are for expositional purposes. The former ensures tractability, while the latter abstracts from finite-sample issues and focuses attention on the systematic implications of investors' statistical models. Our main results remain valid under other estimation approaches, such as maximum likelihood or Bayesian learning. As discussed below, the key assumption is on the restrictions imposed on the coefficient matrices, not on how these matrices are estimated.

as a covariate.[19] The following result shows that, within this class, investors' expectations coincide with rational expectations whenever their statistical models are correctly specified, in the sense that the true data-generating process falls within the set of statistical models they consider.

Proposition 3. *Suppose the short rate follows an* $\mathrm{AR}(p)$ *process with* $p \leq k$*. Then, investors' expectations at all maturities and forecast horizons coincide with rational expectations.*

This result implies that, when investors' model is correctly specified, OLS estimation recovers the true dynamics of the yield curve, and as a result, investors' forecast errors are unpredictable at all maturities and forecast horizons. Note that this holds irrespective of which additional covariates are included in the coefficient matrices $\mathbf{A}_s$ beyond the short rate itself: once investors' models include enough lags to recover the short rate's dynamics exactly, any further covariates provide no additional predictive power for forecasting yields, as all yields are functions of the last p realizations of the short rate through (18). More importantly for our purposes, since investors' expectations coincide with rational expectations whenever $p \leq k$, realized and expected interest rates satisfy the expectations hypothesis, and hence the consistency restrictions in (15) are automatically satisfied—even though (18) is not explicitly incorporated into investors' model. Thus, a correctly specified statistical model, estimated on data that satisfy (18), is enough for investors to uncover the cross-maturity dynamics of the yield curve, including (18) itself and all consistency relationships implied by it.

We next turn to the case in which investors have a misspecified model of yield curve dynamics.

Proposition 4. *Suppose the short rate does not follow an* $\mathrm{AR}(p)$ *process with* $p \leq k$*.*

(a) If only the first columns of $\mathbf{A}_s$ *are nonzero, then the consistency restrictions in* (15) *are satisfied.*

(b) If $\mathbf{A}_s$ *are lower triangular and otherwise unrestricted, the consistency restrictions in* (15) *are generically violated.*

Part (a) establishes that when investors forecast each yield using only the short rate as a covariate, their expectations satisfy the consistency restrictions in (15) even when their model of the short rate is misspecified. This is because, since all yields are linked to the short rate through (18), projecting each yield onto the history of the short rate is sufficient to uncover the structural relationships linking yields of different maturities, without ever explicitly imposing them on the statistical model. As a result, investors' expectations respond to changes in the short rate in a manner consistent with (15), regardless of whether their model of the short rate itself is correctly specified or not.

It is important to emphasize, however, that satisfying (15) does not mean that investors' expectations are consistent with the true data-generating process. As discussed in Section 4, the consistency restrictions in (15) are purely internal restrictions on investors' subjective expectations: they require that expectations across different maturities and forecast horizons respond to changes in the short rate in a mutually consistent manner, but place no requirement on how those expectations relate to realized outcomes. Indeed, since the true data-generating process does not fall within the class of $\mathrm{VAR}(k)$ models that investors consider, their forecast errors will be predictable—not only for the

[19] Restricting attention to lower triangular matrices avoids the severe multicollinearity that would otherwise arise from including longer-maturity yields as covariates.

short rate but potentially across all maturities—manifesting as violations of the kind documented by Coibion and Gorodnichenko (2015), Cieslak (2018), d'Arienzo (2020), and Wang (2021). Yet, as part (a) establishes, this misspecification does not prevent investors' expectations from satisfying (15): their forecasts remain biased relative to the true data-generating process, but mutually consistent across maturities and forecast horizons.

Part (b) of Proposition 4 shows that this conclusion no longer holds once investors use a larger set of covariates—specifically, the past k lags of all yields up to and including the maturity being forecast—to form their expectations: despite using a richer statistical model than in part (a), investors' expectations generically violate the consistency restrictions in (15). The intuition is that while including more covariates allows investors to form more accurate forecasts for each maturity in isolation, this comes at the cost of violating (15). In other words, these violations emerge as a byproduct of forming optimal forecasts using a misspecified model, with nothing in that process disciplining the resulting expectations to remain mutually consistent across maturities.

Taken together, our results in this section illustrate that subjective expectations may be inconsistent with the equilibrium restrictions implied by the bond market whenever investors' models (i) do not explicitly incorporate those relationships and (ii) are misspecified relative to the true data-generating process. Crucially, such inconsistencies can arise even though expectations are endogenously determined and optimally estimated from data that satisfy those very relationships. In fact, as Proposition 4(b) illustrates, they can emerge precisely because investors' models are optimally estimated, not despite it.

6 Conclusion

The sensitivity of long-term interest rates to movements in the short rate is central to the transmission of monetary policy. Yet, decomposing this sensitivity into contributions from expectations and risk premia requires taking a stand on how expectations are formed. In this paper, we develop an approach that sidesteps this requirement. By treating expectations as model primitives rather than imposing a particular expectation formation process, we derive the unique relationship that solely follows from individual optimization and market clearing: the yield on a long-term bond is determined by the current short-term interest rate, a weighted average of investors' expectations of the bond's resale value—with weights reflecting their risk-bearing capacities—and a residual that, irrespective of how expectations are formed, coincides with the bond's risk premium. Our theoretical results also clarify that expressing long-term yields as functions of subjective forecasts of future shorter rates requires investors to understand the equilibrium relationships linking different interest rates, a condition that can be tested directly.

Using survey data from professional forecasters, we find that movements in expectations, rather than risk premia, account for the bulk of the sensitivity of short- and medium-maturity yields to the short rate. We also document that subjective expectations extracted from survey data systematically violate the cross-maturity restrictions implied by bond market equilibrium: short-horizon expectations of long-term yields overreact relative to long-horizon expectations of short-term yields, when benchmarked against those restrictions. This departure from rational expectations is distinct from the predictability of forecast errors emphasized elsewhere in the literature: it concerns the internal consistency of investors'

beliefs across maturities and forecast horizons, not their accuracy relative to the true data-generating process. Because any decomposition that imposes cross-maturity restrictions on expectations will misattribute this form of overreaction to risk premia, the decomposition we derive from individual optimization and market clearing is not merely an alternative to existing approaches, but is the appropriate one given our empirical findings.

To illustrate how such cross-maturity inconsistencies can arise endogenously, we develop a behavioral model of expectation formation in which investors forecast future yields using a statistical model that is estimated from yield curve data but does not explicitly encode the structural relationships implied by bond market equilibrium. We show that when the true data-generating process falls outside the class of models investors consider, the resulting expectations generically violate these relationships, even when the model is estimated optimally from data that satisfies them. Our model is deliberately stylized to isolate the role of misspecification in generating such violations. A richer framework that accounts for these patterns together with the systematic variation in forecast error predictability across maturities and forecast horizons documented elsewhere in the literature could help distinguish among competing theories of expectation formation and identify which features are essential for explaining the dynamics of yields and expectations, a question we leave for future work.

We believe our approach of taking subjective expectations as model primitives, which is rooted in the temporary equilibrium tradition of Grandmont (1977), speaks to a broader methodological point of independent interest.[20] A growing literature in macroeconomics and finance relaxes the full-information rational expectations benchmark by replacing it with alternative, but still structured, models of belief formation, often tied to a specific behavioral bias or friction—such as diagnostic expectations (Bordalo, Gennaioli, and Shleifer, 2018) or rational inattention (Sims, 2003). Our analysis highlights the value of a complementary approach: imposing no structure on how expectations are formed and characterizing which relationships among economic variables hold under any expectation formation mechanism. While this strategy yields fewer such relationships than structural alternatives, those that survive are correspondingly more robust. We view it as especially valuable when expectations are observable, either directly from surveys or indirectly through asset prices, as it allows researchers to separate the structural content of asset pricing models from the content of any particular theory of expectation formation.

Our behavioral model also points to a simple but, in our view, important observation. Even when economic agents are sophisticated statistical forecasters with access to rich information about the environment, their expectations may fail to satisfy the structural relationships linking the variables they forecast if those relationships are not explicitly encoded in their models. Moreover, enriching the statistical model to improve forecast accuracy need not resolve this tension; it may instead sharpen it by improving predictions variable by variable while making the resulting forecasts less consistent with the cross-equation restrictions that link them. This suggests that forecast accuracy and structural consistency are distinct dimensions along which expectations can be evaluated, and that the latter cannot be inferred from the former. Evaluating whether expectations satisfy equilibrium restrictions therefore requires tests explicitly designed to assess consistency across variables, rather than forecast

[20] Despite its long tradition, this approach has not been widely adopted in modern macroeconomics and finance, with García-Schmidt and Woodford (2019) and Werning (2022) being notable recent exceptions.

accuracy one variable at a time.

Appendix

A Proofs

Throughout, we adopt the standard notation that $z = o(\sigma^k)$ if $\lim_{\sigma\to 0} z/\sigma^k = 0$ in probability, where z may be a random variable that depends on σ.

