

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 December 2021

We are grateful to Anna Cieslak (discussant), Ian Dew-Becker, Michael Gällmeyer, Marco Giacomelli, Mathieu Gomez, Valentin Haddad, Christian Heyerdahl-Larsen, Lars Lochstoer, Ricardo Reis (discussant), Avanidhar Subrahmanyam, Andreea Vladu (discussant), as well as seminar and conference participants at the Bank of Canada, Boston Fed, Frankfurt University, Erasmus University, Gothenburg University, UCLA Anderson, the 2021 Bundesbank Term Structure Workshop, the 2021 EFA meeting, the 2021 IAAE meeting, the 2021 SoFiE meeting, and the 2021 conference of the Canadian Derivatives Institute. The latest version is available here. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 December 2021
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

What predicts changes in interest rates? The literature has come a long way from the expectations hypothesis distinguishing between statistical, risk-adjusted, and subjective expectations. Regardless of the channel, the shape of the yield curve has emerged as a key source of information (e.g., [Campbell and Shiller, 1991](#); [Cochrane and Piazzesi, 2005](#); [Piazzesi et al., 2015](#)). The Global Financial Crisis of 2008/2009 led to uncharted territory for the yield curve, with short-term rates constrained by the zero lower bound (ZLB) and long-term rates affected by unconventional monetary policies. In such an environment, interest rate risk was generally perceived to be tilted to the upside, but the slope of the yield curve did not capture this. By contrast, interest rate skewness implied by Treasury options signaled substantial upside risk to yields, suggesting that implied skewness might be a useful forward looking measure to assess interest rate risk.

In this paper we argue that skewness is crucial for understanding yield dynamics and bond risk premia over the last three decades, and not just during ZLB episodes. Because conditional skewness measures the asymmetry in the distribution of future rate changes, it is a natural candidate for capturing changes in the balance of interest rate risks. By contrast, measures of interest rate volatility or uncertainty in general do not contain directional information. We study model-free measures of the skewness of Treasury yields during the period from 1990 to 2021, and document substantial and persistent time variation that is closely related to the business cycle and the stance of monetary policy. Importantly, skewness contains useful information about the outlook for interest rates: it helps predict excess bond excess returns, high-frequency interest rate changes around FOMC announcements, and survey forecast errors for interest rates. Incorporating the information in conditional yield skewness reveals an important role for expectational errors in driving measures of statistical bond risk premia. Our evidence is consistent with a theoretical environment where some economic agents have biased beliefs.

A large literature has documented pronounced negative skewness of stock returns and the implications for asset pricing and investment management (see, e.g., [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. Even papers on non-normality of

interest rates and its links to monetary policy generally assume zero conditional skewness and treat the effects of FOMC announcements as symmetric (e.g., [Johannes, 2004](#); [Piazzesi, 2001](#)). There is a sound empirical reason for the extant perspective: the sample skewness of Treasury yield changes is essentially zero. That is, the unconditional distribution of interest rates is symmetric.

Our first question is whether and when asymmetries are important for the conditional distribution of future interest rates. Specifically, we investigate the time variation and drivers of conditional skewness, which we measure in two different ways. The first measure, realized skewness, is based on realized moments of yield changes that are calculated from daily data. The second, implied skewness, uses risk-neutral moments of yields as implied by Treasury option prices. The two measures are qualitatively similar in our case, their difference reflecting a skewness risk premium and measurement noise. For most of our empirical analysis, we focus on option-implied skewness, which has several advantages including its forward-looking nature, daily availability, and high statistical precision.

We document pronounced and economically interesting time variation in conditional skewness over the past 30 years. The variation is persistent and, in stark contrast to unconditional yield skewness, indicates extended periods of both substantial upward and downward skew in the balance of interest rate risk. Option-implied skewness predicts a substantial share of the variation in realized skewness, establishing formal statistical evidence for the time variation in conditional skewness. The variation in conditional skewness is strongly cyclical and driven by macroeconomic state variables: skewness is closely related to the shape of the yield curve, the stance of monetary policy, and the business cycle. In particular, 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 been tightening and the yield curve is flat or inverted.

