STAFF REPORTS

NO. 934
JULY 2020

REVISED
AUGUST 2021

Fundamental Disagreement about Monetary Policy and the Term Structure of Interest Rates

Shuo Cao | Richard K. Crump | Stefano Eusepi | Emanuel Moench

FEDERAL RESERVE BANK *of* NEW YORK

Fundamental Disagreement about Monetary Policy and the Term Structure of Interest Rates
Shuo Cao, Richard K. Crump, Stefano Eusepi, and Emanuel Moench
Federal Reserve Bank of New York Staff Reports, no. 934
July 2020; revised August 2021
JEL classification: D83, D84, E43, G10, G12

Abstract

Using a unique data set of individual professional forecasts, we document disagreement about the future path of monetary policy, particularly at longer horizons. The stark differences in short rate forecasts imply strong disagreement about the risk-return trade-off of longer-term bonds. Longer-horizon short rate disagreement co-moves with term premiums. We estimate an affine term structure model in which investors hold heterogeneous beliefs about the long-run level of rates. Our model fits Treasury yields and the short rate paths predicted by different groups of investors and thus matches the observed differences in expected return profiles. Investors who correctly anticipated the secular decline in rates became increasingly important for the marginal pricing of risk in the Treasury market. Accounting for heterogeneity in investment performance eliminates the downward trend in the term premium.

Key words: disagreement, heterogeneous beliefs, noisy information, speculation, survey forecasts, yield curve, term premium

---

Crump (corresponding author): Federal Reserve Bank of New York (email: richard.crump@ny.frb.org). Cao: Risk Institute, Shenzhen Stock Exchange (email: shuo.cao@outlook.com). Eusepi: University of Texas at Austin (email: stefano.eusepi@austin.utexas.edu). Moench: Deutsches Bundesbank (email: emanuel.moench@bundesbank.de). The authors thank Pierre Collin-Dufresne, Michael Gallmeyer (WFA discussant), Zongbo Huang (CIFFP discussant), Anh Le (EFA discussant), Tianyu Wang (CICF discussant), Wei Xiong, and seminar participants at the Annual Meeting of the Western Finance Association 2017, 2018 China International Forum on Finance and Policy, 2018 China International Conference in Finance, the Annual Meeting of the European Finance Association 2018, the University of Lausanne, Nankai University, and Peking University for helpful comments and suggestions.

1 Introduction

Bond yields reflect investors' expectations about the future path of short rates as well as their attitudes toward risk. Most term structure models specify these two components of interest rates for a representative investor. While this provides a reasonable starting point for many analyses, it may mask important dynamics among investors and thus fail to provide a complete account of the driving forces behind bond yields. In this paper, we propose and estimate a term structure model which explicitly incorporates differences in beliefs about future short rates.

The paper makes two contributions. We start by using a unique and novel dataset of professional forecasters' individual longer-run expectations from Blue Chip Economic Indicators (BCEI) to document the following facts. First, we confirm that disagreement about future short rates is substantial, particularly at intermediate to long horizons. Second, since term premiums are defined as the difference between observed yields and average expected future short rates, the stark differences in short rate forecasts imply strong disagreement about the risk-return tradeoff of longer-term bonds. Third, disagreement about short rates comoves strongly with the term premium of the consensus forecaster as well as with estimates of the term premium from a reduced-form no-arbitrage model.

In a next step, we build a term structure model that can match these facts. Our model features two investors who disagree about the future path of short rates. We assume that the comovement of bond yields is fully captured by three factors: level, slope, and curvature. While the slope and curvature factors are stationary, the level factor has a time-varying long-run mean which itself follows a random walk. Both investors perfectly observe the three yield curve factors, know the parameters of their data-generating process, have identical preferences and perceive the same volatility of shocks. However, in line with the documented evidence about short-rate disagreement, we assume that they hold different beliefs about the long-run mean of the level factor. As they trade bonds at equilibrium prices, their pricing kernels and hence their perceived risk-return trade-off of longer-term bonds differ. We follow the term structure literature and assume that bond prices are a time-invariant function of the three pricing factors which follow a stationary VAR with a constant mean under the risk-neutral measure. We then show that no-arbitrage restrictions imply that each investor's price of level risk moves proportionally with her belief about the future level of rates.

We estimate our model using zero coupon Treasury yields as well as the term structure of survey forecasts of the federal funds rate for two different hypothetical investors: the top-10 and bottom-10 average responses of the Blue Chip Financial Forecasts (BCFF) survey. Our model fits yields and the two survey forecast paths of the short rate very well in our sample starting in 1983. In the estimated model, investors expecting higher future short rates expect excess bond returns to be negative on average for most maturities. In contrast, investors predicting short rates to be lower perceive average term premiums to range from about 50 basis points at the one-year to about two percent at the ten-year maturity. This suggests that the secular decline in term premiums that typically arises in representative agent term structure models captures expected returns of those investors that better anticipated the observed persistent decline of short rates over the sample.

To study the implications of heterogenous beliefs for the pricing of risk in the Treasury market, we measure the aggregate term premium as the weighted average of expected returns of the individual investors. We obtain the weights by solving a simple portfolio allocation problem under the assumption of mean-variance preferences. Intuitively, an investor expecting higher (lower) future short rates assigns a lower (higher) expected return to longer-term bonds and will tilt her portfolio accordingly. As the weights track the relative wealth of each investor and since short-rates were on a secular downward trend throughout our sample period, this implies that the investor expecting lower future short rates becomes more important for the marginal pricing of risk in the economy over time.

Strikingly, the aggregate term premium does not display any trend unlike common estimates of the term premium in the literature. The ten-year term premium is essentially rangebound between 0% and 2% throughout our sample. The lack of a clear trend is the result of a composition effect. First, the investor who expects short rates to be relatively high perceives low excess bond returns throughout the sample. Second, the investor with a lower expected short rate path anticipates bond returns that feature a strong downward trend in the first part of the sample. Third, this investor with higher expected returns starts with a lower wealth share and thus has little impact on the pricing of risk initially. While the wealth share of this investor is increasing over time, her expected returns continue to decline. This results in an aggregate term premium that does not feature a secular decline.

Our paper contributes to the small but growing literature on bond pricing with heterogeneous beliefs. For an excellent recent discussion of this literature, see Singleton (2021). Our approach is inspired by Xiong and Yan (2010) who theoretically show that heterogeneous beliefs about the long-run mean of fundamentals can be an important driver of bond returns. While in their model individual expected bond returns are constant the economy-wide term premium fluctuates because of time-varying disagreement. Here we propose an affine term structure model with time-varying prices of risk which additionally embeds heterogeneity in beliefs about long-run fundamentals. We estimate the model using a wealth of survey data on short rate expectations and show that disagreement is an important determinant of the price of risk. We further deviate from their analysis by modeling disagreement about the nominal short rate instead of the inflation target. Indeed, as shown in Andrade, Crump, Eusepi, and Moench (2016), disagreement about long-run inflation is not sufficient to account for the sizable long-run disagreement about the short rate.

Other authors have also studied bond pricing with heterogeneous beliefs. Ehling, Gallmeyer, Heyerdahl-Larsen, and Illeditsch (2018) consider a model in which investors with habit formation utility disagree about the distribution of inflation, not just expected inflation. This disagreement induces heterogeneity in investors' consumption and investment decisions and, on average, raises real and nominal bond yields. They further document empirically that inflation disagreement has a strong effect on real and nominal bond yields over and above the impact of expected inflation, consistent with their theoretical model. Buraschi and Whelan (2016) study the interactions between risk aversion and disagreement. In their model heterogeneous beliefs arise because agents have different views about the (constant) long-run growth rate of consumption and because their perceptions of the correlation of shocks differs. They find that disagreement has larger effects on equilibrium bond prices when risk aversion is low. More recently, Buraschi, Piatti, and Whelan (2021) aggregate individual expected excess bond returns based on forecasters' past accuracy in predicting interest rates. In line with our findings, they document that disagreement about bond risk premia is time-varying and persistent. While they show that their measure of aggregate expected bond returns is correlated with disagreement about future real growth and inflation, they do not study its comovement with disagreement about nominal short rates that is the focus of our analysis.

Barillas and Nimark (2019) build a model of the term structure in which investors with heterogeneous information sets form higher-order expectations about the beliefs of all other investors. Equilibrium bond prices then reflect a speculative component which depends on investors' beliefs about the error that the average investor makes when predicting future short rates. Their model suggests that the speculative component explains a sizable fraction of the variation in U.S. Treasury yields. Barillas and Nimark (2017) generalize this model to allow for richer price of risk specifications as used in the empirical term structure literature. In their model, investors observe heterogeneous signals of the state variables driving bond yields. They forecast the forecasts of other investors and engage in speculative trading. In equilibrium, individual investors' prices of risk then reflect idiosyncratic signals, higher order expectations of the true state variables, as well as investor-specific expectations of maturity-specific shocks. Importantly, in their model the pricing factors follow stationary vector autoregressions under both the risk-neutral and the physical measure implying that investors do not disagree about short rates in the long-run, in contrast with our empirical evidence.

Our paper is also related to the term structure literature using survey information in the model estimation. For example, Kim and Wright (2005) and Piazzesi, Salomao, and Schneider (2015) use consensus survey forecasts to discipline the time-series dynamics under the physical measure. Wright (2011) shows that term premiums implied by a standard affine model and model-free term premiums implied by consensus survey expectations of future short rates have seen a secular decline across ten developed economies since the 1990s. He documents that survey-based measures of inflation disagreement show similar dynamics as these estimated term premiums. Giacoletti, Laursen, and Singleton (2021) build a dynamic term structure model in which a representative investor updates her beliefs about future bond yields. They find that when this updating is conditioned on the dispersion in bond yield forecasts, the model produces substantially smaller forecast errors. We provide a structural interpretation to their findings by explicitly relating term premium dynamics to investors' differences in beliefs and resulting relative wealth fluctuations. Finally, Crump, Eusepi, and Moench (2018) use the universe of surveys of professional forecasters to infer the consensus expected path of future Treasury bill rates. They show that although these short rate expectations show sizable variation, term premiums obtained as the simple difference between yields and expected short rates account for the bulk of yield variation at high and medium-term frequencies.

The remainder of this paper is structured as follows. Section 2 documents some novel facts about short rate disagreement and term premiums. In Section 3, we describe our affine term structure model with heterogeneous beliefs about the long-run level of rates. Section 4 presents the estimation results and model fit. In Section 5 we solve the investor portfolio problem and derive its implications for relative wealth dynamics and the aggregate term premium. Section 6 concludes.

2 Disagreement and Expected Returns: Stylized Facts

In this section, we motivate our subsequent analysis by providing some novel stylized facts on disagreement about future policy rates and expected Treasury returns. Our results are based on the Blue Chip Economic Indicators (BCEI) and the Blue Chip Financial Forecasts (BCFF) surveys. The Blue Chip (BC) surveys have been conducted monthly since the early 1980s. They ask two partly overlapping panels, of about 40 professional forecasters, to provide forecasts of the quarterly average of a variety of economic and financial variables. Since the mid 1980s, the surveys have also biannually been collecting forecasts from two years as far as 7-to-11 years ahead. While the BCFF survey publishes the individual forecasts for horizons up to six quarters into the future at a monthly frequency, they only report three quantities for the biannual forecasts of horizons of two years and above. These are the average across all forecasters, which we label the "consensus forecast," as well as the average of the top-10 and bottom-10 responses for a given forecasted variable at a given horizon. While the estimated term structure model with disagreement in Section 4 relies on these three forecast series for the federal funds rate from the BCFF survey, we set the stage in this section by providing information also on individual longer term three-month Treasury bill forecasts from the BCEI survey. To the best of our knowledge, no such individual longer-term forecast data has previously been studied in the literature.