Proof of Proposition 1

Consider the optimization problem (3) of investor i at time t with wealth w_{it}. Denoting the Lagrange multiplier corresponding to the investor's time-τ budget constraint by $(1-\rho)\rho^{\tau-t}\lambda_{i\tau+1}$, the first-order optimality conditions are given by

$$\lambda_{i\tau+1} = u_i'(w_{i\tau+1}) + \rho\mathbb{E}_{i\tau+1}[\lambda_{i\tau+2}]/p_{\tau+1}^{(1)} \tag{A.1}$$

and

$$\mathbb{E}_{i\tau}[\lambda_{i\tau+1}(p_{\tau+1}^{(n-1)} - p_\tau^{(n)}/p_\tau^{(1)})] = 0 \tag{A.2}$$

for all $\tau \geq t$. Solving for $\lambda_{i\tau+1}$ from (A.1) recursively and using the transversality condition, we have

$$\lambda_{i\tau+1} = \sum_{h=0}^{\infty} \rho^h \mathbb{E}_{i\tau+1}\left[\frac{u_i'(w_{i\tau+h+1})}{p_{\tau+1}^{(1)}\cdots p_{\tau+h}^{(1)}}\right]. \tag{A.3}$$

Furthermore, multiplying the budget constraint in (3) by $\lambda_{i\tau+1}$ and using (A.2) implies that

$$\mathbb{E}_{i\tau}[\lambda_{i\tau+1}w_{i\tau+1}] = \mathbb{E}_{i\tau}[\lambda_{i\tau+1}]w_{i\tau}/p_\tau^{(1)}. \tag{A.4}$$

We use the above equations to derive a second-order approximation to log bond prices as $\sigma \to 0$. But first, we characterize the steady-state values of w_{it+h} and $\mathbb{E}_{it}[\lambda_{it+1}]$. Recall from (4) that when $\sigma = 0$, investor i faces no subjective uncertainty regarding the future and believes that the short-term interest rate is held constant and equal to zero in all future periods $\tau \geq t+1$. Thus, (A.4) implies that $\lim_{\sigma\to 0} w_{it+h+1} = w_{it}$ for all $h \geq 0$. This, together with (A.3), then implies that

$$\lim_{\sigma\to 0} \mathbb{E}_{it}[\lambda_{it+1}] = \frac{u_i'(w_{it})}{1-\rho}. \tag{A.5}$$

With the above in hand, we next obtain a second-order approximation to (A.2) as $\sigma \to 0$. Note that setting $\tau = t$ in (A.2) implies that

$$\mathrm{cov}_{it}(\lambda_{it+1}, p_{t+1}^{(n-1)}) + \mathbb{E}_{it}[\lambda_{it+1}](\mathbb{E}_{it}[p_{t+1}^{(n-1)}] - p_t^{(n)}/p_t^{(1)}) = 0. \tag{A.6}$$

Equation (A.3), combined with the assumption that investors' expectations satisfy the law of iterated expectations, leads to

$$\mathrm{cov}_{it}(\lambda_{it+1}, p_{t+1}^{(n-1)}) = \sum_{h=0}^{\infty} \rho^h \mathrm{cov}_{it}\left(\frac{u_i'(w_{it+h+1})}{p_{t+1}^{(1)}\cdots p_{t+h}^{(1)}}, p_{t+1}^{(n-1)}\right),$$

and as a result,

$$\begin{aligned}\mathrm{cov}_{it}(\lambda_{it+1}, p_{t+1}^{(n-1)}) = u_i''(w_{it}) \sum_{h=0}^{\infty} \rho^h \mathrm{cov}_{it}(w_{it+h+1}, \log p_{t+1}^{(n-1)}) \\ - u_i'(w_{it}) \sum_{h=1}^{\infty} \rho^h \sum_{r=1}^{h} \mathrm{cov}_{it}(\log p_{t+r}^{(1)}, \log p_{t+1}^{(n-1)}) + o(\sigma^2)\end{aligned} \tag{A.7}$$

to a second-order approximation as $\sigma \to 0$, where we are using the fact that when $\sigma = 0$, investor i believes that $\lim_{\sigma \to 0} w_{it+h+1} = w_{it}$. Next, note that setting $\tau = t + h + 1$ in (A.4) implies that

$$w_{it+h+1} = \frac{\mathbb{E}_{it+h+1}[\lambda_{it+h+2} w_{it+h+2}] p_{t+h+1}^{(1)}}{\mathbb{E}_{it+h+1}[\lambda_{it+h+2}]}$$

for all $h \geq 0$. Consequently,

$$\mathrm{cov}_{it}(w_{it+h+1}, \log p_{t+1}^{(n-1)}) = \mathrm{cov}_{it}(w_{it+h+2}, \log p_{t+1}^{(n-1)}) + w_{it} \mathrm{cov}_{it}(\log p_{t+h+1}^{(1)}, \log p_{t+1}^{(n-1)}) + o(\sigma^2).$$

Using this recursion to solve for $\mathrm{cov}_{it}(w_{it+h+1}, \log p_{t+1}^{(n-1)})$ in terms of $\mathrm{cov}_{it}(w_{it+1}, \log p_{t+1}^{(n-1)})$, we obtain

$$\mathrm{cov}_{it}(w_{it+h+1}, \log p_{t+1}^{(n-1)}) = \mathrm{cov}_{it}(w_{it+1}, \log p_{t+1}^{(n-1)}) - w_{it} \sum_{r=0}^{h-1} \mathrm{cov}_{it}(\log p_{t+r+1}^{(1)}, \log p_{t+1}^{(n-1)}) + o(\sigma^2)$$

for all $h \geq 1$. Plugging the above into the first term on the right-hand side of (A.7) implies that

$$\begin{aligned}\mathrm{cov}_{it}(\lambda_{it+1}, p_{t+1}^{(n-1)}) = -\frac{u_i'(w_{it})}{1-\rho} \Bigg(\theta_{it} \mathrm{cov}_{it}(w_{it+1}, \log p_{t+1}^{(n-1)}) \\ + (1 - \gamma_{it}) \sum_{h=1}^{\infty} \rho^h \mathrm{cov}_{it}(\log p_{t+h}^{(1)}, \log p_{t+1}^{(n-1)}) \Bigg) + o(\sigma^2),\end{aligned}$$

where θ_{it} and γ_{it} denote, respectively, investor i's coefficients of absolute and relative risk aversion at time t, defined in (6). Combining the above equation with (A.6) and substituting for w_{it+1} from the budget constraint in (3), we obtain

$$\begin{aligned}\mathbb{E}_{it}[p_{t+1}^{(n-1)}] - p_t^{(n)}/p_t^{(1)} = \theta_{it} \sum_{m=2}^{\infty} x_{it}^{(m)} \mathrm{cov}_t(\log p_{t+1}^{(m-1)}, \log p_{t+1}^{(n-1)}) \\ + (1 - \gamma_{it}) \sum_{h=1}^{\infty} \rho^h \mathrm{cov}_t(\log p_{t+h}^{(1)}, \log p_{t+1}^{(n-1)}) + o(\sigma^2),\end{aligned} \tag{A.8}$$

where we are also using the assumption that $\mathrm{cov}_{it}(\log p_{t+h}^{(n)}, \log p_{t+h}^{(m)})$ is independent of i (equation (4)), and the fact that $\lim_{\sigma \to 0} \mathbb{E}_{it}[\lambda_{it+1}] = u_i'(w_{it})/(1-\rho)$ (equation (A.5)). Dividing both sides of the above equation by θ_{it}, integrating over all i, and imposing market clearing, we obtain

$$\begin{aligned}p_t^{(n)}/p_t^{(1)} = \int_0^1 (\theta_{it}/\bar{\theta}_t)^{-1} \mathbb{E}_{it}[p_{t+1}^{(n-1)}] \mathrm{d}i - \bar{\theta}_t \sum_{m=2}^{\infty} s_t^{(m)} \mathrm{cov}_t(\log p_{t+1}^{(m-1)}, \log p_{t+1}^{(n-1)}) \\ - \left(\int_0^1 (\theta_{it}/\bar{\theta}_t)^{-1} (1 - \gamma_{it}) \mathrm{d}i \right) \sum_{h=1}^{\infty} \rho^h \mathrm{cov}_t(\log p_{t+h}^{(1)}, \log p_{t+1}^{(n-1)}) + o(\sigma^2),\end{aligned}$$

where $\bar{\theta}_t$ is the harmonic average of investors' coefficients of absolute risk aversion at time t. Taking logarithms of both sides of the above equation leads to

$$\begin{aligned}\log p_t^{(n)} = \log p_t^{(1)} &+ \int_0^1 (\theta_{it}/\bar{\theta}_t)^{-1}\mathbb{E}_{it}[\log p_{t+1}^{(n-1)}]\mathrm{d}i + \frac{1}{2}\mathrm{var}_t(\log p_{t+1}^{(n-1)}) - \bar{\theta}_t \sum_{m=2}^{\infty} s_t^{(m)}\mathrm{cov}_t(\log p_{t+1}^{(m-1)}, \log p_{t+1}^{(n-1)}) \\ &- \left(\int_0^1 (\theta_{it}/\bar{\theta}_t)^{-1}(1-\gamma_{it})\mathrm{d}i\right)\sum_{h=1}^{\infty}\rho^h \mathrm{cov}_t(\log p_{t+h}^{(1)}, \log p_{t+1}^{(n-1)}) \\ &+ \frac{1}{2}\int_0^1 (\theta_{it}/\bar{\theta}_t)^{-1}\mathbb{E}_{it}^2[\log p_{t+1}^{(n-1)}]\mathrm{d}i - \frac{1}{2}\left(\int_0^1 (\theta_{it}/\bar{\theta}_t)^{-1}\mathbb{E}_{it}[\log p_{t+1}^{(n-1)}]\mathrm{d}i\right)^2 + o(\sigma^2).\end{aligned}$$

The last two terms on the right-hand side of the above equation reflect the cross-sectional covariance of investors' forecasts of next-period prices, with higher weights assigned to investors with smaller coefficients of absolute risk aversion. However, recall from (5) that the disagreement between investors shrinks to zero at a rate faster than σ. As a result, this cross-sectional variance term decays to zero at a rate faster than σ^2, thus establishing (7).