Our second question is whether measured conditional skewness contains useful information about the future outlook for interest rates, in addition to the macroeconomic and financial variables that drive skewness or those that are commonly used as predictors in this context. We find that indeed it contains substantial additional information. Skewness exhibits highly significant predictive power for excess returns on Treasury bonds, and this finding is robust

to controlling for the shape of the yield curve and a wide range of other predictors. 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 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 money market futures rate changes around upcoming FOMC announcements, that is, it correctly predicts a substantial portion of “monetary policy surprises.” This finding is mainly driven by the fact that negative skewness captures downside risk early on in monetary easing cycles and successfully predicts the large dovish surprises that typically occur during those phases. The risk of monetary policy surprises is asymmetric, and skewness correctly captures this asymmetry.

The evidence on FOMC surprises suggests that the predictive power of skewness is highest at monetary policy turning points, a time when disagreement among forecasters is particularly high. Because predictability of asset returns may arise from time-varying risk premia or expectational errors in survey forecasts, we next study the relation between these errors and conditional skewness. Using the Blue Chip Financial Forecasts and Survey of Professional Forecasters, we document that skewness measured at the time of the surveys is highly informative for the forecast errors for future interest rates. 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. One needs a framework where some agents have biased beliefs. We adopt the heterogeneous beliefs two-agent framework (Basak, 2005; Ehling et al., 2018) by assuming that one agent knows the true distribution of the state, while the other one has erroneous beliefs. As a result, the usual measure of disagreement in these models becomes a measure of bias in beliefs. The model directly speaks to our evidence and is consistent with our findings on time-varying skewness, predictability, and expectational errors. The economic mechanism is straightforward: as investors differ in beliefs about future consumption they take different bond positions, with the more optimistic investor selling bonds to the pessimistic investor. Ex ante, each investor expects to capture wealth from the other investor and, hence, both expect future consumption to be higher than without disagreement. As a result, interest rates depends nonlinearly on the dispersion in beliefs between the agents, and are non-normally distributed even though the state variables are Gaussian. Yield skewness arises endogenously and varies over time with changes in beliefs. Bond yields do not linearly span expected excess returns, leaving an important role for expectational errors and skewness to explain statistical bond risk premia. Skewness is negative when belief dispersion is high, as is the case around turning points of monetary policy, when low skewness tends to anticipate easing surprises.

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 (e.g., Campbell and Shiller, 1991, Cochrane and Piazzesi, 2005, Ludvigson and Ng, 2009, and Cieslak and Povala, 2015, among many others). Recent work has revisited this predictability allowing for deviations from the benchmark of full information rational expectations (FIRE). In particular, Piazzesi et al. (2015), Cieslak (2018), Schmeling et al. (2021), Buraschi et al. (2021) and Nagel and Xu (2021) use survey forecasts of interest rates to study the contributions of subjective bond risk premia and expectational errors to the predictability of excess returns, i.e., to variation in statistical risk premia. We document that conditional yield skewness captures systematic errors in interest rate expectations (and indeed may be their source), and in this way

substantially contributes to return predictability.

[Giacoletti et al. \(2021\)](#) (GLS) show that survey disagreement in yield forecasts predicts excess bond returns, and attribute this result to learning about interest rate dynamics. Our predictability results are robust to controlling for the GLS measure of disagreement. GLS argue that their evidence is inconsistent with heterogeneous beliefs about macro fundamentals, since they find no relationship between disagreement about inflation and yields (see also [Singleton, 2021](#)). The mechanism that we emphasize with our heterogeneous beliefs model is distinct from the one explored by GLS. In our model, disagreement is closely connected to biases in beliefs about fundamentals. We document supportive evidence in the form of predictability of survey forecast errors for GDP growth using conditional skewness. Thus, in the data skewness is connected to biased beliefs about both fundamental and financial variables, consistent with the heterogeneous beliefs framework.

