

NBER WORKING PAPER SERIES

INTEREST RATE SKEWNESS AND BIASED BELIEFS

Michael D. Bauer
Mikhail Chernov

Working Paper 28954
<http://www.nber.org/papers/w28954>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021, Revised January 2023

We thank the editor Stefan Nagel, the associate editor, and two referees for their valuable feedback. We are also grateful to Anna Cieslak (discussant), Ian Dew-Becker, Greg Duffee, Michael Gallmeyer, Marco Giacoletti, Mathieu Gomez, Valentin Haddad, Christian Heyerdahl-Larsen, Lars Lochstoer, Ricardo Reis (discussant), Avanidhar Subrahmanyam, Andreea Vladu (discussant), as well as participants at various research seminars and conferences. Bauer acknowledges funding by the German Research Foundation (Deutsche Forschungsgemeinschaft), project no. 425909451. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Michael D. Bauer and Mikhail Chernov. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Interest Rate Skewness and Biased Beliefs
Michael D. Bauer and Mikhail Chernov
NBER Working Paper No. 28954
June 2021, Revised January 2023
JEL No. E43,E44,E52,G12

ABSTRACT

The conditional skewness of Treasury yields is an important indicator of the risks to the macroeconomic outlook. Positive skewness signals upside risk to interest rates during periods of accommodative monetary policy and an upward-sloping yield curve, and vice versa. Skewness has substantial predictive power for future bond excess returns, high-frequency interest rate changes around FOMC announcements, and survey forecast errors for interest rates. The estimated expectational errors, or biases in beliefs, are quantitatively important for statistical bond risk premia. These findings are consistent with a heterogeneous-beliefs model where one of the agents is wrong about consumption growth.

Michael D. Bauer
Universität Hamburg
Department of Economics
Von-Melle-Park 5
20146 Hamburg
Germany
michael.d.bauer@gmail.com

Mikhail Chernov
Anderson School of Management
University of California, Los Angeles
110 Westwood Plaza, Suite C-417
Los Angeles, CA 90095
and CEPR
and also NBER
mikhail.chernov@anderson.ucla.edu

1 Introduction

Predicting bond returns and changes in interest rates is a key challenge for financial economists, investors, and policymakers. Conditional skewness of interest rates is a natural candidate for a useful predictor. In contrast to measures of uncertainty or volatility, skewness contains information about the likely direction of future rate changes, because it measures asymmetry and the balance of interest rate risks. This paper demonstrates that conditional skewness of Treasury yields is important for understanding yield dynamics and bond risk premia. We study model-free measures of skewness, focusing on forward-looking skewness implied by Treasury option prices. Over the last three decades, conditional interest rate skewness has exhibited pronounced cyclical variation. Importantly, it contains additional predictive power for future interest rates, above and beyond the information in commonly used predictors or survey forecasts. Our evidence is consistent with a theoretical environment where some economic agents have biased beliefs.

A large literature has documented the negative skewness of stock returns and its implications for asset pricing and investment management (see [Neuberger, 2012](#), and the references therein). In stark contrast, there is a dearth of evidence on the asymmetry of the distribution of interest rates and bond returns. In general, yield-curve models assume symmetric distributions.¹ This modeling choice may be justified by the apparent symmetry of the unconditional distribution of interest rates. The sample skewness coefficient of changes in Treasury yields is essentially zero. But important asymmetries arise once the focus is shifted to the conditional distribution of Treasury yields.

We measure conditional skewness in two different ways, using prices of Treasury futures and options. Realized skewness is based on realized moments of yield changes. Implied skewness is based on risk-neutral moments implied by Treasury options. While both measures are highly correlated, for most of our empirical analysis we focus on option-implied skewness, which is forward-looking, less noisy, and available at a daily frequency.

Our evidence shows substantial cyclical variation in conditional skewness over the past 30 years. The variation is persistent, meaning that the balance of interest rate risk is characterized by substantial upward or downward skew over extended periods of time. Option-implied

¹See, for example, [Ang and Piazzesi \(2003\)](#), [Rudebusch and Wu \(2008\)](#), and [Bauer and Rudebusch \(2020\)](#), among many others. Even papers on non-normality of interest rates and its links to monetary policy generally, such as [Piazzesi \(2001\)](#) and [Johannes \(2004\)](#), assume zero conditional skewness and treat the effects of FOMC announcements as symmetric.

skewness predicts a substantial share of the variation in realized skewness, establishing formal statistical evidence for the time variation in conditional skewness. The cyclical variation in conditional skewness is related to macroeconomic conditions, such as the shape of the yield curve, the stance of monetary policy, and the business cycle. Skewness tends to be high when the Fed has been easing the stance of monetary policy and the yield curve is upward-sloping, and low when the Fed has recently been tightening and the yield curve is flat or inverted. Intuitively, after a series of rate hikes, investors are becoming increasingly concerned about an impending turning point for monetary policy and downside risk to rates.

Changes in conditional skewness reflect shifting perceptions of interest rate risk and indeed often correctly anticipate turning points of monetary policy and changes in interest rates. We document that conditional skewness contains useful information for the future course of interest rates. Skewness significantly predicts excess returns on Treasury bonds, and this finding is robust to controlling for other commonly used predictors such as the shape of the yield curve or macroeconomic trends (Cochrane and Piazzesi, 2005; Cieslak and Povala, 2015; Bauer and Rudebusch, 2020). In fact, skewness is particularly informative about future bond returns when considered jointly with the yield curve, in violation of the spanning hypothesis for bond markets (Duffee, 2011; Bauer and Hamilton, 2018). The COVID-19 episode is a powerful illustration: implied skewness signaled the increased downside risk to rates at the onset of the pandemic and then anticipated the steep rise in long-term Treasury yields starting in mid-2021.

Monetary policy appears to be an important factor underlying the predictive power of skewness, which on the whole is even more pronounced for changes in short-term rates. Skewness is highly informative about high-frequency changes in money market futures rates around upcoming FOMC announcements. In other words, it correctly predicts a substantial portion of “monetary policy surprises.” This finding is mainly driven by the fact that negative skewness captures downside risk to rates during easing cycles.

Because predictability of asset returns may arise from both time-varying subjective risk premia and expectational errors, we next study the relation between survey forecast errors and conditional skewness. Using surveys of professional forecasters, we document that skewness measured at the time of the surveys is highly informative for interest rate forecast errors. This finding, which extends across all forecast horizons and is robust to controlling for the shape of the yield curve, suggests that expectational errors may be important for the predictability of excess bond returns.

To quantify this channel we use a decomposition of statistical risk premia into (survey-based) subjective risk premia and expectational errors. Conditioning only on information in the yield curve leads to the mistaken conclusion that subjective risk premia are the main source of variation in statistical bond risk premia. But conditioning on yield skewness uncovers the quantitatively important role of variation in expectational errors, which explains more than half of the variance of statistical bond risk premia for bonds with maturities between one and five years. This evidence supports the view that time-varying differences between statistical and subjective expectations, or biased beliefs, play an important role for explaining the predictive power of skewness for interest rates. Skewness appears to be a proxy for the bias in beliefs about future interest rates.

That observation offers a clue about a possible economic mechanism behind the evidence: biased beliefs. We adopt the heterogeneous beliefs two-agent framework ([Basak, 2005](#); [Xiong and Yan, 2010](#); [Ehling et al., 2018](#)) by assuming that one agent knows the endowment growth process, while the other agent has to form beliefs.² As a result, the usual measure of disagreement in these models becomes a measure of bias in beliefs, which is the key state variable. Interest rates depend nonlinearly on disagreement, and are non-normally distributed even though the state variables are Gaussian. Yield skewness arises endogenously and varies over time with changes in beliefs. We use simulations to show that this simple, almost “off-the-shelf” model is qualitatively consistent with our findings on time-varying skewness, predictability, and expectational errors.

Our paper is related to several strands of the macro-finance literature. It builds on a long tradition of research on predictability in bond markets, including [Campbell and Shiller \(1991\)](#), [Cochrane and Piazzesi \(2005\)](#), [Ludvigson and Ng \(2009\)](#), [Cieslak and Povala \(2015\)](#) and many others. Conditional yield skewness is a powerful predictor in this context that has not previously been studied. Recent work has revisited this predictability allowing for deviations from the benchmark of full information rational expectations (FIRE). [Bacchetta et al. \(2009\)](#), [Piazzesi et al. \(2015\)](#), [Cieslak \(2018\)](#), [Schmeling et al. \(2022\)](#), [Buraschi et al. \(2022\)](#) and [Nagel and Xu \(2022\)](#) use survey forecasts of interest rates to study subjective bond risk premia and expectational errors. Our evidence supports the view that bond risk premia and interest rate expectations deviate from the FIRE assumption, and it suggests that changes

²Existing work in macroeconomics and finance has proposed various types of belief formation and deviations from rational expectations; prominent examples include “natural expectations” ([Fuster et al., 2010](#)), “diagnostic expectations” ([Bordalo et al., 2018](#)), and, more generally, extrapolative expectations ([Barberis et al., 2015](#)). In our model specification, we do not take a stand on the precise nature of the belief formation.

in conditional skewness capture expectational errors for interest rates.

[Giacoletti et al. \(2021\)](#) show that disagreement about future yields across forecasters predicts excess bond returns, and attribute this result to learning about interest rate dynamics. Our predictability results are robust to controlling for their measure of disagreement. [Giacoletti et al.](#) argue that their evidence is inconsistent with heterogeneous beliefs about macro fundamentals, since they find no relationship between intra-survey disagreement about inflation and yields (see also [Singleton, 2021](#)). The mechanism that we emphasize with our heterogeneous beliefs model is a different one, as we are focusing on the consensus forecasts. In our model and evidence, disagreement between consensus forecast and the truth is closely connected to biases in beliefs about fundamentals. We document that conditional skewness predicts survey forecast errors for GDP growth, indicating that skewness is connected to biased beliefs about both fundamental and financial variables, consistent with the heterogeneous beliefs framework.

Previous research has extensively studied the interaction of monetary policy and bond markets (for a survey, see [Gürkaynak and Wright, 2012](#)). Since [Kuttner \(2001\)](#), many studies have studied the effects of surprises in monetary policy announcements on asset prices (e.g., [Gürkaynak et al., 2005](#), [Nakamura and Steinsson, 2018](#)). Our results show that high-frequency rate changes around FOMC announcements are predictable with the information in conditional skewness, questioning their use as exogenous monetary policy surprises. Our findings are consistent with related evidence on this type of predictability using information in macroeconomic and financial variables ([Cieslak, 2018](#); [Bauer and Swanson, 2022a,b](#)).

Our paper also connects to studies on the importance of asymmetries in the macro-financial outlook (e.g. [Barro, 2006](#); [Conrad et al., 2013](#); [Adrian et al., 2019](#)). Recent work on the “Fed put” by [Cieslak and Vissing-Jorgensen \(2021\)](#) documents the predictive power of downside risk in the stock market for Fed easing actions, which is related to our finding that asymmetric risk perceptions in the bond market predict market moves around FOMC announcements.

Our methodology for using options to measure skewness follows that of [Bakshi and Madan \(2000\)](#) and [Neuberger \(2012\)](#). An early paper on time-varying skewness in interest rates is [Durham \(2008\)](#), who shows that skewness implied by Eurodollar options is related to estimates of the term premium in long-term yields. [Trolle and Schwartz \(2014\)](#) study swaption-implied skewness and document some variation over their 2002-2009 sample, but do not relate skewness to bond returns or survey forecasts of yields. Some other empirical work has

considered interest rate skewness in other contexts, such as the impact of unconventional monetary policy on tail risks (Hattori et al., 2016), the distinction between different equilibria based on skewness at the ZLB (Mertens and Williams, 2021), the connection between skewness and recessions (Li, 2021), or the skewness in subjective interest rate distributions in surveys (Diercks et al., 2022).

Asset pricing with heterogeneous beliefs has been studied since Harrison and Kreps (1978) and Detemple and Murthy (1994), and this theory is reviewed in Basak (2005). Xiong and Yan (2010) and Ehling et al. (2018) use heterogeneous beliefs models to price Treasury bonds. Buraschi and Whelan (2022) analyze speculation and hedging in a bond pricing model with disagreement. Cao et al. (2021) and Crump et al. (2021) study the effects of disagreement about long-run trends in no-arbitrage term structure models. Our contribution to this literature is to uncover a link between biases in beliefs about fundamentals and skewness of interest rates, a connection we document both in theory and in the data.

2 Cyclical variation in interest rate skewness

Interest rate skewness captures the degree of asymmetry in the probability distribution of changes in interest rates. Given that average interest rate changes are close to zero, positive skewness indicates that large rate increases are more likely than large rate declines, so that the balance of risk is tilted to the upside, and vice versa. In this section, we show that unconditional skewness of Treasury yields has been essentially zero over recent decades, explain how we measure conditional skewness from prices of Treasury futures and options, and then investigate its properties, cyclical behavior, and underlying economic drivers.

2.1 Data and sample skewness

We use prices of Treasury futures and options from CME group. We focus on the ten-year T-Note (TY) futures and option contracts, because these Treasury derivatives are the most liquid (in terms of open interest and volume) and have the longest available price history, compared to derivatives on Treasuries with shorter or longer maturities. In addition, the ten-year Treasury yield is a key benchmark interest rate in global financial markets.

We use end-of-day prices for the TY futures and options. The deliverable securities for the TY futures contract are Treasury bonds with maturities between 6.5 and 10 years. Changes

in futures prices are closely associated with negative yield changes in the cheapest-to-deliver (CTD) Treasury security. Because skewness is scale invariant, we can take the negative of the skewness of futures price changes as a measure of skewness of yield changes for the CTD bond. Details are in Appendix A.

Treasury yield data are the smoothed zero-coupon yield curves from [Gürkaynak et al. \(2007\)](#). To construct yield curve factors—level, slope and curvature of the yield curve—we take the first three principal components of these yields up to a maturity of ten years. For monthly samples we use end-of-month yields.

The sample period is from the beginning of January 1990 to the end of May 2021. The starting date is dictated by the availability of options data. While the CME Treasury options data starts in May 1985, there are only few contracts and prices available during the early years.

Table 1 reports unconditional sample skewness coefficients for changes in the ten-year yield at different frequencies, ranging from one-month to twelve-month changes. It also shows 90%-confidence intervals that are calculated using a simple bootstrap (since yields are highly persistent and the serial correlation of their changes is close to zero). In addition to yield changes, we also consider negative futures price changes.

Table 1: Unconditional sample skewness for changes in ten-year Treasury yield

Months	Yield changes		Neg. fut. price changes	
$m = 1$	0.03	(-0.42, 0.51)	-0.14	(-0.62, 0.40)
$m = 2$	-0.41	(-1.28, 0.39)	-0.20	(-0.65, 0.30)
$m = 3$	-0.24	(-0.57, 0.05)	-0.16	(-0.39, 0.06)
$m = 12$	0.36	(-0.09, 0.84)	-0.05	(-0.28, 0.17)

Notes: Sample skewness coefficient for m -month changes of the ten-year Treasury yield and negative changes of futures prices. Sample period: January 1990 to May 2021.

The skewness coefficients in Table 1 are all close to zero, and we generally cannot reject the hypothesis that in population the skewness is zero. This full-sample, unconditional perspective suggests that interest rate changes are essentially symmetric. It is missing important features of the distribution of interest rates, which only become evident once we consider conditional skewness.

2.2 Realized and implied skewness

To measure the skewness of the conditional distribution of yields,

$$\frac{E_t(y_T - E_t y_T)^3}{(Var_t y_T)^{3/2}},$$

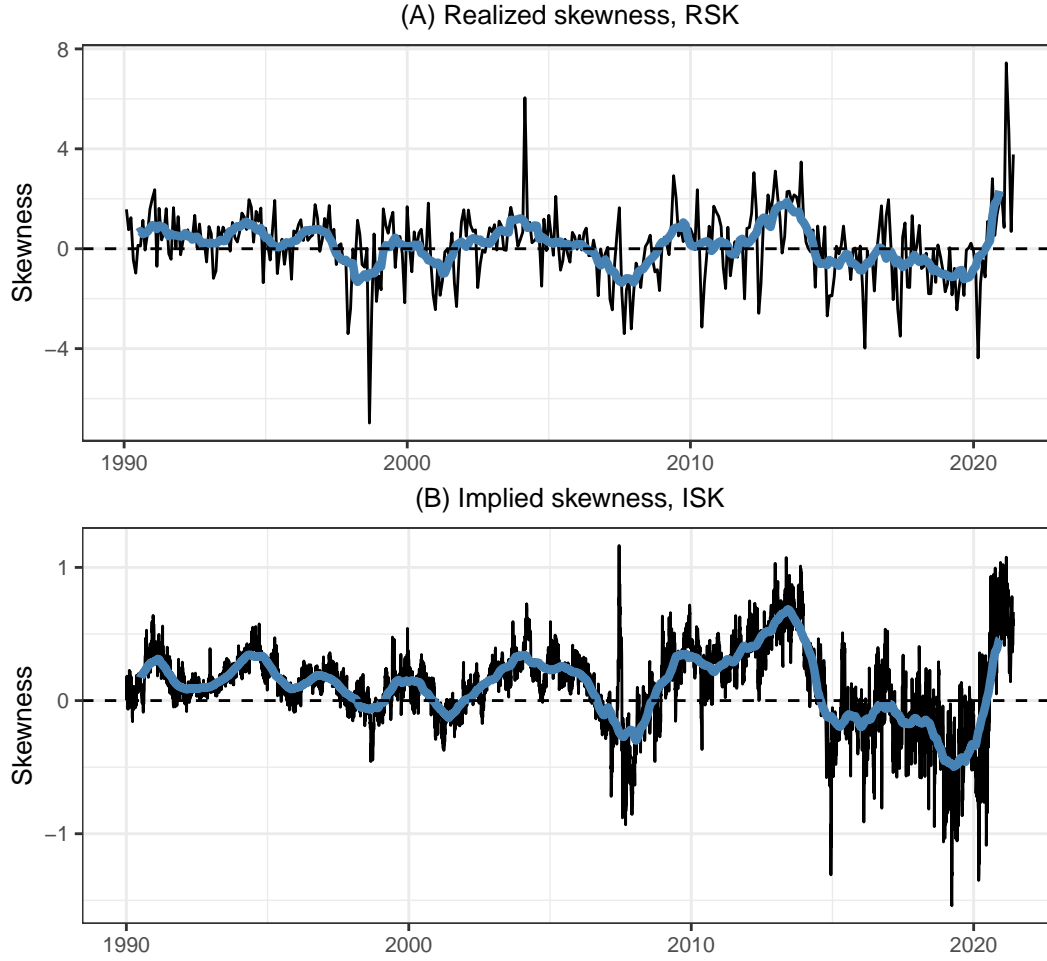
estimates of second and third conditional moments are required. Following the literature on skewness in stock returns, we use realized and option-implied moments for Treasury futures price changes, and we measure yield skewness as negative price skewness.