To complete the proof, we verify that equilibrium quantities $x_{it}^{(m)}$ remain bounded for all investors i as $\sigma \to 0$. To this end, note that equation (A.8) implies that

$$\begin{aligned}\mathbb{E}_{it}[p_{t+1}^{(n-1)}] - \mathbb{E}_{jt}[p_{t+1}^{(n-1)}] = &\sum_{m=2}^{\infty}(\theta_{it}x_{it}^{(m)} - \theta_{jt}x_{jt}^{(m)})\mathrm{cov}_t(\log p_{t+1}^{(m-1)}, \log p_{t+1}^{(n-1)}) \\ &+ (\gamma_{jt}-\gamma_{it})\sum_{h=1}^{\infty}\rho^h \mathrm{cov}_t(\log p_{t+h}^{(1)}, \log p_{t+1}^{(n-1)}) + o(\sigma^2)\end{aligned}$$

for all $i \neq j$. Recall from (4) that $\mathrm{cov}_t(\log p_{t+h}^{(n)}, \log p_{t+s}^{(m)}) = \sigma^2\Sigma_{t,h,s}^{(n,m)}$. Therefore, if we let $\mathbf{\Sigma}_t$ denote the matrix with (n,m) element given by $\Sigma_{t,1,1}^{(n,m)}$, the above equation implies that

$$\begin{aligned}\theta_{it}x_{it}^{(m)} - \theta_{jt}x_{jt}^{(m)} = &\frac{1}{\sigma^2}\sum_{n=1}^{\infty}(\mathbf{\Sigma}_t^{-1})_{m-1,n}(\mathbb{E}_{it}[p_{t+1}^{(n)}] - \mathbb{E}_{jt}[p_{t+1}^{(n)}]) \\ &- (\gamma_{jt}-\gamma_{it})\frac{1}{\sigma^2}\sum_{h=1}^{\infty}\rho^h\sum_{n=1}^{\infty}(\mathbf{\Sigma}_t^{-1})_{m-1,n}\mathrm{cov}_t(\log p_{t+h}^{(1)}, \log p_{t+1}^{(n)}) + o(1)\end{aligned}$$

for all $m \geq 2$. Observe that $\mathbb{E}_{it}[p_{t+1}^{(n)}] = 1 + \mathbb{E}_{it}[\log p_{t+1}^{(n)}] + \frac{1}{2}\mathbb{E}_{it}[(\log p_{t+1}^{(n)})^2] + o(\sigma^2)$ as $\sigma \to 0$. Also recall that, by assumption, disagreement about any random variable that decays to zero at rate σ^k vanishes at rate σ^{k+1} or faster. Since $\log p_{t+1}^{(n)}$ is of order σ, it follows that $\mathbb{E}_{it}[(\log p_{t+1}^{(n)})^2] - \mathbb{E}_{jt}[(\log p_{t+1}^{(n)})^2] = o(\sigma^2)$. Consequently, we can rewrite the above equation as

$$\begin{aligned}\theta_{it}x_{it}^{(m)} - \theta_{jt}x_{jt}^{(m)} = &\frac{1}{\sigma^2}\sum_{n=1}^{\infty}(\mathbf{\Sigma}_t^{-1})_{m-1,n}(\mathbb{E}_{it}[\log p_{t+1}^{(n)}] - \mathbb{E}_{jt}[\log p_{t+1}^{(n)}]) \\ &- (\gamma_{jt}-\gamma_{it})\frac{1}{\sigma^2}\sum_{h=1}^{\infty}\rho^h\sum_{n=1}^{\infty}(\mathbf{\Sigma}_t^{-1})_{m-1,n}\mathrm{cov}_t(\log p_{t+h}^{(1)}, \log p_{t+1}^{(n)}) + o(1).\end{aligned}$$

Note that assumptions (4) and (5) guarantee that the right-hand side of the above equation remains finite as $\sigma \to 0$. Hence,

$$\limsup_{\sigma\to 0}|\theta_{it}x_{it}^{(m)} - \theta_{jt}x_{jt}^{(m)}| < \infty$$

for all i and j. Dividing by θ_{jt}, integrating over all investors j, and using market clearing then guarantees that $\limsup_{\sigma\to 0}|x_{it}^{(m)}|<\infty$ for all $m\geq 2$. □

Proof of Proposition 2

Recall from Proposition 1 that realized and expected bond yields satisfy

$$y_t^{(n)}=\frac{1}{n}y_t^{(1)}+\frac{n-1}{n}\overline{\mathbb{E}}_t[y_{t+1}^{(n-1)}]+\mathrm{rp}_t^{(n)}+o(\sigma^2), \tag{A.9}$$

as $\sigma\to 0$, where $\mathrm{rp}_t^{(n)}$ denotes the bond's risk premium. Since investor i understands market clearing and other investors' portfolio optimization, the investor also understands the equilibrium mapping from primitives to bond prices, including the fact that all bond yields must satisfy the above equation at any given future date $t+h$. As a result,

$$\mathbb{E}_{it}[y_{t+h}^{(n)}]=\frac{1}{n}\mathbb{E}_{it}[y_{t+h}^{(1)}]+\frac{n-1}{n}\mathbb{E}_{it}\overline{\mathbb{E}}_{t+h}[y_{t+h+1}^{(n-1)}]+\mathbb{E}_{it}[\mathrm{rp}_{t+h}^{(n)}]+o(\sigma^2). \tag{A.10}$$

Next, note that since investors believe that risk premia are constant, it follows that $\mathbb{E}_{it}[\mathrm{rp}_{t+h}^{(n)}]=\overline{\mathrm{rp}}_i^{(n)}+o(\sigma^2)$ for all t and all $h\geq 1$, where $\overline{\mathrm{rp}}_i^{(n)}$ is a constant that does not depend on the time index. Substituting this expression into (A.10), multiplying both sides by $\bar{\theta}_t/\theta_{it}$, integrating over all i, and taking time differences, we get

$$\Delta\overline{\mathbb{E}}_t[y_{t+h}^{(n)}]=\frac{1}{n}\Delta\overline{\mathbb{E}}_t[y_{t+h}^{(1)}]+\frac{n-1}{n}\Delta\overline{\mathbb{E}}_t\overline{\mathbb{E}}_{t+h}[y_{t+h+1}^{(n-1)}]+o(\sigma^2)$$

for all maturities n and all forecast horizons $h\geq 1$. If the consensus expectations operator satisfies the law of iterated expectations as $\sigma\to 0$ to a second-order approximation—a claim that we prove below—the above equation simplifies to

$$n\Delta\overline{\mathbb{E}}_t[y_{t+h}^{(n)}]-(n-1)\Delta\overline{\mathbb{E}}_t[y_{t+h+1}^{(n-1)}]=\Delta\overline{\mathbb{E}}_t[y_{t+h}^{(1)}]+o(\sigma^2) \tag{A.11}$$

for all $n\geq 1$ and $h\geq 1$. Summing both sides of the above equation over n from $n=1$ to q with h replaced by $h+q-n$, we obtain

$$\Delta\overline{\mathbb{E}}_t[y_{t+h}^{(q)}]=\frac{1}{q}\sum_{s=0}^{q-1}\Delta\overline{\mathbb{E}}_t[y_{t+h+s}^{(1)}]+o(\sigma^2).$$

Regressing both sides of the above equation on $\Delta y_t^{(1)}$ implies that

$$\gamma_h^{(q)}=\frac{1}{q}\sum_{s=0}^{q-1}\gamma_{h+s}^{(1)}+o(\sigma), \tag{A.12}$$

where $\gamma_h^{(q)}$ denotes the sensitivity coefficient defined in (14) and we are using the fact that $\Delta y_t^{(1)}=O(\sigma)$ and $\mathrm{var}(\Delta y_t^{(1)})=O(\sigma^2)$. Substituting for $\gamma_{h+n_1+\cdots+n_{j-1}}^{(n_j)}$ from the above into the right-hand side of (15) leads to

$$\frac{1}{n}\sum_{j=1}^{\ell}n_j\gamma_{h+n_1+\cdots+n_{j-1}}^{(n_j)}=\frac{1}{n}\sum_{j=1}^{\ell}\sum_{s=0}^{n_j-1}\gamma_{h+n_1+\cdots+n_{j-1}+s}^{(1)}+o(\sigma)=\frac{1}{n}\sum_{s=0}^{n-1}\gamma_{h+s}^{(1)}+o(\sigma)=\gamma_h^{(n)}+o(\sigma)$$

for all $h \geq 1$, where the last equality uses (A.12) applied to $\gamma_h^{(n)}$. This establishes (15).

The proof is therefore complete once we show that consensus expectations operator satisfies the law of iterated expectations to a second-order approximation, that is,

$$\overline{\mathbb{E}}_t \overline{\mathbb{E}}_{t+s-1}[y_{t+s}^{(n)}] = \overline{\mathbb{E}}_t[y_{t+s}^{(n)}] + o(\sigma^2) \tag{A.13}$$

as $\sigma \to 0$ for all $s \geq 1$. To this end, observe that

$$\overline{\mathbb{E}}_t \overline{\mathbb{E}}_{t+s}[y_{t+s+1}^{(n)}] = \int_0^1 \omega_{it} \mathbb{E}_{it} \overline{\mathbb{E}}_{t+s}[y_{t+s+1}^{(n)}] \mathrm{d}i = \int_0^1 \int_0^1 \omega_{it} \mathbb{E}_{it}[\omega_{jt+s} \mathbb{E}_{jt+s}[y_{t+s+1}^{(n)}]] \mathrm{d}j \mathrm{d}i$$

and

$$\overline{\mathbb{E}}_t[y_{t+s+1}^{(n)}] = \int_0^1 \omega_{it} \mathbb{E}_{it} \mathbb{E}_{it+s}[y_{t+s+1}^{(n)}] \mathrm{d}i = \int_0^1 \int_0^1 \omega_{it} \mathbb{E}_{it}[\omega_{jt+s} \mathbb{E}_{it+s}[y_{t+s+1}^{(n)}]] \mathrm{d}j \mathrm{d}i,$$

where $\omega_{it} = \bar{\theta}_t / \theta_{it}$ and we are using the fact that $\int_0^1 \omega_{jt+s} \mathrm{d}j = 1$. Therefore,

$$\overline{\mathbb{E}}_t \overline{\mathbb{E}}_{t+s}[y_{t+s+1}^{(n)}] - \overline{\mathbb{E}}_t[y_{t+s+1}^{(n)}] = \int_0^1 \int_0^1 \omega_{it} \mathbb{E}_{it} \left[\omega_{jt+s} \left(\mathbb{E}_{jt+s}[y_{t+s+1}^{(n)}] - \mathbb{E}_{it+s}[y_{t+s+1}^{(n)}] \right) \right] \mathrm{d}j \mathrm{d}i.$$

By a simple change of variables,

$$\begin{aligned} \overline{\mathbb{E}}_t \overline{\mathbb{E}}_{t+s}[y_{t+s+1}^{(n)}] - \overline{\mathbb{E}}_t[y_{t+s+1}^{(n)}] = & \int_0^1 \int_i^1 \omega_{it} \mathbb{E}_{it} \left[\omega_{jt+s} \left(\mathbb{E}_{jt+s}[y_{t+s+1}^{(n)}] - \mathbb{E}_{it+s}[y_{t+s+1}^{(n)}] \right) \right] \mathrm{d}j \mathrm{d}i \\ & - \int_0^1 \int_i^1 \omega_{jt} \mathbb{E}_{jt} \left[\omega_{it+s} \left(\mathbb{E}_{jt+s}[y_{t+s+1}^{(n)}] - \mathbb{E}_{it+s}[y_{t+s+1}^{(n)}] \right) \right] \mathrm{d}j \mathrm{d}i. \end{aligned}$$