Our analysis is motivated by Andrade, Crump, Eusepi, and Moench (2016) who show that the term structure of disagreement about future short rates is upward sloping. While forecasters agree to a large extent about monetary policy in the near-term, they have strongly opposing views about the medium and longer-term policy outlook, which at least partially reflects disagreement about the fundamentals of the economy. Here, we expand on these results by showing that fundamental or long-term disagreement about short term interest rates *i*) is not driven by outlier predictions; *ii*) is a persistent phenomenon in the sense that individual forecasters tend to see high or low future short rates across all forecast horizons; *iii*) implies sizable fundamental disagreement about bond risk premia; and *iv*) is strongly correlated with the term premium of the consensus forecaster and that implied by an affine term structure model.

Figure 1 shows the individual predictions for the three-month Treasury bill at horizons of two and

7-11 years into the future. The left-hand chart shows that while individual longer-term forecasts broadly move together there is a considerable degree of disagreement among forecasters already at the two-year horizion. They disagree by as much as six percentage points about the level of the three-month TBill. The strong disagreement is particularly pronounced just after the start of the large-scale asset purchase programs by the Federal Reserve in 2009, but drops considerably when calendar-based forward guidance was introduced in the summer of 2011 (Crump, Eusepi, and Moench (2013)). The chart also shows that the width of the forecast distribution as measured by the difference between the top-10 and bottom-10 average responses is wide and varies considerably over time.

The right-hand chart of Figure 1 provides the predictions of the three-month TBill of the same individuals in the long-run. As expected, there is less of a cyclical element in these forecasts. That said, the chart also shows that the entire distribution of long-run forecasts of the short rate has trended down over the sample. This strongly suggests that the long-run level of the short rate is perceived to vary over time. This feature of forecasters' beliefs will be a central element of our modeling strategy. Interestingly, while there clearly is a strong common element in the individual forecasts, the distribution at this very long horizon is also quite wide. This indicates that forecasters disagree to a considerable extent about the long-run value of the short rate. Quite strikingly, at the end of our sample some forecasters believe the long-run value of the TBill will remain below two percent while others see it go back to a level of around four percent. These heterogeneous assessments likely reveal sharply different views of the steady state of the economy.

As it is inherently difficult to predict far into the future, one might worry that individual forecasters' responses are to some extent arbitrary and do not necessarily reflect their views of the world. While we do not observe the names of individual forecasters in our sample of long-term predictions, we are able to trace their forecast paths at any given point in time. We can thus check whether the individual medium to long-run predictions are consistent in the sense that they reveal a particular forecaster expecting higher or lower future short rates. We rank the individual forecasts at all medium to long-run horizons and compute the rank correlation between two adjacent horizons. This measures the probability that a forecaster who expects low rates at, say, the four-year ahead horizon also expects low rates at the five-year horizon.

Figure 2 shows the rank correlations across forecasters and their 90 percent confidence interval for adjacent medium to long-term forecast horizons. At all horizons, these correlations are large and precisely estimated. The rank correlations are somewhat lower, around 70 percent, at medium-term horizons suggesting that individual forecasts are to some degree driven by different views about the state of the business cycle and the corresponding monetary stance at these horizons. That said, for longer forecast horizons the rank correlations increase further and reach almost 90 percent at the six year and 7-11 year ahead horizon. This implies that individual forecasts are highly consistent across horizons and likely reflect fundamentally different views about the economy.

The term premium is defined as the difference between the yield on a government bond and the average short rate expected to prevail over the life of the bond. Since we observe survey participants' individual forecast paths for the short rate, we can compute the term premiums consistent with their short-rate beliefs for various forward horizons. Figure 3 displays the evolution of forward term premiums implied by the individual TBill forecasts for the one-to-two year and 7-11 year forward horizons.

The figure shows that the assessment of the compensation that long-term bond investors command differs widely across forecasters. Moreover, especially in the latter part of the sample quite a few survey participants see term premiums in negative territory, possibly suggesting that they view longer term Treasuries as hedges against adverse states of the economy. While individuals' views about term premiums are quite heterogeneous, the top and bottom-10 average predictions appear to represent well the dispersion of beliefs across the forecaster distribution.

Figure 4 displays the one-two year and the 7-11 year forward term premium implied by the consensus forecaster, which we obtain as the difference between observed forward yields and the average short rate expected by the consensus forecast over the same maturity. We superimpose the difference between the top-10 and bottom-10 average forecasts of the short rate for the same horizon. The charts show that the two measures comove at both horizons. The model presented later in the paper rationalizes this comovement.

The positive correlation between a measure of the disagreement about future short rates and the term premium shown above is not restricted to the term premium implied by the consensus forecast. The upper panel of Figure 5 shows the time series of the two and ten-year Treasury term premium obtained from the Adrian, Crump, and Moench (2013) (ACM) model.[1] This no-arbitrage term structure model uses the first five principal components of Treasury yields as pricing factors and does not include survey forecasts in the estimation. We superimpose the difference between the top and bottom-10 average responses at the two-year horizon and the five and 7-11 year ahead horizons, respectively. The charts show a strong co-movement of the ACM term premium and short rate disagreement at both horizons. This becomes even more apparent when considering scatter plots of the same series in the bottom panel of Figure 5. Both survey-based and statistical term premiums are strongly correlated with measures of longer-term disagreement about the short rate.

3 A GATSM with Fundamental Disagreement

In this section, we introduce a Gaussian affine term structure model (GATSM) with two investors who hold different beliefs about the future level of rates. As is common in GATSM models, we assume that yields are affine functions of three pricing factors which follow a stationary vector autoregression under the pricing measure. The pricing factors are portfolios of yields and can be interpreted as level, slope, and curvature. We extend the standard affine model in two ways. First, we assume different dynamics of the pricing factors under the physical measure. Specifically, while

[1] See https://www.newyorkfed.org/research/data_indicators/term_premia.html.

the slope and curvature factors are stationary, we assume that the level has a time-varying long-run mean which itself follows a random walk. Second, we allow investors to have different beliefs about that long-run level of rates. This assumption is motivated by the evidence presented in the previous section. In the following, we introduce the individual pieces of this specification which will form the basis for our empirical analysis in Section 4.

3.1 The Model

We assume that there are two investors $i \in \{A, B\}$ who interact in a complete market setting, trading bonds at different maturities. Three factors determine the evolution of bond prices, which we label 'L', 'S' and 'C'. We assume that these factors evolve according to the following data generating process:

$$\begin{aligned} x_t^L &= \left(1-\beta^L\right)\mu_{t-1}^L + \beta^L x_{t-1}^L + \sigma_L \epsilon_t^L, \\ x_t^S &= \left(1-\beta^S\right)\mu^S + \beta^S x_{t-1}^S + \sigma_S \epsilon_t^S, \\ x_t^C &= \left(1-\beta^C\right)\mu^C + \beta^C x_{t-1}^C + \sigma_C \epsilon_t^C, \\ \mu_t^L &= \mu_{t-1}^L + \sigma_\mu \eta_t^L. \end{aligned}$$

We can write the system more compactly as

$$\mathbf{X}_t = (I-\beta)\,\mu_{t-1} + \beta \mathbf{X}_{t-1} + \Sigma_\epsilon^{1/2}\epsilon_t, \tag{1}$$

where $\mathbf{X}_t = \left(x_t^L, x_t^S, x_t^C\right)'$, $\mu_t = \left(\mu_t^L, \mu^S, \mu^C\right)'$, $\beta = \text{diag}\left(\beta^L, \beta^S, \beta^C\right)$, $\epsilon_t = \left(\epsilon_t^L, \epsilon_t^S, \epsilon_t^C\right)'$, and $\Sigma_\epsilon = \text{diag}\left(\sigma_L^2, \sigma_S^2, \sigma_C^2\right)$. The short rate is assumed to be linear in the pricing factors:

$$r_t = \delta_0 + \delta_1' \mathbf{X}_t. \tag{2}$$

Both investors perfectly observe the three pricing factors $\mathbf{X}_t$. They also know the data generating process and the parameters (β^j, σ_j) with $j \in \{L, S, C\}$ as well as $\{\mu^S, \mu^C, \sigma_\mu\}$. However, we assume that the two investors have different beliefs about the drift μ_t^L which determines the long-run mean of the level factor. Both investors believe μ_t^L evolves according to a random walk.[2] Their estimates of the drift evolve according to

$$\mu_t^i = \mu_{t-1}^i + \eta_t^i, \quad \text{for } i \in \{A, B\}. \tag{3}$$

This equation parsimoniously captures the updating of investors' beliefs about slow-moving trends and reflects investors' attempts to disentangle short-run macroeconomic developments from long-term changes in the economy. This difference in beliefs across the two investors could be driven,

[2] Bauer and Rudebusch (2020) introduce a similar model in a representative-agent framework.

for example, by idiosyncratic signals or some other form of informational friction.[3] Importantly we assume that the innovations η_t^i are correlated among each other and with the innovations η_t^L. This correlation structure embeds the assumption that both investors' private signals about the innovation to the drift are noisy measures of its true innovation. As such, their signals will also be correlated with each other.

We now turn to equilibrium bond prices. Absence of arbitrage implies that

$$P_t^{(n)} = E_t^i \left(M_{t+1}^i P_{t+1}^{(n-1)} \right), \quad \text{for } i \in \{A, B\}.$$

where $P_t^{(n)}$ is the price of an n-period bond in time t and M_t^i denotes the stochastic discount factor of investor i. As is common in the affine term structure literature we assume an exponentially affine functional form:

$$M_{t+1}^i = \exp\left(-r_t - \frac{1}{2}\lambda_t^{i\prime}\Sigma_\epsilon\lambda_t^i - \lambda_t^{i\prime}\Sigma_\epsilon^{1/2}\epsilon_{t+1} \right), \tag{4}$$

where ϵ_t are the innovations to the pricing factors and λ_t^i denotes the investor-specific vector of market prices of risk associated with these pricing factors. Following Duffee (2002) and many others, we assume that prices of risk are linear in the pricing factors:

$$\lambda_t^i = \lambda_{0,t}^i + \Lambda_1^i \mathbf{X}_t, \tag{5}$$

where $\lambda_{0,t}^i$ is a 3×1 vector and Λ_1^i a 3×3 matrix. We assume a time-varying intercept in the market prices of risk to embed the notion that investors' perceived risk-return tradeoff can fluctuate with their assessment of the long-run mean of the risk factors.

Both investors trade at equilibrium bond prices and conjecture that these are exponentially affine in the observed pricing factors:[4]

$$\ln P_t^{(n)} = \mathcal{A}_n + \mathcal{B}_n' \mathbf{X}_t$$

We can then use this in investor i's pricing condition

$$\begin{aligned} P_t^{(n)} &= E_t^i \left(M_{t+1}^i P_{t+1}^{(n-1)} \right) \\ &= E_t^i \left(M_{t+1}^i \exp\left(\mathcal{A}_{n-1} + \mathcal{B}_{n-1}' \mathbf{X}_{t+1} \right) \right) \quad \text{for } i = \{A, B\}. \end{aligned}$$

[3]For example, Andrade, Crump, Eusepi, and Moench (2016) show that the observed term structures of disagreement can be rationalized in a model with informational frictions where the state variables follow a VAR with slow-moving long-run means which investors filter from the imperfectly observed data.