We calculate realized skewness (RSK) at a monthly frequency using daily changes in prices and implied volatilities for Treasury futures (Neuberger, 2012, equation 5). Figure 1A plots RSK and its 12-month moving average. RSK is volatile and on average close to zero, but exhibits some persistent variation. During three episodes it was markedly negative: the dot-com bubble 1998-2000, the financial crisis of 2007-2009, and the period since 2015 when the Fed lifted its policy rate off the ZLB. Skewness declines sharply in the wake of the COVID-19 pandemic in early 2020 but then reaches historical high level in the wake of global fiscal and monetary stimulus.

RSK is a simple and useful measure, but it is quite noisy, available only at lower frequencies, and backward-looking. Implied skewness (ISK) does not suffer from these drawbacks. This measure of conditional yield skewness uses the risk-neutral moments implied by options on Treasury futures. Appendix A provides details for how we construct these moments. On every trading day, we calculate ISK for the most active option contract, the shortest quarterly contract, which generally has a maturity between 1 and 3 months.

Figure 1B shows a time series of implied skewness. Over the full sample, the average level of ISK is slightly positive, with a mean of 0.10 and a standard deviation of 0.30. There is substantial, persistent variation in conditional skewness, indicating shifting risk perceptions about future yields. Similar to realized skewness, ISK has exhibited pronounced cyclical swings over the course of our sample. But its variation is much more persistent: the first-order autocorrelation of monthly ISK is 0.95, whereas for RSK it is only 0.42. Particularly striking is the behavior during the first ZLB episode, when the Fed's policy rate was near zero. Between 2009 and 2014, ISK averaged 0.35, while outside of this period the average was only 0.04. Just before liftoff from the ZLB in 2015, skewness shifted markedly negative,

Figure 1: Treasury yield skewness



Panel (A): monthly realized skewness, calculated from changes in daily Treasury futures prices and implied volatilities, with a 12-month moving average (blue line). Panel (B): daily implied skewness, calculated from options on Treasury futures and interpolated to a constant horizon of 0.2 years (2.4 months, the average horizon of all option contracts), with a 250-day moving average (blue line). Sample period: January 1990 to May 2021.

averaging -0.2 between 2015 and the end of the sample. Similar to RSK, it then turned positive during the COVID-19 episode.

While ISK is much less erratic than RSK at the monthly frequency, it does have significant day-to-day volatility, which might be partly due to measurement error. For example, option prices in the tails affect measured skewness but may be more volatile and less reliable due to low trading volumes or rounding to the tick size. In predictive regressions we smooth out some of the day-to-day movements by averaging ISK over several days.

Table 2: Predicting realized skewness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
RSK	0.43*** (0.04)		0.19*** (0.06)		0.40*** (0.05)		0.17*** (0.06)
ISK		2.14*** (0.21)	1.60*** (0.27)			2.17*** (0.23)	1.65*** (0.27)
Level				−0.02 (0.06)	−0.02 (0.04)	−0.03 (0.03)	−0.03 (0.03)
Slope				0.27*** (0.08)	0.14** (0.06)	0.04 (0.06)	0.03 (0.06)
Curvature				0.08 (0.55)	0.15 (0.33)	−0.57* (0.31)	−0.38 (0.29)
Constant	0.08 (0.07)	−0.10 (0.07)	−0.07 (0.07)	−0.39 (0.47)	−0.18 (0.32)	0.14 (0.28)	0.11 (0.25)
Observations	376	376	376	376	376	376	376
R ²	0.18	0.24	0.26	0.05	0.19	0.24	0.26


Predictive regressions for one-month realized skewness (RSK). *ISK* is option-implied yield skewness; *RSK* is realized yield skewness based on daily changes in futures prices and implied volatilities; *Level*, *Slope* and *Curvature* are the first three principal components of Treasury yields. All predictors are measured at the end of the previous month. Sample: monthly observations from January 1990 to May 2021. Newey-West standard errors with automatic bandwidth selection are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

To assess whether the time variation in conditional skewness is statistically and economically significant, we check whether option-implied skewness predicts realized skewness. Table 2 presents results for various predictive regressions using one-month lags of RSK itself, ISK, and yield-curve factors. For ISK, we use the average over the last five days of the previous month. We report Newey-West standard errors with automatic bandwidth selection. As we noted before, RSK exhibits some autocorrelation, and regressions on its own lag yield an R^2 of 0.18. Lagged values of ISK have slightly stronger predictive power, and in a joint regression with lagged RSK and ISK, both are strongly significant. The slope of the yield curve has some predictive power for RSK, but this is driven out once we control for lagged RSK and ISK. These regressions provide evidence that conditional yield skewness varies over time, and that ISK is a useful forward-looking measure of this conditional skewness.

2.3 Skewness, the yield curve, and the business cycle

We now turn to the cyclical nature of the variation in skewness and its macro-financial drivers. Figure 2 plots annual moving averages of implied yield skewness and the slope of the Treasury yield curve. The two series are clearly positively related: Skewness tends to increase when the yield curve is steepening, and vice versa. This pattern is most pronounced during the monetary easing episodes in 2002-2003 and 2008-2013, but the positive correlation is evident throughout most of the sample.

Figure 2: Skewness and interest rates



Option-implied yield skewness (left axis) and the slope of the Treasury yield curve (the second principal component of Treasury yields, right axis). Annual moving averages of daily values. Green/orange shaded areas indicate monetary policy easing/tightening cycles (based on changes in the fed funds rate). Sample period: January 1990 to May 2021.

Table 3 shows regression evidence on the relationship between conditional skewness and macro-financial variables. The dependent variable in these regressions is ISK, the sample frequency is monthly, and Newey-West standard errors are reported in parentheses. A regression on the level and slope of the yield curve, in the first column, confirms that the slope is important for skewness: an upward-sloping yield curve is associated with high skewness.

The yield curve’s level, by contrast, does not have a statistically significant relationship with skewness. The second column shows that a negative interaction effect between level and slope adds substantial explanatory power: the slope exhibited a stronger relationship with skewness when the level of yields was low (later in the sample, in particular the 2008-2015 ZLB period) than when it was high (early in the sample). This specification explains about 20 percent of the variation in skewness. But the unexplained portion still has substantial cyclical variation, which we show below in Section 3 to contain relevant information about future yields.

Table 3: Explaining the level of conditional yield skewness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Level	0.004 (0.02)	0.06 (0.04)		0.02 (0.02)		0.02 (0.02)	0.08*** (0.03)
Slope	0.10*** (0.03)	0.20*** (0.06)		−0.01 (0.04)		0.03 (0.04)	0.14** (0.06)
Level*Slope		−0.03** (0.01)					−0.03** (0.01)
Easing			0.23*** (0.08)	0.29*** (0.10)			
Tightening			−0.16** (0.08)	−0.13 (0.08)			
Unemployment rate					0.08*** (0.02)	0.08** (0.03)	0.08*** (0.03)
Constant	−0.11 (0.14)	−0.32 (0.20)	0.07 (0.06)	−0.01 (0.15)	−0.36*** (0.13)	−0.50*** (0.16)	−0.71*** (0.18)
R ²	0.15	0.19	0.31	0.33	0.23	0.27	0.31

Regressions for the level of option-implied yield skewness of the ten-year Treasury yield, using monthly data from January 1990 to May 2021 ($N = 377$). *Level* and *Slope* are the first two principal components of Treasury yields; *Easing* and *Tightening* are dummy variables indicating whether the Federal Reserve was easing or tightening monetary policy one year ago (based on observed changes in the policy rate). The unemployment rate is the real-time observation for the previous month. Newey-West standard errors (with automatic bandwidth selection) are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

Shifts in monetary policy are a key driver of shifts in the yield curve, and in particular of its slope (Piazzesi, 2005; Rudebusch and Wu, 2008). Figure 2 shades monetary policy easing and tightening cycles, identified as the months from the first to the last successive decrease and increase, respectively, in the federal funds rate target. During easing cycles, as the Fed lowers the policy rate, the yield curve steepens and implied skewness rises. Intuitively, investors are starting to become more concerned about upside risk to rates after a series of rate cuts,

since the next tightening cycle may begin soon. As an example, during the monetary easing episode 2001-2003 following the dot-com bust, implied skewness rose steadily as option prices reflected rising risk of an impending turning point for monetary policy. During monetary tightening episodes, the fed funds rate rises as both the slope of the yield curve and skewness tend to decline, likely because investors are increasingly turning their attention to downside risk in the rate outlook.

These observations suggest that the positive correlation between conditional skewness and the slope of the yield curve arises in part from the monetary policy cycle and its effects on the balance of interest rate risk. Regressions with indicator variables for monetary easing and tightening cycles confirm this impression. Since the slope and implied skewness change in response to shifts in monetary policy, their *level* lags the cycle. For both variables, and for both tightening and easing indicators, the correlation is generally strongest for a lag of about one year, so we include one-year lags of the cyclical indicators in our regressions. The estimates in Table 3 confirm a statistically strong and economically intuitive relationship with the monetary policy cycle. A regression with only the two (lagged) cyclical indicators, shown in column 3, explains almost a third of the variation in conditional yield skewness. In a regression with both the cyclical indicators and yield curve factors, the coefficient on the slope is statistically insignificant, as shown in column 4. The two cyclical indicators apparently provide a more nuanced measure of the monetary policy cycle, driving out the relationship between skewness and the slope.

Of course, the stance of monetary policy is ultimately determined by macroeconomic conditions. In particular, cyclical indicators like the output gap or the unemployment rate tend to be strongly correlated with the slope of the yield curve (Rudebusch and Wu, 2008). Consistent with this evidence, ISK is also closely related to such cyclical indicators. Table 3 shows that variation in the unemployment rate explains about 20 percent of the variation in ISK. The coefficient on the unemployment rate remains statistically significant even after adding yield curve variables.³

Overall, we find strong correlations of conditional skewness with the slope of the yield curve, the stance of monetary policy, and the business cycle. In particular, when the yield curve is upward-sloping as a result of monetary easing during economic downturns, then ISK tends

³Table 3 uses real-time vintages for the unemployment rate in the previous month. We have also found evidence for the cyclical behavior of skewness using various business cycle variables, including NBER recession dummies, industrial production growth, the output gap, and the Chicago Fed National Activity Index. We omit these results for the sake of brevity.

to be high, and vice versa. Changes in conditional skewness appear to capture the changing balance of interest rate risk resulting from shifts in monetary policy over the business cycle. In Section 3 we will investigate whether the shifting perceptions signaled by conditional skewness indeed anticipate changes in interest rates and monetary policy, controlling for the cyclical drivers of skewness.

2.4 Skewness and the ZLB

Over most of the two ZLB periods, implied skewness was significantly elevated. One might hypothesize that skewness generally tends to be high at the ZLB because the left tail is truncated and skewness mechanically raised. Such truncation and asymmetry is certainly an important issue for conditional distributions of future short-term rates, which during ZLB periods have only one way to go, but the following observations suggest that it was not quantitatively important for Treasury yield skewness.


First, if conditional yield skewness depended on the proximity of yields to the ZLB, then this should imply a negative relationship between skewness and the level of the yield curve. By contrast, the estimates in Table 3 suggest either a non-existent or positive relationship, depending on the specification. Furthermore, skewness is time-varying and sign-switching in our sample prior to 2008, without any apparent time trend despite the secular downward trend in interest rates (Bauer and Rudebusch, 2020). And yields were lower in 2016 than during most of the first ZLB period, yet skewness was mainly negative in 2016.

Second, skewness behaved very differently during the two ZLB episodes in our sample. Skewness turned positive when the ZLB was reached in 2008 and remained mainly positive for several years, but then switched to negative in 2014, more than a year before the Fed lifted its policy rate off the ZLB. During the most recent period, skewness remained negative for several months after the ZLB was reached, but then turned positive in mid-2020 before long-term Treasury yields started to increase, as discussed in Section 3.5.

Third, a closer investigation of the entire option-implied density of future yields during the two ZLB episodes shows no evidence of a mechanical ZLB explanation of positive skewness. Figure 3 plots, for two different ZLB dates, the option-implied conditional densities for the future CTD Treasury yield. The densities are based on normal-inverse-gamma distributions (Eriksson et al., 2009) that match the first four implied conditional moments for future yields, which are calculated from (i) the Treasury futures and options prices, (ii) an approximate

mapping from bond price changes to yield changes (see Appendix A), and (iii) the current CTD yield. December 31, 2012, was a day with a particularly low yield level (1.14 percent) and a high level of skewness (0.8), but even for this extreme example the 1st percentile of the distribution is comfortably above the ZLB, at 0.6 percent. This suggests that the left tail is not thinner because it is cut off by the ZLB, but instead because investors perceived an upward tilted balance of risk and large right tail. On June 16, 2020, yields were even lower but ISK was actually *negative*, a pattern that persisted for some time during this period (see Figure 5 below). The 1st percentile of the option-implied conditional distribution was deeply negative, at -0.6 percent, indicating that the ZLB does not mechanically eliminate left tails.

Figure 3: Densities for future yields at ZLB



Option-implied probabilities densities for future CTD bond yield given market prices on December 31, 2013 (left panel) and June 16, 2020 (right panel). Red shaded area indicates 1st percentile.

The presence of the ZLB certainly affects some aspects of the distribution of interest rates including its skewness. But the ZLB, as well as the Fed's forward guidance on its likely duration, mainly affect the behavior of short-term interest rates (Swanson and Williams, 2014). Our evidence here and the predictability results in Section 3 suggest that proximity to the ZLB in itself is not an important driver of skewness of Treasury yields, and that high conditional skewness during ZLB episodes was not due to a mechanical truncation effect but changing perceptions of interest rate risk.

3 The information in conditional skewness

Our evidence so far has established the cyclical variation in skewness and linked it to economic driving forces and the business cycle. We now turn to the question whether yield skewness contains useful forward-looking information for interest rates. Consider the link between expected bond returns and risk premia:

$$E_t(RX_{t+1}^{(n)}) = E_t(RX_{t+1}^{(n)}) - E_t^s(RX_{t+1}^{(n)}) - \frac{Cov_t^s(M_{t+1}, RX_{t+1}^{(n)})}{E_t^s(M_{t+1})}, \quad (1)$$

where $RX_{t+1}^{(n)} = P_{t+1}^{(n-1)}/P_t^{(n)} - 1/P_t^{(1)}$ is the one-period excess gross return on an n -period bond with price $P_t^{(n)}$, M_{t+1} is the stochastic discount factor (SDF), and the superscript s refers to subjective probability.⁴ This representation is helpful because it demonstrates that predictability of excess returns can arise from variation in subjective risk premia and from a time-varying bias in beliefs, i.e., changes in systematic expectational errors. Variation in risk premia is the traditional explanation for bond return predictability. Under the assumption of full information rational expectations (FIRE), this is the only possible explanation, since statistical/rational expectations and subjective expectations coincide. More recent work, however, has emphasized the possibility of a failure of the FIRE assumption, allowing for biases in beliefs (e.g. [Cieslak, 2018](#)). Equation (1) provides a framework that guides our empirical work. We first establish whether skewness predicts excess returns. We implement that analysis in two different ways. We consider conventional predictive regressions for excess returns on Treasury bonds, similar to [Cochrane and Piazzesi \(2005\)](#) and many others. In addition, we analyze high-frequency interest rate changes around FOMC announcements, i.e., monetary policy surprises, because high-frequency interest changes closely correspond to (negative) excess bond returns. Next, we disentangle the source of predictability by using survey forecasts as a proxy for subjective expectations.

3.1 Bond returns

We begin with conventional predictive regressions for excess returns on Treasury bonds, similar to [Cochrane and Piazzesi \(2005\)](#) and many others. We work with monthly data and follow common practice by using end-of-month interest rates. The holding period is one

⁴Excess gross returns allow for the cleanest decomposition, while our empirical analysis below uses excess log returns, as is common in this literature. The two are very similar in the data, and our empirical results are essentially unchanged if we use excess gross returns.

quarter, because ISK is based on derivative contracts with an expiration of 2.4 months on average, and it is natural to match the horizons of option contracts and bond returns. Log excess returns are calculated as $rx_{t,t+3}^{(n)} = p_{t+3}^{(n-3)} - p_t^{(n)} - y_t^{(3)}$ where $p_t^{(n)} = -ny_t^{(n)}$ is the log price of a zero-coupon bond with n months to maturity and continuously compounded yield $y_t^{(n)}$, and the risk-free rate $y_t^{(3)}$ is taken to be the three-month T-bill rate. Our main predictive regressions are:

$$\overline{rx}_{t,t+3} = \beta' X_t + \varepsilon_{t,t+3}, \quad (2)$$

where $\overline{rx}_{t,t+3} = \sum_{j=1}^{10} rx_{t,t+3}^{(12j)} / (10j)$ is the weighted average log excess return across maturities from one to ten years, and $\varepsilon_{t,t+3}$ is the serially correlated prediction error. As in [Cieslak and Povala \(2015\)](#), we scale excess returns by maturity so that all excess returns have the same duration and similar volatility.