As a result,

$$\overline{\mathbb{E}}_t \overline{\mathbb{E}}_{t+s}[y_{t+s+1}^{(n)}] - \overline{\mathbb{E}}_t[y_{t+s+1}^{(n)}] = \int_{j>i} (\omega_{it} \mathbb{E}_{it}[\omega_{jt+s} \Delta_{ij,t+s}^{(n)}] - \omega_{jt} \mathbb{E}_{jt}[\omega_{it+s} \Delta_{ij,t+s}^{(n)}]) \mathrm{d}j \mathrm{d}i,$$

where $\Delta_{ij,t}^{(n)} = \mathbb{E}_{jt}[y_{t+1}^{(n)}] - \mathbb{E}_{it}[y_{t+1}^{(n)}]$ denotes the disagreement between investors i and j at time t about bond yields in the next period. Dividing both sides by σ^2, and taking the limit as $\sigma \to 0$, we obtain

$$\lim_{\sigma \to 0} \frac{1}{\sigma^2} \left(\overline{\mathbb{E}}_t \overline{\mathbb{E}}_{t+s}[y_{t+s+1}^{(n)}] - \overline{\mathbb{E}}_t[y_{t+s+1}^{(n)}] \right) = \lim_{\sigma \to 0} \frac{1}{\sigma^2} \int_{j>i} (\omega_{it} \mathbb{E}_{it}[\omega_{jt+s} \Delta_{ij,t+s}^{(n)}] - \omega_{jt} \mathbb{E}_{jt}[\omega_{it+s} \Delta_{ij,t+s}^{(n)}]) \mathrm{d}j \mathrm{d}i$$

Recall from (4) that when $\sigma = 0$, investor i faces no uncertainty regarding the future and believes that the short-term interest rate is held constant and equal to zero in all future periods $\tau \geq t+1$ (i.e., $\lim_{\sigma \to 0} p_t^{(1)} = 1$). Therefore, (A.4) implies that $\lim_{\sigma \to 0} w_{it} = \widetilde{w}_i$ for all t. This, together with inequality (5), which maintains that the disagreement term $\Delta_{ij,t+s}^{(n)} = O(\sigma^2)$ as $\sigma \to 0$, implies that

$$\lim_{\sigma \to 0} \frac{1}{\sigma^2} \left(\overline{\mathbb{E}}_t \overline{\mathbb{E}}_{t+s}[y_{t+s+1}^{(n)}] - \overline{\mathbb{E}}_t[y_{t+s+1}^{(n)}] \right) = \lim_{\sigma \to 0} \frac{1}{\sigma^2} \int_{j>i} \widetilde{\omega}_i \widetilde{\omega}_j (\mathbb{E}_{it}[\Delta_{ij,t+s}^{(n)}] - \mathbb{E}_{jt}[\Delta_{ij,t+s}^{(n)}]) \mathrm{d}j \mathrm{d}i.$$

Recall that, by assumption, for any random variable z that decays to zero at rate σ^k, the disagreement $\mathbb{E}_{it}[z] - \mathbb{E}_{jt}[z]$ between any two investors $i \neq j$ about that variable decays to zero at rate σ^{k+1} or faster. Setting $z = \Delta_{ij,t+s}^{(n)}$ and once again using the fact that $\Delta_{ij,t+s}^{(n)} = O(\sigma^2)$ as $\sigma \to 0$ implies that the right-hand side of the above equation is equal to zero, thus establishing (A.13). □

Proof of Proposition 3

We use an inductive argument to establish that for every $n \geq 1$ the yield on the n-period bond admits a representation of the form

$$y_t^{(n)} = c_0^{(n)} + \sum_{s=1}^{p} c_s^{(n)} y_{t-s+1}^{(1)} \tag{A.14}$$

for some constants $c_0^{(n)}, \ldots, c_p^{(n)}$, and that

$$\overline{\mathbb{E}}_t[y_{t+1}^{(n)}] = \mathbb{E}_t^*[y_{t+1}^{(n)}], \tag{A.15}$$

where $\mathbb{E}_t^*[\cdot]$ denotes the expectation operator under the true data-generating process.

We start by establishing the base of the induction. Observe that (A.14) holds trivially for $n = 1$, where $c_1^{(1)} = 1$ and $c_s^{(1)} = 0$ for all $s \neq 1$. To establish (A.15) for $n = 1$, note that the assumption that the short rate follows an $\mathrm{AR}(p)$ process implies that

$$y_{t+1}^{(1)} = a_0^* + \sum_{s=1}^{p} a_s^* y_{t-s+1}^{(1)} + \varepsilon_{t+1}, \tag{A.16}$$

for some constants $a_0^*, \ldots, a_p^*$, where ε_{t+1} satisfies $\mathbb{E}_t^*[\varepsilon_{t+1}] = 0$ and is therefore orthogonal to all date-t variables. Since the coefficient matrices $\mathbf{A}_s$ in (19) are lower triangular, investors form their subjective expectation of the short rate, $\overline{\mathbb{E}}_t[y_{t+1}^{(1)}]$, by regressing $y_{t+1}^{(1)}$ on its own past k lags $y_t^{(1)}, \ldots, y_{t-k+1}^{(1)}$. Since $p \leq k$, every variable on the right-hand side of (A.16) is among these regressors, all of which are orthogonal to ε_{t+1}. Thus, with an infinitely large sample, OLS recovers the conditional mean of $y_{t+1}^{(1)}$, implying that $\overline{\mathbb{E}}_t[y_{t+1}^{(1)}] = \mathbb{E}_t^*[y_{t+1}^{(1)}]$, and hence establishing (A.15) for $n = 1$.

As the induction hypothesis, suppose that (A.14) and (A.15) hold for all maturities $m \leq n-1$. Using (A.15) at maturity $n-1$ to replace $\overline{\mathbb{E}}_t[y_{t+1}^{(n-1)}]$ on the right-hand side of (18) with $\mathbb{E}_t^*[y_{t+1}^{(n-1)}]$, we obtain

$$y_t^{(n)} = \frac{1}{n} y_t^{(1)} + \frac{n-1}{n} \mathbb{E}_t^*[y_{t+1}^{(n-1)}] + \text{constant}.$$

Furthermore, the induction hypothesis (A.14) at maturity $n-1$ implies that

$$\mathbb{E}_t^*[y_{t+1}^{(n-1)}] = c_0^{(n-1)} + c_1^{(n-1)} \mathbb{E}_t^*[y_{t+1}^{(1)}] + \sum_{s=2}^{p} c_s^{(n-1)} y_{t-s+2}^{(1)}.$$

Combining these two equations and using (A.16) to replace for $\mathbb{E}_t^*[y_{t+1}^{(1)}]$ establishes (A.14) for maturity n.

It remains to establish (A.15) for maturity n. Substituting for $y_{t+1}^{(1)}$ from (A.16) on the right-hand side of (A.14) at maturity n and time $t+1$, we obtain

$$y_{t+1}^{(n)} = c_0^{(n)} + c_1^{(n)} a_0^* + \sum_{s=1}^{p} \left(c_1^{(n)} a_s^* + c_{s+1}^{(n)} \right) y_{t-s+1}^{(1)} + c_1^{(n)} \varepsilon_{t+1}, \tag{A.17}$$

with the convention that $c_{p+1}^{(n)} = 0$. Recall that investors form their subjective expectation $\overline{\mathbb{E}}_t[y_{t+1}^{(n)}]$ by regressing $y_{t+1}^{(n)}$ on the covariates dictated by the restrictions on the coefficient matrices $\mathbf{A}_s$ in (19), with the last k lags of the short rate always included among the covariates. We make two observations.

First, since $p \leq k$, all lags of the short rate on the right-hand side of (A.17) are among these covariates. Second, since $\mathbb{E}_t^*[\varepsilon_{t+1}] = 0$ and all regressors in investors' estimation problem are dated t or earlier, ε_{t+1} is orthogonal to all such regressors. Together, these observations imply that, with an infinitely large sample, OLS recovers the conditional mean of $y_{t+1}^{(n)}$ in (A.17). Hence $\overline{\mathbb{E}}_t[y_{t+1}^{(n)}] = \mathbb{E}_t^*[y_{t+1}^{(n)}]$, establishing (A.15) for maturity n and completing the inductive argument.