[4]In line with the existing affine term structure literature, we only consider a solution with a constant intercept under the pricing measure. While alternative solutions with a time-varying interecept are conceivable, specific parametric assumptions about the intercept would be required to ensure a closed-form system of Ricatti equations as the one we derive below.

This yields the typical Ricatti equations:[5]

$$\mathcal{A}_n = \mathcal{A}_{n-1} + \mathcal{B}'_{n-1}\left(\left(I-\beta\right)\mu_t^i - \Sigma_\epsilon \lambda_{0,t}^i\right) + \frac{1}{2}\mathcal{B}'_{n-1}\Sigma_\epsilon \mathcal{B}_{n-1} - \delta_0 \tag{6}$$

$$\mathcal{B}'_n = \mathcal{B}'_{n-1}\left(\beta - \Sigma_\epsilon \Lambda_1^i\right) - \delta'_1 \quad \text{for } i \in \{A, B\}. \tag{7}$$

An equilibrium where these pricing equations hold and where both investors trade therefore requires that

$$\left(I-\beta\right)\mu_t^i - \Sigma_\epsilon \lambda_{0,t}^i = \Psi_0 \quad i \in \{A, B\} \quad \forall t \in \{1, \ldots, T\} \tag{8}$$

$$\text{and } \beta - \Sigma_\epsilon \Lambda_1^i = \Psi_1 \quad i \in \{A, B\}. \tag{9}$$

Hence we can express

$$\ln P_t^{(n)} = \mathcal{A}_n\left(\Psi_0\right) + \mathcal{B}'_n\left(\Psi_1\right)\mathbf{X}_t \tag{10}$$

where Ψ_0 is a 3×1 vector and Ψ_1 a 3×3 matrix denoting the intercept and autoregressive coefficient of $\mathbf{X}_t$ under the pricing measure.

The restrictions in equation (8) imply that

$$\sigma_L^2 \lambda_{0,t}^{i,L} = \left(1-\beta^L\right)\mu_t^i - \Psi_0^L \quad \text{for } i \in \{A, B\} \quad \forall t \in \{1, \ldots, T\},$$

where $\lambda_{0,t}^{i,L}$ is the first element of the vector $\lambda_{0,t}^i$ and Ψ_0^L is the first element of Ψ_0. Hence, it follows that

$$\mu_t^A - \mu_t^B = \frac{\sigma_L^2}{1-\beta^L}\left(\lambda_{0,t}^{A,L} - \lambda_{0,t}^{B,L}\right) \quad \forall t \in \{1, \ldots, T\}. \tag{11}$$

Conversely, the restrictions in equation (9) imply that

$$\begin{aligned} \beta - \Sigma_\epsilon \Lambda_1^A &= \beta - \Sigma_\epsilon \Lambda_1^B = \Psi_1 \\ \text{or} \quad \Lambda_1^A &= \Lambda_1^B = \Lambda_1 \end{aligned} \tag{12}$$

Combined, we therefore obtain

$$\frac{\sigma_L^2}{1-\beta^L}\left(\lambda_t^{A,L} - \lambda_t^{B,L}\right) = \mu_t^A - \mu_t^B. \tag{13}$$

In other words, any disagreement about the long-run mean of the level factor translates one-for-one into a difference in risk attitudes. This is intuitive. Assume investor A anticipates a higher

[5]Note that to simplify the analysis we abstract from individual investors' uncertainty about the estimate μ_t. Here we assume pricing occurs under anticipated utility as in Kreps (1998): investors take their current estimate as if they knew it with certainty.

long-run level of rates than investor B. This implies that A expects the future level of short rates to be higher than B which, in turn, makes longer-term bonds a relatively less attractive investment for A than for B and thus requires a higher expected compensation for bearing level risk. Expected one-month excess holding period returns across maturities of each investor are given by

$$E_t^i\left[rx_{t+1}\right] = \mathcal{B}_X \lambda_t^i - \frac{1}{2}\Sigma, \tag{14}$$

where $\mathcal{B}_X$ is the matrix of corresponding loadings of the log price of a bond with maturity n on the pricing factors $\mathbf{X}_t$ derived above and Σ is the model-implied conditional variance-covariance matrix of excess returns given by

$$\Sigma = \mathcal{B}_X \Sigma_\epsilon \mathcal{B}_X'.$$

Inspecting Equation (14), a higher estimate μ_t^i results in a higher value of the market price of risk λ_t^i. Since the bond pricing parameters $\mathcal{B}_X$ are negatively related to the yield loadings on the pricing factors, an increase in λ_t^i implies lower expected returns.

3.2 The State-Space Model

With these ingredients we now write our model in state-space form. In addition to data on observed bond yields, $\mathbf{y}_t^o$, we use observations of short rate forecasts of two investors, labeled $\mathbf{y}_t^{E,A}$and $\mathbf{y}_t^{E,B}$ to identify the parameters of the model. The observation equation is given by:

$$\begin{bmatrix} \mathbf{y}_t^o \\ \mathbf{y}_t^{E,A} \\ \mathbf{y}_t^{E,B} \end{bmatrix} = \begin{bmatrix} \mathbf{A}^X \\ \mathbf{H}_0^A(\beta,\mu) \\ \mathbf{H}_0^B(\beta,\mu) \end{bmatrix} + \begin{bmatrix} \mathbf{B}^X & 0 & 0 & 0 \\ \mathbf{H}_X^A(\beta) & \mathbf{H}_\mu^A(\beta) & 0 & 0 \\ \mathbf{H}_X^B(\beta) & 0 & \mathbf{H}_\mu^B(\beta) & 0 \end{bmatrix} \times \begin{bmatrix} \mathbf{X}_t \\ \mu_t^A \\ \mu_t^B \\ \mu_t^L \end{bmatrix} + e_t, \tag{15}$$

where the first few entries of e_t are yield pricing errors $e_t^{(n)}$ and e_t has a diagonal variance-covariance matrix. Moreover, $\mathbf{H}_0^A(\beta,\mu)$, $\mathbf{H}_0^B(\beta,\mu)$, $\mathbf{H}_X^A(\beta)$, $\mathbf{H}_\mu^A(\beta)$, $\mathbf{H}_X^B(\beta)$, $\mathbf{H}_\mu^B(\beta)$ determine the model-implied loadings of future short rate expectations of investors A and B on the observed model factors $\mathbf{X}_t$ as well as their respective updates of the drift of the level factor, μ_t^A and μ_t^B. These loadings are known, nonlinear transformations of the model parameters, see also Crump, Eusepi, and Moench (2018).

Combining (1) and (3), the transition equation is given by

$$\begin{bmatrix} \mathbf{X}_t \\ \mu_t^A \\ \mu_t^B \\ \mu_t^L \end{bmatrix} = \begin{bmatrix} \boldsymbol{\alpha}^X \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \beta & 0 & 0 & l \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} \mathbf{X}_{t-1} \\ \mu_{t-1}^A \\ \mu_{t-1}^B \\ \mu_{t-1}^L \end{bmatrix} + u_t. \tag{16}$$

where $\boldsymbol{\alpha}^X = (0, \left(1-\beta^S\right)\mu^S, \left(1-\beta^C\right)\mu^C)'$ and $l = (1-\beta^L, 0, 0)$ and where

$$u_t = \begin{pmatrix} \Sigma_\epsilon^{1/2} & \Gamma \\ 0 & \Sigma_\eta^{1/2} \end{pmatrix} \begin{pmatrix} \epsilon_t \\ \eta_t \end{pmatrix}, \qquad \Gamma = \begin{pmatrix} 0 & 0 & 1-\beta^L \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

with $u_t = (\epsilon_t^L, \epsilon_t^S, \epsilon_t^C, \eta_t^A, \eta_t^B, \eta_t^L)'$. Hence, the innovations of different beliefs about the level factor can be correlated with each other, while the innovations of the pricing factors X_t are uncorrelated with the innovations of these beliefs.

The vector of parameters to be estimated is given by

$$\Theta = \begin{pmatrix} \beta^L & \beta^S & \beta^C & \mu^S & \mu^C & \operatorname{diag}(\Sigma_\epsilon)' & \operatorname{vech}(\Sigma_\eta)' & \boldsymbol{\Psi}_0' & \operatorname{diag}(\boldsymbol{\Psi}_1)' & \sigma_y & \sigma_{short}' & \sigma_{long}' \end{pmatrix}$$

where Σ_η is a variance-covariance matrix capturing the correlation between innovations to the drift and investors' beliefs, and the last three elements of Θ are the variances of the *i.i.d.* measurement error on yields as well as short and long-horizon survey forecasts of the short rate. We estimate the parameters via maximum likelihood and filter the perceived drifts of the level factor μ_t^A and μ_t^B, and as a result, estimates of the prices of risk λ_t^A and λ_t^B as per Equation (13). Note that while the level factor is fully spanned by the cross-section of bond yields, its time-varying long-run mean μ_t as perceived by different investors is not. At the same time, the different beliefs about the long-run mean bear information about future expected short rates and expected bond returns. Hence, we can interpret the long-run means perceived by the two investors as "hidden" or "unspanned" factors in the spirit of Duffee (2011b) and Joslin, Priebsch, and Singleton (2014). As disagreement about future short rates in our model is fully captured by the difference in beliefs about the long-run mean of the level factor, disagreement is itself unspanned in our model.

4 Bringing the Model to the Data

In this section, we bring our two-investor GATSM with disagreement to the data. We first describe the data used and the estimation approach. We then discuss the model fit and provide decompositions of Treasury yields into expected short rates and term premiums. In Section 5, we then use the model to measure the aggregate term premium.

4.1 Data and Estimation

We jointly estimate our model using zero-coupon Treasury yields as well as survey expectations of short rates for two different groups of investors. We obtain the latter from the *Blue Chip Financial Forecasts* (BCFF) survey. Specifically, we use two- and four-quarter-ahead, one-to-two, four-to-five year-ahead, and long-term forecasts which cover horizons between six and ten or seven and eleven years into the future, depending on when the survey was taken. The short-term forecasts are observed monthly and this is our frequency of observation also for Treasury yields. The medium-term and long-term forecasts are observed biannually. The missing monthly observations in between biannual survey observations can easily be accommodated in our state-space framework. The BCFF provides medium and long-term forecasts for three different cross-sectional averages of the forecaster distribution: the average across all responses (the "consensus" forecast), the average of the top-10 responses and the average of the bottom-10 responses. We employ the top-10 and bottom-10 average responses as proxies representing two investors at the opposite spectrum of the belief distribution about future short rates. As discussed in Section 2, the difference between the top-10 and bottom-10 average responses is closely correlated with common measures of forecaster disagreement, such as the cross-sectional standard deviation or the interquartile range.

We obtain zero coupon Treasury yields from Gürkaynak, Sack, and Wright (2007) (GSW henceforth).[6] The GSW zero coupon yields are based on fitted Nelson-Siegel-Svensson curves, the parameters of which are published along with the estimated zero coupon curve. We use these parameters to back out the cross-section of zero-coupon yields for maturities up to ten years, using end-of-month values. In our estimation, we use $N = 8$ Treasuries with maturities $n = 3, 6, 12, 24, 36, 60, 84, 120$ months. Our sample period is $1983:03 - 2015:08$ for a total of 390 monthly observations. We employ the normalization scheme in Joslin, Singleton, and Zhu (2011) to estimate our model and impose additional zero restrictions in the physical dynamics to ensure a parsimonious model (see Appendix for details).