A very large empirical literature studies the interaction of monetary policy and bond markets, including [Balduzzi et al. \(1997\)](#), [Ang and Piazzesi \(2003\)](#), and [Piazzesi \(2005\)](#). Since [Kuttner \(2001\)](#) an important focus has been on the effects of surprises in monetary policy announcements on asset prices (e.g., [Gürkaynak et al., 2005](#), [Nakamura and Steinsson, 2018](#)). [Brooks et al. \(2020\)](#) document that FOMC announcement surprises cause a persistent drift in Treasury yields. 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, 2021](#)).

Our paper also connects to studies on the importance of asymmetries in the macro-financial outlook ([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.

A large literature in macroeconomics and finance studies 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](#)). Our evidence and model supports the view that interest rate expectations deviate from the FIRE assumption, but we do not take a stand on the precise nature of this deviation or on a specific type of belief formation. Our theoretical results show the implications of a deviation from FIRE, and thus biases in beliefs, for interest rates and bond risk premia.

In terms of using options to measure skewness, the methodology follows that of [Bakshi and Madan \(2000\)](#) and [Neuberger \(2012\)](#). [Trolle and Schwartz \(2014\)](#) is a paper close to ours as it also measures skewness in fixed-income markets using the Neuberger approach and documents some time-variation in swaption-implied conditional skewness. However, the sample period is relatively short, 2002-2009, and the authors do not relate skewness to bond returns or survey-based forecasts of yields.

Some other empirical work looks at interest rate skewness from different perspectives: [Hattori et al. \(2016\)](#) demonstrate that unconventional monetary policy has reduced option-implied tail risks for the stock and bond markets. [Mertens and Williams \(2021\)](#) use option-implied distributions during the 2008-2015 ZLB period to distinguish between the constrained and unconstrained monetary policy equilibria. [Li \(2021\)](#) documents the connection between option-implied Treasury skewness and recessions between 2000 and 2018. [Diercks et al. \(2021\)](#) measure the skewness in subjective distributions for the future federal funds rate derived from the New York Fed’s Survey of Primary Dealers. Their series, which starts in 2011, exhibits broadly similar patterns as option-implied skewness for Treasury yields.

The early work on heterogeneous beliefs is [Harrison and Kreps \(1978\)](#) and [Detemple and Murthy \(1994\)](#). Heterogeneous beliefs-based asset pricing applications are reviewed in [Basak \(2005\)](#). [Xiong and Yan \(2010\)](#) is the first application of heterogeneous beliefs to Treasury bonds. Our contribution to this literature is to uncover a link between dispersion in beliefs and non-normality of yields. Disagreement and biased beliefs about fundamentals generates skewness in interest rates, a connection we document both in theory and in the data.

2 Time 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, which implies that the balance of risk is tilted to the upside, and vice versa.

Unconditional interest rate skewness—the sample skewness of Treasury yield changes over long periods of time—has been essentially zero. In this section we document that this contrasts with pronounced shifts and large cyclical swings over the last three decades in conditional skewness, measured either as realized skewness (using short rolling windows of Treasury yield changes) or option-implied skewness (using model-free moments implied by Treasury options). The statistical evidence for this time variation is that option-implied skewness strongly predicts realized skewness. The variation is strongly cyclical, and skewness is closely related to the slope of the yield curve and stance of monetary policy. Periods when the slope is high (low) or when the Fed has been easing (tightening) the stance of monetary policy are characterized by high (low) option-implied skewness. The orthogonal component of skewness still exhibits cyclical variation, which our subsequent analysis will show to contain substantial additional predictive information.

2.1 Data

The data we use in this analysis are Treasury yields, as well as Treasury futures and options. Our Treasury yields are the daily “GSW” smoothed Treasury yield curves from [Gürkaynak et al. \(2007\)](#). When we need a monthly data frequency we take monthly averages. Most of our analysis focuses on the 10-year yield.

The Treasury derivative prices are from CME group. In particular, we use end-of-day prices of the 10-year T-Note futures contract, and options written on this contract. Both types of contracts are among the most actively traded Treasury derivatives, with high liquidity (in terms of open interest and volume) and long available price histories. The deliverable maturities for this futures contract are 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](#).

Our 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 allowing us to consistently calculate option-based moments using prices across many contracts/strikes. While the historical Treasury options data starts in May 1985, there are only few contracts and prices available during the early years.