The inductive argument above establishes (A.15) for all maturities $n \geq 1$; that is, investors' subjective one-step-ahead forecasts coincide with the corresponding forecasts under the true data-generating process. At the same time, recall that forecasts at horizons $h > 1$ are obtained by iterating the one-step forecasting rule. Since the subjective and objective one-step forecasts agree at every history, starting from the same date-t history results in the same forecasts at every horizon. That is, $\overline{\mathbb{E}}_t[y_{t+h}^{(n)}] = \mathbb{E}_t^*[y_{t+h}^{(n)}]$ for all maturities n and all $h \geq 1$. □

Proof of Proposition 4

Proof of part (a). Since all columns other than the first column of matrices $\mathbf{A}_s$ are restricted to be zero, investors' one-step-ahead forecasts in (19) are given by

$$\overline{\mathbb{E}}_t[y_{t+1}] = \sum_{s=0}^{k-1} a_s \, y_{t-s}^{(1)} + \text{constant},$$

where $a_s \in \mathbb{R}^{N\times 1}$ denotes the first (and only nonzero) column of $\mathbf{A}_s$. Defining $z_t = [y_t^{(1)} \; \dots \; y_{t-k+1}^{(1)}]' \in \mathbb{R}^k$, we can rewrite the above equation as

$$\overline{\mathbb{E}}_t[y_{t+1}] = \tilde{\mathbf{A}} z_t + \text{constant}, \tag{A.18}$$

where $\tilde{\mathbf{A}}$ is an $N \times k$ matrix with $a_0, \dots, a_{k-1}$ as its columns. The assumption that bond risk premia are constant implies that realized and expected yields are related through (18) for all $n \geq 1$. Together with (A.18), this implies that realized yields at any given time t are given by

$$y_t = \tilde{\mathbf{B}} z_t + \text{constant}, \tag{A.19}$$

where $\tilde{\mathbf{B}}$ is an $N \times k$ matrix that is related to $\tilde{\mathbf{A}}$ via

$$\tilde{\mathbf{B}} = \mathbf{V} + \mathbf{Q}\tilde{\mathbf{A}}, \tag{A.20}$$

and $\mathbf{V} \in \mathbb{R}^{N\times k}$ and $\mathbf{Q} \in \mathbb{R}^{N\times N}$ have entries $v_{n,1} = 1/n$ and $q_{n,n-1} = (n-1)/n$ for all n, and zero otherwise. By assumption, investors construct their subjective models by regressing realized yields at time $t+1$ on the past k realizations of the short rate, that is, on z_t. Thus, advancing (A.19) one period, regressing y_{t+1} on z_t, and matching coefficients with (A.18) implies that

$$\tilde{\mathbf{A}} = \tilde{\mathbf{B}}\mathbf{\Phi}, \tag{A.21}$$

where $\mathbf{\Phi}$ is the coefficient matrix from the population regression of z_{t+1} on z_t. Combining (A.20) and (A.21) results in the fixed point equation $\tilde{\mathbf{B}} = \mathbf{V} + \mathbf{Q}\tilde{\mathbf{B}}\mathbf{\Phi}$ in $\tilde{\mathbf{B}}$ in terms of model primitives, whose solution is given by

$$\tilde{\mathbf{B}} = \sum_{s=0}^{\infty} \mathbf{Q}^s \mathbf{V} \mathbf{\Phi}^s = \sum_{s=0}^{N-1} \mathbf{Q}^s \mathbf{V} \mathbf{\Phi}^s, \tag{A.22}$$

where the second equality follows from the observation that $\mathbf{Q}^N = 0$. Since $\overline{\mathbb{E}}_t[z_{t+h}] = \mathbf{\Phi}^h z_t$, equations (A.18), (A.19), and (A.21) then imply that investors' subjective forecasts of yields at any horizon $h \geq 1$ are

$$\overline{\mathbb{E}}_t[y_{t+h}] = \sum_{s=0}^{N-1} \mathbf{Q}^s \mathbf{V} \mathbf{\Phi}^{s+h} z_t + \text{constant}.$$

Multiplying both sides of the above equation by the n-th unit vector, e_n', from the left, taking time differences, and using the observation that $e_n' \mathbf{Q}^s \mathbf{V} = (1/n) e_1'$ for $s \leq n-1$ and zero otherwise, we obtain

$$\Delta \overline{\mathbb{E}}_t[y_{t+h}^{(n)}] = \frac{1}{n} \sum_{s=0}^{n-1} e_1' \mathbf{\Phi}^{s+h} \Delta z_t.$$

Since the above equality holds for all maturities n and all forecast horizons $h \geq 1$, it follows that

$$\Delta \overline{\mathbb{E}}_t[y_{t+h}^{(n)}] - \frac{1}{n} \sum_{r=0}^{n-1} \Delta \overline{\mathbb{E}}_t[y_{t+h+r}^{(1)}] = \frac{1}{n} \sum_{s=0}^{n-1} e_1' \mathbf{\Phi}^{s+h} \Delta z_t - \frac{1}{n} \sum_{r=0}^{n-1} e_1' \mathbf{\Phi}^{h+r} \Delta z_t = 0.$$

Hence, regressing both sides of the above equation on $\Delta y_t^{(1)}$ implies that

$$\gamma_h^{(n)} = \frac{1}{n} \sum_{r=0}^{n-1} \gamma_{h+r}^{(1)}$$

for all $n, h \geq 1$. Fix $h \geq 1$, $\ell \geq 1$, and positive integers $n_1, \ldots, n_\ell$ such that $n = n_1 + \cdots + n_\ell$. Since the above equality holds for all maturities and forecast horizons,

$$\gamma_h^{(n)} - \frac{1}{n} \sum_{s=1}^{\ell} n_s \gamma_{h+n_1+\cdots+n_{s-1}}^{(n_s)} = \frac{1}{n} \sum_{r=0}^{n-1} \gamma_{h+r}^{(1)} - \frac{1}{n} \sum_{s=1}^{\ell} \sum_{r=0}^{n_s-1} \gamma_{h+n_1+\cdots+n_{s-1}+r}^{(1)} = 0,$$

thus establishing (15). □

Proof of part (b). We show that when the coefficient matrices $\mathbf{A}_s$ are lower triangular but otherwise unrestricted, $\delta_1\langle 1,1\rangle \neq 0$ generically, where $\delta_1\langle 1,1\rangle$ is given by (15) corresponding to $h = 1$, $\ell = 2$, and $n_1 = n_2 = 1$.

Let $a_{ij}^{(s)}$ denote the (i,j) entry of $\mathbf{A}_s$. The assumption that the coefficient matrices $\mathbf{A}_s$ are lower triangular implies that investors' one-step-ahead forecast of the short rate is given by

$$\overline{\mathbb{E}}_t[y_{t+1}^{(1)}] = \sum_{s=0}^{k-1} a_{11}^{(s)} y_{t-s}^{(1)} + \text{constant},$$

where investors estimate the coefficients $a_{11}^{(s)}$ by regressing $y_{t+1}^{(1)}$ on the past k lags of the short rate. Since the short rate is not an $\mathrm{AR}(p)$ process with $p \leq k$, the coefficient on the most distant lag in this $\mathrm{AR}(k)$ projection is generically nonzero; that is,

$$a_{11}^{(k-1)} \neq 0. \tag{A.23}$$

Next, recall that realized and expected yields are related via (18). Setting $n = 2$ and substituting investors' one-step-ahead forecast of the short rate implies that the yield on the 2-period bond is given by

$$y_t^{(2)} = \frac{1}{2} y_t^{(1)} + \frac{1}{2} \sum_{s=0}^{k-1} a_{11}^{(s)} y_{t-s}^{(1)} + \text{constant}. \tag{A.24}$$

Turning to investors' one-step-ahead forecast of the yield on the 2-period bond, the assumption that the coefficient matrices $\mathbf{A}_s$ are lower triangular implies that

$$\overline{\mathbb{E}}_t[y_{t+1}^{(2)}] = \sum_{s=0}^{k-1} a_{21}^{(s)} y_{t-s}^{(1)} + \sum_{s=0}^{k-1} a_{22}^{(s)} y_{t-s}^{(2)} + \text{constant}, \tag{A.25}$$

where the coefficients on the right-hand side are obtained by regressing $y_{t+1}^{(2)}$ on the past k lags of the one-period and two-period yields. We next characterize how the coefficient $a_{22}^{(k-1)}$ on the most distant lag of the two-period yield depends on the autocovariances of the short rate. To this end, we use (A.24) to express both the regressand and the regressors in terms of the short rate. For the regressand, advancing (A.24) one period implies that

$$y_{t+1}^{(2)} = \frac{1}{2}\big(1 + a_{11}^{(0)}\big) y_{t+1}^{(1)} + \frac{1}{2} \sum_{s=0}^{k-2} a_{11}^{(s+1)} y_{t-s}^{(1)} + \text{constant}, \tag{A.26}$$

while for the regressors $\{y_{t-s}^{(2)}\}_{s=0}^{k-1}$ in (A.25), we have

$$y_{t-s}^{(2)} = \frac{1}{2} \left(y_{t-s}^{(1)} + \sum_{r=0}^{k-1} a_{11}^{(r)} y_{t-s-r}^{(1)} \right) + \text{constant}. \tag{A.27}$$

We make two observations. First, among the regressors in (A.25), only $y_{t-k+1}^{(2)}$ depends on $y_{t-2k+2}^{(1)}$. The coefficient on this lag is $\frac{1}{2} a_{11}^{(k-1)}$, which is generically nonzero by (A.23). All other regressors depend only on $y_t^{(1)}, \ldots, y_{t-2k+3}^{(1)}$. Consequently, $y_{t-k+1}^{(2)}$ does not lie in the span of the remaining regressors, which guarantees that the coefficient $a_{22}^{(k-1)}$ is uniquely determined despite potential collinearity of the regressors. Second, by (A.26), the regressand $y_{t+1}^{(2)}$ differs from $\frac{1}{2}(1 + a_{11}^{(0)}) y_{t+1}^{(1)}$ only by terms that are already included among the regressors in (A.25). Therefore, by the Frisch–Waugh–Lovell theorem,

$$a_{22}^{(k-1)} = \pi \, (1 + a_{11}^{(0)}) / a_{11}^{(k-1)}, \tag{A.28}$$

where $\pi = \operatorname{cov}^*(y_{t+1}^{(1)}, u_{t-2k+2}) / \operatorname{var}^*(u_{t-2k+2})$ is the coefficient of the projection of $y_{t+1}^{(1)}$ on the component u_{t-2k+2} of $y_{t-2k+2}^{(1)}$ that is orthogonal to $y_t^{(1)}, \ldots, y_{t-2k+3}^{(1)}$, and var^* and cov^* denote the variance and covariance under the true data-generating process.