4.2 Model Fit and Individual Expected Excess Returns

Our model fits both yields and survey forecasts of the short rate very precisely, as displayed in Figure 6 which shows the time series and cross section of observed and model-implied yields. All parameter estimates and associated standard errors are provided in the Appendix. The average of yield pricing errors is no more than 5 basis points in absolute value and is thus well in line with previous studies. The bottom two panels of Figure 6 displays the unconditional mean and standard deviation of yields across maturities as observed and fitted by the model. The charts show that the model fits both moments well.

We next turn to the model fit of survey forecasts of the short rate. The top two charts in Figure 7 show the observed and fitted top-10 and bottom-10 average survey forecasts of the federal funds rate, where actual values are plotted by solid lines. These two charts document that with only the perceived long-run mean of the level factor being different across investors, our model is able to capture the substantial time variation in investors' disagreement about future short rates. The bottom four panels of Figure 7 provide a plot of unconditional first and second moments of the two groups' survey forecasts as observed and fitted by the model, again documenting that the model

[6] See http://www.federalreserve.gov/econresdata/researchdata.htm

fits survey forecasts at all horizons quite precisely.

We use the model to study the entire forecast path of individual short rate expectations. Figure 8 displays the evolution of the short rate along with the expected short rate paths ten years into the future. The model generates these paths for every month in the sample, but for ease of exposition we only show them at five year intervals.

The following observations are worth highlighting. First, for both investors, the longer-run expected short rate has steadily declined over time. At the same time, the short rate disagreement at the ten year horizon has declined from more than four percentage points in the mid 1980s to around one percentage point at the end of the sample. Disagreement about the long term also results in very different expected paths in the short and medium term. For example, in the mid 1980s and mid 2000s the top-10 investor saw a steepening path of the short rate while the bottom-10 expected a declining path throughout the forecasting horizon. While both investors were too optimistic about policy normalization in the zero lower bound period after the Great Financial Crisis, the bottom-10 investor consistently expected a flatter short rate path. For example, during the period 2011-2013 when the Federal Reserve was heavily relying on forward guidance to steer short-term rate expectations at the zero lower bound, near-term rate expectations of the bottom-10 investor were compressed to zero. In contrast, the top-10 investor saw short rates rise above zero even in the short term.

Having discussed the evolution of expected short rates, Figure 9 shows the corresponding term premiums which – as discussed in Section 3 – are equivalent to expected excess one-month returns over the life of the bond. As shown in the upper panel, the top-10 investor expecting high future short rates implicitly expected returns on two-year notes to hover between minus one and one percent over the last 30 years. Similarly, this investor expects ten-year Treasury term premiums to fluctuate around zero in a somewhat wider range. In stark contrast, the bottom-10 investor expecting low future policy rates implicitly saw positive expected excess returns across time and maturity, declining from about three (six) percent at the two-year (ten-year) maturity in the early 1980s to just below one percent for both maturities at the end of the sample.

The bottom-left panel of Figure 9 provides the time series average of the implied term premiums across maturities for the two investors. While the bottom-10 investor has an upward sloping term structure of term premiums ranging between 50 basis points at the one-year maturity and 200 basis points at the ten-year maturity, the top-10 investor essentially sees term premiums on average slightly negative across all except the very long maturities. In other words, to this investor, Treasuries provide insurance for which she is willing to pay a premium. The bottom-right chart presents the differences between the two term premium estimates across time and maturities, showing that these differences have been much less pronounced since around the year 2000.

5 Portfolio Allocation and Aggregate Term Premium

We have documented the evolution of individual short rate expectations and their implications for subjective term premiums. In this section we attempt at measuring an aggregate term premium of this heterogeneous-investor economy. To do so, we have to evaluate how important these two investors have been over time. We proceed in two steps. First, we evaluate the portfolio performance (and therefore the hypothetical wealth) of the two investors over the sample. Second, we construct an aggregate term premium by appropriately weighting the individual term premiums using investors' wealth.

5.1 Portfolio Allocation and Investment Performance

What do the different short rate forecasts imply for investors' bond portfolio allocation? Intuitively, the investor predicting higher future short rates would go long in short-term maturities while the investor predicting lower future short rates should tilt her portfolio towards longer maturities. Given the secular decline of interest rates over the sample period, one would thus expect the investor with a larger portfolio weight on longer maturities to accumulate higher returns. While our reduced-form approach does not allow to derive portfolio allocations from explicit micro foundations, we can use the model-implied expected excess returns to compute hypothetical mean-variance portfolio weights.

Specifically, we assume that each investor constructs a portfolio by choosing the weights that minimize a quadratic utility function subject to the constraint that they add up to one. We follow Carriero and Giacomini (2011) and obtain the optimal weights by solving in every period

$$w_{i,t}^{*} = \arg\min_{w} \left\{ w' E_t \left[rx_{t+1}^{i} \right] - \frac{\gamma}{2} w' \Sigma w \right\}. \tag{17}$$

Recall that the expected excess return of investor i is

$$E_t \left[rx_{t+1}^{i} \right] = \mathcal{B}_X \lambda_t^{i} - \frac{1}{2}\Sigma,$$

As shown in Carriero and Giacomini (2011), the solution to this problem is

$$\begin{aligned} w_{i,t}^{*} &= \phi + \Phi E_t \left[rx_{t+1}^{i} \right] \\ \text{where } \phi &= \frac{\Sigma^{-1}\iota}{\iota'\Sigma^{-1}\iota} \text{ and } \Phi = \frac{1}{\gamma}\left(\Sigma^{-1} - \frac{\Sigma^{-1}\iota\iota'\Sigma^{-1}}{\iota'\Sigma^{-1}\iota} \right) \end{aligned}$$

The two investors's expected returns are a key input to the optimal portfolio weights in Equation 18. The top panel of Figure 10 shows the expected excess one-month return for a ten-year Treasury note. The top-10 investor's expected return (solid blue line) has been slightly negative throughout the sample. In contrast, the bottom-10 investor (dashed green line) has been expecting positive but declining returns over time. This pattern of return expectations is the mirror image of the short rate expectations of the two investors. While the difference in expected returns was sizable at the beginning of the sample in the mid 1980s, it has been shrinking over the later part of the sample. Of note, even the bottom-10 investor who has predicted lower future short rates implicitly expected negative returns since the mid 2000s.

The bottom panel of Figure 10 displays the cumulative realized returns from both investors' portfolios, where we obtain the realized returns as $\hat{w}^{*\prime}_{i,t} \times \left(rx^i_{t+1} + r_t\right)$. Not surprisingly, the bottom-10 investor strongly outperforms the top-10 investor over the sample period. The reason is that the bottom-10 investor is tilting her portfolio towards longer maturities, in line with the higher expected returns. Given the observed secular decline of rates over the sample, the price of longer-term bonds appreciated substantially, delivering higher returns compared to short-term maturities. This can be seen by the additional two lines in the figure which display the cumulative returns from rolling over a three-month Tbill (solid gray) and a ten-year Treasury (dashed gray), respectively. Given his portfolio allocation tilted towards shorter maturities, the top-10 investor's cumulative return is comparable to that of the Tbill. Conversely, the bottom-10 investor even outperforms the ten-year Treasury as she combines the holdings of longer maturities with a short position in short-term maturities.

5.2 Implications for the Term Premium

In structural asset pricing models with heterogeneous agents such as Xiong and Yan (2010) the evolution of relative wealth is an important determinant of asset prices. Commonly, in these models the investor with a higher share of the total wealth has a greater impact on the marginal valuation of risky assets. Specifically, a central result in heterogeneous asset pricing models with general preferences is that the discount factor and beliefs of the representative investor can be represented as weighted averages of the discount factors and beliefs of the individual agents (see, e.g., Jouini and Napp (2007) and Bhamra and Uppal (2014)). Applied to a bond-pricing model Xiong and Yan (2010) show that the bond prices in a representative economy equal a wealth-weighted average of individual investors' bond prices. Interestingly, while in their model individual term premiums are constant due to specific assumptions about preferences, the aggregate term premium fluctuates over time because of changes in relative wealth and disagreement.

Inspired by this result, we provide a measure of the aggregate term premium that would arise if the Treasury market was only populated by the top-10 and bottom-10 investor. We begin by approaching this problem in a more agnostic way by utilizing a nonparametric approach to estimating the weights. The nonparametric approach relies on two assumptions. The first, as discussed above, is that we can write the expected returns for the marginal investor in the economy as a weighted average of the top 10 and bottom-10 investor. Second, we assume that these weights vary smoothly over time in a way that is made precise in the Appendix. With these two assumptions, we can then use realized returns over our sample, to estimate the corresponding weights on each of our investors.

Figure 11 plots the estimated weight of the bottom-10 investor in the economy-wide term premium. In particular, the red line is the point estimate, with 68% confidence intervals represented by the black lines, which corresponds to an initial weight of 20%. Over the course of the sample, this weight rises monotonically, ending the sample a bit above 90%. The blue line in Figure 11 represents the wealth share derived from the cumulative returns for each investor shown in Figure 10. As initial condition, we use the 20% initial weight obtained from the nonparametric approach. Both the nonparametric estimated weights and the wealth share for the bottom-10 investor follow broadly similar dynamics, rising steadily across our sample. Consequently, we proceed by estimating the economy-wide term premium as the wealth-weighted term premium of each of our individual investors.

The bottom panel of Figure 11 shows the wealth-weighted term premium obtained using a range of initial wealth shares. These initial wealth shares correspond to the two standard deviation interval around the 20% mean initial wealth weight from the non-parametric approach. For comparison, we superimpose a term premium estimate obtained as the difference between the 10-year Treasury yield and a consensus survey-based forecast of the average short rate over the next ten years from Crump, Eusepi, and Moench (2018). A key difference is that this "consensus term premium" is substantially higher than the wealth-weighted term premium until the mid 1990s. This is driven by two related effects. First, the top-10 investor has substantially higher wealth weights at the beginning of the sample which reduces the wealth-weighted term premium as this investor expected lower future excess returns. Second, disagreement about future short rates is high before the late 1990s which magnifies the effect of the differential wealth weights on the term premium. Conversely, later in the sample the bottom-10 investor increases her relative wealth share and contemporaneously the expected paths of future short rates converge. As a result, the wealth-weighted term premium appears stationary over the sample, in contrast to the pronounced downward trend of the consensus term premium at the beginning of the sample. As the wealth-weighted term premium reflects the required compensation of the marginal investor for bearing interest rate risk, the differential behavior of the two term premium measures at the beginning of the sample implies a large gap between the beliefs of the marginal and the average investor. Our results thus show that the distribution of short-rate forecasts is highly informative about the pricing of risk.

6 Conclusion

Bond investors disagree about the future path of policy rates, and particularly about their long-run level. Accordingly, they disagree about the risk-return tradeoff of longer-term bonds and engage in speculative trading. This induces shifts in their relative wealth which, in turn, affects the marginal pricing of risk in the economy. Hence, the term premium of an econometrician observing only yields partly reflects disagreement-driven changes in the marginal pricing of risk.

In this paper, we have formalized this intuition in an affine term structure model with heterogeneous beliefs. In our model investors perfectly observe the level, slope, and curvature of the yield curve but have different beliefs about the long-run level of rates. Our model fits yields and survey forecasts of future short rates very well. It implies that investors who projected a faster decline in the short rate expected positive and sizable excess returns across maturities, while investors who anticipated a more modest decline expected near zero excess bond returns. With realized yields falling since the early 1980s, this implies that the first group of investors would have consistently accumulated more wealth. As a result, we show that the wealth-weighted term premium displays essentially no trend, in sharp contrast to term premiums based on a representative investor. Therefore, one broad conclusion we reach is that the measured secular decline in term premiums obtained from representative models can be sourced to the group of investors that correctly forecast a downward trend in interest rates and therefore expected positive expected returns from holding longer maturity bonds.