A vector of predictors, X_t , is observable at the end of month t . We are interested in the predictive power of option-implied yield skewness, ISK. To smooth out high-frequency movements in ISK and reduce measurement error, we use the average over the last five business days of the month. The predictors X_t also typically include the level, slope and curvature of the yield curve, the first three principal components of Treasury yields. Controlling for the shape of the yield curve is important since a natural null hypothesis is that the yield curve reflects all the information that is relevant for expectations of future interest rates; this “spanning hypothesis” has been investigated by [Duffee \(2011\)](#), [Bauer and Hamilton \(2018\)](#), and many others. For more reliable statistical inference in this setting with multi-period overlapping returns we calculate standard errors using the reverse regression delta method of [Hodrick \(1992\)](#) and [Wei and Wright \(2013\)](#).

Table 4 reports estimates of Equation (2) for different sets of predictors. The first column displays results for level, slope, and curvature alone. The coefficients on level and slope are statistically significant, although for the slope only marginally so. The slope coefficient is positive, in line with previous work that found a high slope to predict falling long-term yields and high bond returns ([Campbell and Shiller, 1991](#), [Cochrane and Piazzesi, 2005](#)).

Adding ISK roughly doubles the predictive power, measured by R^2 , relative to the yields-only specification. The coefficient on ISK is negative and highly significant, indicating that high skewness predicts low bond returns and rising yields. Compared to the yields-only specification, the coefficient on the slope is larger and more strongly significant. The fact that conditional yield skewness has significant predictive power even controlling for the information in yields, i.e., that the predictive power of ISK is not subsumed by the shape of the yield

Table 4: Predicting excess returns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Level	0.003** (0.002)	0.004** (0.002)		0.004** (0.001)	0.004** (0.002)		0.01*** (0.01)	0.01*** (0.002)
Slope	0.02* (0.01)	0.03*** (0.01)		0.03** (0.01)	0.03*** (0.01)		0.05*** (0.01)	0.05*** (0.01)
Curvature	0.03 (0.06)	0.004 (0.06)		0.04 (0.06)	0.01 (0.06)		-0.05 (0.07)	-0.02 (0.06)
ISK		-0.34*** (0.12)	-0.19* (0.10)		-0.28** (0.13)	-0.27*** (0.10)	-0.32*** (0.12)	-0.29** (0.14)
RSK				-0.06*** (0.02)	-0.02 (0.02)			
CF						1.12*** (0.30)		
i^*							-0.26** (0.12)	
GLS								-0.39** (0.17)
Constant	-0.07 (0.10)	-0.15* (0.09)	0.19*** (0.04)	-0.11 (0.09)	-0.15 (0.09)	0.01 (0.06)	0.26 (0.18)	-0.03 (0.14)
Observations	374	374	374	374	374	374	374	347
R ²	0.06	0.11	0.02	0.09	0.11	0.11	0.13	0.14

Predictive regressions for three-month excess bond returns (average of duration-normalized excess returns on Treasury bonds with one to ten years maturity) using monthly data from January 1990 to May 2021. Predictors: *Level*, *Slope* and *Curvature* are the first three principal components of end-of-month Treasury yields; *ISK* is option-implied yield skewness averaged over the last five business days of the month; *RSK* is monthly realized yield skewness based on daily changes in futures prices and implied volatilities; *CF* is the cycle factor of Cieslak and Povala (2015); i^* is an estimate of the trend component of nominal interest rates from Bauer and Rudebusch (2020); *GLS* is survey disagreement about future ten-year yields from Giacomelli et al. (2021). Standard errors based on the reverse regression delta method of Wei and Wright (2013) are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

curve, indicates a violation of the spanning hypothesis.⁵ To account for small-sample problems in predictive regressions for bond returns, we also test the spanning hypothesis using the bootstrap method proposed by Bauer and Hamilton (2018), which yields a small-sample p -value for ISK below 1%.

⁵Cieslak and Povala (2015) and Bauer and Rudebusch (2017) point out that apparent violations of the spanning hypothesis could be due to measurement error. However, theoretically, conditional skewness is non-linearly related to yields. Since the yield curve cannot span conditional skewness, the predictive power of skewness represents strong evidence for a violation of the spanning hypothesis.

The third column reports estimates for a univariate specification with only ISK. The coefficient is still negative and marginally significant, but the predictive power is weaker than for the specification which includes information in current yields. Comparing the results in columns (2) and (3) strongly supports the case that ISK contains additional information about future returns relative to the current yield curve: including yield-curve variables improves the predictive power of the regression, and at the same time raises the absolute magnitude of the coefficient on ISK and its t -statistic. In addition, the whole (the R^2 for the joint specification in column 2) is larger than the sum of its parts (the R^2 's in columns 1 and 3), which shows that *combining* the information in the yield curve and in option-implied yield skewness is key to taking full advantage of these two sources of forward-looking information.

The additional predictive power of skewness appears to mainly stem from its ability to capture downside risk to rates, i.e., positive surprises in excess returns. This becomes evident in a univariate regression of excess returns on ISK orthogonalized with respect to information in the yield curve (which in line with columns (1) and (2) of Table 4 yields an R^2 of 0.05). In this regression, the seven most influential observations are positive excess returns.

Column (4) adds RSK to the yield curve variables, instead of ISK. Realized skewness also has economically and statistically significant predictive power for bond returns, but it is somewhat less powerful than implied skewness, as evidenced by the slightly lower R^2 in column (4) than in column (2). To investigate whether both ISK and RSK are important for return predictions, we estimate a regression that includes both of them, shown in column (5). In this specification, ISK exhibits significant predictive power, but not RSK, confirming the notion that the former is a more accurate measure of conditional yield skewness than the latter.

Long-run trends in interest rates are an important issue for estimation of bond risk premia, as first documented by Cieslak and Povala (2015). They detrended yields using a slow-moving average of past inflation, and showed that a linear combination of detrended yields, which they called a “cycle factor,” is an excellent predictor of excess bond returns. In column (6), we control for this cycle factor, which optimally combines the information in both the current yield curve and the underlying inflation trend.⁶ The coefficient on ISK remains highly significant, and the R^2 is similar to the specification including the standard yield curve predictors in column (2).

⁶We estimate the cycle factor by detrending one- through ten-year yields with a moving average of core CPI inflation, predicting $\bar{r}x_{t,t+3}$ using the one-year and the average yield cycles, and calculating the fitted values.

Another way to account for long-run rate trends in bond return predictions is to include the trend proxy as an additional regressor. Column (7) controls for an estimate for the trend component of nominal interest rates, or i^* , suggested by [Bauer and Rudebusch \(2020\)](#). This trend variable includes proxies for both the inflation trend emphasized by [Cieslak and Povala \(2015\)](#) and the trend in real interest rates that has been the focus of much macroeconomic research since [Laubach and Williams \(2003\)](#), consistent with a long-run Fisher equation $i_t^* = \pi_t^* + r_t^*$. Accounting for the slow-moving interest rate trend in this way further raises the predictive power for future bond returns relative to the specification with only the yield curve. Importantly, the coefficient on ISK remains highly statistically significant.

The last column explores the relation between ISK and the survey disagreement about the ten-year yield that is measured as the interdecile range of BCFF forecasts and advocated by [Giacoletti et al. \(2021\)](#). The sample for this regression is shorter because their disagreement series ends in November of 2018. The estimates show that ISK continues to be significant when combined with yield disagreement, and that both variables add forecasting power for future bond returns.

Estimates for individual excess bond returns are shown in Appendix Table [B.1](#). ISK contains additional information about future returns for all bond maturities from one to ten years. Interestingly, the predictive power of ISK for excess returns is even stronger for short than for long bond maturities.

Our main results use quarterly holding periods, whereas empirical work on bond risk premia more commonly considers annual excess bond returns. The results in Appendix Table [B.2](#) show that for one-year holding periods the predictive power of conditionsl skewness is slightly weaker but remains statistically significant for the average excess return and most individual bond maturities.

An important question is how robust these results are across different sample periods. Appendix Table [B.3](#) shows estimates of our baseline predictive regression—including yield curve predictors and ISK—for various sample periods and demonstrates that the predictive power is robust. Of particular interest is the role of the first ZLB episode, when yield skewness was elevated and Treasury yields increased after the Fed lifted the policy rate off the ZLB, at least for some time. Table [B.3](#) shows that excluding this episode, by using a pre-ZLB sample period that ends in November 2008, leads to even stronger predictive power of ISK, so the ZLB episode does not seem to play a unique role in explaining our results.

Another way to consider robustness of predictive models over time is to assess out-of-sample (OOS) forecast accuracy. Appendix Table B.4 shows the results of an OOS analysis, which are consistent with our conclusion that ISK contains useful information for future bond returns. The improvements in OOS forecast accuracy are, however, small and not statistically significant. This lack of significance is likely due, at least in part, to the relatively small sample and the lower statistical power of OOS methods (Inoue and Kilian, 2005).

We have also estimated predictive regressions for excess bond returns using unsmoothed Fama-Bliss Treasury yields. In this dataset Cochrane and Piazzesi (2005) documented that a single linear combination of forward rates captures the predictive power of the yield curve for excess returns across all bond maturities. Appendix Table B.5 shows that ISK robustly predicts future bond returns in the Fama-Bliss data, and that it remains important when we control for the Cochrane-Piazzesi factor.⁷

3.2 Monetary policy surprises

We now zoom in on an important source of new information for bond markets: FOMC announcements. Going back to Kuttner (2001), an extensive literature has studied the reaction of interest rates to the surprise change in short-term interest rates. Such monetary policy surprises are typically calculated based on intraday changes in money market futures rates over a tight window around the FOMC announcements (Gürkaynak et al., 2005). Several recent papers have found that these high-frequency rate changes are predictable using publicly available macroeconomic data, possibly due to incomplete information of market participants about the Fed’s implicit policy reaction function (Cieslak, 2018; Bauer and Swanson, 2022a,b).⁸

We measure the policy “surprise” as the first principal component of intraday rate changes around the announcement that are derived from changes in fed funds and Eurodollar futures prices, following Nakamura and Steinsson (2018). This surprise, denoted below by s_t , is a univariate summary of the shift in short- and medium-term interest rates—the change in

⁷In addition, we have verified that the predictive power of ISK is robust to controlling for (i) the other drivers of skewness in Table 3, (ii) option-implied skewness in equity markets, measured by the CBOE skew index (Crump and Gospodinov, 2019), and (iii) measures of option-implied variance of interest rates (Choi et al., 2017).

⁸Monetary policy surprises may also contain so-called “information effects” and directly impact beliefs about macroeconomic fundamentals, as argued by Campbell et al. (2012), Nakamura and Steinsson (2018), Cieslak and Schrimpf (2019), and others. In contrast, the evidence in Bauer and Swanson (2022a,b) suggests that information effects are likely to be small.

the expected path of future policy rates, up to a term premium component. Appendix B.2 contains additional results for other measures of monetary policy surprises used in the literature, including the target and path factors of [Gürkaynak et al. \(2005\)](#). We report estimates for predictive regressions

$$s_t = \beta' X_{t-1} + \varepsilon_t, \quad (3)$$

where t are days with FOMC announcements, X_{t-1} are predictors observed on the day before the announcement and ε_t is a prediction error. We report White heteroskedasticity-robust standard errors, since ε_t is not serially correlated between FOMC announcements. Our sample contains 213 FOMC announcements from the beginning of 1994 (when the FOMC started making regular post-meeting announcements) to June 2019 (where our data for intradaily policy surprises ends). The sample includes both scheduled and unscheduled FOMC announcements, but our results are not sensitive to the exclusion of unscheduled announcements.

Table 5 shows the results. Information in the yield curve alone does not predict monetary policy surprises, as evident from the specification in column (1) where X_t contains only the level, slope and curvature of the yield curve. Column (2) shows estimates of a univariate regression. ISK alone has statistically significant predictive power and explains three percent of the variance in the monetary policy surprise.⁹ When combining the information in ISK and the yield curve, both the slope and ISK are highly statistically significant, and the R^2 is about 0.05, as shown in column (3). The slope's coefficient is negative, while the coefficient of ISK is positive, mirroring the findings for the return regressions in Table 4.

If we include RSK instead of ISK as a predictor, its coefficient is found to be statistically significant at the five-percent level. However, the predictive power is somewhat lower as for ISK, as shown in column (4). In a regression with both ISK and RSK, in addition to the yield curve factors, only the slope and ISK are statistically significant, as shown in column (5). Similarly to return regressions, the information in implied skewness appears more relevant for predicting future rate changes than the information in realized skewness.

Several previous studies have documented predictability of FOMC policy surprises. Useful predictors include the federal funds rate and employment growth ([Cieslak, 2018](#)), as well

⁹Consistent with our other predictive regressions, we use a five-day average for ISK. We note, however, that the predictive power of ISK is increased quite a bit if the moving average is calculated with 22 trading days, similar to the window length for the calculation of RSK.

Table 5: Predicting FOMC surprises with ISK and RSK


	(1)	(2)	(3)	(4)	(5)
Level	−0.0002 (0.001)		−0.001 (0.001)	−0.001 (0.001)	−0.001 (0.001)
Slope	−0.001 (0.002)		−0.005** (0.002)	−0.003 (0.002)	−0.005** (0.002)
Curvature	−0.016 (0.013)		−0.019 (0.013)	−0.015 (0.013)	−0.019 (0.013)
ISK		0.021*** (0.008)	0.029*** (0.009)		0.032*** (0.012)
RSK				0.004** (0.002)	−0.001 (0.003)
Constant	0.010 (0.007)	−0.002 (0.003)	0.018** (0.008)	0.015* (0.008)	0.017** (0.008)
Observations	213	213	213	213	213
R ²	0.009	0.030	0.054	0.027	0.055

Predictive regressions for the monetary policy surprise around FOMC announcements from January 1994 to June 2019. The dependent variable is the first principal component of 30-minute futures rate changes around the announcement for five different contracts with up to about one year maturity. *Level*, *Slope* and *Curvature* are the first three principal components of Treasury yields on the day before the announcement; *ISK* is a five-day average of option-implied yield skewness, *RSK* is realized yield skewness based on daily changes in futures prices and implied volatilities over the previous month (22 trading days). White heteroskedasticity-robust standard errors are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

as macroeconomic news, such as the surprise component in the monthly nonfarm payrolls number (Bauer and Swanson, 2022a). Appendix B.2 shows that ISK usually retains its predictive power when we control for these other variables.

These results raise the question what type of monetary policy decisions the information in ISK helps predict. The scatter plot in Figure 4 reveals that most of the predictive power of ISK arises from its correlation with upcoming easing surprises. It plots the monetary policy surprise against the residual ISK after projecting out the level, slope and curvature of the yield curve. The correlation in this plot captures the information that ISK adds to the information in the yield curve. A univariate regression for these observations yields an R^2 of 4.5 percent—consistent with the additional explanatory power in column (3) vs. column (1) in Table 5—and a t -statistic on the slope coefficient of 3.2. The eight most influential observations, as measured by the contribution to the slope coefficient, are labeled in Figure 4. They are all dovish surprises that occurred during monetary easing cycles—the U.S. economy

Figure 4: Skewness and monetary policy surprises



Horizontal axis: residual of ISK (averaged over month preceding each FOMC announcement) after regressing on level, slope and curvature of the yield curve. Vertical axis: monetary policy surprise. Sample: 213 FOMC announcements from January 1994 to June 2019. Labels indicate top eight most influential observations.

was typically in a recession—and they were preceded by unusually low levels of implied skewness. The ability of conditional skewness to anticipate dovish surprises during Fed easing cycles, with downside risk to rates, appears to be an important part of its predictive power.

3.3 Survey forecast errors

The evidence on interest-rate predictability in Sections 3.1 and 3.2 speaks to time variation in expected excess bond returns, that is, in the left-hand side of equation (1). We now investigate the role of expectational errors play for this predictability—the first term on the right-hand side of equation (1)—using survey forecasts as proxies for subjective expectations.

We calculate survey forecast errors for interest rates using the Blue Chip Financial Forecasts (BCFF). This is a monthly survey that contains interest rate forecasts for the current quarter (nowcasts) and each of the next five quarters.¹⁰ We consider forecasts for the ten-year

¹⁰The surveys conducted before 1997 extend out only four quarters.

Treasury yield, since ISK pertains to Treasury maturities between 6.5 and 10 years, and the federal funds rate, motivated by influential previous work by [Cieslak \(2018\)](#). The specific forecast targets are the quarterly averages of the daily ten-year constant-maturity Treasury yield and the effective federal funds rate, which we obtain from FRED. We calculate forecast errors as the difference between the quarterly average daily value and the consensus forecast, which is the arithmetic mean of the individual forecasts.

We estimate monthly predictive regressions for the forecast errors of the form

$$y_{q(t,h)} - \hat{y}_t^{(h)} = \beta' X_t + \varepsilon_t^{(h)},$$

where t indexes the month of the survey forecast, h is the forecast horizon from 0 to 5 quarters, $y_{q(t,h)}$ is the average interest rate over quarter $q(t, h)$ that contains the month $t + 3h$, $\hat{y}_t^{(h)}$ is the forecast for the average in quarter $q(t, h)$, X_t are predictors observable at the time the survey forecasts are made, and $\varepsilon_t^{(h)}$ is a forecast error.¹¹ We measure the predictors on the day of the BCFF survey deadline to ensure that they are observable at the time the forecast is made.¹² As before, we take a five-day average of ISK to reduce measurement error. The forecast errors of these regressions are necessarily serially correlated due to both the monthly frequency and also overlapping observations. Because of the latter, Hansen-Hodrick standard errors are preferable to Newey-West standard errors ([Cochrane and Piazzesi, 2005](#)), and we use $3(h + 1)$ lags.