2.2 Sample statistics and unconditional skewness

The top panel of Table 1 reports summary statistics for quarterly changes (using the last month of the quarter) in the 10-year yield, including the mean, median, variance and third central moment. We report the statistics for the full sample and the first and second half of the sample. In addition to sample statistics, we also calculate 90%-confidence intervals using a simple bootstrap, since yields are highly persistent and the serial correlation of their changes is close to zero. The mean and median are negative and, like the variance, have changed little between the first and second sub-sample. By contrast, the third sample moment of quarterly yield changes shifted: over the first half of the sample, the third moments is zero, while over the second half it has turned negative.

The middle and bottom panels of Table 1 report the sample skewness coefficient for yield changes and (negative) futures price changes, respectively, and we report skewness for different frequencies, ranging from one-month to twelve-month changes. The results for bond yields and futures are similar: For the full sample, sample skewness of interest rate changes is statistically close to zero. Its value in split samples depends on the frequency, but it is typically higher and positive over the first half of the sample and negative over the second half of the sample. The magnitude ranges between -1 and 0.5, depending on the specific sample and frequency, is comparable to the skewness estimates for foreign currency and equity index returns reported in the literature (e.g., [Chernov et al., 2018](#), Table 1).

Thus, while the mean and variance of yield changes have not changed, the shape of the asymmetry has shifted noticeably. While the skew of the distribution generally appears slightly positive from 1990 to 2004, it has shifted significantly negative for the period from 2005 to 2021. This empirical pattern suggests that an unconditional, full-sample perspective on skewness may miss interesting features of the distribution of interest rates. Therefore we

Table 1: Summary statistics for changes in 10-year Treasury yield

	Full sample		1990-2004		2005-2021	
<i>Summary statistics of quarterly yield changes</i>						
Mean	-0.05	(-0.13, 0.01)	-0.07	(-0.18, 0.03)	-0.04	(-0.14, 0.05)
Median	-0.02	(-0.10, 0.04)	-0.04	(-0.18, 0.13)	-0.02	(-0.14, 0.06)
Variance	0.23	(0.18, 0.28)	0.23	(0.18, 0.29)	0.23	(0.16, 0.31)
Third moment	-0.03	(-0.06, 0.01)	0.00	(-0.04, 0.04)	-0.05	(-0.11, 0.01)
<i>Skewness of m-month yield changes</i>						
$m = 1$	0.03	(-0.41, 0.52)	0.52	(0.12, 0.96)	-0.49	(-1.20, 0.16)
$m = 2$	-0.41	(-1.27, 0.36)	0.46	(0.15, 0.79)	-1.05	(-2.52, -0.20)
$m = 3$	-0.24	(-0.56, 0.05)	0.01	(-0.35, 0.34)	-0.46	(-1.05, -0.02)
$m = 6$	0.03	(-0.30, 0.37)	0.27	(-0.20, 0.77)	-0.27	(-0.71, 0.19)
$m = 12$	0.36	(-0.10, 0.84)	0.40	(-0.30, 1.10)	0.37	(-0.19, 1.17)
<i>Skewness of m-month (negative) futures price changes</i>						
$m = 1$	-0.14	(-0.63, 0.41)	0.40	(0.02, 0.81)	-0.69	(-1.47, 0.04)
$m = 2$	-0.20	(-0.66, 0.30)	0.32	(0.05, 0.61)	-0.72	(-1.55, -0.01)
$m = 3$	-0.16	(-0.39, 0.06)	0.15	(-0.07, 0.40)	-0.45	(-0.84, -0.14)

Notes: Summary statistics for changes in 10-year Treasury yield and futures prices. Sample period: January 1990 to May 2021.

next turn to conditional yield skewness.