Figure 1: Disagreement about Short Rates at Medium and Long Horizons

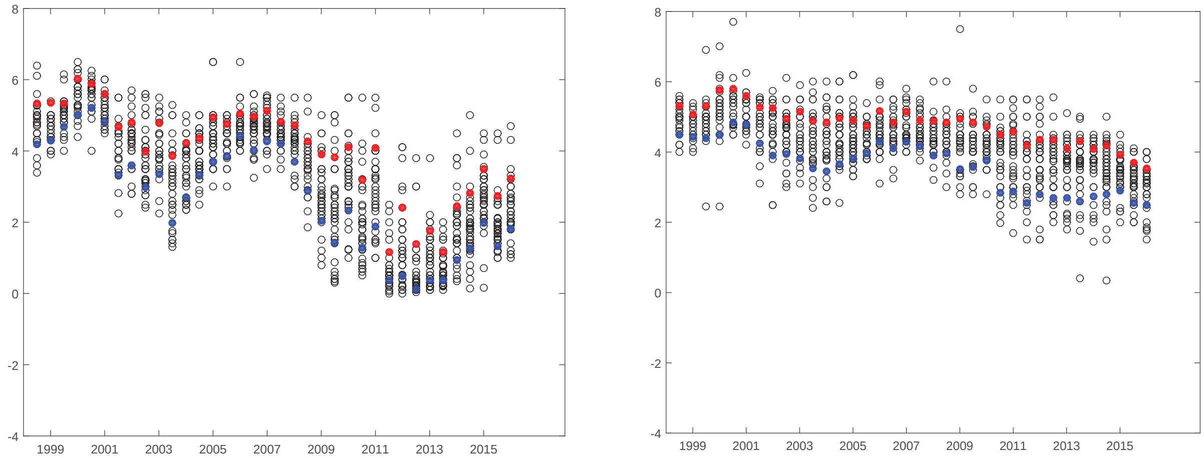

Notes:
This figure plots individual forecasts from the Blue Chip Economic Indicators survey for the three-month Treasury Bill at forecast horizons of two (left chart) and 7-11 years (right chart) into the future. The red and blue dots represent the top and bottom-10 average responses, respectively. The sample is from 1999-2016.

Figure 2: Consistency in Individual Beliefs Across Horizons

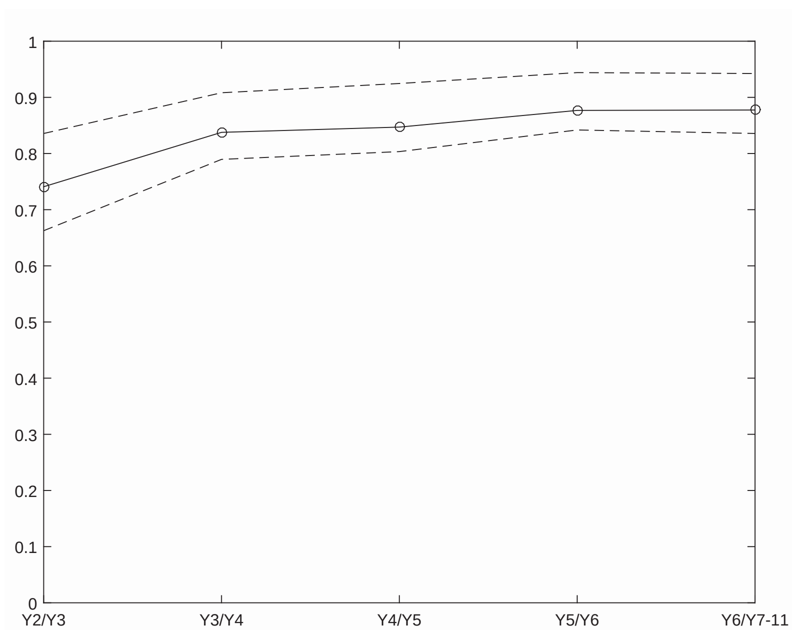

Notes:
This figure plots the rank correlation among individual forecasts from the Blue Chip Economic Indicators (BCEI) survey for the three-month Treasury Bill at adjacent forecast horizons between two and 7-11 years into the future. The dashed lines provide the 5th and 95th pointwise confidence intervals. The sample is from 1999-2016.

Figure 3: Disagreement About Term Premiums at Medium and Long Horizons

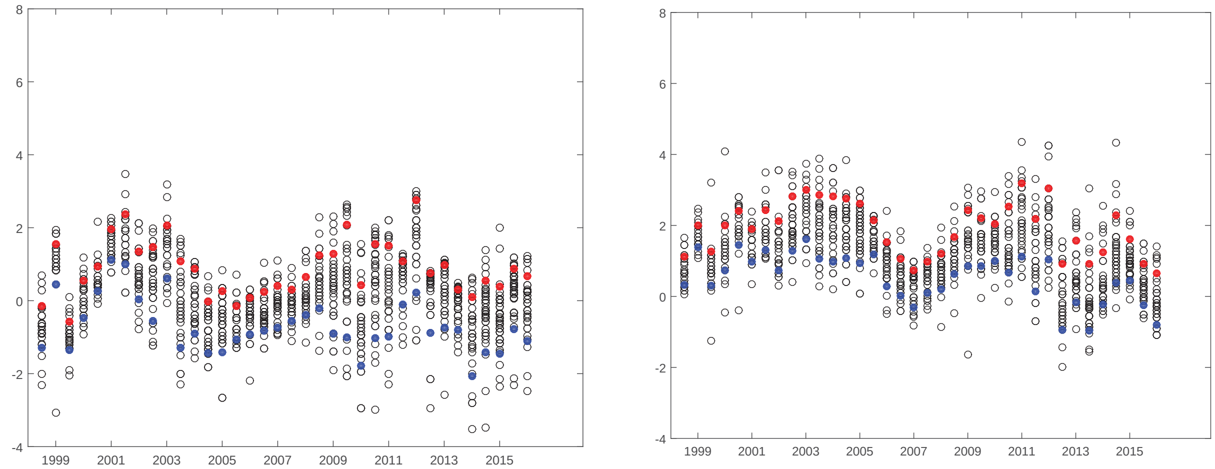

Notes:
This figure plots forward term premiums implied by individual forecasts of the three-month Treasury Bill from the Blue Chip Economic Indicators (BCEI) survey at forecast horizons 1-2 and 7-11 years into the future. The red and blue dots represent the top and bottom-10 average responses, respectively. The sample is from 1999-2016.

Figure 4: Consensus Term Premium and Disagreement About Short Rates

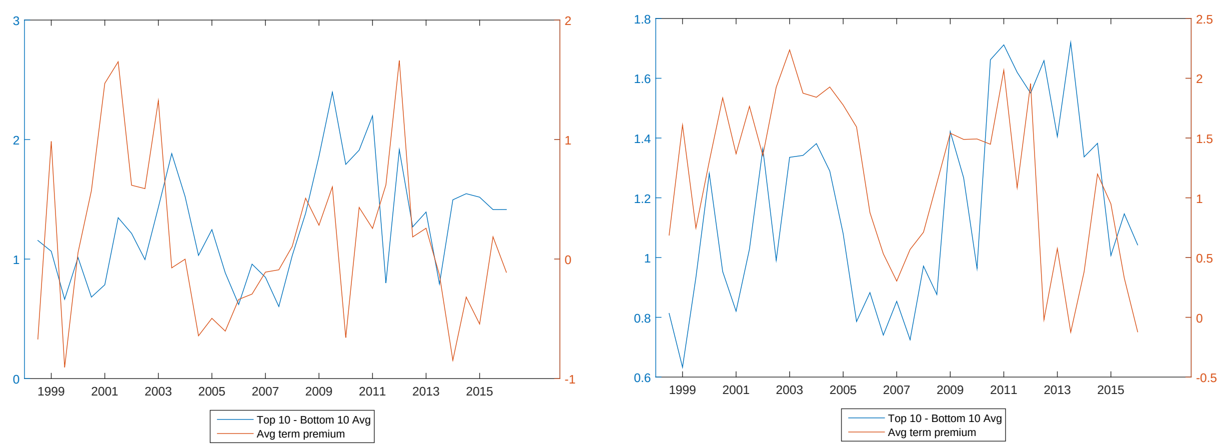

Notes:
This figure plots forward term premiums implied by the consensus forecast along with the difference between the top and bottom-10 average forecasts for the three-month Treasury Bill from the Blue Chip Economic Indicators (BCEI) survey at forecast horizons 1-2 and 7-11 years into the future. The sample is from 1999-2016.

Figure 5: Disagreement About Short Rates and Term Premiums

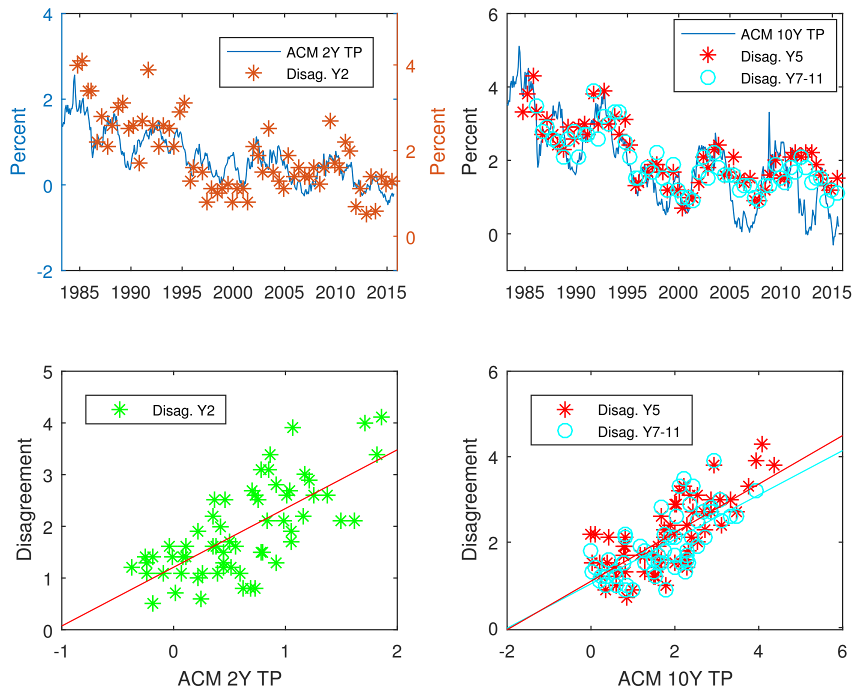

Notes:
This figure plots disagreement measures calculated using survey forecasts and ACM term premiums obtained from the model described in Adrian, Crump, and Moench (2013). The upper two charts display the difference between top-10 and bottom-10 average forecasts of the federal funds rate obtained from the Blue Chip Financial Forecasts (BCFF) survey. One- to four-quarter ahead, five-year ahead and long-horizon (7-11 years) survey forecasts are used. The lower two charts compare two-year and ten-year ACM term premiums with disagreement over similar horizons. Asterisks and circles in the charts are long-horizon forecasts from surveys conducted biannually.