Table 6 shows the results. When we predict the forecast error for the ten-year yield with ISK alone, the coefficient is positive and statistically significant for short forecast horizons out to two quarters. In predictive regressions that also include yield factors, the coefficient on ISK is larger than in the univariate specification, and statistically significant at the five-percent level for all horizons. As before, the slope of the yield curve has a coefficient with the opposite sign, which is marginally statistically significant for horizons of three quarters and beyond. The R^2 of these regressions increases with the horizon from four to 18 percent. Appendix B.3 presents analogous results for the Survey of Professional Forecasters, confirming the strong predictive power of ISK for interest rate forecast errors.

Previous work has documented that systematic expectational errors are particularly impor-

¹¹For example, in January, February, and March, forecasts for $h = 1$ are for the average over the second quarter (April to June).

¹²The surveys are conducted between the 23rd and 26th of the preceding month; the January survey is conducted between the 17th and the 21st of December.

Table 6: Predicting Blue Chip forecast errors

(A) <i>Ten-year yield</i>	Current	1Q ahead	2Q ahead	3Q ahead	4Q ahead	5Q ahead
ISK	0.14*** (0.05)	0.33*** (0.12)	0.34* (0.19)	0.30 (0.25)	0.33 (0.28)	0.33 (0.29)
Constant	−0.07*** (0.02)	−0.23*** (0.05)	−0.39*** (0.08)	−0.54*** (0.11)	−0.70*** (0.13)	−0.93*** (0.16)
R ²	0.03	0.04	0.03	0.02	0.01	0.02
ISK	0.15*** (0.05)	0.38*** (0.12)	0.45*** (0.17)	0.49*** (0.18)	0.58*** (0.19)	0.74*** (0.20)
Level	0.004 (0.01)	0.004 (0.02)	0.004 (0.04)	0.01 (0.04)	0.01 (0.04)	−0.03 (0.06)
Slope	−0.01 (0.02)	−0.05 (0.04)	−0.11* (0.06)	−0.16** (0.08)	−0.21** (0.10)	−0.32*** (0.12)
Curvature	−0.02 (0.10)	0.11 (0.22)	0.27 (0.28)	0.28 (0.38)	0.31 (0.50)	0.24 (0.49)
Constant	−0.07 (0.05)	−0.18 (0.13)	−0.28 (0.20)	−0.35 (0.25)	−0.45 (0.33)	−0.33 (0.49)
R ²	0.04	0.05	0.06	0.08	0.10	0.18
(B) <i>Federal funds rate</i>	Current	1Q ahead	2Q ahead	3Q ahead	4Q ahead	5Q ahead
ISK	0.07** (0.03)	0.33*** (0.12)	0.61*** (0.21)	0.87*** (0.31)	1.10** (0.46)	1.36*** (0.51)
Constant	−0.05*** (0.01)	−0.18*** (0.05)	−0.37*** (0.12)	−0.58*** (0.19)	−0.81*** (0.26)	−1.01*** (0.35)
R ²	0.03	0.07	0.08	0.08	0.08	0.11
ISK	0.11*** (0.04)	0.45*** (0.12)	0.86*** (0.20)	1.31*** (0.24)	1.71*** (0.30)	1.98*** (0.40)
Level	0.001 (0.003)	−0.02 (0.02)	−0.06 (0.04)	−0.09 (0.07)	−0.14 (0.09)	−0.15 (0.16)
Slope	−0.01 (0.01)	−0.04 (0.03)	−0.10 (0.07)	−0.15 (0.11)	−0.20 (0.14)	−0.25 (0.16)
Curvature	−0.14** (0.07)	−0.44** (0.17)	−0.70* (0.36)	−1.15** (0.53)	−1.61*** (0.58)	−1.75*** (0.63)
Constant	0.03 (0.04)	0.13 (0.10)	0.28 (0.21)	0.50 (0.32)	0.70* (0.42)	0.59 (0.53)
R ²	0.08	0.14	0.16	0.20	0.21	0.22
Observations	372	372	371	368	365	279

Predictive regressions for Blue Chip forecast errors for the ten-year Treasury yield (panel A) and the federal funds rate (panel B), for monthly surveys from January 1990 to January 2021. Horizons are quarterly from 0 to 5. *ISK* is a five-day average of option-implied yield skewness, and *Level*, *Slope* and *Curvature* are the first three principal components of Treasury yields, all measured on the day of the survey deadline. Hansen-Hodrick standard errors with $3(h + 1)$ lags are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

tant for short-term yields (Cieslak, 2018; Schmeling et al., 2022). The results for the federal funds rate in the bottom panel of Table 6 confirm this finding. They show that ISK is a very powerful predictor of short rate forecast errors. Even univariate regressions have R^2 ranging from seven to eleven percent for forecasts beyond the current quarter, and adding information in the yield curve raises this to 14 to 22 percent. In the multivariate regressions, the coefficient on the slope has the opposite sign as the coefficient on ISK, as usual. Curiously, the coefficient on the curvature is quite strongly significant, although this factor typically explains only a small fraction of yield variation. In any event, ISK has strong predictive power for the fed funds rate in both univariate and multivariate regressions. Our R^2 are even larger than in Cieslak (2018), who predicted fed funds rate forecast errors using the level of the funds rate and employment growth and found R^2 from three to 18 percent. Cieslak used quarterly observations and a different sample period, from 1984:Q3 to 2011:Q3. We cannot start before 1990 when our ISK series starts, but using a quarterly sample that ends in 2011:Q3 our R^2 range from 16 to 22 percent for univariate predictions and from 24 to 30 percent for predictions including the yield curve.¹³

Under the FIRE hypothesis, forecast errors should be unpredictable using information that was publicly available at the time the forecast was made. Our results indicate that this hypothesis is unlikely to hold for interest rate forecasts, and that the predictive evidence in Sections 3.1–3.2 is likely driven, as least in part, by a relationship between conditional skewness and expectational errors.

Better forecasts of monetary policy appear to be an important source of the predictive power of conditional skewness for interest rates. From Table 6 it is clear that ISK has more information about the future federal funds rate than about long-term Treasury yields.¹⁴ Similarly, we found that the predictive power of skewness for excess bond returns is larger for shorter bond maturities (see Appendix Table B.1). The predictability of monetary policy surprises around FOMC announcements also supports this view.

Finally, a large part of the predictive power of skewness seems to be due to its ability

¹³Following Coibion and Gorodnichenko (2015) we have also estimated predictive regressions that include ex-ante survey forecast revisions, which we found to have some explanatory power for fed funds forecast errors, but not for ten-year yield forecast errors (consistent with the suggestion by Coibion and Gorodnichenko that informational rigidities play a larger role for less persistent time series, like short-term rates). Importantly, including ex-ante survey revisions as additional predictors leaves our results on the information in ISK materially unchanged. Relatedly, Cieslak (2018) also found that the “predictability of FFR forecast errors cannot be explained with information rigidities such as sticky or noisy information” (fn. 14).

¹⁴Similar analysis for yields of shorter maturities (1y, 2y, 5y), omitted for the sake of brevity, shows that the predictive power of ISK systematically increases when the maturity shortens.

to capture downside risk and negative surprises in interest rates. Similar to Figure 4 for FOMC surprises, we found the most influential observations in the regressions for Table 6 are typically negative survey forecast errors.¹⁵ To put this finding into perspective, we note that the majority—between 60 and 80 percent—of interest rate forecast errors are negative, and the average forecast error is negative as well.

Overall, our evidence suggests that ISK has strong predictive power for interest rate forecast errors, in particular for short-term interest rates. This predictive power is mainly due to its ability to capture downside risk and anticipate declines in short-term rates and dovish monetary policy surprises by the Fed.

3.4 Subjective risk premia vs. expectational errors

Since we have documented predictability of interest rate forecast errors, the logical next question is whether this is quantitatively important for measured bond risk premia. Specifically, to which extent is the variation in statistical bond risk premia driven by changes in subjective risk premia and shifting bias in beliefs?

The expected h -period log excess return on a bond with maturity n , $E_t r_{t,t+h} = E_t p_{t+h}^{(n-h)} - p_t^{(n)} - y_t^{(h)}$, can be written as

$$E_t r_{t,t+h}^{(n)} = E_t^s r_{t,t+h}^{(n)} + \underbrace{(n-h) \left(E_t^s y_{t+h}^{(n-h)} - E_t y_{t+h}^{(n-h)} \right)}_{ee_t},$$

where the first term on the right hand side corresponds to subjective risk premia, and the second term to expectational errors, denoted by ee_t for short (see also Piazzesi et al., 2015). This decomposition corresponds to a log version of equation (1). To quantify the importance of these two terms, we use the following variance decomposition:

$$Var(E_t r_{t,t+h}^{(n)}) = Cov(E_t^s r_{t,t+h}^{(n)}, E_t r_{t,t+h}^{(n)}) + Cov(ee_t, E_t r_{t,t+h}^{(n)}).$$

For our empirical implementation of this decomposition, we use quarterly data, one-quarter holding periods, and one-quarter-ahead yield expectations from the BCFF consensus forecast, so t indexes quarters and $h = 1$. The maturities n we consider are 5, 9, 21 and 41

¹⁵For example, in regressions for fed funds forecast errors—using non-overlapping observations—seven to nine out of the ten most influential observations have negative forecast errors, depending on the horizon.

quarters, so that we can use the survey forecasts of 1, 2, 5 and 10 year yields for $E_t^s y_{t+1}^{(n-1)}$. Subjective expected excess returns are calculated using survey forecasts as $E_t^s(rx_{t+1}^{(n)}) = -(n-1)E_t^s(y_{t+1}^{(n-1)}) + ny_t^{(n)} - y_t^{(1)}$.¹⁶ To estimate statistical expectations $E_t y_{t+1}^{(n-1)}$, which are used in the calculations both of expectational errors ee_t and of statistical bond risk premia $E_t rx_{t+1}^{(n)}$, we estimate predictive regressions with information available at time t . Our interest naturally lies in the role played by ISK and we therefore compare results for regression specifications with time- t yield curve information that either include or exclude this additional predictor.

Table 7: Risk premia and expectational errors

Mat.	Yields only			With ISK		
	$Var(E_t^{(n)} rx_{t+1}^{(n)})$	RP (%)	EE (%)	$Var(E_t rx_{t+1}^{(n)})$	RP (%)	EE (%)
1y	0.022	91	9	0.035	40	60
2y	0.074	80	20	0.117	34	66
5y	0.469	88	12	0.718	45	55
10y	1.510	112	-12	2.324	72	28

Variance of statistical bond risk premia, i.e., of expected excess returns $E_t rx_{t+1}^{(n)}$, and relative contributions, in percent, of survey-based subjective risk premia (RP), $E_t^s rx_{t+1}^{(n)}$, and expectational errors (EE), $ee_t = (n-1) \left(E_t^s y_{t+h}^{(n-h)} - E_t y_{t+h}^{(n-h)} \right)$. Relative contributions are calculated as ratios of covariances to the variance of expected excess returns. In the “yields only” case, statistical expectations of future yields, $E_t y_{t+1}^{(n-1)}$, are calculated using predictive regressions with only time- t yields, whereas under “with ISK” we also include implied yield skewness as a predictor. Data are quarterly and the holding period is one quarter.

Table 7 shows the results of our variance decomposition. When only information in current yields is used to calculate statistical expectations and risk premia, it would appear that expectational errors account for very little of the variation in bond risk premia. The largest contribution is 20 percent, estimated for two-year yields (i.e., for expected excess returns on bonds with initial maturity of nine quarters). However, when ISK is added to the information set, expectational errors contribute a much larger share of the variation: The fraction of the variance of statistical risk premia explained by expectational errors is larger than 50% for maturities up to five years. In other words, more than half of the variation in these expected excess returns is explained by shifting biases in beliefs. To uncover this important role for expectational errors, it is crucial to condition on implied skewness of interest rates.


¹⁶This calculation ignores the fact that excess returns are based on zero-coupon yields while survey forecasts pertain to (constant-maturity) yields on coupon bonds, but the two are highly correlated in the data. Nagel and Xu (2022) use a different methodology in which they calculate subjective expected returns from survey-implied zero-coupon yields.

3.5 Skewness during the COVID-19 pandemic

As we noted earlier, skewness has reached all-time high values in the wake of the COVID-19 pandemic. A natural question is whether the unusual circumstances have disrupted the properties of skewness established in this paper. The answer is no. In fact, the COVID-19 episode serves as a helpful illustration of our findings.

Figure 5 displays ISK, the ten-year Treasury yield and its survey forecasts from BCFF, and the slope of the yield curve, measured here as the difference between the ten-year and three-month Treasury yields. We see that skewness was negative throughout 2019, and, in fact, sharply dropped to -1.5 in early 2020 as the pandemic was taking hold. After that, coincident with aggressive monetary and fiscal stimulus, it started climbing back and ultimately reached historically high values around 1.0 in the second half of 2020. Was skewness helpful in predicting ten-year yields during this period? Was it related to expectational errors of forecast surveys?

Figure 5: Skewness and interest rates since 2019



Ten-year Treasury yield, yield forecasts from Blue Chip Financial Forecasts, option-implied yield skewness, and slope of the yield curve (measured as difference between ten-year and three-month yield). Sample period: January 2, 2019, to May 28, 2021.

Early 2019 and late 2020 are two episodes where the slope was close to zero in both cases, predicting low bond returns and rising interest rates. But the level of skewness differed substantially, negative during the first and positive during the second episode. In 2019 the signal from the slope turned out to be incorrect, as yields dropped precipitously. This was correctly anticipated by implied skewness. In 2020 the prediction of the flat yield curve for rising long-term rates turned out to be correct, but the slope was essentially unchanged over most of the year, so that it was of little use as a timely indicator of interest rate risk. Skewness, by contrast, all of a sudden rose substantially in the middle of the year, correctly anticipating the rising long-term yields. Both of these episodes highlight the extra information in skewness that is not present in the current yield curve.

Large swings in skewness during this period indicate large expectational errors. Consistent with our regression results in Table 6, forecasters were overshooting yields in the beginning of the pandemic and then undershooting later in this episode. In fact, in late 2020 expectational errors were large as Treasury yields began a sustained ascent from historical lows (from 0.5% to 1.5%). At the time market observers were surprised by the development. Again, this was correctly predicted by skewness, which started rising in advance of the rise in yields.

The COVID-19 episode was unique in many respects including extreme rate volatility. However, the information content of conditional skewness remained intact, and correctly anticipated both the dramatic decline in long-term Treasury yields in 2019 and early 2020, as well as their pronounced increase starting in the middle of the pandemic.

4 A potential explanation: heterogeneous beliefs

We have documented that conditional skewness of Treasury yields varies over time and predicts future interest rates and bond returns. The predictive power is particularly strong when interest rates surprise on the downside—during monetary easing cycles and recessions—and in regressions with yield curve factors, in violation of the spanning hypothesis. Skewness predicts survey forecast errors for interest rates, and such expectational errors are quantitatively important for statistical bond risk premia, in violation of the FIRE hypothesis.

A natural question is whether all this evidence is consistent with an economic mechanism. We use a standard two-agent heterogeneous-beliefs model to show that this is indeed the case. Like [Basak \(2005\)](#) we consider a continuous-time economy featuring agents with CRRA preferences and an unknown mean of the aggregate endowment. We derive bond pricing

implications following [Xiong and Yan \(2010\)](#) and [Ehling et al. \(2018\)](#).¹⁷ Going beyond these earlier studies, we show that time-varying interest rate skewness arises endogenously, and we analyze its relation to bond risk premia and forecast errors. Disagreement is the key state variable and it drives the dynamics of all the variables that we have analyzed empirically: skewness, bond returns, and agents' expectational errors.

In what follows we describe the main assumptions of the model and key results, with details and derivations provided in [Appendix C.1](#). Consumption is exogenous and follows a geometric Brownian motion,

$$dC_t/C_t = \mu dt + \sigma dz_t. \quad (4)$$

As in [Basak \(2005\)](#), the two agents in the economy disagree about the true dynamics of consumption growth, specifically about the mean growth rate μ . To connect the model more closely to our evidence on expectational errors, we assume that agent 1 knows the true value of μ while agent 2 does not know it and has to form beliefs μ_t^s . Our interpretation is that agent 1 represents true, statistical expectations, while agent 2 represents (consensus) survey forecasts which deviate from FIRE. We assume that subjective beliefs follow a linear process

$$d\mu_t^s = \kappa(\mu - \mu_t^s)dt + \delta dz_t. \quad (5)$$

Such a specification encompasses many mechanisms of beliefs formation explored in the literature, including Bayesian learning (e.g., [Basak, 2005](#), in which case κ and δ are deterministic functions of time), sentiment (e.g., [Dumas et al., 2009](#)), diagnostic expectations (e.g., [Bordalo et al., 2018](#)), and other forms of extrapolative expectations. The exact form of belief formation is not important for our results, so we do not take a stand on how the beliefs of agent 2 deviate from FIRE.

The key state variable of the model is disagreement, which is measured as

$$\Delta_t \equiv \frac{\mu - \mu_t^s}{\sigma}. \quad (6)$$

The Gaussian dynamics of μ_t^s translate into Gaussian dynamics of Δ_t . Because agent 1 has full information, disagreement Δ_t is directly linked to the belief bias of agent 2, $\mu_t^s - \mu$.