2.3 Realized and implied skewness

In order to measure the skewness of the conditional distribution of yields,

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

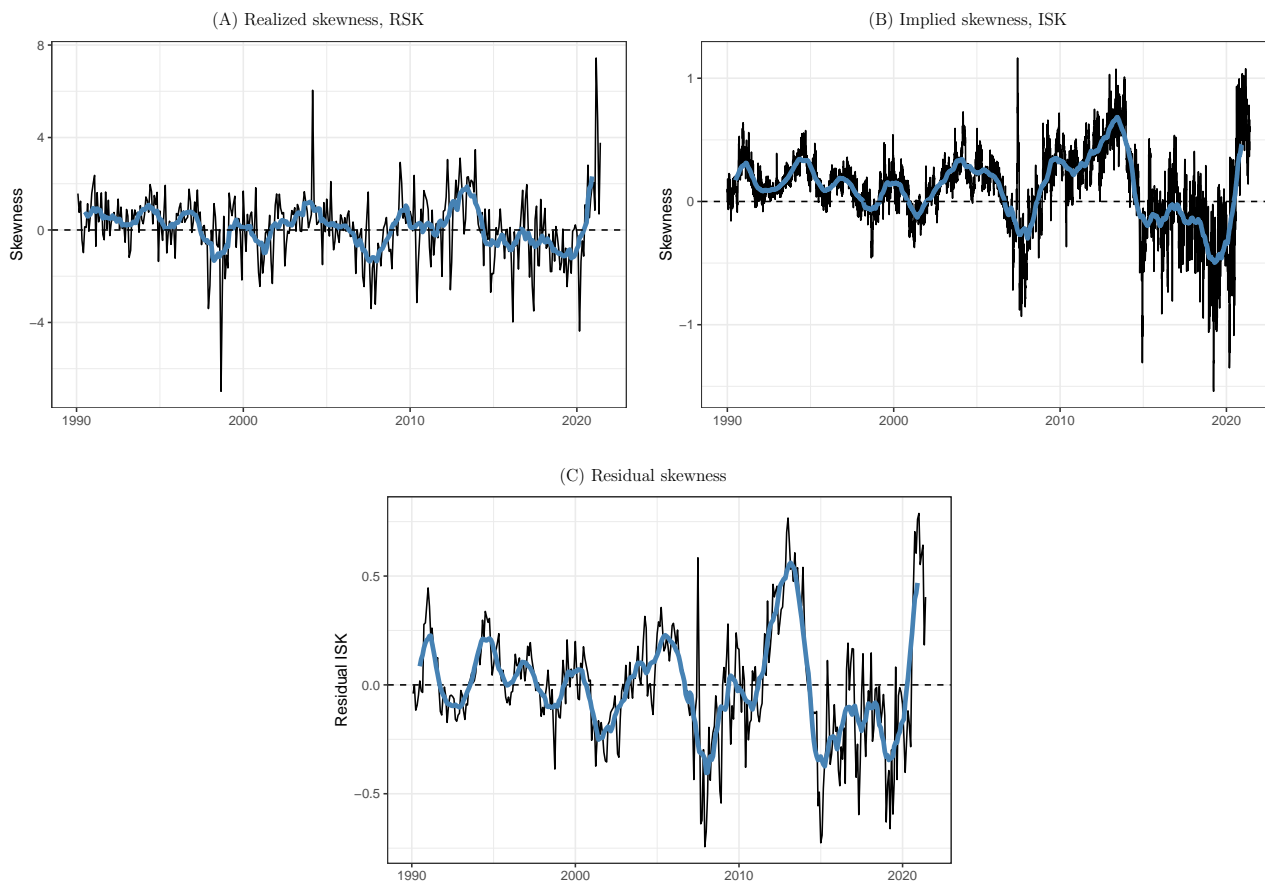
we require estimates of second and third conditional moments. We follow the literature on skewness in stock returns and use both realized and option-implied moments for Treasury futures price changes, using the negative of price skewness to measure yield skewness. We

calculate realized skewness (RSK) at a monthly frequency using daily changes in prices and implied volatility for Treasury futures (Neuberger, 2012, equation 5). Figure 1A plots this time series of RSK, as well as a 12-month moving average. Monthly realized yield skewness is volatile and on average close to zero, but exhibits some persistence and pronounced time variation. During three episodes skewness 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.