Figure 6: Time-series and Sross-sectional Fit of Yields

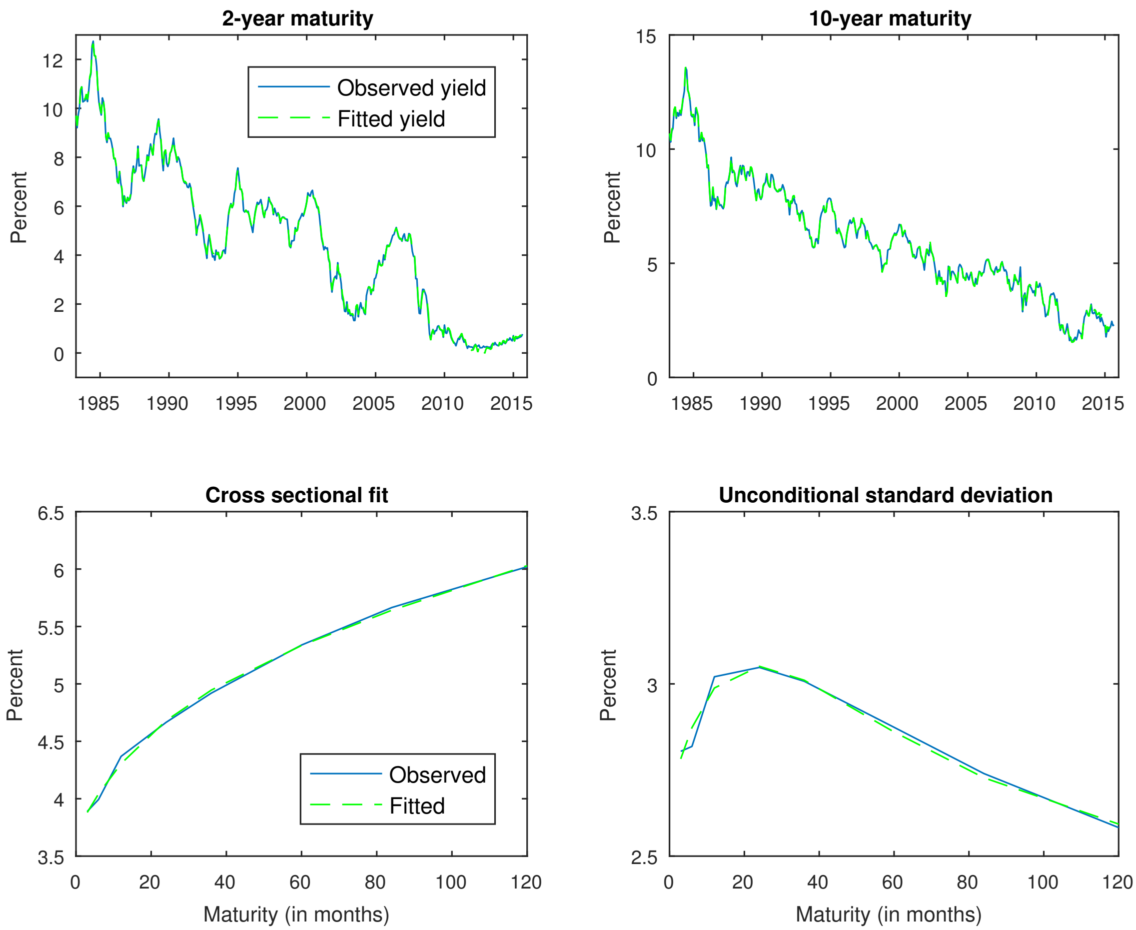

Notes:
This figure provides plots of observed and model-implied yields for the two- and ten-year maturities in the upper two charts. Observed yields are displayed by solid lines, dashed lines correspond to model-implied yields. The bottom panels plot unconditional averages and standard deviations of observed yields against those implied by the model.

Figure 7: Time-series and Cross-sectional Fit of Survey Forecasts

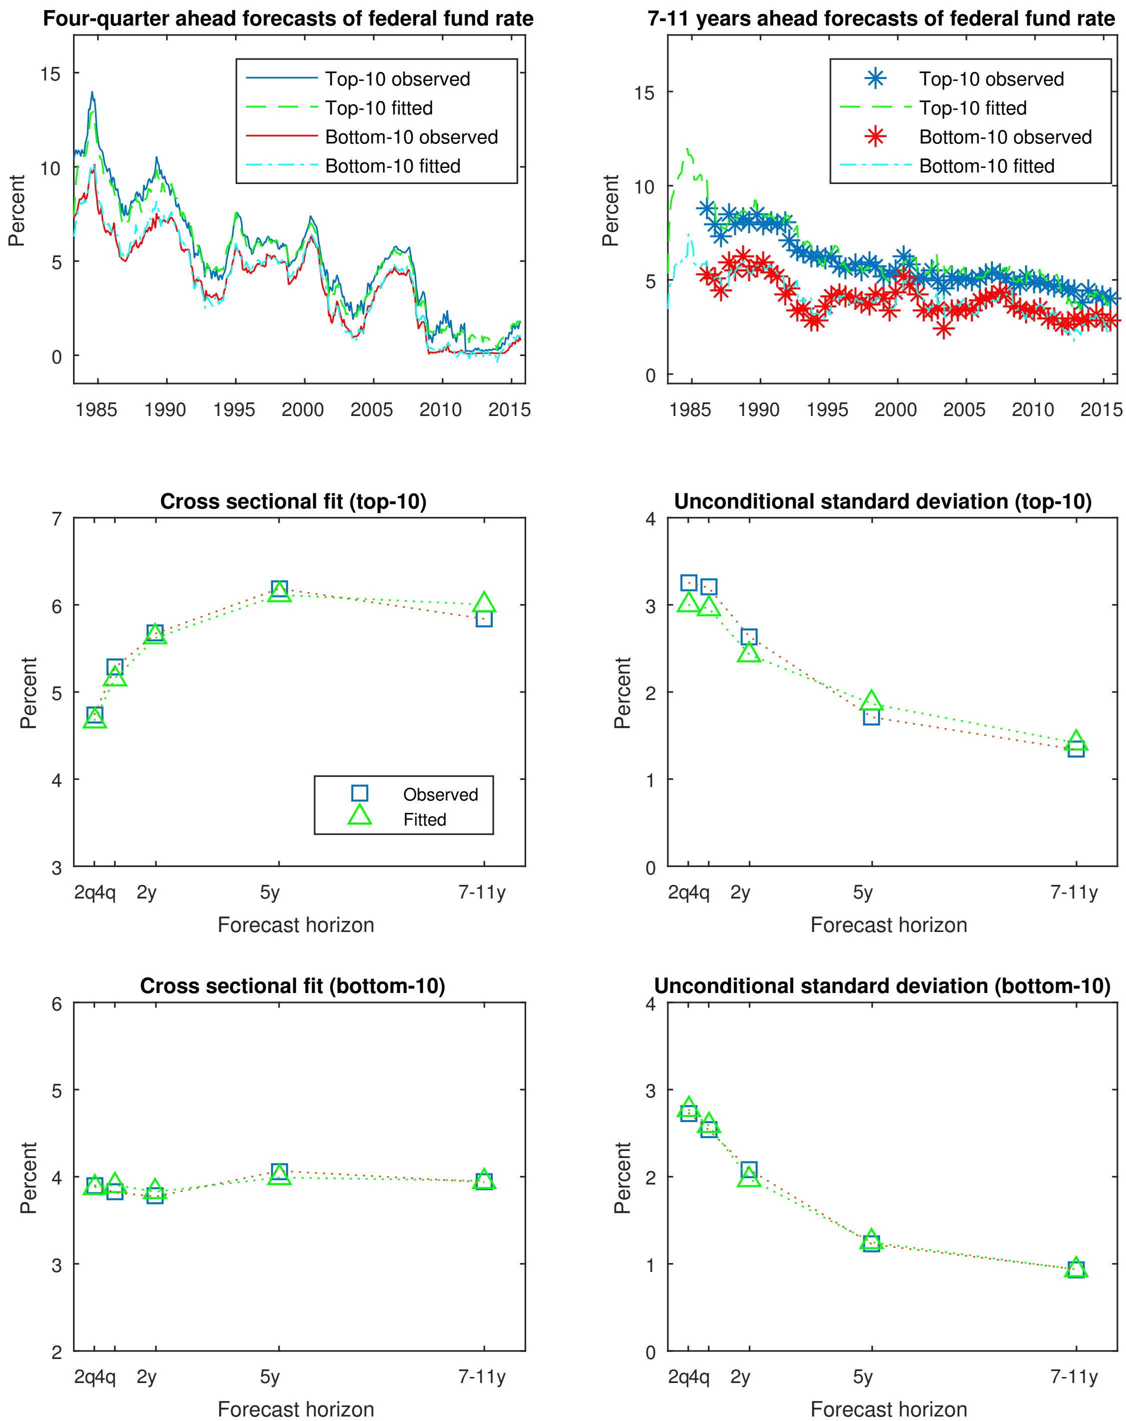

Notes:

This figure provides plots of observed and model-implied survey forecasts of the fed funds rate. Observed survey forecasts are displayed as solid lines, dashed lines correspond to model-implied survey forecasts. The top two charts show the Blue Chip Financial Forecasts (BCFF) four-quarter ahead and 7-11 years ahead top-10 and bottom-10 average forecasts of the federal funds rate. Asterisks in the top right chart are long-term forecasts which are observed biannually. The bottom four panels plot unconditional means and standard deviations of survey forecasts of the top-10 and bottom-10 average responses against those implied by the model.

Figure 8: Expected Short Rate Paths

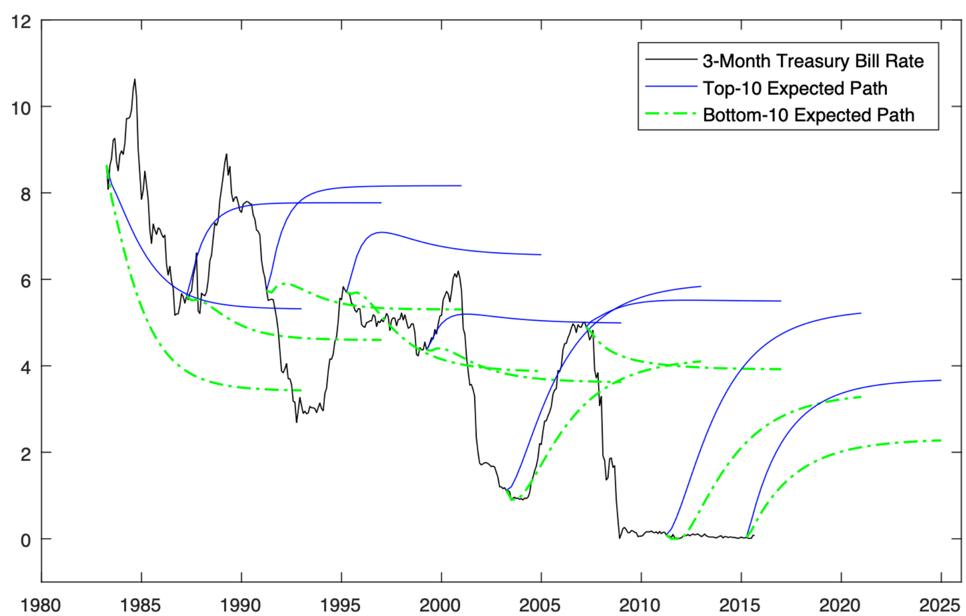

Notes:

This figure plots the realized three-month Treasury Bill along with the expected Treasury Bill paths for the top-10 and bottom-10 investors.