We assume that agents have power utility with risk aversion γ . [Appendix C.1](#) derives the

¹⁷[Xiong and Yan \(2010\)](#) is a special case of logarithmic preferences. Both [Xiong and Yan \(2010\)](#) and [Ehling et al. \(2018\)](#) focus on unknown expected inflation, but we follow [Basak \(2005\)](#) who considers unknown mean consumption growth.

equilibrium consumption allocations and solves for the real interest rate, bond yields, and risk premia. The expression for the real short rate provides some intuition:

$$r_t = \rho + \gamma [f(\lambda_t)\mu + (1 - f(\lambda_t))\mu_t^s] - \frac{1}{2}\gamma(\gamma + 1)\sigma^2 + \frac{1}{2}(1 - \gamma^{-1})f(\lambda_t)(1 - f(\lambda_t))\Delta_t^2. \quad (7)$$

The short rate is driven by (i) the rate of time preference, (ii) consumption smoothing with consumption-share-weighted average perceived consumption growth (instead of true expected growth), (iii) a standard precautionary savings effect, and (iv) an “investor disagreement” effect, arising whenever $\gamma \neq 1$.¹⁸ Terms (ii) and (iv) involve the likelihood ratio of objective and subjective probabilities λ_t and the nonlinear consumption sharing rule $f(\lambda)$, both of which are defined in the appendix. In the case of log-utility, $\gamma = 1$, the investor-disagreement term disappears from equation (7) and r_t is solely driven by consumption-weighted average beliefs of the two agents in the economy. When there is no disagreement and the two agents are symmetric ($\Delta_0 = 0, \lambda_0 = 1$), the risks that one of the agents gets a larger consumption share balance out and skewness is zero; once an agent becomes more dominant, the interest rate dynamics become skewed. Simulation results in Appendix C.2 confirm that for log-utility skewness is zero when $\Delta_0 = 0$. Moving away from log-preferences, the investor disagreement effect comes into play. In this case, risks do not balance out and skewness is nonzero even when $\Delta_0 = 0$. In general, the distribution of r_t is non-Gaussian, and its conditional skewness is related to Δ_0 and expectational errors.

The model assumes biased beliefs about expected endowment growth and links them to interest rate skewness and bond risk premia. Before evaluating whether the model can capture our evidence, we check whether its premise is in fact supported by the data. Table 8 shows predictive regressions for GDP forecast errors, calculated as the difference between four-quarter real GDP growth and the consensus (median) expectation in the Survey of Professional Forecasters.¹⁹ These forecast errors are predictable with ISK, suggesting that expectational errors about macroeconomic fundamentals, and real growth in particular, are related to conditional yield skewness.

This evidence represents a sharp departure from the findings of [Giacoletti et al. \(2021\)](#), which is motivated by the heterogeneous beliefs framework as well. Their measure, which predicts

¹⁸Using equation (6), term (ii) can be rewritten as $\gamma[\mu - \sigma(1 - f(\lambda_t))\Delta_t]$, which shows that for $\gamma \neq 1$, r_t is a linear-quadratic function of Δ_t .

¹⁹The predictors are measured on the day of the survey deadline. For ISK we use a five-day average, as in our analysis of interest rate forecast errors in Section 3.3. Following [Coibion and Gorodnichenko \(2015\)](#) we use real-time data (from ALFRED) as estimated in the vintage one year after the period being predicted.

bond excess returns together with ISK in Table 4, reflects the intra survey disagreement about Treasury yields. In the model such disagreement can be captured by two agents, both of which are biased about the GDP prospects. Thus, under the null of the model, the [Giacoletti et al. \(2021\)](#) measure should be predictive about intra survey disagreement about future GDP. [Giacoletti et al. \(2021\)](#) and [Singleton \(2021\)](#) find no such predictive relation, which leads them to conclude that the heterogeneous beliefs framework cannot capture their evidence. We follow [Reis \(2021\)](#), who posits that bond traders are better informed than the general public proxied by surveys, and reinterpret the same model as disagreement between one correct and one biased agent. That perspective qualitatively reconciles our evidence about bond and GDP predictability.

While we cannot directly observe $\mu - \mu_t^s$ or Δ_t in the data, we can consider sample statistics of GDP forecast errors. Both the mean and median are negative, and close to 60% of forecast errors are negative. In addition, the predictive power of ISK in Table 8 largely stems from its ability to capture downside risks, similar to our findings for bond returns, FOMC surprises, and interest rate forecast errors.²⁰ To summarize, ISK has a positive relation with future forecast errors, and this relation is primarily driven by negative forecast errors, which represent the majority of all observations.

We use simulations to study the model’s implications and its ability to match our empirical evidence. In our calibration, we choose annualized volatility of consumption growth to be 3.5%, $\sigma = 0.035$. The persistence of biased beliefs, captured by the mean-reversion coefficient, is $\kappa = 0.1$. We assume that $\delta = \kappa\sigma$, corresponding to constant-gain learning about μ . We select the smallest value of risk aversion, which allows for the investor disagreement effect, that is, we set $\gamma = 2$.

We simulate our economy over a range for the initial disagreement, or belief bias, Δ_0 . For each value of Δ_0 we simulate 100,000 paths for Δ_t , λ_t , spot interest rates and bond prices over a period of one quarter. We use sample moments across these paths to obtain conditional skewness of the ten-year bond yield, expected excess returns for the ten-year bond, and mean forecast errors for real GDP growth and the short rate. Appendix C.2 provides details of our simulation setup.

As anticipated by equation (7), interest rate skewness depends on Δ_0 non-linearly. Appendix

²⁰Specifically, in a regression of GDP growth forecast errors on residual ISK (orthogonalized with respect to level, slope and curvature of the yield curve), the top six (and eight out of the ten) most influential observations are negative forecast errors.

Table 8: Predicting real GDP forecast errors

	(1)	(2)	(3)
Level	0.06 (0.12)	0.05 (0.12)	
Slope	0.19 (0.22)	0.06 (0.16)	
Curvature	-1.15 (0.87)	-1.43* (0.76)	
ISK		1.84** (0.89)	1.80* (1.02)
Constant	-0.39 (0.86)	-0.16 (0.60)	-0.38 (0.31)
Observations	123	123	123
R ²	0.03	0.09	0.06

Regressions of forecast errors for four-quarter real GDP growth, based on real-time GDP data and the Survey of Professional Forecasters from January 1990 to May 2021. *ISK* is option-implied yield skewness, averaged over five trading days; *Level*, *Slope* and *Curvature* are the first three principal components of Treasury yields; all predictors are measured on the day of the survey deadline. Hansen-Hodrick standard errors with 4 lags are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

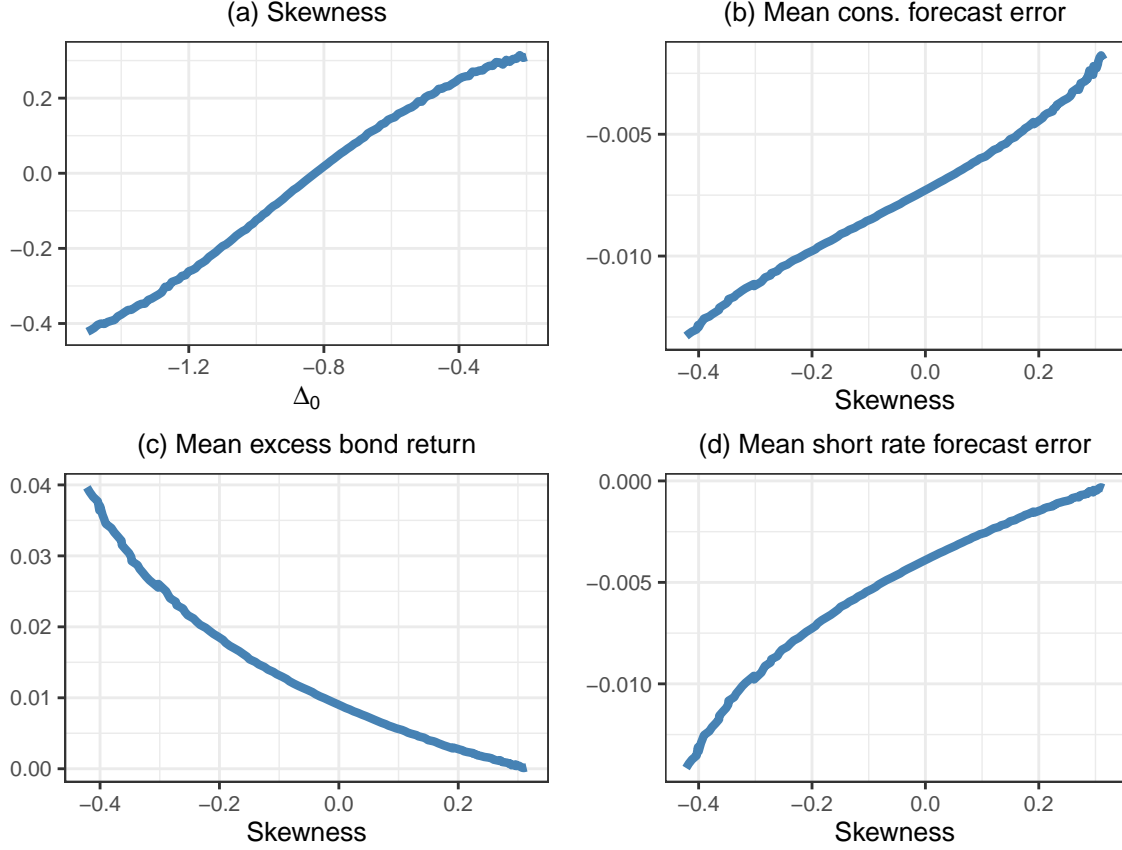
Figure C.1 shows that skewness increases with Δ_0 to a peak of about 0.3 for Δ_0 close to zero, and then declines. As a result, the signs of the key model-implied relationships generally depend on the sign of Δ_0 . The evidence on GDP forecast errors suggests that $\Delta_0 < 0$ (an overly optimistic consensus forecast/agent 2) is the empirically more relevant case, so this is the case we focus on.²¹

Figure 6 plots the simulation results. Panel (a) displays one-quarter conditional skewness of the ten-year yield, similar to the ISK measure computed in the data. We see that conditional skewness depends on current beliefs and disagreement, as captured by Δ_0 . Since beliefs vary over time, skewness does as well. The range of values for conditional skewness is roughly in line with the values for ISK shown in Figures 1B and 2.

The model has implications for the relationships of various objects with disagreement, measured by the current value of the state variable, Δ_0 , while our empirical analysis focuses on the relationships with conditional skewness. Because of the one-to-one relationship shown

²¹Appendix C.2 discusses the alternative case for completeness. It also plots conditional skewness for the entire range of Δ_0 , and for alternative choices of γ , including log-utility.

Figure 6: Yield skewness in the heterogeneous beliefs model



Simulation results for heterogeneous-beliefs model. For each starting value of disagreement, Δ_0 , we simulate 100,000 paths over a period of one quarter, using $\gamma = 2$, $\sigma = 0.035$, and $\kappa = 0.1$. Panel (a) shows the sample skewness for the distribution of the ten-year yield. Panels (b)-(d) plot mean forecast errors for consumption, expected excess returns on the ten-year bond, and mean forecast errors for the real short rate against skewness of the ten-year yield.

in panel (a) of Figure 6, we can use skewness as a “sufficient statistic” for Δ_0 .

Panel (b) shows the relationship between skewness and mean consumption forecast errors, which are obtained as the difference between conditional expectations of future consumption under the true/objective probability measure and under the subjective measure, representing survey forecasts. The relationship is positive, consistent with the coefficients on ISK in Table 8. The sign of the model-implied forecast errors is negative, due to the choice of a negative range for Δ_0 , consistent with most observations and all of the most influential observations in the regressions for GDP growth forecast errors.

Panel (c) displays the relationship between skewness and the mean excess return on the ten-

year bond. Expected excess returns decline as conditional skewness rises, consistent with the evidence in Section 3.1. Over the range we focus on, model-implied expected excess bond returns are positive. In the data, the mean and median excess return, 65% of the observations and all of the most influential observations are positive as well.²²

In this model, bond risk premia depend on disagreement non-linearly, as shown in Appendix C.1. In addition, yields are non-linear functions of disagreement. As a consequence, linear combinations of yields cannot span disagreement or bond excess returns. In other words, the spanning hypothesis is violated in this model, and skewness captures unspanned risks. This feature of the model is consistent with the empirical finding in Table 4 that adding skewness to yield curve variables substantially increases predictive power for bond returns, in violation of the spanning hypothesis.

Panel (d) relates skewness to the mean forecast error for the real short rate. This expectational error is calculated as the difference between the conditional mean short rate under the objective and subjective probability measures. The calibrated model implies a positive relationship, consistent with our evidence in Section 3.3. Over the range we focus on, mean interest rate forecast errors are negative. In the data, negative rate forecast errors are more prevalent and empirically most important. For the ten-year yield and the federal funds rate, mean and median forecast errors, the majority of observations (65-90%, depending on the series and forecast horizon) and the most influential observations are negative forecast errors.

Overall, this essentially “off-the-shelf” theoretical framework is consistent with our empirical evidence, and provides a coherent explanation for the role of conditional skewness.²³ We have shown that for broadly plausible parameter choices and ranges, the model can generate patterns that are similar to the empirical findings. Nevertheless, some caveats should be acknowledged. First, we have only shown that the model properties are qualitatively consistent with the data, without any quantification of the effects. Second, the relationships depend on the model calibration and the range of the initial disagreement, Δ_0 . Finally, while our evidence is based on nominal interest rates and FOMC surprises, we use a real model without any explicit role for inflation or monetary policy.

²²Analogously, mean and median FOMC surprises are negative, and all of the top 10 most influential FOMC announcements have dovish surprises. FOMC surprises are essentially negative excess bond returns. See the discussion in Section 3.2.

²³Other explanations for the predictability of short-rate forecast errors and monetary policy surprises involve incomplete information about the Fed’s reaction to economic conditions (Cieslak, 2018; Bauer and Swanson, 2022a). Our model does not contradict those explanations, as erroneous beliefs about fundamentals, as in our model, naturally translate into erroneous beliefs about monetary policy (e.g., via a policy rule).

5 Conclusion

Our paper makes three contributions to the macro-finance literature. First, we document novel empirical patterns for the conditional skewness of Treasury yields, including a tight empirical relationship between conditional skewness and the shape of the yield curve, the business cycle, and the stance of monetary policy. Second, we show that option-based yield skewness contains useful forward-looking information for interest rates, including predictive power for survey forecast errors. The evidence suggests that conditional skewness captures biased beliefs about future interest rates. Third, we argue that our empirical findings can be rationalized by a standard theoretical framework with heterogeneous beliefs.

Our results have implications for asset pricing, macroeconomic forecasting, and investment practice. Forecasters and investors can benefit from implied yield skewness as a helpful additional indicator of the balance of interest rate risk, which is available in real time, can be calculate at a daily frequency, and does not require any model or estimation. Our evidence also raises additional questions and suggests other avenues for investigation of the role of interest rate skewness. For example, skewness calculated from money market futures and options could provide important forward-looking information, because of the central role of monetary policy in determining the balance of interest rate risk. And going beyond individual interest rate maturities, an entire term structure of skewness could be constructed. More generally, higher-order moments derived from interest rate derivatives are a rich source of information, with much promise for a better understanding of risks and beliefs in financial markets.

References

- Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone (2019) “Vulnerable Growth,” *American Economic Review*, 109 (4), 1263–89.
- Ahn, Dong-Hyun, Robert F. Dittmar, and A. Ronald Gallant (2002) “Quadratic Term Structure Models: Theory and Evidence,” *Review of Financial Studies*, 15 (1), 243–288.
- Ang, Andrew and Monika Piazzesi (2003) “A No-Arbitrage Vector Regression of Term Structure Dynamics with Macroeconomic and Latent Variables,” *Journal of Monetary Economics*, 50 (4), 745–787.

- Bacchetta, Philippe, Elmar Mertens, and Eric van Wincoop (2009) “Predictability in financial markets: What do survey expectations tell us?” *Journal of International Money and Finance*, 28 (1), 406–426.
- Bakshi, Gurdip and Dilip Madan (2000) “Spanning and derivative-security valuation,” *Journal of Financial Economics*, 55 (2), 205–238.
- Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer (2015) “X-CAPM: An extrapolative capital asset pricing model,” *Journal of Financial Economics*, 115 (1), 1–24.
- Barro, Robert (2006) “Rare disasters and asset markets in the twentieth century,” *Quarterly Journal of Economics*, 121 (3), 823–867.
- Basak, Suleyman (2005) “Asset pricing with heterogenous beliefs,” *Journal of Banking and Finance*, 29 (11), 2849–2881.
- Bauer, Michael D. and James D. Hamilton (2018) “Robust Bond Risk Premia,” *Review of Financial Studies*, 31 (2), 399–448.
- Bauer, Michael D. and Glenn D. Rudebusch (2017) “Resolving the Spanning Puzzle in Macro-Finance Term Structure Models,” *Review of Finance*, 21 (2), 511–553.
- (2020) “Interest Rates Under Falling Stars,” *American Economic Review*, 110 (5), 1316–1354.
- Bauer, Michael D. and Eric T. Swanson (2022a) “An Alternative Explanation for the ‘Fed Information Effect’,” forthcoming, *American Economic Review*.
- (2022b) “A Reassessment of Monetary Policy Surprises and High-Frequency Identification,” *NBER Macroeconomics Annual*, 37, forthcoming.
- Beber, Alessandro and Michael W. Brandt (2006) “The Effect of Macroeconomic News on Beliefs and Preferences: Evidence from the Options Market,” *Journal of Monetary Economics*, 53 (8), 1997–2039.
- Bikbov, Ruslan and Mikhail Chernov (2009) “Unspanned stochastic volatility in affine models: evidence from eurodollar futures and options,” *Management Science*, 55 (8), 1292–1305.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018) “Diagnostic expectations and credit cycles,” *Journal of Finance*, 73 (1), 199–227.
- Buraschi, Andrea, Ilaria Piatti, and Paul Whelan (2022) “Subjective Bond Returns and Belief Aggregation,” *Review of Financial Studies*, 35 (8), 3710–3741.
- Buraschi, Andrea and Paul Whelan (2022) “Speculation, Sentiment, and Interest Rates,” *Management Science*, 68 (3), 2308–2329.