Realized skewness allows us to gauge time variation but it is noisy, available only at lower frequencies, and backward-looking. Implied skewness (ISK) does not suffer from these drawbacks. It measures the conditional, risk-neutral skewness of future yield changes based on the conditional moments implied by options on Treasury futures. Details of how we construct ISK are in Appendix A. On every trading day we calculate ISK for each futures contract expiration. For most of our analysis, we then use ISK for the most active option contract, namely the shortest quarterly contract, which has a maturity between about 1 and 3 months. Figure 1B shows a time series of implied skewness that is linearly interpolated to a constant horizon of 2.4 months (the average horizon of all option contracts). Over the full sample, the average level of ISK is positive, with a mean of 0.10 that is significantly different from zero, and a standard deviation of 0.30. But this average level of skewness masks substantial variation in risk perceptions about for future Treasury yields. Similar to realized skewness, ISK has exhibited pronounced cyclical swings over the course of our sample, but the variation is more cyclical and more persistent—the first-order autocorrelation of ISK is 0.95. 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. Before liftoff from the ZLB in 2015 skewness shifted markedly negative, and it averaged -0.21 between 2015 and the end of the sample. As is the case with RSK, the COVID stimulus has changed that with ISK reaching levels of around 1.0 in mid-2020.

More formal statistical analysis is helpful to better understand the time variation that is visually evident in these figures. Specifically, we want to test whether time variation in

Figure 1: Yield skewness



Notes: Panel (A) displays monthly realized Treasury yield skewness, calculated from changes in daily Treasury futures prices and implied volatilities, with a 12-month moving average (blue line). Panel (B) plot daily implied Treasury yield skewness, calculated from options on Treasury futures and interpolated to a constant horizon of 0.2 years, with a 250-day moving average (blue line). Panel (C) shows residual skewness from a regression of monthly implied skewness on yield-curve factors (specification 2 in Table 3), with a 12-month moving average (blue line). Sample period: January 2, 1990, to May 28, 2021.

option-implied skewness is statistically and economically significant. One straightforward way to do so is to assess whether it predicts realized skewness. In Table 2 we present results for various regression specifications, predicting monthly realized skewness with its

Table 2: Predicting realized skewness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
RSK	0.43*** (0.04)		0.20*** (0.06)		0.40*** (0.05)		0.19*** (0.06)
ISK		2.05*** (0.22)	1.51*** (0.27)			2.04*** (0.23)	1.52*** (0.27)
Level				-0.02 (0.06)	-0.02 (0.04)	-0.02 (0.04)	-0.02 (0.03)
Slope				0.27*** (0.08)	0.14** (0.06)	0.05 (0.06)	0.04 (0.06)
Curvature				0.08 (0.55)	0.15 (0.33)	-0.48 (0.34)	-0.31 (0.31)
Constant	0.08 (0.07)	-0.09 (0.08)	-0.06 (0.07)	-0.39 (0.47)	-0.18 (0.32)	0.07 (0.29)	0.06 (0.26)
Observations	376	376	376	376	376	376	376
R ²	0.18	0.23	0.26	0.05	0.19	0.24	0.26

Notes: 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, following [Neuberger \(2012\)](#); *Level*, *Slope* and *Curvature* are the first three principal components of Treasury yields from one to ten years maturity (appropriately scaled). 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.

own lag, option-implied skewness or the shape of the yield curve, all measured at the end of the previous month. To guard against serial correlation due to the persistence of RSK we report Newey-West standard errors (with automatic bandwidth selection). RSK exhibits significant autocorrelation, but lagged values of ISK have even stronger predictive power than lagged values of RSK itself. In a regression that includes both predictors, both are strongly significant (see column 3). The slope of the yield curve has some explanatory power for future RSK, but once we include ISK this is driven out and the information in the yield curve becomes irrelevant. In sum, there is strong evidence for time variation in

the conditional expectation of RSK, and that much of this variation is captured by ISK. This establishes that conditional yield skewness varies over time, and that ISK is a useful forward-looking measure of this conditional skewness.

2.4 Skewness and the yield curve

We now turn to the cyclical nature of the variation in skewness and its macro-financial drivers. First, to provide a visual impression of the relationship between interest rates and skewness, Figure 2 plots annual moving averages of implied skewness and the slope of the Treasury yield curve (calculated as explained below). Skewness tends to increase when the yield curve is steep or steepening. This pattern is most striking during the episodes in 2002-2003 and 2008-2013. In other words, the slope of the yield curve is positively related to yield skewness.