With the above in hand, we now turn to establishing that $\delta_1\langle 1,1\rangle \neq 0$. Recall from (17) that $\delta_1\langle 1,1\rangle$ is the slope coefficient in the regression

$$\Delta\overline{\mathbb{E}}_t[y_{t+1}^{(2)}] - \frac{1}{2}\Delta\overline{\mathbb{E}}_t[y_{t+1}^{(1)}] - \frac{1}{2}\Delta\overline{\mathbb{E}}_t[y_{t+2}^{(1)}] = \alpha + \delta_1\langle 1,1\rangle \Delta y_t^{(1)} + \varepsilon_t. \tag{A.29}$$

Setting $n = 2$ and $n = 3$ in (18) implies that

$$\Delta y_t^{(3)} - \frac{2}{3}\Delta y_t^{(2)} - \frac{1}{3}\Delta\overline{\mathbb{E}}_t[y_{t+2}^{(1)}] = \frac{2}{3}\left(\Delta\overline{\mathbb{E}}_t[y_{t+1}^{(2)}] - \frac{1}{2}\Delta\overline{\mathbb{E}}_t[y_{t+1}^{(1)}] - \frac{1}{2}\Delta\overline{\mathbb{E}}_t[y_{t+2}^{(1)}] \right).$$

Thus, it is sufficient to show that the slope coefficient from regressing the left-hand side of the above equation on $\Delta y_t^{(1)}$ is generically nonzero. Since investors forecast the short rate using its past k realizations, (A.24), (A.25), and (A.27) imply that this expression can be written as

$$\Delta y_t^{(3)} - \frac{2}{3}\Delta y_t^{(2)} - \frac{1}{3}\Delta\overline{\mathbb{E}}_t[y_{t+2}^{(1)}] = \sum_{\ell=0}^{2k-2} \beta_\ell \Delta y_{t-\ell}^{(1)}, \tag{A.30}$$

where the coefficients β_ℓ are uniquely determined, as the lags of the short rate are linearly independent. In particular,

$$\beta_{2k-2} = \frac{1}{3}\, a_{11}^{(k-1)} a_{22}^{(k-1)} = \frac{1}{3}\big(1 + a_{11}^{(0)}\big) \frac{\operatorname{cov}^*(y_{t+1}^{(1)}, u_{t-2k+2})}{\operatorname{var}^*(u_{t-2k+2})}, \tag{A.31}$$

where the second equality follows from (A.28). Consequently, the regression coefficient $\delta_1\langle 1,1\rangle$ satisfies

$$\frac{2}{3}\delta_1\langle 1,1\rangle \cdot \operatorname{var}^*(\Delta y_t^{(1)}) = \sum_{\ell=0}^{2k-2} \beta_\ell \cdot \operatorname{cov}^*(\Delta y_t^{(1)}, \Delta y_{t-\ell}^{(1)}). \tag{A.32}$$

It is therefore sufficient to establish that the right-hand side of the above equation is generically nonzero. Note that the right-hand side of (A.32) depends on the short rate process only through $\xi_0^*, \ldots, \xi_{2k-1}^*$, where $\xi_\ell^* = \operatorname{cov}^*(y_t^{(1)}, y_{t-\ell}^{(1)})$ denotes the lag-ℓ autocovariance of the short rate. In what follows, we establish that the right-hand side of (A.32) is not identically zero as a function of these autocovariances. To this end, we fix $\xi_0^*, \ldots, \xi_{2k-2}^*$ and vary only ξ_{2k-1}^*. We make three observations.

First, the covariance matrix of the regressors in (A.25) depends only on autocovariances up to lag $2k-2$. Hence, varying ξ_{2k-1}^* affects the regression coefficients only through the covariance between the regressand $y_{t+1}^{(2)}$ and the most distant regressor $y_{t-k+1}^{(2)}$. This implies that every coefficient β_ℓ in (A.30) is an affine function of ξ_{2k-1}^*.

Second, since neither u_{t-2k+2} nor $\operatorname{var}^*(u_{t-2k+2})$ depends on ξ_{2k-1}^*, while $\operatorname{cov}^*(y_{t+1}^{(1)}, u_{t-2k+2})$ is affine in ξ_{2k-1}^* with slope one, (A.31) implies that β_{2k-2} is affine in ξ_{2k-1}^* with slope $(1 + a_{11}^{(0)})/(3\operatorname{var}^*(u_{t-2k+2}))$, which is generically nonzero.

Third, note that $\operatorname{cov}^*(\Delta y_t^{(1)}, \Delta y_{t-\ell}^{(1)}) = 2\xi_\ell^* - \xi_{\ell-1}^* - \xi_{\ell+1}^*$. As a result, among the covariances on the right-hand side of (A.32), the one corresponding to $\ell = 2k-2$ is an affine function of ξ_{2k-1}^* with slope -1, while all other covariances do not depend on ξ_{2k-1}^*.

Putting the above three observations together implies that the right-hand side of (A.32) is a quadratic function of ξ_{2k-1}^*, with a coefficient on the quadratic term that is generically nonzero. Hence, the right-hand side of (A.32) is not identically zero as a function of the autocovariances of the short rate. □

B Blue Chip Financial Forecasts

Blue Chip Financial Forecasts is a monthly survey of interest rates and macroeconomic variables dating back to the 1980s, with an average of approximately 60 forecasters providing estimates each month. BCFF is particularly suited for our purposes for two reasons. First, forecasters are identified by name, allowing us to match their forecasts with balance-sheet information. Second, forecasters consistently provide projections for the whole term structure of interest rates at several horizons. Forecasters consist of asset managers, banks, brokers, insurance and mortgage brokers, as well as research institutions. In our empirical analysis we focus on primary dealers in the US Treasury market because we have detailed balance sheet data on those institutions. In our sample from 2001 to 2024, we identify 24 primary dealers. Not all primary dealers are present throughout the sample, and on average, we have 12 forecasters per quarter.

A first concern maybe that primary dealers are different from other forecasters. In the following, we first investigate whether primary dealers' forecasts differ systematically from those of other forecasters.

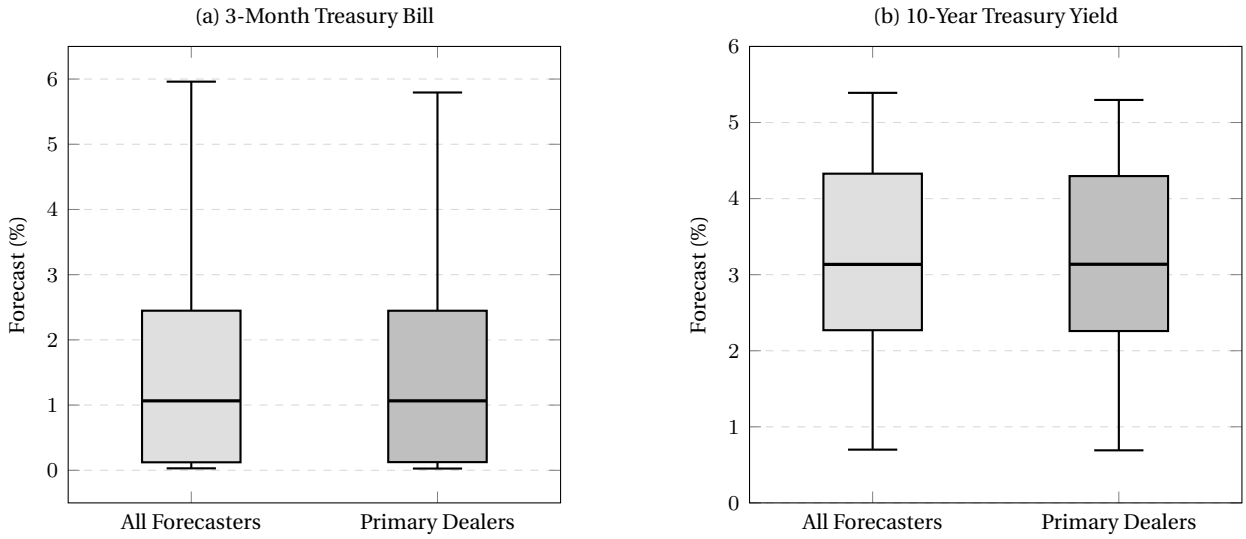


Figure B.1. Box Plots of Interest Rate Forecasts

Notes: This figure presents box plots of one-quarter-ahead forecasts for the 3-month Treasury bill rate (left panel) and the 10-year Treasury yield (right panel), separately for all BCFF forecasters and for primary dealers. The sample period runs from 2001:Q1 to 2024:Q4.

Figure B.1 presents box plots of one-quarter-ahead forecasts for the 3-month Treasury bill and the 10-year Treasury yield, separately for all forecasters and for primary dealers. The distributions are nearly identical across the two groups, and the time-series correlation between the two consensus measures exceeds 99% for both series.

A second question is whether BCFF forecasts differ from those of other surveys. To this end, we compare the BCFF forecasts to those in the Survey of Professional Forecasters (SPF) provided by the Federal Reserve Bank of Philadelphia. SPF forecasters are anonymous, though the survey provides an industry classification. On average, around 40 forecasters contribute projections each quarter, with more than half working in non-financial industries such as academic or other research institutions. The SPF covers a broad range of macroeconomic variables and asset prices, including the 3-month Treasury bill rate and the 10-year Treasury yield.

We use the unweighted average of SPF forecasts as our benchmark. Comparing forecasts across all BCFF forecasters, primary dealers, and SPF, we find a time-series correlation of 99% at all horizons for both the 3-month Treasury bill and the 10-year Treasury yield. Forecast accuracy is also similar across groups: Table B.1 reports root mean square errors (RMSE) and mean absolute errors (MAE) for each forecast series, showing that all three measures perform comparably and are nearly unbiased.

Overall, these comparisons suggest that our results are unlikely to be sensitive to the choice of forecaster group or survey.

C Robustness Checks

The empirical analysis in the main text uses data on primary dealer forecasts from the BCFF survey, covering the period 2001:Q1–2024:Q4. The starting date of the sample is dictated by the availability of Treasury holdings data from the EDGAR database, which we use to construct a wealth-weighted

Table B.1. One-Year-Ahead Forecast Accuracy

	3-Month Treasury		10-Year Treasury	
	RMSE	MAE	RMSE	MAE
SPF	1.37	1.02	0.94	0.78
All Forecasters	1.43	1.01	0.92	0.76
Primary Dealers	1.42	0.99	0.91	0.75

Notes: This table reports root mean square errors (RMSE) and mean absolute errors (MAE) for one-year-ahead forecasts of the 3-month Treasury bill rate and the 10-year Treasury yield, in percent, for three forecast series: the Survey of Professional Forecasters (SPF), all BCFF forecasters, and primary dealers only.

consensus forecast as the empirical counterpart to the consensus expectations in (8), as characterized in Proposition 1.