- Campbell, Jeffrey, Charles Evans, Jonas Fisher, and Alejandro Justiniano (2012) “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Brookings Papers on Economic Activity*, 1–54.
- Campbell, John Y. and Robert J. Shiller (1991) “Yield Spreads and Interest Rate Movements: A Bird’s Eye View,” *Review of Economic Studies*, 58 (3), 495–514.
- Cao, Shuo, Richard K. Crump, Stefano Eusepi, and Emanuel Moench (2021) “Fundamental disagreement about monetary policy and the term structure of interest rates,” Staff Report 934, Federal Reserve Bank of New York.
- Choi, Hoyong, Philippe Mueller, and Andrea Vedolin (2017) “Bond Variance Risk Premiums,” *Review of Finance*, 21 (3), 987–1022.
- Cieslak, Anna (2018) “Short-Rate Expectations and Unexpected Returns in Treasury Bonds,” *Review of Financial Studies*, 31 (9), 3265–3306.
- Cieslak, Anna and Pavol Povala (2015) “Expected Returns in Treasury Bonds,” *Review of Financial Studies*, 28 (10), 2859–2901.
- Cieslak, Anna and Andreas Schrimpf (2019) “Non-Monetary News in Central Bank Communication,” *Journal of International Economics*, 118, 293–315.
- Cieslak, Anna and Annette Vissing-Jorgensen (2021) “The Economics of the Fed Put,” *Review of Financial Studies*, 34 (9), 4045–4089.
- Cochrane, John H. and Monika Piazzesi (2005) “Bond Risk Premia,” *American Economic Review*, 95 (1), 138–160.
- Coibion, Olivier and Yuriy Gorodnichenko (2015) “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 105 (1), 2644–2678.
- Conrad, Jennifer, Robert Dittmar, and Eric Ghysels (2013) “Ex Ante Skewness and Expected Stock Returns,” *Journal of Finance*, 68 (1), 85–124.
- Crump, Richard K., Stefano Eusepi, Emanuel Moench, and Bruce Preston (2021) “The Term Structure of Expectations,” Staff Report 992, Federal Reserve Bank of New York.
- Crump, Richard K. and Nikolay Gospodinov (2019) “Deconstructing the Yield Curve,” Staff Report 884, Federal Reserve Bank of New York.
- Detemple, Jerome and Shashidhar Murthy (1994) “Intertemporal asset pricing with heterogeneous beliefs,” *Journal of Economic Theory*, 62 (2), 294–320.
- Diercks, Anthony M., Hiroatsu Tanaka, and Paul Cordova (2022) “Asymmetric Monetary Policy Expectations,” September, unpublished working paper.

- Duffee, Gregory R. (2011) “Information In (and Not In) the Term Structure,” *Review of Financial Studies*, 24 (9), 2895–2934.
- Dumas, Bernard, Alexander Kurshev, and Raman Uppal (2009) “Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility,” *Journal of Finance*, 64 (2), 579–629.
- Durham, J. Benson (2008) “Implied Interest Rate Skew, Term Premiums, and the ”Conundrum”,” *Journal of Fixed Income*, 17 (4), 88–99.
- 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.
- Eriksson, Anders, Eric Ghysels, and Fangfang Wang (2009) “The normal inverse Gaussian distribution and the pricing of derivatives,” *Journal of Derivatives*, 16 (3), 23–37.
- Fama, Eugene F. and Robert R. Bliss (1987) “The Information in Long-Maturity Forward Rates,” *American Economic Review*, 77 (4), 680–692.
- Fuster, Andreas, David Laibson, and Brock Mendel (2010) “Natural expectations and macroeconomic fluctuations,” *Journal of Economic Perspectives*, 24 (4), 67–84.
- Gertler, Mark and Peter Karadi (2015) “Monetary policy surprises, credit costs, and economic activity,” *American Economic Journal: Macroeconomics*, 7 (1), 44–76.
- Giacoletti, Marco, Kristoffer Laursen, and Kenneth Singleton (2021) “Learning From Disagreement in the U.S. Treasury Bond Market,” *Journal of Finance*, 76 (1), 395–441.
- Gürkaynak, Refet S., Brian P. Sack, and Eric T. Swanson (2005) “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 1 (1), 55–93.
- Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright (2007) “The U.S. Treasury Yield Curve: 1961 to the Present,” *Journal of Monetary Economics*, 54 (8), 2291–2304.
- Gürkaynak, Refet S. and Jonathan H. Wright (2012) “Macroeconomics and the term structure,” *Journal of Economic Literature*, 50 (2), 331–67.
- Harrison, J. Michael and David Kreps (1978) “Speculative investor behavior in a stock market with heterogenous expectations,” *Quarterly Journal of Economics*, 92 (2), 323–336.
- Hattori, Masazumi, Andreas Schrimpf, and Vladyslav Sushko (2016) “The Response of Tail Risk Perceptions to Unconventional Monetary Policy,” *American Economic Journal: Macroeconomics*, 8 (2), 111–136.
- Hodrick, Robert J. (1992) “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement,” *Review of Financial Studies*, 5 (3), 357–386.

- Inoue, Atsushi and Lutz Kilian (2005) “In-sample or out-of-sample tests of predictability: Which one should we use?” *Econometric Reviews*, 23 (4), 371–402.
- Jiang, George and Yisong Tian (2005) “Model-Free Implied Volatility and Its Information Content,” *Review of Financial Studies*, 18 (4), 1305–1342.
- Johannes, Michael (2004) “The statistical and economic role of jumps in continuous-time interest rate models,” *Journal of Finance*, 59 (1), 227–260.
- Kuttner, Kenneth N. (2001) “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market,” *Journal of Monetary Economics*, 47 (3), 523–544.
- Laubach, Thomas and John C. Williams (2003) “Measuring the Natural Rate of Interest,” *Review of Economics and Statistics*, 85, 1063–1070.
- Li, Xinyang (2021) “Bond Implied Risks Around Macroeconomic Announcements,” October, Working Paper.
- Ludvigson, Sydney and Serena Ng (2009) “Macro Factors in Bond Risk Premia,” *Review of Financial Studies*, 22 (12), 5027–5067.
- Mertens, Thomas M. and John C. Williams (2021) “What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices,” *American Economic Review*, 111 (8), 2473–2505.
- Nagel, Stefan and Zhengyang Xu (2022) “Dynamics of Subjective Risk Premia,” Working Paper 29803, National Bureau of Economic Research.
- Nakamura, Emi and Jón Steinsson (2018) “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect,” *Quarterly Journal of Economics*, 133 (3), 1283–1330.
- Neuberger, Anthony (2012) “Realized Skewness,” *Review of Financial Studies*, 25 (11), 3423–3455.
- Piazzesi, Monika (2001) “An econometric model of the yield curve with macroeconomic jump effects,” Working Paper w8246, National Bureau of Economic Research.
- (2005) “Bond Yields and the Federal Reserve,” *Journal of Political Economy*, 113 (2), 311–344.
- Piazzesi, Monika, Juliana Salomao, and Martin Schneider (2015) “Trend and Cycle in Bond Premia,” March, unpublished working paper.
- Reis, Ricardo (2021) “The People versus the Markets: A Parsimonious Model of Inflation Expectations,” Discussion Paper 15624, CEPR.
- Rudebusch, Glenn D. and Tao Wu (2008) “A Macro-Finance Model of the Term Structure, Monetary Policy, and the Economy,” *Economic Journal*, 118 (530), 906–926.

- Schmeling, Maik, Andreas Schrimpf, and Sigurd Steffensen (2022) “Monetary Policy Expectation Errors,” *Journal of Financial Economics*, 146 (3), 841–858.
- Singleton, Kenneth J. (2021) “Presidential Address: How Much ‘Rationality’ Is There in Bond-Market Risk Premiums?” *Journal of Finance*, 76 (4), 1611–1654.
- Swanson, Eric T. and John C. Williams (2014) “Measuring the effect of the zero lower bound on medium-and longer-term interest rates,” *American Economic Review*, 104 (10), 3154–85.
- Trolle, Anders B. and Eduardo S. Schwartz (2014) “The Swaption Cube,” *Review of Financial Studies*, 27 (8), 2307–2353.
- Wei, Min and Jonathan H. Wright (2013) “Reverse Regressions And Long-Horizon Forecasting,” *Journal of Applied Econometrics*, 28 (3), 353–371.
- Xiong, Wei and Hongjun Yan (2010) “Heterogenous expectations and bond markets,” *Review of Financial Studies*, 23 (4), 1433–1466.

Appendix

A Option-implied moments of Treasury yields

Our Treasury derivatives data are end-of-day prices of Treasury futures and options from CME.²⁴ We focus on the ten-year T-note futures contract (or “TY”). The deliverable securities for the TY contract are “U.S. Treasury notes with a remaining term to maturity of at least six and a half years, but not more than 10 years” (according to the CME contract specifications). The contract expirations are at the end of each calendar quarter, and at each point in time three consecutive quarterly contracts are available; the exact delivery date is roughly in the third week of the month. The first quarterly contract is the most active, until about 2-3 weeks before expiration when trading in the subsequent quarterly contract becomes more active. Therefore, when working with futures prices (e.g., for calculating sample moments or realized moments of price changes), we always use the first quarterly expiration that is not in the current calendar month (e.g., we use the March contract until the end of February, and the June contract starting in the beginning of March).

The options on the TY contract are available for three quarterly and three serial (monthly) expirations, and they each exercise into the next futures contract. For example, February and March options exercise into the March futures contract, and April options exercise into June futures contract. The last trading day for each options contract is the “2nd last business day of the month prior to the contract month” so that trading for the March options ends at the end of February. We denote by t the current trading day and by T the last trading day (or expiration date) of an options contract. For most of our analysis we focus on the first quarterly option expiration. For Figure 1B we linearly interpolate option-implied moments to a constant horizon of 0.2 years (about the average maturity across all expirations, strikes and put/call prices).

Based on option prices on day t for the contract expiration T we can calculate conditional market-based/risk-neutral moments for the price of the underlying futures contract at the time of the option expiration, F_T . The implied risk-neutral variance is

$$\begin{aligned} Var_t F_T &= E_t(F_T - F_t)^2 = 2 \left[\int_{F_t}^{\infty} C(K) dK + \int_0^{F_t} P(K) dK \right] \\ &= 2 \int_0^{\infty} C(K) - \max(0, F_T - K) dK \end{aligned}$$

where all moments are under the time- T forward measure, we treat F_t as the forward price for simplicity, and the forward call and put prices for options with strike K are $C(K)$

²⁴For details see <https://www.cmegroup.com/trading/interest-rates/us-treasury.html>.

and $P(K)$. Because expectations are under the T -forward measure, $E_t F_T = F_t$, $C(K) = E_t \max(0, F_T - K)$ and $P(K) = E_t \max(0, K - F_T)$. The second line follows from put call parity, $C(K) - P(K) = F_T - K$. The implied third moment is

$$\begin{aligned} E_t(F_T - F_t)^3 &= 6 \left[\int_{F_t}^{\infty} (K - F_t)C(K)dK - \int_0^{F_t} (F_t - K)P(K)dK \right] \\ &= 6 \int_0^{\infty} (K - F_t)(C(K) - \max(0, F_T - K))dK. \end{aligned}$$

See also [Trolle and Schwartz \(2014\)](#) who use similar formulas for calculating swaption-implied moments for future swap yields. The implied skewness coefficient is

$$skew_{t,T}^F = \frac{E_t(F_T - F_t)^3}{(Var_t F_T)^{3/2}}$$

We now describe how we implement these measures empirically. In what follows σ is the normal implied volatility (IV) for at-the-money options. Normal IV, the most common way to measure IV in bond markets, is based on the Bachelier model and measures the volatility of future price changes under the assumption that they are Gaussian. First, we filter our options data to reduce the impact of measurement error and eliminate data errors, similar to [Beber and Brandt \(2006\)](#). Specifically, we exclude options that

- have maturity of at most two weeks
- have prices of at most two ticks (2/64)
- have moneyness greater than 15 in absolute value (moneyness is $(F_t - K)/\sqrt{(T - t)\sigma^2}$ for calls and $(K - F_t)/\sqrt{(T - t)\sigma^2}$ for puts),
- have absolute moneyness of less than -15 (the absolute moneyness is $F - K$ for calls and $K - F$ for puts),
- have distinct duplicate prices for the same strike (using the IVs and other prices we can eliminate the erroneous price by hand),
- have prices which are not monotone across strikes, or
- violate the no-arbitrage condition that the price is no lower than the intrinsic value.

Then we calculate implied moments for each pair (t, T) if we observe at least five option prices (puts and calls across all strikes) in the following way:

1. We select all option prices that are ATM/OTM

2. We calculate the normal IVs for these observed prices.
3. We fit a curve in strike-IV space by linearly interpolating IVs and, outside the range of observed prices, using IVs at the endpoints of the range.
4. We obtain a continuous price function $C(X)$ by mapping the IVs back to call prices using the Bachelier pricing formula.
5. We approximate the required integrals using trapezoid rule for grid of strike prices from $F_t - 10$ to $F_t + 10$ with 200 grid points (see also [Jiang and Tian, 2005](#)).

As a result, we have conditional model-implied variances and skewness coefficients for the change in the futures price between t and T .

With the moments for futures prices in hand we can also calculate certain moments for changes in the yields of the cheapest-to-deliver (CTD) bond. The reason is that for small changes, the relationship between changes in futures prices and changes in the CTD yield is approximately linear. The “dollar value of a basis point” (DV01) is the negative sensitivity of the futures price (in points) to a change in the CTD yield (in basis points). Denoting the change in the futures price as ΔF and the change in the CTD yield by Δy , we have

$$\Delta y \approx -\frac{\Delta F}{DV01}.$$

Under the assumption that the change in the CTD yield until expiration, $y_T - y_t$, is small, and that $DV01$ remains approximately unchanged between t and T , we can obtain risk-neutral moments for future yields as

$$Var_t y_T \approx \frac{Var_t F_T}{(DV01)^2}, \quad E_t(y_T - y_t)^3 \approx -\frac{E_t(F_T - F_t)^3}{(DV01)^3}, \quad skew_{t,T}^y \approx -skew_{t,T}^F$$

The $DV01$ data, as well as any information about the CTD bonds, becomes available on Bloomberg in 2004. But this information is not required for the skewness coefficient, since skewness of yield changes is approximately equal simply to the negative of the skewness of futures price changes.

Our derivation and implementation abstract from the fact that Treasury options are American options on futures contracts, and not, as assumed, European options on forward contracts. Existing results suggest that accounting for early exercise would lead to only minor adjustments; see [Bikbov and Chernov \(2009\)](#) and [Choi et al. \(2017\)](#). In addition, since we only use out-of-the-money options any adjustment for early exercise would be minimal, as there are no dividends and the early-exercise premium increases with the moneyness of options.

B Additional results for Section 3

B.1 Additional results for Treasury bond returns

Table B.1 shows predictive regressions for quarterly excess returns on bonds of different maturities. The predictive power is a substantially higher for shorter than for longer maturities. For example, the R^2 for the one-year maturity is about twice as large as for the ten-year maturity. In additional analysis we have found that the predictive power of univariate regressions with only ISK also decreases with maturity. The fact that ISK is more powerful for short maturities is somewhat surprising, given that the underlying securities for our implied skewness measure are futures on ten-year Treasury bonds (with maturity of the cheapest-to-deliver bonds ranging between 7 and 10 years).

Table B.1: Predicting excess returns: individual bond maturities

	1y	2y	3y	5y	7y	10y
ISK	−0.38*** (0.14)	−0.37*** (0.14)	−0.36*** (0.14)	−0.35*** (0.13)	−0.34*** (0.12)	−0.30*** (0.11)
Level	0.01*** (0.001)	0.01*** (0.002)	0.004** (0.002)	0.003** (0.002)	0.003* (0.002)	0.002 (0.002)
Slope	0.02** (0.01)	0.03** (0.01)	0.03** (0.01)	0.03*** (0.01)	0.03*** (0.01)	0.03*** (0.01)
Curvature	−0.08 (0.06)	−0.02 (0.07)	0.01 (0.07)	0.03 (0.07)	0.02 (0.07)	0.01 (0.06)
Constant	−0.25*** (0.10)	−0.18* (0.10)	−0.14 (0.10)	−0.12 (0.10)	−0.13 (0.10)	−0.14 (0.10)
Observations	374	374	374	374	374	374
R^2	0.18	0.12	0.10	0.09	0.09	0.09

Predictive regressions for three-month excess returns on Treasury bonds with different maturities, using monthly data from January 1990 to May 2021. Predictors: *ISK* is option-implied yield skewness averaged over the last five business days of the month; *Level*, *Slope* and *Curvature* are the first three principal components of end-of-month Treasury yields. Standard errors based on the reverse regression delta method of [Wei and Wright \(2013\)](#) are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

In the paper, we focus on three-month holding periods because this aligns with the horizon of the option contracts used to calculate implied skewness. The recent empirical literature on risk premia in Treasury bonds instead usually considers excess returns over annual holding periods (e.g. [Cochrane and Piazzesi, 2005](#); [Ludvigson and Ng, 2009](#); [Cieslak and Povala, 2015](#)). In Table B.2 we report additional results for one-year holding periods. As expected, the predictive power of ISK is somewhat weaker, but it remains statistically significant both

for the average excess return (across bonds with 2-10 years maturities) and individual excess returns for bonds with short and medium maturities.