We formalize the evidence on the relationship between skewness and other macro-financial variables by a series of regressions reported in Table 3. In all specifications, the dependent variable is conditional implied skewness, ISK.¹ We use monthly averages of skewness and interest rates. Newey-West standard errors (with automatic bandwidth selection) are reported in parentheses.

The first column in Table 3 reports estimates for a regression on the level and slope of the yield curve. These are calculated as the first two principal components of yields from one to ten years maturity, normalized such that high level/slope are associated with high yields/an upward sloping yield curve. The numbers confirm that the slope is important for skewness: an upward-sloping yield curve is associated with high skewness. Level, by contrast, does not have a statistically significant relationship with skewness.

Including an interaction between level and slope adds substantial explanatory power. The reason is that the slope exhibited a stronger relationship with skewness when the level of yields was low (later in the sample) than when it was high (early in the sample). This

¹To save space, we do not report results for specifications including contemporaneous RSK. Since RSK and ISK are highly correlated, including RSK substantially raises R^2 , but it does not materially affect the statistical relationships documented in Table 3.

Figure 2: Skewness and interest rates



Notes: 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.

pattern is partly driven by the 2008-2015 ZLB period, when the level of yields was low and both skewness and the slope of the yield curve were generally high.

Figure 1C plots the residual from a regression of ISK on level, slope and a level-slope interaction, that is, from the regression specification in column 2 of Table 3. Accounting for the shifts in the yield curve dampens the large cyclical swings, compared to the original series. However, this unexplained portion still has interesting cyclical variation, which we show below to contain substantial 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)		0.04 (0.04)	0.15** (0.07)
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)	0.07** (0.03)	0.07** (0.03)
Constant	-0.11 (0.14)	-0.32 (0.19)	0.07 (0.06)	-0.02 (0.15)	-0.33** (0.14)	-0.45*** (0.16)	-0.68*** (0.19)
Observations	377	377	377	377	377	377	377
R ²	0.15	0.19	0.31	0.33	0.20	0.25	0.29

Notes: Regressions for the level of option-implied yield skewness of the ten-year Treasury yield, using monthly data from January 1990 to May 2021. *Level* and *Slope* are the first two principal components of Treasury yields from one to ten years maturity, scaled to correspond to level and slope of the yield curve; *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). 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.

2.5 Skewness and the ZLB

Over most of the ZLB period beginning in 2008, and since mid-2020, implied skewness was significantly higher than over the rest of the sample. One might hypothesize that skewness generally tends to be high when interest rates are low and close to the ZLB, because lower bound might truncate the left tail, thus making the right tail comparatively longer and

skewness positive. But the following observations suggest that this kind of ZLB effect was not quantitatively important for variation in 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). For example, yields were lower in 2016 than during most of the first ZLB period, yet skewness was mainly negative in 2016.