As a robustness check, we reestimate our main results using the full BCFF sample, which covers the period 1988:Q1–2024:Q4. Because the earlier part of this sample predates the availability of Treasury holdings data from the EDGAR database, we cannot construct a wealth-weighted consensus forecast over the entire period. We therefore use an equal-weighted average of forecasts across all available BCFF respondents rather than restricting the sample to primary dealers. While forecasts at horizons $h = 1, \ldots, 5$ are generally available throughout the full sample, six-quarter-ahead forecasts are unavailable during the earlier survey years. Consequently, after taking first differences, the full-sample regressions generally use 147 quarterly observations, while those requiring six-quarter-ahead forecasts begin in 1997:Q2 and use 111 observations.

This extended sample includes two episodes during which the short-term interest rate was constrained near the zero lower bound (ZLB). The short rate exhibits substantially less variation during these periods, which may limit the precision with which sensitivities can be estimated. Furthermore, the relationship between movements in the short rate and movements in yields and expectations may differ when the short rate is constrained near its lower bound. As an additional robustness check, we therefore repeat the analysis while excluding the ZLB quarters. We classify a quarter as a ZLB quarter if the 1-quarter Treasury yield $y_t^{(1)}$ falls below 0.25 percentage points. This criterion identifies 35 quarters, comprising 2009:Q1–2015:Q4 and 2020:Q2–2021:Q4.

C.1 Sensitivity Decomposition

We begin by reestimating the sensitivity decomposition in Table 1 for each of the two robustness checks. Table C.1 reports the results: Panel A uses the full 1988:Q1–2024:Q4 sample with an equal-weighted consensus forecast across all available forecasters, while Panel B excludes the ZLB quarters. The two sets of estimates are strikingly similar to one another, both in point estimates and in statistical significance, and are also similar to our baseline results reported in Table 1 based on the wealth-weighted consensus forecast over the shorter 2001:Q1–2024:Q4 sample restricted to primary dealers.

Consider first the overall sensitivity $\hat{\beta}^{(n)}_{\text{HLW}}$, reported in the first columns of the two panels. The same monotone decline with maturity appears in all three specifications. At $n = 2$ quarters, the estimate is

Table C.1. Sensitivity Decomposition: Robustness

Panel A: Full Sample

n	$\hat{\beta}^{(n)}_{\text{HLW}}$	$\beta^{(n)}_{\text{short}}$	$\hat{\beta}^{(n)}_{\text{exp}}$	$\hat{\beta}^{(n)}_{\text{rp}}$
2	0.99^{***}	0.50	0.46^{***}	0.03
	(0.01)		(0.03)	(0.03)
3	0.97^{***}	0.33	0.65^{***}	-0.01
	(0.02)		(0.04)	(0.05)
5	0.91^{***}	0.20	0.74^{***}	-0.03
	(0.04)		(0.05)	(0.06)
9	0.78^{***}	0.11	0.76^{***}	-0.09
	(0.06)		(0.05)	(0.07)
21	0.53^{***}	0.05	0.64^{***}	-0.15^{**}
	(0.07)		(0.05)	(0.06)
41	0.32^{***}	0.02	0.50^{***}	-0.20^{***}
	(0.06)		(0.06)	(0.04)

Panel B: ZLB-Excluded Sample

n	$\hat{\beta}^{(n)}_{\text{HLW}}$	$\beta^{(n)}_{\text{short}}$	$\hat{\beta}^{(n)}_{\text{exp}}$	$\hat{\beta}^{(n)}_{\text{rp}}$
2	0.99^{***}	0.50	0.45^{***}	0.04
	(0.01)		(0.02)	(0.03)
3	0.97^{***}	0.33	0.63^{***}	0.00
	(0.02)		(0.04)	(0.05)
5	0.91^{***}	0.20	0.72^{***}	-0.01
	(0.04)		(0.05)	(0.06)
9	0.78^{***}	0.11	0.74^{***}	-0.08
	(0.06)		(0.04)	(0.07)
21	0.52^{***}	0.05	0.62^{***}	-0.14^{**}
	(0.07)		(0.05)	(0.06)
41	0.31^{***}	0.02	0.48^{***}	-0.19^{***}
	(0.06)		(0.05)	(0.04)

Notes: This table reports the estimated slope coefficients from regressions (10), (12), and (13), with $\Delta y^{(1)}_t$ as the regressor, where by construction $\beta^{(n)}_{\text{short}}+\hat{\beta}^{(n)}_{\text{exp}}+\hat{\beta}^{(n)}_{\text{rp}} = \hat{\beta}^{(n)}_{\text{HLW}}$. The mechanical component $\beta^{(n)}_{\text{short}} = 1/n$ is known analytically. Consensus forecasts $\overline{\mathbb{E}}_t[\cdot]$ are constructed as the equal-weighted average across all available BCFF respondents. Panel A uses the full sample. Panel B excludes observations dated t for which $y^{(1)}_t < 0.25\%$, comprising 2009:Q1–2015:Q4 and 2020:Q2–2021:Q4. The numbers in parentheses report EWC-HAC standard errors according to Lazarus et al. (2018). Panel A uses $T = 147$ and $\nu = 11$ cosine basis functions; Panel B uses $T = 112$ and $\nu = 9$. Stars indicate statistical significance for the null hypothesis that the coefficient equals zero, with *, **, and *** denoting significance at the 10%, 5%, and 1% levels, respectively. The underlying data run from 1988:Q1–2024:Q4.

0.97 in the baseline and 0.99 in both the full-sample and ZLB-excluded robustness checks. It falls to 0.46, 0.53, and 0.52, respectively, at $n = 21$ quarters, and to 0.26, 0.32, and 0.31 at $n = 41$. Thus, although the equal-weighted estimates are somewhat larger than the baseline estimates at longer maturities, all three specifications imply that the sensitivity declines from near unity at short maturities to roughly one-quarter to one-third at the 10-year horizon.

The decomposition of yield sensitivities into expectations and risk-premium components is also similar across the three specifications. The expectations component $\hat{\beta}^{(n)}_{\text{exp}}$ accounts for the bulk of the total sensitivity at all maturities, while the risk-premium component $\hat{\beta}^{(n)}_{\text{rp}}$ is statistically indistinguishable from zero at short and medium maturities. For instance, at $n = 5$ quarters, the expectations and risk-premium components are 0.79 and -0.14 in the baseline, 0.74 and -0.03 in the full-sample robustness check, and 0.72 and -0.01 after excluding the ZLB. The risk-premium component is statistically insignificant in all three specifications. At longer maturities, the risk-premium component becomes more negative, with the longer samples providing greater precision. At $n = 21$, the estimates are -0.18 in the baseline and -0.15 and -0.14 in the two robustness checks; the baseline estimate is statistically insignificant, while the robustness estimates are significant at the 5% level. At $n = 41$, the corresponding estimates are -0.19, -0.20, and -0.19, with the baseline estimate significant at the 5% level and both robustness estimates significant at the 1% level.

C.2 Sensitivity of Yield Expectations

Next, we reestimate the sensitivity of yield expectations reported in Table 2 for each of the two robustness checks. Table C.2 reports the estimated sensitivity coefficients $\hat{\gamma}_h^{(n)}$ from regression (14) across all six maturities $n \in \{1, 2, 4, 8, 20, 40\}$ and forecast horizons $h \in \{1, \ldots, 6\}$. Panel A uses the full sample, while Panel B excludes the ZLB quarters. The estimates are similar across the two panels: the largest absolute difference is 0.04, the standard errors are comparable, and the significance classifications are identical.

Both robustness checks also reproduce the two main patterns in the baseline results. First, at a given forecast horizon, the sensitivity is relatively stable across shorter maturities and declines toward the long end of the yield curve. At $h = 1$, for example, the estimates for the 1-quarter yield are 0.93 in the baseline, 0.93 in the full equal-weighted sample, and 0.90 after excluding the ZLB. For the 10-year yield, the corresponding estimates are 0.44, 0.51, and 0.49. Second, at a given maturity, the sensitivity generally declines with the forecast horizon. For $n = 1$, the estimates decline from 0.93 to 0.64 between $h = 1$ and $h = 6$ in the baseline, from 0.93 to 0.61 in the full equal-weighted sample, and from 0.90 to 0.57 after excluding the ZLB. For $n = 40$, the corresponding declines are from 0.44 to 0.22, from 0.51 to 0.25, and from 0.49 to 0.23. Thus, the general decline in the sensitivity of yield expectations with both maturity and forecast horizon is preserved across the two robustness checks.

C.3 Tests of Structural Consistency Restrictions

Finally, we examine whether the rejection of the structural-consistency restrictions in (15) is robust to the two alternative specifications. Table C.3 compares the baseline results in Section 4 with those obtained using the full equal-weighted sample and the sample excluding the ZLB quarters. The direction of the estimated deviations is remarkably stable: 33 of the 34 estimates are positive in all three specifications. The average deviation is 0.077 in the baseline, 0.051 in the full equal-weighted sample, and 0.053 after excluding the ZLB. Thus, although the average deviation is somewhat smaller under equal weighting, it is virtually unchanged when the ZLB quarters are excluded.

Both tests of the joint null continue to reject in the robustness samples. The average deviation is statistically significant at the 5% level in the baseline and at the 1% level in each of the two robustness checks. The Bonferroni-adjusted p-values are 0.042, 0.002, and 0.011, respectively. Both tests therefore reject at the 5% level in all three specifications. The conclusion that the structural-consistency restrictions are jointly violated is thus robust to using an equal-weighted consensus forecast over the longer sample and to excluding the ZLB quarters.