Table B.2: Predicting excess returns: one-year holding periods

	avg.	2y	3y	5y	7y	10y
ISK	−0.72** (0.35)	−0.97** (0.42)	−0.90** (0.40)	−0.80** (0.37)	−0.66* (0.36)	−0.46 (0.36)
Level	0.01* (0.01)	0.02*** (0.01)	0.02** (0.01)	0.01* (0.01)	0.01 (0.01)	0.01 (0.01)
Slope	0.09* (0.05)	0.07 (0.06)	0.08 (0.06)	0.09* (0.05)	0.10** (0.05)	0.10** (0.04)
Curvature	−0.04 (0.18)	−0.20 (0.20)	−0.11 (0.20)	−0.03 (0.19)	0.005 (0.18)	0.04 (0.17)
Constant	−0.45 (0.34)	−0.81* (0.42)	−0.59 (0.40)	−0.42 (0.36)	−0.37 (0.33)	−0.32 (0.31)
Observations	365	365	365	365	365	365
R ²	0.17	0.22	0.17	0.16	0.17	0.17

Predictive regressions for annual excess returns on Treasury bonds, using monthly data from January 1990 to May 2021. The first column shows results for the average excess return across 2-10 years maturities, the remaining columns show results for individual bond maturities. Predictors: *ISK* is option-implied yield skewness averaged over the last five business days of the month; *Level*, *Slope* and *Curvature* are the first three principal components of end-of-month Treasury yields. Standard errors based on the reverse regression delta method of [Wei and Wright \(2013\)](#) are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

Table B.3 reports estimates of our baseline specification—including yield curve predictors and ISK—over a variety of different samples. The dependent variable is the same as in Table 4, the three-month excess return on bonds with maturities from 1-10 years. The predictive power of ISK appears to be very robust across different samples.

Table B.4 reports out-of-sample (OOS) forecast accuracy for excess bond returns of three different linear models, using only information in the yield curve (level, slope and curvature of Treasury yields), using only ISK, or combining the yield curve and ISK as predictors. The dependent variable is the three-month excess bond return, either averaged across 1-10 year maturities or for individual bond maturities. The first forecast is made in January 1995, once five years of data are available, and the models are then re-estimated every month using expanding estimation windows. We measure forecast accuracy as the root-mean-squared forecast errors. The two models that include ISK always have higher prediction accuracy than the yields-only model. Forecasts using only ISK are the most accurate; the difference with in-sample regressions, where combining yield factors and ISK leads to the highest predictive power, is due to the increased importance of parsimony in OOS analysis. However, these differences in forecast accuracy are not statistically significant, using standard Diebold-

Table B.3: Predicting excess returns: different sample periods

	Full	Pre-2000	Post-2000	Pre-ZLB	Post-crisis	Pre-2018
ISK	−0.34*** (0.12)	−0.98** (0.44)	−0.33*** (0.12)	−0.75*** (0.24)	−0.24** (0.10)	−0.30** (0.15)
Level	0.004** (0.002)	0.02*** (0.01)	0.01** (0.003)	0.01*** (0.003)	0.01** (0.01)	0.004*** (0.002)
Slope	0.03*** (0.01)	0.03 (0.02)	0.03** (0.01)	0.03** (0.02)	0.03* (0.02)	0.04*** (0.01)
Curvature	0.004 (0.06)	0.23* (0.13)	−0.02 (0.07)	0.09 (0.11)	−0.10 (0.08)	−0.01 (0.07)
Constant	−0.15* (0.09)	−1.09** (0.53)	−0.22** (0.10)	−0.27 (0.26)	−0.31** (0.13)	−0.27** (0.11)
Observations	374	120	254	227	134	336
R ²	0.11	0.25	0.16	0.14	0.20	0.11

Predictive regressions for three-month excess bond returns (average of duration-normalized excess returns on Treasury bonds with one to ten years maturity) using different monthly sub-samples: *Full* is Jan-1990 to May-2021, *Pre-2000* is Jan-1990 to Dec-1999, *Post-2000* is Jan-2000 to May-2021, *Pre-ZLB* is Jan-1990 to Nov-2008, *Post-crisis* is Jan-2010 to May-2021, *Pre-2018* is Jan-1990 to Dec-2017. Predictors: *ISK* is option-implied yield skewness averaged over the last five business days of the month; *Level*, *Slope* and *Curvature* are the first three principal components of end-of-month Treasury yields. Standard errors based on the reverse regression delta method of [Wei and Wright \(2013\)](#), are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

Mariano tests. The lack of statistical significance is likely due to the relatively small sample and lower power of OOS comparisons.

Table B.4: Predicting excess returns: out-of-sample analysis

	Avg.	2y	3y	5y	7y	10y
Yield curve only	0.470	0.504	0.522	0.513	0.493	0.466
ISK only	0.447	0.478	0.493	0.487	0.472	0.450
ISK & yield curve	0.463	0.498	0.518	0.508	0.487	0.461

Root-mean-squared forecast errors for out-of-sample predictions of three-month excess bond returns, averaged across 1-10 year maturities or for individual bond maturities. Estimation uses expanding windows, starting in January 1995. The sample is monthly, from January 1990 to May 2021. Predictors: *ISK* is option-implied yield skewness averaged over the last five business days of the month; *Yield curve* corresponds to the first three principal components of end-of-month Treasury yields.

Finally, Table B.5 shows results for predictive regressions using Fama-Bliss yields. As opposed to the smoothed Treasury yields we have used in the rest of our analysis, here we use the unsmoothed zero-coupon yields of [Fama and Bliss \(1987\)](#), obtained from Wharton

Research Data Services, that have been used in much previous work on risk premia in Treasury bonds, including [Cochrane and Piazzesi \(2005\)](#) and many other studies. As in our main analysis, we predict bond returns in excess of the three-month T-bill rate. The Fama-Bliss data contain annual maturities from one to five years, and we predict both excess returns on individual zero-coupon bonds as well as the average excess return across all five bonds. The first four columns of Table B.5 show results for the average excess return across the five bond maturities. Adding skewness to the yield curve factors more than doubles the predictive power. As before, ISK has less power in a univariate regression than in combination with the yield curve, but its coefficient nevertheless is statistically significant at the 5 percent level. The same is true in a regression that includes a “CP” factor, the linear combination of forward rates that best predicts the average excess return, following [Cochrane and Piazzesi \(2005\)](#). We have found similar results when controlling for all five forward rates that go into the CP factor (not shown). For completeness, the remaining columns of Table B.5 show regressions for each of the five individual bond maturities. ISK is strongly statistically significant in each case.

Table B.5: Predicting excess returns: Fama-Bliss yields

	Average excess return				Bond maturity:				
					1y	2y	3y	4y	5y
Level	0.02*** (0.01)	0.02** (0.01)			0.02*** (0.01)	0.02*** (0.01)	0.02** (0.01)	0.01** (0.01)	0.01* (0.01)
Slope	−0.06 (0.07)	−0.11* (0.06)			−0.03 (0.06)	−0.09 (0.07)	−0.13* (0.07)	−0.15** (0.07)	−0.15** (0.07)
Curvature	0.09 (0.22)	0.11 (0.22)			0.33 (0.23)	0.13 (0.23)	0.07 (0.24)	0.03 (0.23)	0.003 (0.22)
ISK		−0.36*** (0.12)	−0.27** (0.12)	−0.27** (0.12)	−0.33** (0.13)	−0.35*** (0.13)	−0.37*** (0.13)	−0.38*** (0.12)	−0.36*** (0.12)
CP				1.00*** (0.32)					
Constant	−0.07 (0.09)	−0.09 (0.08)	0.18*** (0.04)	0.03 (0.05)	−0.06 (0.08)	−0.10 (0.09)	−0.11 (0.09)	−0.11 (0.09)	−0.10 (0.09)
Observations	363	363	363	363	363	363	363	363	363
R ²	0.04	0.09	0.03	0.09	0.12	0.09	0.09	0.09	0.08

Predictive regressions for three-month excess bond returns using unsmoothed Fama-Bliss Treasury yield data. The sample is monthly from January 1990 to June 2020. Predictors: *ISK* is option-implied yield skewness averaged over the last five business days of the month; *Level*, *Slope* and *Curvature* are the first three principal components of end-of-month yields; *CP* is the linear combination of forward rates that best predicts the average excess return, following [Cochrane and Piazzesi \(2005\)](#). Standard errors based on the reverse regression delta method of [Wei and Wright \(2013\)](#), are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

B.2 Additional results for FOMC announcement surprises

Since [Gürkaynak et al. \(2005\)](#) the literature on high-frequency event studies of FOMC announcements has focused on two measures of the policy surprises: a target surprise which, similar to the original measure proposed by [Kuttner \(2001\)](#), measures the surprise change in the federal funds rate, and a path surprise which captures the change in the expected policy path that is orthogonal to the target surprise. The two surprises are the first two principal components of the high-frequency changes in different money market futures rates, appropriately rotated and scaled (for details see [Gürkaynak et al., 2005](#)). More recently, [Gertler and Karadi \(2015\)](#) used changes in individual money market futures rates, which they found to be powerful instruments in monetary VARs. We consider these measures as well. Specifically, we include $FF1$, $FF4$, and $ED4$ as measures of monetary policy surprises: changes in the rates on the current-month fed funds futures contract, the three-month-ahead fed funds futures contract, and the four-quarter Eurodollar futures contract. Table [B.6](#) shows estimates of predictive regressions for these different monetary policy surprises. In all cases, the dependent variable is based on the changes in the 30-minute window around FOMC announcements. The top panel shows results for a univariate specification using only ISK, and the bottom panel for regressions that also add the usual yield-curve variables. The predictive power of ISK is weaker for target-rate surprises but generally quite robust across different measures of monetary policy surprises.

Table [B.7](#) considers specifications with macroeconomic variables that have been found to predict FOMC surprises in previous studies. In column (1) we include the predictors considered by [Cieslak \(2018\)](#): the average federal funds rate over the month preceding the FOMC meeting, and annual employment growth, measured as the 12-month log-change in total nonfarm payroll employment, appropriately lagged so that it is known by the day before the FOMC announcement. In this specification, employment growth but not the Federal Funds rate exhibits predictive power. The lack of predictive power of the funds rate is partly due to our different sample period and partly due to the different policy surprise measure than in the estimates of [Cieslak \(2018\)](#). Using Cieslak’s exact sample and regression specification we are able to replicate her results, and we still find that when we add ISK to the regression it significantly raises the predictive power. Columns (2) to (4) add some of the macroeconomic variables considered by [Bauer and Swanson \(2022a\)](#): the Brave-Butters-Kelley business cycle indicator produced by the Chicago Fed, the change in nonfarm payroll employment in the previous month (again appropriately accounting for the publication lag), and the return of the S&P 500 stock index over the three months preceding the FOMC announcement. Except for the last case that includes the past stock return, both predictors exhibit statistically significant explanatory power for the FOMC policy surprise.

B.3 Additional results for SPF forecast errors

Here we present additional evidence using the quarterly Survey of Professional Forecasters (SPF). As in the BCFF, we consider two forecast targets, the ten-year yield and a short-

Table B.6: Predicting different monetary policy surprises

	PC1	Target	Path	FF1	FF4	ED4
ISK	0.02*** (0.01)	0.02 (0.01)	0.03** (0.01)	0.02* (0.01)	0.02** (0.01)	0.04*** (0.01)
Constant	−0.002 (0.003)	−0.002 (0.01)	−0.003 (0.004)	−0.01** (0.003)	−0.01*** (0.004)	−0.02*** (0.005)
R ²	0.03	0.01	0.02	0.01	0.02	0.04
ISK	0.03*** (0.01)	0.02 (0.02)	0.04*** (0.01)	0.02* (0.01)	0.03** (0.01)	0.06*** (0.01)
Level	−0.001 (0.001)	−0.003 (0.002)	0.002 (0.002)	−0.001 (0.001)	−0.002 (0.002)	0.0000 (0.002)
Slope	−0.01** (0.002)	−0.002 (0.004)	−0.01** (0.004)	0.0002 (0.003)	−0.004 (0.003)	−0.01** (0.005)
Curvature	−0.02 (0.01)	−0.04 (0.02)	−0.003 (0.02)	−0.03* (0.02)	−0.02 (0.02)	−0.03 (0.02)
Constant	0.02** (0.01)	0.03** (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.02 (0.01)
R ²	0.05	0.03	0.06	0.04	0.04	0.07

Predictive regressions for alternative measures of monetary policy surprises: *PC1* is the same measure used in Table 5, the first principal component of five futures rate changes; *Target* and *Path* are the target and path factors of the policy surprise from [Gürkaynak et al. \(2005\)](#); *FF1* is the change in the current-month fed funds futures rate; *FF4* is the change in the three-month-ahead fed funds futures rate; *ED4* is the change in the four-quarter Eurodollar futures rate. The sample contains 213 FOMC announcements from January 1994 to June 2019. For a description of the predictors see the notes to Table 5. White heteroskedasticity-robust standard errors are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

term interest rate. Because the SPF does not include forecasts of the federal funds rate, the short-term rate is the three-month T-Bill rate. Both interest rates are quarterly averages of the daily values in the Fed’s H.15 statistical release. Forecasts are reported for the current quarter (nowcasts) and each of the subsequent four quarters. We use the median as the consensus forecast. We regress forecast errors on yield curve factors and ISK. To ensure that the predictors are observable at the time the forecast is made, we use observations on the day of the response deadline of the survey. Because the regression residuals are serially correlated we use Hansen-Hodrick standard errors with the number of lags corresponding to the forecast horizon.

Table B.8 shows the results. For each forecast horizon, we estimate two specifications, one with ISK only, and one that also includes yield factors. For the ten-year yield, we find that ISK has statistically significant predictive power for most forecast horizons, and that this is increased when yield curve factors are added to the regression, although those are themselves rarely statistically significant. For the T-bill rate, the predictive power is quite a bit stronger,

Table B.7: Predicting FOMC surprises with ISK and macro variables

	(1)	(2)	(3)	(4)
ISK	0.023*** (0.007)	0.020*** (0.008)	0.024*** (0.008)	0.011 (0.007)
FFR	−0.001 (0.001)			
Annual employment growth	0.460*** (0.175)			
BBK index		0.010** (0.004)		
Change in employment			0.057*** (0.018)	
S&P 500 return				0.143*** (0.053)
Constant	−0.004 (0.004)	−0.0004 (0.003)	−0.009** (0.004)	−0.003 (0.003)
Observations	213	213	213	213
R ²	0.059	0.084	0.103	0.108

Predictive regressions for the monetary policy surprise around FOMC announcements from January 1994 to June 2019. The dependent variable is the first principal component of 30-minute futures rate changes around the announcement for five different contracts with up to about one year maturity. *ISK* is option-implied yield skewness averaged over the month (22 trading days) before the FOMC announcement; *FFR* is the average federal funds rate over the calendar month preceding the meeting, and *Annual employment growth* is the 12-month log-change in total nonfarm payroll employment (appropriately lagged), as used by [Cieslak \(2018\)](#); *BBK index* is the Brave-Butters-Kelley business cycle indicator from the Chicago Fed, *Change in employment* is the change in non-farm payrolls released in the most recent employment report, and *S&P 500 return* is the stock return over the three months (65 days) up to the day before the FOMC announcement, as used by [Bauer and Swanson \(2022a\)](#). White heteroskedasticity-robust standard errors are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

with higher levels of significance and higher R^2 .

C A model of biased beliefs

C.1 Details and derivations

The subjective innovation to the beliefs of agent 2 is

$$dz_t^s = \frac{1}{\sigma} (dC_t/C_t - \mu_t^s dt) = dz_t + \frac{\mu - \mu_t^s}{\sigma} dt.$$

which yields

$$dz_t^s - dz_t = \Delta_t dt.$$

Table B.8: Predicting SPF forecast errors

(A) 10y yield	Current		1Q ahead		2Q ahead		3Q ahead		4Q ahead	
ISK	0.18***	0.20***	0.53***	0.60***	0.57***	0.68***	0.48	0.63**	0.54	0.74**
	(0.04)	(0.05)	(0.13)	(0.14)	(0.20)	(0.21)	(0.32)	(0.30)	(0.34)	(0.29)
Level		0.001		−0.01		−0.01		0.003		0.02
		(0.01)		(0.02)		(0.05)		(0.06)		(0.06)
Slope		−0.01		−0.05		−0.09		−0.10		−0.14
		(0.01)		(0.04)		(0.06)		(0.09)		(0.10)
Curve		−0.16*		−0.09		0.08		0.03		0.11
		(0.09)		(0.23)		(0.33)		(0.39)		(0.48)
Constant	−0.08***	−0.01	−0.26***	−0.10	−0.41***	−0.23	−0.54***	−0.35	−0.69***	−0.53
	(0.02)	(0.05)	(0.04)	(0.11)	(0.08)	(0.22)	(0.12)	(0.31)	(0.14)	(0.40)
R ²	0.09	0.14	0.08	0.10	0.05	0.08	0.03	0.05	0.03	0.07
(B) 3m TB	Current		1Q ahead		2Q ahead		3Q ahead		4Q ahead	
ISK	0.14***	0.15***	0.51***	0.53***	0.85***	0.89***	1.20***	1.29***	1.49***	1.63***
	(0.05)	(0.05)	(0.16)	(0.14)	(0.32)	(0.28)	(0.42)	(0.35)	(0.54)	(0.44)
Level		−0.01		−0.02		−0.03		−0.04		−0.06
		(0.004)		(0.02)		(0.04)		(0.06)		(0.08)
Slope		0.01		0.01		0.01		0.01		0.01
		(0.01)		(0.02)		(0.06)		(0.10)		(0.15)
Curve		−0.22***		−0.51***		−0.76*		−1.13*		−1.37*
		(0.05)		(0.16)		(0.39)		(0.61)		(0.81)
Constant	−0.07***	0.02	−0.21***	0.01	−0.37***	−0.004	−0.54***	0.01	−0.72***	−0.02
	(0.01)	(0.03)	(0.04)	(0.09)	(0.10)	(0.17)	(0.17)	(0.26)	(0.24)	(0.37)
R ²	0.11	0.29	0.12	0.20	0.12	0.19	0.12	0.20	0.12	0.19
Observations	118	118	117	117	116	116	115	115	114	114

Predictive regressions for forecast errors in the h -quarter ahead median forecast for the ten-year Treasury yield (panel A) and the three-month T-Bill rate (panel B) in the Survey of Professional Forecasters, using quarterly surveys from 1992:Q1 to 2020:Q3. The forecast horizon h ranges from 0 (current/nowcast) to 4 quarters. *ISK* is a five-day average of option-implied yield skewness, *Level*, *Slope* and *Curvature* are the first three principal components of Treasury yields. All predictors are measured on the day of the survey deadline. Hansen-Hodrick standard errors with h lags are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

The dynamics of the expectational error Δ_t are:

$$d\Delta_t = -\kappa\Delta_t dt - \frac{\delta}{\sigma} dz_t. \quad (\text{C.1})$$

In particular the equation implies that Δ_t is a Gaussian variable.