Figure 9: Short Rate Expectations and Term premiums

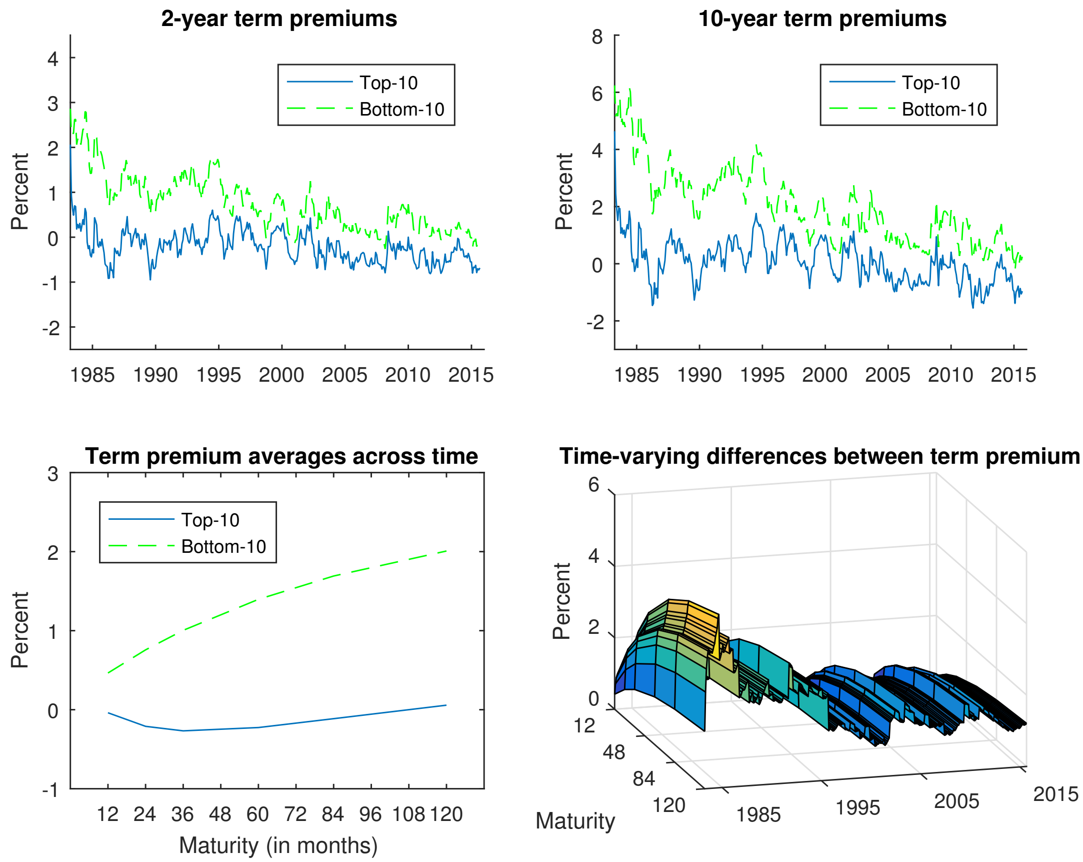

Notes:
This figure plots the term premiums implied by the top-10 and bottom-10 investors' beliefs about future short rates. The upper panels plot the the term premium estimates for two- and ten-year treasury notes. The lower left panel plots the sample averages of term premium estimates in the investors' beliefs for different maturities. The lower right panel displays time-varying differences across maturities between the two investors' beliefs.

Figure 10: Expected Returns and Investment Performance

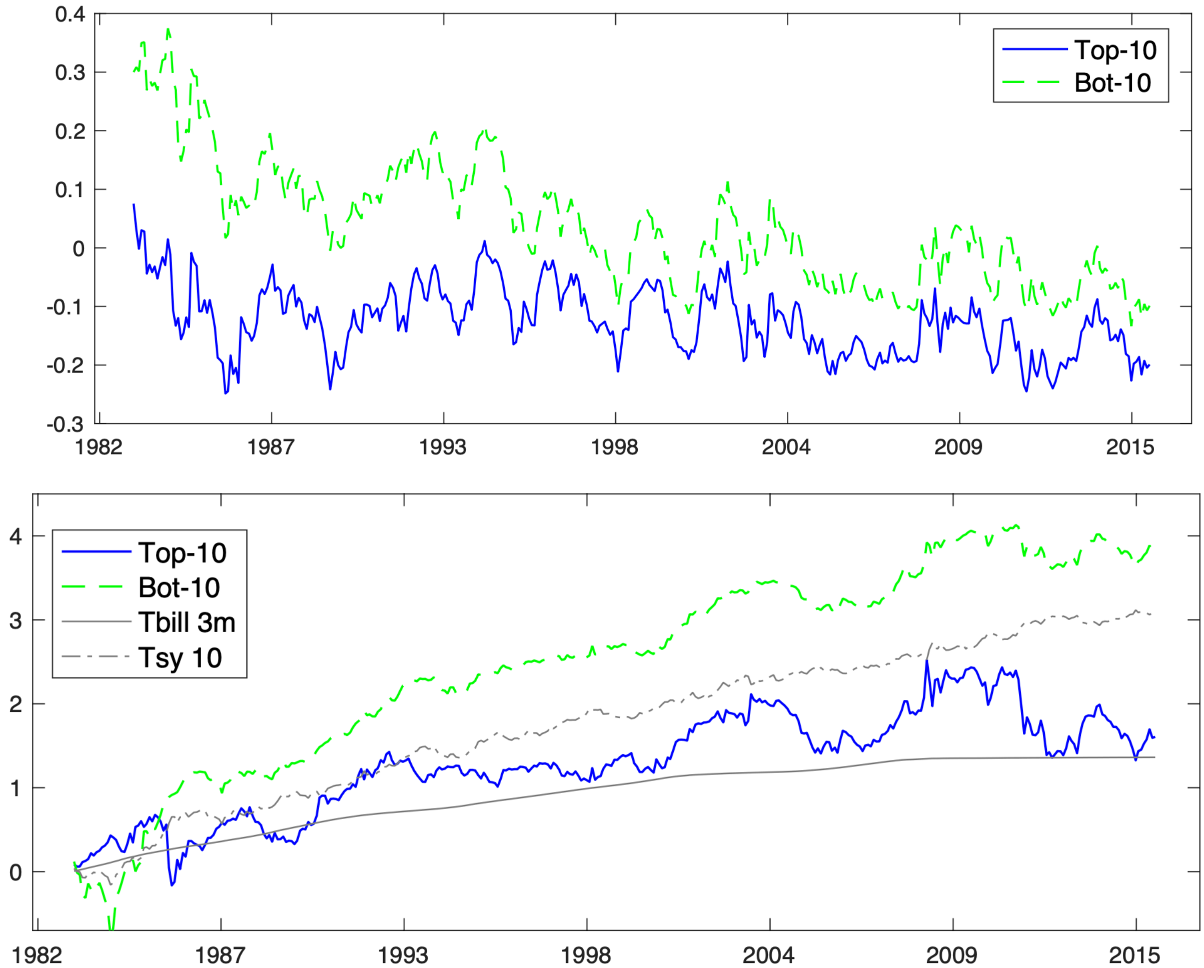

Notes:
The upper panel of this figure plots the top-10 and bottom-10 investor's expected excess return on the ten-year Treasury. The bottom panel shows the cumulative returns on the mean-variance portfolios chosen by the two investors, along with the cumulative return of rolling over the three-month TBill and that of rolling over the ten-year Treasury.

Figure 11: Wealth-Weighted Term Premium

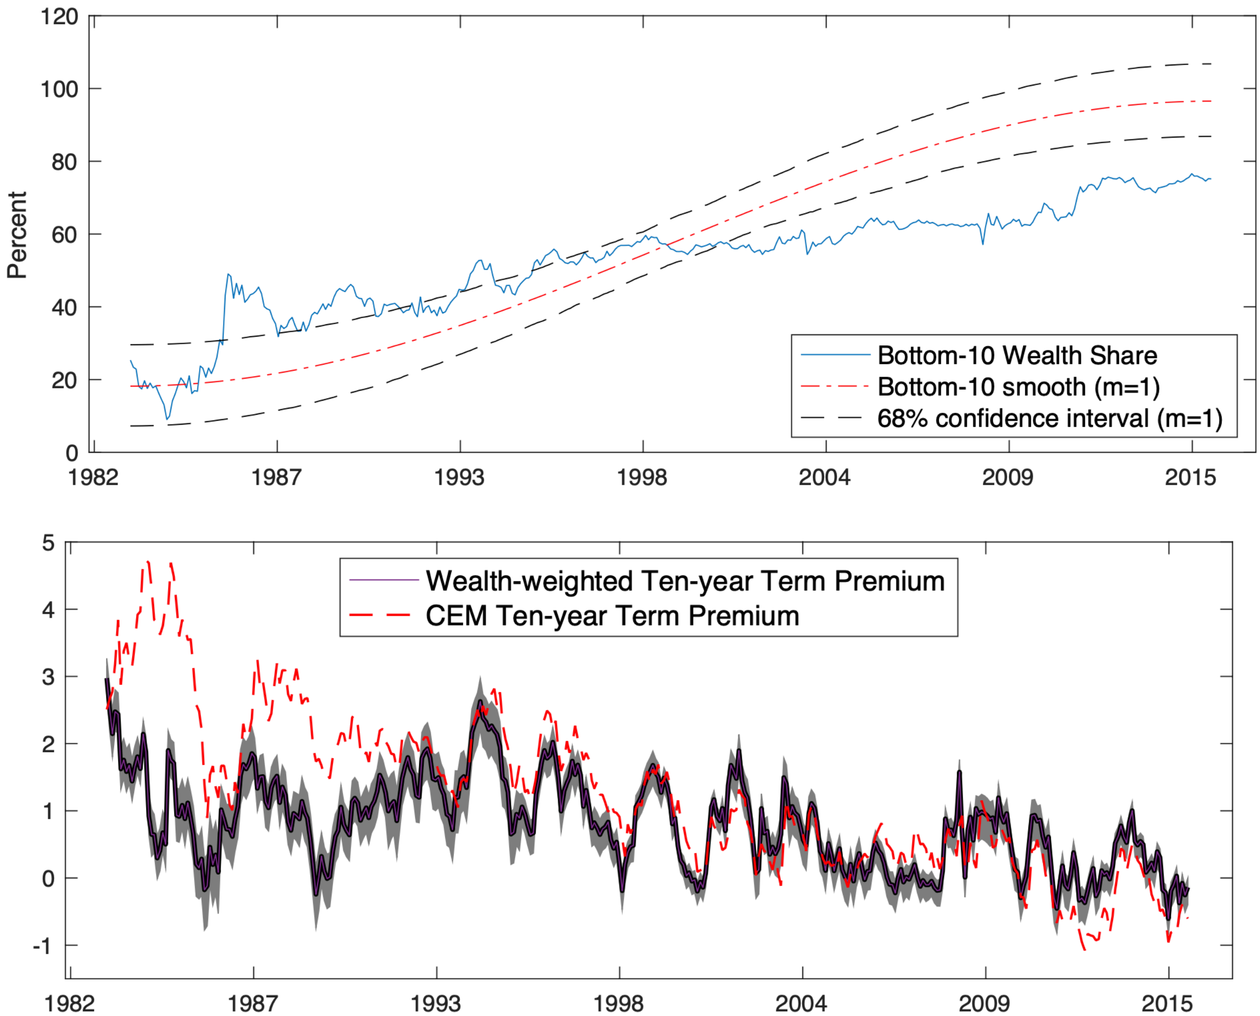

Notes:
The upper panel of this figure plots the estimated weight of the bottom-10 investor in the economy-wide term premium (red line) with associated 68% confidence intervals (black lines) along with the wealth share (blue line) derived from the cumulative returns for each investor. The bottom panel shows the wealth-weighted term premium based on a 20% initial wealth share (black line) and based on a two standard deviation interval around 20% (shaded region), along with the survey-based term premium of Crump, Eusepi, and Moench (2018) (red line).