Table C.2. Sensitivity of Yield Expectations to Changes in the Short Rate: Robustness

maturity n (in quarters)	forecast horizon h (in quarters)					
	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
Panel A: Full Sample						
1	0.93^{***}	0.96^{***}	0.90^{***}	0.80^{***}	0.71^{***}	0.61^{***}
	(0.05)	(0.06)	(0.05)	(0.05)	(0.06)	(0.05)
2	0.97^{***}	0.96^{***}	0.90^{***}	0.80^{***}	0.70^{***}	0.59^{***}
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.05)
4	0.92^{***}	0.93^{***}	0.86^{***}	0.77^{***}	0.68^{***}	0.56^{***}
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.05)
8	0.86^{***}	0.85^{***}	0.79^{***}	0.68^{***}	0.60^{***}	0.50^{***}
	(0.05)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
20	0.67^{***}	0.66^{***}	0.62^{***}	0.55^{***}	0.45^{***}	0.38^{***}
	(0.05)	(0.05)	(0.06)	(0.06)	(0.06)	(0.05)
40	0.51^{***}	0.49^{***}	0.45^{***}	0.39^{***}	0.33^{***}	0.25^{***}
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.05)
Panel B: ZLB-Excluded Sample						
1	0.90^{***}	0.94^{***}	0.88^{***}	0.78^{***}	0.69^{***}	0.57^{***}
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.04)
2	0.95^{***}	0.93^{***}	0.88^{***}	0.78^{***}	0.68^{***}	0.56^{***}
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.04)
4	0.90^{***}	0.91^{***}	0.84^{***}	0.75^{***}	0.66^{***}	0.53^{***}
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.04)
8	0.84^{***}	0.83^{***}	0.77^{***}	0.66^{***}	0.58^{***}	0.47^{***}
	(0.05)	(0.06)	(0.06)	(0.06)	(0.06)	(0.05)
20	0.65^{***}	0.65^{***}	0.60^{***}	0.53^{***}	0.43^{***}	0.35^{***}
	(0.05)	(0.05)	(0.06)	(0.06)	(0.06)	(0.04)
40	0.49^{***}	0.47^{***}	0.43^{***}	0.37^{***}	0.32^{***}	0.23^{***}
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.04)

Notes: This table reports the estimated slope coefficients $\hat{\gamma}_h^{(n)}$ from regression (14) for maturities $n \in \{1,2,4,8,20,40\}$ quarters and forecast horizons $h \in \{1,\ldots,6\}$ quarters. Consensus forecasts $\overline{\mathbb{E}}_t[\cdot]$ are constructed as the equal-weighted average across all available BCFF respondents. Panel A uses the full sample. Panel B excludes observations dated t for which $y_t^{(1)} < 0.25\%$. The numbers in parentheses report EWC-HAC standard errors according to Lazarus et al. (2018). In Panel A, the $h = 6$ estimates use $T = 111$ and $\nu = 9$; the remaining estimates use $\nu = 11$, with $T = 147$ except for $(n,h) = (8,5)$, where $T = 146$. In Panel B, the $h = 6$ estimates use $T = 76$ and $\nu = 7$; the remaining estimates use $\nu = 9$, with $T = 112$ except for $(n,h) = (8,5)$, where $T = 111$. Stars indicate statistical significance for the null hypothesis that the coefficient equals zero, with *, **, and *** denoting significance at the 10%, 5%, and 1% levels, respectively. The underlying data run from 1988:Q1–2024:Q4.

Table C.3. Tests of Structural Consistency Restrictions: Robustness

	Baseline	Equal-weighted full sample	Equal-weighted ZLB-excluded sample
Positive estimates	$33/34$	$33/34$	$33/34$
Average deviation	0.077^{**}	0.051^{***}	0.053^{***}
	(0.024)	(0.008)	(0.008)
Bonferroni-adjusted p-value	0.042	0.002	0.011

Notes: This table summarizes tests of the structural-consistency restrictions in (15). The baseline uses the wealth-weighted consensus forecast across primary dealers and the main-text sample. The two robustness columns construct consensus forecasts $\overline{\mathbb{E}}_t[\cdot]$ as the equal-weighted average across all available BCFF respondents. The equal-weighted full sample uses all available observations. The equal-weighted ZLB-excluded sample omits observations dated t for which $y_t^{(1)} < 0.25\%$. The average deviation is estimated by regressing the unweighted average of the 34 left-hand-side variables in (17) on $\Delta y_t^{(1)}$. The numbers in parentheses report EWC-HAC standard errors according to Lazarus et al. (2018). Stars on the average deviation indicate statistical significance based on the fixed-smoothing t_ν reference distribution, with *, **, and *** denoting significance at the 10%, 5%, and 1% levels, respectively. The average-deviation test uses $T = 95$ and $\nu = 8$ in the baseline, $T = 111$ and $\nu = 9$ in the equal-weighted full sample, and $T = 76$ and $\nu = 7$ in the equal-weighted ZLB-excluded sample. The Bonferroni-adjusted p-value equals 34 times the smallest two-sided EWC-HAC p-value among the individual restrictions. The data for the two equal-weighted specifications run from 1988:Q1–2024:Q4.

References

Abrahams, Michael, Tobias Adrian, Richard K. Crump, Emanuel Moench, and Rui Yu (2016), "Decomposing real and nominal yield curves." *Journal of Monetary Economics*, 84, 182–200.

Adam, Klaus and Albert Marcet (2011), "Internal rationality, imperfect market knowledge and asset prices." *Journal of Economic Theory*, 146(3), 1224–1252.

Barillas, Francisco and Kristoffer P. Nimark (2017), "Speculation and the term structure of interest rates." *Review of Financial Studies*, 30(11), 4003–4037.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam (2024), "Perceptions about monetary policy." *Quarterly Journal of Economics*, 139(4), 2227–2278.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020), "Overreaction in macroeconomic expectations." *American Economic Review*, 110(9), 2748–2782.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018), "Diagnostic expectations and credit cycles." *Journal of Finance*, 73(1), 199–227.

Buraschi, Andrea and Paul Whelan (2022), "Speculation, sentiment, and interest rates." *Management Science*, 68(3), 2308–2329.

Campbell, John Y. and Luis M. Viceira (2001), "Who should buy long-term bonds?" *American Economic Review*, 91(1), 99–127.

Cieslak, Anna (2018), "Short-rate expectations and unexpected returns in treasury bonds." *Review of Financial Studies*, 31(9), 3265–3306.

Coibion, Olivier and Yuriy Gorodnichenko (2015), "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review*, 105(8), 2644–2678.

Crump, Richard K., Stefano Eusepi, and Emanuel Moench (2026), "Is there hope for the expectations hypothesis?" Federal Reserve Bank of New York Staff Reports, No. 1098.

d'Arienzo, Daniele (2020), "Maturity increasing overreaction and bond market puzzles." Working paper.

Ehling, Paul, Michael Gallmeyer, Christian Heyerdahl-Larsen, and Philipp Illeditsch (2018), "Disagreement about inflation and the yield curve." *Journal of Financial Economics*, 127(3), 459–484.

Eyster, Erik and Michele Piccione (2013), "An approach to asset pricing under incomplete and diverse perceptions." *Econometrica*, 81(4), 1483–1506.

Farmer, Leland, Emi Nakamura, and Jon Steinsson (2024), "Learning about the long run." *Journal of Political Economy*, 132(10), 3334–3377.

Friedman, Benjamin M. (1979), "Interest rate expectations versus forward rates: Evidence from an expectations survey." *Journal of Finance*, 34(4), 965–973.

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

Fuster, Andreas, Benjamin Hébert, and David Laibson (2012), "Natural expectations, macroeconomic dynamics, and asset pricing." *NBER Macroeconomics Annual*, 26, 1–48.

Gabaix, Xavier (2014), "A sparsity-based model of bounded rationality." *Quarterly Journal of Economics*, 129(4), 1661–1710.

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.

Gertler, Mark and Peter Karadi (2015), "Monetary policy surprises, credit costs, and economic activity." *American Economic Journal: Macroeconomics*, 7(1), 44–76.

Giglio, Stefano and Bryan Kelly (2018), "Excess volatility: Beyond discount rates." *Quarterly Journal of Economics*, 133(1), 71–127.

Gilchrist, Simon, David López-Salido, and Egon Zakrajšek (2015), "Monetary policy and real borrowing costs at the zero lower bound." *American Economic Journal: Macroeconomics*, 7(1), 77–109.

Grandmont, Jean Michel (1977), "Temporary general equilibrium theory." *Econometrica*, 45(3), 535–572.

Hanson, Samuel G., David O. Lucca, and Jonathan H. Wright (2021), "Rate-amplifying demand and the excess sensitivity of long-term rates." *Quarterly Journal of Economics*, 136(3), 1719–1781.

Hanson, Samuel G. and Jeremy C. Stein (2015), "Monetary policy and long-term real rates." *Journal of Financial Economics*, 115(3), 429–448.

Hubert de Fraisse, Antoine, Peter Maxted, Kunal Sangani, and David Sraer (2026), "Pre-trends ahead of monetary policy surprises: Recovering the response of corporate investment." Working paper.

Kekre, Rohan, Moritz Lenel, and Federico Mainardi (2026), "Monetary policy, segmentation, and the term structure." NBER Working Paper No. 32324.

Lazarus, Eben, Daniel J. Lewis, James H. Stock, and Mark W. Watson (2018), "HAR inference: Recommendations for practice." *Journal of Business and Economic Statistics*, 36, 541–559.

Leombroni, Matteo, Carolin E. Pflueger, and Adi Sunderam (2026), "Risk and return in government bonds." Working paper, Boston College.

Liu, Yan and Cynthia Jing Wu (2021), "Reconstructing the Yield Curve." *Journal of Financial Economics*, 142(3), 1395–1425.

Molavi, Pooya (2026), "Simple models and biased forecasts." Working Paper, Northwestern University.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Andrea Vedolin (2024), "Model complexity, expectations, and asset prices." *The Review of Economic Studies*, 91(4), 2462–2507.

Nagel, Stefan and Zhengyang Xu (2023), "Dynamics of subjective risk premia." *Journal of Financial Economics*, 150(2), 103713.

Piatti, Ilaria and Paul Whelan (2026), "Subjective risk premia in bond and FX markets." Working paper, Queen Mary University.

Piazzesi, Monika, Juliana Salomao, and Martin Schneider (2015), "Trend and cycle in bond premia." Working Paper, Stanford University.

Shue, Kelly, Richard Townsend, and Chen Wang (2025), "Long-short interest rate confusion." Working paper, Yale University.

Sims, Christopher A. (2003), "Implications of rational inattention." *Journal of Monetary Economics*, 50(3), 665–690.

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

Wang, Chen (2021), "Under- and overreaction in yield curve expectations." Working paper, University of Notre Dame.

Werning, Iván (2022), "Expectations and the rate of inflation." Massachusetts Institute of Technology Working Paper.

Xiong, Wei and Hongjun Yan (2010), "Heterogeneous expectations and bond markets." *Review of Financial Studies*, 23(4), 1433–1466.