Let \mathcal{P} and \mathcal{P}^s denote the true and subjective probability measures, respectively. Let ξ_t and ξ_t^s denote the state-price density (SPD) under \mathcal{P} and \mathcal{P}^s , respectively. E^s is the expecta-

tion taken under \mathcal{P}^s . Agents 1 and 2 solve their consumption-savings problems given by, respectively,

$$\begin{aligned} \max E \left(\int_0^T e^{-\rho t} u(C_t^1) dt \right) \quad \text{s.t.} \quad E \left(\int_0^T \xi_t / \xi_0 \cdot C_t^1 dt \right) &\leq w_0^1, \\ \max E^s \left(\int_0^T e^{-\rho t} u(C_t^2) dt \right) \quad \text{s.t.} \quad E^s \left(\int_0^T \xi_t^s / \xi_0^s \cdot C_t^2 dt \right) &\leq w_0^2, \end{aligned}$$

where it is assumed that the agents have identical power utility functions

$$u(C) \equiv C^{1-\gamma} / (1-\gamma).$$

Consumption allocations and state price densities. Denote the likelihood ratio by $\lambda_t = d\mathcal{P}/d\mathcal{P}^s = y^{-1} \xi_t / \xi_t^s$, where $y = y^2 / y^1$, and y^i is the constant Lagrange multiplier from the respective budget constraint. Optimal consumption allocations are

$$C_t^1 = f(\lambda_t) C_t, \quad C_t^2 = (1 - f(\lambda_t)) C_t, \quad f(\lambda_t) = (1 + (y\lambda_t)^{1/\gamma})^{-1}.$$

The state price densities are:

$$\begin{aligned} \xi_t &= (y^1)^{-1} e^{-\rho t} C_t^{-\gamma} f(\lambda_t)^{-\gamma} \\ &= (y^1)^{-1} e^{-\rho t} C_t^{-\gamma} (1 + (y\lambda_t)^{1/\gamma})^\gamma \\ &= \sum_{k=0}^{\gamma} \binom{\gamma}{k} (y^1)^{-1} e^{-\rho t} C_t^{-\gamma} (y\lambda_t)^{k/\gamma}, \\ \xi_t^s &= (y^2)^{-1} e^{-\rho t} C_t^{-\gamma} (1 - f(\lambda_t))^{-\gamma}. \end{aligned}$$

Note that y^i and y cancel out in the SDF, ξ_T^i / ξ_t^i . Lastly,

$$d\lambda_t = -\Delta_t \lambda_t dz_t. \tag{C.2}$$

To derive equation (7) simply apply Itô's lemma to ξ_t . The real short rate is the negative of the drift of $d\xi_t / \xi_t$.

Bond pricing. Set $y = 1$ w.l.o.g. The real bond price is, for integer γ ,

$$\begin{aligned} B_{t,T} &= E_t(\xi_T / \xi_t) = \sum_{k=0}^{\gamma} w_t^{(k)} E_t \left[e^{-\rho(T-t)} \left(\frac{C_T}{C_t} \right)^{-\gamma} \left(\frac{\lambda_T}{\lambda_t} \right)^{k/\gamma} \right] = \sum_{k=0}^{\gamma} w_t^{(k)} E_t \left[\frac{\xi_T^{(k)}}{\xi_t^{(k)}} \right], \\ w_t^{(k)} &= \binom{\gamma}{k} \lambda_t^{k/\gamma} (1 + \lambda_t^{1/\gamma})^{-\gamma}, \quad \sum_{k=0}^{\gamma} w_t^{(k)} = 1, \\ \xi_t^{(k)} &= e^{-\rho t} C_t^{-\gamma} \lambda_t^{k/\gamma}. \end{aligned}$$

Then

$$\begin{aligned} d\xi_t^{(k)} / \xi_t^{(k)} &= -r_t^{(k)} dt - [\gamma\sigma + \gamma^{-1}k\Delta_t] dz_t, \\ r_t^{(k)} &= \rho + \gamma\mu - \frac{1}{2}\gamma(\gamma+1)\sigma^2 - \sigma k\Delta_t + \frac{1}{2}\frac{k}{\gamma} \left(1 - \frac{k}{\gamma}\right) \Delta_t^2. \end{aligned} \quad (\text{C.3})$$

These expressions imply that the “artificial” bond prices $B_{t,T}^{(k)}$, corresponding to the SPDs $\xi_t^{(k)}$, are exponentially quadratic functions of the state variable Δ_t . Following [Ahn et al. \(2002\)](#), they are given by

$$\begin{aligned} \log B_{t,T}^{(k)} &= A^{(k)}(T-t) + B^{(k)}(T-t) \cdot \Delta_t + C^{(k)}(T-t) \cdot \Delta_t^2, \\ \frac{dC^{(k)}(\tau)}{d\tau} &= 2\frac{\delta^2}{\sigma^2} (C^{(k)}(\tau))^2 - 2\left(\kappa - \frac{k}{\gamma}\frac{\delta}{\sigma}\right) C^{(k)}(\tau) - \frac{1}{2}\frac{k}{\gamma} \left(1 - \frac{k}{\gamma}\right), \quad C^{(k)}(0) = 0, \\ \frac{dB^{(k)}(\tau)}{d\tau} &= 2\frac{\delta^2}{\sigma^2} C^{(k)}(\tau) B^{(k)}(\tau) - \left(\kappa - \frac{k}{\gamma}\frac{\delta}{\sigma}\right) B^{(k)}(\tau) + 2\gamma\delta C^{(k)}(\tau) + \sigma k, \quad B^{(k)}(0) = 0, \\ \frac{dA^{(k)}(\tau)}{d\tau} &= \frac{\delta^2}{\sigma^2} C^{(k)}(\tau) + \frac{1}{2}\frac{\delta^2}{\sigma^2} (B^{(k)}(\tau))^2 + \gamma\delta B^{(k)}(\tau) - \left(\rho + \gamma\mu - \frac{1}{2}\gamma(\gamma+1)\sigma^2\right), \quad A^{(k)}(0) = 0. \end{aligned}$$

The real bond price is

$$B_{t,T} = \sum_{k=0}^{\gamma} w_t^{(k)} B_{t,T}^{(k)}. \quad (\text{C.4})$$

Thus, bond prices are weighted averages of exponentially quadratic functions of the Gaussian state variable Δ_t . The weights $w_t^{(k)}$ are affected by the bias in beliefs via λ_t .

Bond risk premia. Without disagreement ($\mu_t = \mu$) bond risk premia are constant due to the constant volatility of the log SDF, as can be seen from the evolution equation for the SPD $\xi_t^{(k)}$ in (C.3). To see the dependence of the bond risk premium on Δ_t , note that

$$\begin{aligned} E_t \left(\frac{dB_{t,T}}{B_{t,T}} \right) &= \sum_{k=0}^{\gamma} \left[\frac{B_{t,T}^{(k)}}{B_{t,T}} E_t \left(dw_t^{(k)} \right) + w_t^{(k)} E_t \left(\frac{dB_{t,T}^{(k)}}{B_{t,T}^{(k)}} \right) \right] \\ &= \sum_{k=0}^{\gamma} w_t^{(k)} \frac{B_{t,T}^{(k)}}{B_{t,T}} \cdot E_t \left(\frac{dB_{t,T}^{(k)}}{B_{t,T}^{(k)}} \right). \end{aligned}$$

The expected instantaneous bond return in each artificial economy k is equal to the artificial risk-free rate, $r_t^{(k)}$, plus a linear function of the price of risk, with is affine in disagreement, $\gamma\sigma + \gamma^{-1}k\Delta_t$. Continuing the previous expression, one can then write:

$$\frac{1}{dt} E_t \left(\frac{dB_{t,T}}{B_{t,T}} \right) = \sum_{k=0}^{\gamma} \frac{w_t^{(k)} B_{t,T}^{(k)}}{B_{t,T}} \cdot \left(r_t^{(k)} - \frac{\delta}{\sigma} (B^{(k)}(T-t) + 2C^{(k)}(T-t)\Delta_t) (\gamma\sigma + \gamma^{-1}k\Delta_t) \right),$$

see equation (11) in [Ahn et al. \(2002\)](#). Thus, one can obtain the risk premium, $E_t \left(\frac{dB_{t,T}}{B_{t,T}} - r_t dt \right)$.

These expressions have two implications central to our findings. First, they connect non-normality and skewness of bond yields to the bond risk premia via Δ_t . Second, they show that the expected excess return is a non-linear function of Δ_t . Thus, it generally cannot be captured by a linear combination of yields, consistent with the violations of the spanning hypothesis we have documented in Section 3.1.

Dynamics under subjective probability measure. To be able to calculate “consensus forecast errors” we need the dynamics of the state variables under the subjective probability measure. To derive them, we use Girsanov’s theorem and the fact that $dz_t^s = dz_t + \Delta_t dt$. First, we note that the law of motion for C_t under the subjective measure is simply

$$dC_t = \mu_t^s dt + \sigma dz_t^s \quad (\text{C.5})$$

with $\mu_t^s = \mu - \sigma \Delta_t$ by definition of Δ_t . Second, for the S -dynamics of Δ_t we have

$$d\Delta_t = \mu_{\Delta,t}^s dt - \frac{\delta}{\sigma} dz_t^s = \mu_{\Delta,t}^s dt - \frac{\delta}{\sigma} (dz_t + \Delta_t dt) = \left(\mu_{\Delta,t}^s - \frac{\delta}{\sigma} \Delta_t \right) dt - \frac{\delta}{\sigma} dz_t.$$

The term in parentheses needs to equal the drift under the objective measure,

$$\mu_{\Delta,t}^s - \frac{\delta}{\sigma} \Delta_t = -\kappa \Delta_t \quad \Rightarrow \quad \mu_{\Delta,t}^s = -\left(\kappa - \frac{\delta}{\sigma} \right) \Delta_t = -\kappa^* \Delta_t,$$

with $\kappa^* = \kappa - \delta/\sigma$. As a result, the S -dynamics of Δ_t are

$$d\Delta_t = -\kappa^* \Delta_t dt - \frac{\delta}{\sigma} dz_t^s. \quad (\text{C.6})$$

Finally, for λ_t we have

$$d\lambda_t = \mu_{\lambda,t}^s dt - \Delta_t \lambda_t dz_t^s = \mu_{\lambda,t}^s dt - \Delta_t \lambda_t (dz_t + \Delta_t dt) = (\mu_{\lambda,t}^s - \Delta_t^2 \lambda_t) dt - \Delta_t \lambda_t dz_t$$

Since λ_t is a martingale under the objective measure, $\mu_{\lambda,t}^s - \Delta_t^2 \lambda_t = 0$, so that the S -dynamics of λ_t are

$$d\lambda_t = \Delta_t^2 \lambda_t dt - \Delta_t \lambda_t dz_t^s. \quad (\text{C.7})$$

C.2 Model simulations

As noted in the main text, we choose the parameters $\sigma = 0.035$, $\kappa = 0.1$, $\delta = \kappa\sigma$, and $\gamma = 2$. We also need to choose ρ and μ , which we both set to 0.02; alternative reasonable choices lead to essentially identical simulation results.

For each initial value of Δ_0 , we simulate the model $M = 100,000$ times, in each run obtaining simulated paths from $t = 0$ to $t = T = 0.25$ (one quarter). We use time increments of

$dt = 1/400$. The simulations are initialized at $C_0 = 1$ and $\lambda_0 = 1$. The initial wealth share is captured by λ_0 , which implies a 50-50 initial consumption shares; using other values or rules for λ_0 generally gave similar results

We simulate $\log C_t$ and $\log \lambda_t$ to avoid negative values. We use the following discrete-time simulation scheme:

$$\begin{aligned}\log C_{t+dt} &= \log C_t + \left(\mu - \frac{1}{2}\sigma^2\right)dt + \sigma u_{t+dt} \\ \Delta_{t+dt} &= (1 - \kappa dt)\Delta_t - \frac{\delta}{\sigma}u_{t+dt} \\ \log \lambda_{t+dt} &= \log \lambda_t - \frac{1}{2}\Delta_t^2 dt - \Delta_t u_{t+dt}\end{aligned}$$

where u_{t+dt} are independent Gaussian innovations with standard deviation $\sqrt{dt} = 1/20$, corresponding to dz_t in our model formulation. The first equation above follows from (4) and Itô's lemma. The second equation corresponds to (C.1). The third equation follows from (C.2) and Itô's lemma.

For each of the M simulated values for Δ_T and λ_T we can calculate the price of the ten-year bond at $t = T$ using equation (C.4), as well as the corresponding ten-year yield. The sample skewness across the M simulated values of the time- T yield is our simulation estimate of conditional yield skewness, which we plot against Δ_0 in Figure 6, panel (a), as well as in Figure C.1 below. In addition, we calculate the expected excess log return on a bond with maturity 10.25 years as a simulation estimate of the risk premium on long-term bonds.

In order to calculate “survey expectations” we need expectations under the subjective measure. We simulate under this measure as follows:


$$\begin{aligned}\log C_{t+dt} &= \log C_t + \left(\mu - \frac{1}{2}\sigma^2 - \sigma\Delta_t\right)dt + \sigma u_{t+dt} \\ \Delta_{t+dt} &= (1 - \kappa^* dt)\Delta_t - \frac{\delta}{\sigma}u_{t+dt} \\ \log \lambda_{t+dt} &= \log \lambda_t + \frac{1}{2}\Delta_t^2 dt - \Delta_t u_{t+dt}\end{aligned}$$

The first equation above follows from (C.5) and Itô's lemma. The second equation corresponds to (C.6). The third equation follows from (C.7) and Itô's lemma. The average value of C_T across these simulations is our estimate of $E_0^s(C_T)$, and we calculate the mean forecast error as $E_0(C_T) - E_0^s(C_T)$. Using the paths for λ_t and Δ_t simulated under the objective and subjective probabilities in the short rate equation (7), we can calculate mean forecast errors for the short rate as $E_0(r_T) - E_0^s(r_T)$.

In the main text we showed simulation results for a range of Δ_0 where agent 2 is too optimistic about consumption growth, i.e., $\Delta_0 < 0$, and conditional yield skewness is increasing in Δ_0 . Here we present additional simulation results to get a broader impression of the model's

implications for skewness. In particular, we show skewness over a much broader range of Δ_0 from -1.5 to +1.5. We also consider alternative choices for γ , in addition to our main calibration which assumed $\gamma = 2$.

Figure C.1: Model-implied yield skewness for different levels of risk aversion



Conditional skewness of ten-year yield, based on simulations from heterogeneous-beliefs model. For each starting value of disagreement, Δ_0 , we simulate 100,000 paths over a period of one quarter, using $\sigma = 0.035$, $\kappa = 0.1$, and different values of γ .

Figure C.1 shows that the implications of the model for yield skewness depend on the degree of risk aversion. For log-utility, skewness first gradually rises up to about $\Delta_0 = -0.4$, then declines up to about $\Delta_0 = 0.4$, and then rises again. The relationship is point-symmetric around $\Delta_0 = 0$ where skewness is zero, up to simulation error.

For $\gamma > 1$, the behavior of skewness is not symmetric. Skewness first rises up to a point near zero, and then declines. The maximum level of skewness depends on γ . It is gradually moving towards zero as γ increases, and breaches this threshold when $\gamma = 9$. We display the case of $\gamma = 10$ in which skewness remain negative for the entire range of Δ_0 .

Broadly speaking, the model can capture the correct sign of the relationship between skewness and various objects of interest, such as bond excess returns, or forecast errors, whenever skewness increases with Δ_0 . We demonstrate this for the case of $\gamma = 2$ and a negative range of Δ_0 in Figure 6. We see from Figure C.1 that when $\gamma = 1$, skewness is increasing in Δ_0 only in the tails of the distribution for Δ_0 , where the relation between skewness and Δ_0 is nearly flat, whereas for values of Δ_0 between -0.4 and 0.4, skewness is decreasing. Thus, there is little hope that with log utility the model can be consistent with the evidence.

The cases of $\gamma = 2$ and $\gamma = 10$ are qualitatively similar as skewness increases to a point near zero in both cases. High values of γ are inconsistent with the empirically observed sign-switching in skewness, since skewness remains negative for all values of Δ_0 . For an increasing, sign-switching pattern, we require γ greater than one but less than ten.

For the baseline case of $\gamma = 2$, skewness is increasing only over a certain range of Δ_0 , and we focused on this range in Section 4. The behavior of the key objects of interest is reversed as compared to that figure for positive Δ_0 , when skewness is declining with Δ_0 , as is evident from Figure C.1. Over this region the model generates relationships with expected excess bond returns and mean forecast errors for interest rates and consumption that are opposite to those in Figure 6 and the patterns we have documented in the data.