Table 1: Model Parameter Estimates

$k_{\infty}^{\mathbb{Q}}$	2.67E-04	(4.29E-5)				
	X^L		X^S		X^C	
$\lambda^{\mathbb{Q}}$	0.000	(2.82E-6)	-0.029	(1.73E-2)	-0.091	(3.81E-2)
diag(Σ_{XX})	1.16E-05	(1.64E-5)	2.77E-05	(3.32E-5)	3.29E-05	(5.50E-5)
$\boldsymbol{\alpha}^X$			4.65E-04	(3.64E-4)	-3.30E-04	(7.49E-4)
$\boldsymbol{\beta}$	0.941	(3.52E-2)	0.965	(1.36E-2)	0.892	(6.22E-2)
	A		B		L	
σ_{short}	7.49E-03	(9.82E-3)	5.01E-03	(6.00E-3)	6.09E-03	(7.73E-3)
σ_{long}	1.87E-03	(1.03E-3)	1.25E-03	(6.35E-4)	1.52E-03	(1.33E-2)
σ_y	8.93E-04	(2.67E-4)				
chol(Σ_η)	1.20E-05	(1.15E-5)	-1.19E-03	(9.68E-4)	1.78E-03	(4.52E-3)
			7.96E-08	(1.64E-6)	2.28E-03	(2.53E-3)
					8.14E-12	(2.82E-6)

Notes:
This table reports parameter estimates for our affine term structure model. The sample period is 1983:03-2015:08, and standard errors are reported in parentheses. σ_y is the standard deviation of bond yield observational errors. σ_{short} and σ_{long} denote observational error standard deviations of short-horizon forecasts (less than one year) and long-horizon forecasts, respectively. L, A and B denote the econometrician's estimate of μ_t^L, σ_{short}, and σ_{long} and the respective beliefs of the top-10 and bottom-10 investor.

References

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

Andrade, Philippe, Richard K. Crump, Stefano Eusepi, and Emanuel Moench, 2016, Fundamental disagreement, *Journal of Monetary Economics* 83, 106 – 128.

Barillas, Francisco, and Kristoffer Nimark, 2019, Speculation and the bond market: An empirical no-arbitrage framework, *Management Science* 65, 4179–4203.

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

Bauer, Michael D, and Glenn D Rudebusch, 2020, Interest rates under falling stars, *American Economic Review* 110, 1316–54.

Bhamra, Harjoat S, and Raman Uppal, 2014, Asset prices with heterogeneity in preferences and beliefs, *Review of Financial Studies* 27, 519–580.

Bierens, Herman J., and Luis F. Martins, 2010, Time-varying cointegration, *Econometric Theory* 26, 1453–1490.

Buraschi, Andrea, Ilaria Piatti, and Paul Whelan, 2021, Subjective bond returns and belief aggregation, *Saïd Business School WP* 6, 2016–36.

Buraschi, Andrea, and Paul Whelan, 2016, Speculation, hedging, and interest rates, Working paper Imperial College London.

Carriero, Andrea, and Raffaella Giacomini, 2011, How useful are no-arbitrage restrictions for forecasting the term structure of interest rates?, *Journal of Econometrics* 164, 21–34.

Crump, Richard K., Stefano Eusepi, and Emanuel Moench, 2013, Making a statement: How did professional forecasters react to the August 2011 FOMC statement?, Liberty Street Economics Blog.

Crump, Richard K, Stefano Eusepi, and Emanuel Moench, 2018, The term structure of expectations and bond yields, Staff Report 775 Federal Reserve Bank of New York.

Duffee, Gregory R., 2002, Term premia and interest rate forecasts in affine models, *Journal of Finance* 57, 405–443.

Duffee, Gregory R., 2011a, Forecasting with the term structure: The role of no-arbitrage restrictions, Working paper Johns Hopkins University, Department of Economics.

Duffee, Gregory R., 2011b, Information in (and not in) the term structure, *Review of Financial Studies* 24, 2895–2934.

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

Giacoletti, Marco, Kristoffer T. Laursen, and Kenneth J. Singleton, 2021, Learning from disagreement in the U.S. Treasury bond market, *Journal of Finance* 76, 395–441.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The US Treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291–2304.

Joslin, Scott, Marcel Priebsch, and Kenneth J. Singleton, 2014, Risk premiums in dynamic term structure models with unspanned macro risks, *Journal of Finance* 69, 1197–1233.

Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu, 2011, A new perspective on Gaussian dynamic term structure models, *Review of Financial Studies* 24, 926–970.

Jouini, Elyes, and Clotilde Napp, 2007, Consensus consumer and intertemporal asset pricing with heterogeneous beliefs, *Review of Economic Studies* 74, 1149–1174.

Kim, Don H., and Jonathan H. Wright, 2005, An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates, Working paper Board of Governors of the Federal Reserve System.

Kreps, David M, 1998, Anticipated utility and dynamic choice, *Econometric Society Monographs* 29, 242–274.

Piazzesi, Monika, Juliana Salomao, and Martin Schneider, 2015, Trend and cycle in bond premia, Working paper Stanford University.

Singleton, Kenneth J, 2021, Presidential Address: How much "rationality" is there in bond-market risk premiums?, *Journal of Finance*.

Wright, Jonathan H, 2011, Term premia and inflation uncertainty: Empirical evidence from an international panel dataset, *American Economic Review* 101, 1514–34.

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

Appendix

Normalization Scheme and Parameter Estimates

To estimate the model, we employ the normalization scheme proposed by Joslin, Singleton, and Zhu (2011) (henceforth, JSZ). Under the JSZ normalization scheme, we have a risk-neutral parameter set $\Theta^{\mathbb{Q}} \equiv (\Sigma_{XX}, \lambda^{\mathbb{Q}}, k^{\mathbb{Q}}_{\infty})$. Let $\tilde{\mathbf{X}}_t$ denote a set of risk factors with

$$r_t = 1'\tilde{\mathbf{X}}_t,$$
$$\tilde{\mathbf{X}}_{t+1} = C(k^{\mathbb{Q}}_{\infty}) + J(\lambda^{\mathbb{Q}})\tilde{\mathbf{X}}_t + \Sigma^{1/2}_{XX} e^{\mathbb{Q}}_{X,t+1}.$$

JSZ show that there exists a unique rotation of $\tilde{\mathbf{X}}_t$ so that the factors are portfolios (or principal components) of bond yields:

$$\mathcal{P}_t = v(\Theta^{\mathbb{Q}}, W) + L(\lambda^{\mathbb{Q}}, W)\tilde{\mathbf{X}}_t,$$

where W denote weights used to construct factor-mimicking portfolios such that the latent states are portfolios of yields.[7] That is, $\mathcal{P}_t = \mathbf{y}^o_t \cdot W'$. It can be shown that the parameters controlling the risk neutral dynamics (Ψ_0, Ψ_1, δ_0, δ_1 and Σ_ϵ) are all functions of the elements in $\Theta^{\mathbb{Q}}$.

The physical dynamics can be written in a similar normalized form

$$\tilde{\mathbf{X}}_{t+1} = C(\mu_t) + J(\lambda^{\mathbb{P}})\tilde{\mathbf{X}}_t + \Sigma^{1/2}_{XX}\varepsilon_{t+1},$$

where we impose that the coefficient matrix $J(\lambda^{\mathbb{P}})$ is diagonal.

Nonparametric Estimation of Time-Varying Weights

In our model we have that

$$rx^{(n-1)}_{t+1} = \mathcal{B}_{n-1}\Lambda_1\mathbf{X}_t + \mathcal{B}_{n-1}\left(w^A_t\lambda^A_{0,t} + w^B_t\lambda^B_{0,t}\right) + \mathcal{B}_{n-1}v_{t+1},$$

where we have used that $\lambda^R_t = w^A_t\lambda^A_{0,t} + w^B_t\lambda^B_{0,t} + \Lambda_1\mathbf{X}_t$ for some weights $w^A_t + w^B_t = 1$. Importantly we have that $\mathbb{E}_t\left[v_{t+1}\right] = 0$. Suppose that we assume that w^A_t is a smooth process over time. Then it can be well approximated by a linear combination of an appropriate set of basis functions:

$$w^A_t \approx \sum_{j=0}^{m} \zeta_j \cdot P_{j,t},$$

where

$$\begin{aligned} P_{0,t} &= 1 \\ P_{i,t} &= \sqrt{2}\cos\left(i\pi\left(t-.5\right)/T\right), \end{aligned}$$

[7] We choose the portfolio weights following Duffee (2011a). The portfolios can be interpreted as empirical Level, Slope and Curvature.

and m is a user-chosen parameter which governs how much smoothness to impose (see, e.g., Bierens and Martins (2010)). For example if $m = 0$ then we would set w_t^A to be constant over time. Next note that,

$$\begin{aligned} & rx_{t+1}^{(n-1)} \\ = \; & \mathcal{B}_{n-1}\Lambda_1\mathbf{X}_t + \mathcal{B}_{n-1}\lambda_{0,t}^B + \mathcal{B}_{n-1}w_t^A\left(\lambda_{0,t}^A - \lambda_{0,t}^B\right) + \mathcal{B}_{n-1}v_{t+1} \\ = \; & \mathcal{B}_{n-1}\Lambda_1\mathbf{X}_t + \mathcal{B}_{n-1}\lambda_{0,t}^B + \zeta_0 \cdot \mathcal{B}_{n-1}\left(\lambda_{0,t}^A - \lambda_{0,t}^B\right)P_{0,t} + \cdots + \zeta_m \cdot \mathcal{B}_{n-1}\left(\lambda_{0,t}^A - \lambda_{0,t}^B\right)P_{m,t} + \mathcal{B}_{n-1}v_{t+1}. \end{aligned}$$

We observe one-period excess returns across time and maturity. From our model outputs we have $\left(\mathcal{B}_{n-1}, \Lambda_1, \lambda_{0,t}^A, \lambda_{0,t}^B\right)$. However, we do not observe $\mathcal{B}_{n-1}v_{t+1}$ because this is the realized return error relative to the representative investor. Thus, we can estimate the coefficients on each basis function directly by linear regression. To choose m in practice, we use the BIC,

$$BIC = nT\log\left(\sum\nolimits_{n,t}\left(\xi_t^{(n)}\right)^2\right) + m\log\left(nT\right),$$

where

$$\begin{aligned} & \xi_t^{(n-1)} \\ = \; & rx_{t+1}^{(n-1)} - \mathcal{B}_{n-1}\lambda_1\mathbf{X}_t - \mathcal{B}_{n-1}\lambda_{0,t}^B - \\ & \hat{\zeta}_0 \cdot \mathcal{B}_{n-1}\left(\lambda_{0,t}^A - \lambda_{0,t}^B\right)P_{0,t} - \cdots - \hat{\zeta}_m \cdot \mathcal{B}_{n-1}\left(\lambda_{0,t}^A - \lambda_{0,t}^B\right)P_{m,t} - \hat{c} \end{aligned}$$

and $\hat{c}$ is the estimated constant. The red line in Figure 11 shows the nonparametrically estimated weights using the optimal choice based on BIC ($m = 1$). As we observe, these model-free estimated weights are consistent with the weights we use in our baseline analysis.