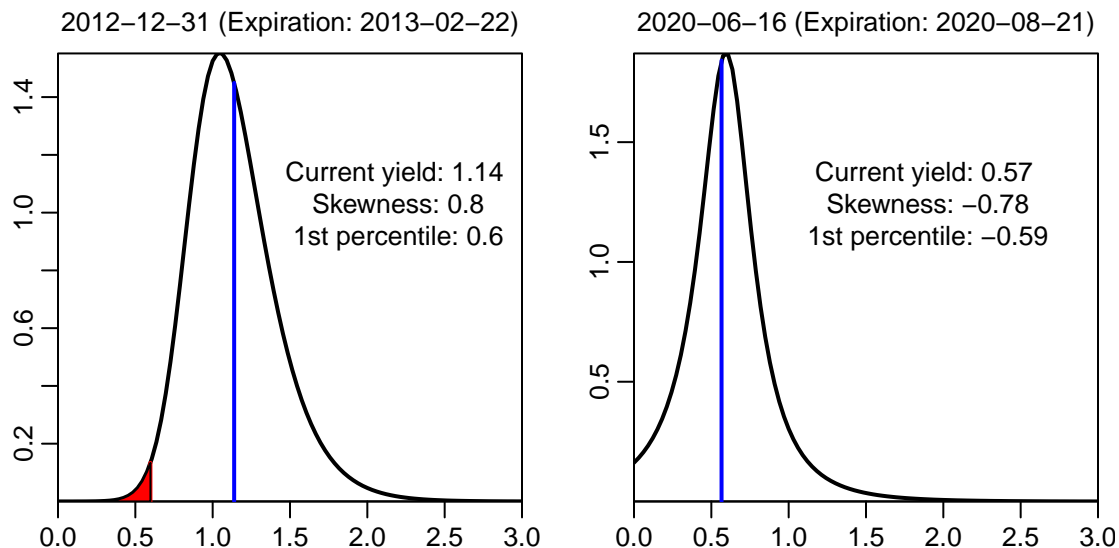
Second, skewness behaved very differently during the two ZLB episodes in our sample. Skewness turned positive when the ZLB was reached in 2008 and remained mainly positive for several years, but then switched to negative in 2014, more than a year before the Fed lifted its policy rate off the ZLB. During 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 commenced a pronounced 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 shows, for two different dates, the implied densities for yields at the option expiration date, obtained from (i) the bond derivative prices, (ii) a normal-inverse-gamma distribution that matches the first four option-implied moments (Eriksson et al., 2009), (iii) an approximate mapping from bond price changes to yield changes (see Appendix A), and (iv) the current CTD bond yield.²

The first date is December 31, 2012, a day with a particularly low yield level (1.14 percent) and a high level of skewness (0.8). The density shows the conditional distribution of yields on February 22, 2013, the expiration date. Even for this extreme example of low yields and high skewness during this episode, 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

²One could also infer the entire risk-neutral distribution from option prices directly using non-parametric methods. That approach, however, is quite sensitive to measurement errors and outliers in the options data, and to the implementation details and specific smoothing techniques (e.g., Ait-Sahalia and Duarte, 2003, Bondarenko, 2003).

Figure 3: Densities for future yields at ZLB



Notes: 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.

by the ZLB, but instead because investors perceived an upward tilted balance of risk and large right tail.³

The second date is June 16, 2020, a day with extremely low yields and *negative* skewness, which was not uncommon during this episode. For this distribution, pertaining to the yield level on August 21, the 1st percentile is deeply negative, at -0.6 percent, suggesting that the ZLB does not eliminate left tails and mechanically lead to positive skewness, at least not during this episode.

The presence of the ZLB certainly affects some aspects of the distribution of interest rates including its skewness. But overall our evidence here and the predictability results in Section 3 suggest that proximity to the ZLB is not an important driver of skewness, and that high conditional skewness during ZLB episodes was not due to a mechanical truncation effect.

³Even a counterfactual distribution with an equally large negative skewness, at -0.8, still has a 1st percentile quite a bit above the ZLB, at 0.3 percent.

Instead, other factors likely affected risk perceptions during this episode, such as perceptions about the likely future course of monetary policy and the risk that long-term yields might return to higher levels. Conditional skewness mainly reflects changes in the outlook of interest rate risk, to which we turn next.

2.6 Monetary policy and the business cycle

Figure 2 also shades monetary policy easing and tightening cycles, which we identify based on changes in the federal funds rate, since shifts in monetary policy are a key driver of shifts in the yield curve (Piazzesi, 2005). This shows that conditional skewness tends to increase during or after monetary easing cycles, most prominently during the easing after the 2000 dot-com bust and during the Great recession and ZLB period. Vice versa, tightening episodes coincide with or precede episodes of falling skewness. We now dig deeper into the relationship between skewness and monetary policy cycles.

To this end, we estimate regressions that include indicator variables for monetary easing and tightening cycles. Cross-correlations reveal that one-year lags of the indicator variables for easing and tightening episodes have the strongest correlation with skewness, so we include these lags instead of contemporaneous indicators in our regressions. Columns 6 and 7 of Table 3 show the results.

Conditional skewness has a statistically strong and economically intuitive relationship with the monetary policy cycle. A regression of ISK on the cyclical indicators, shown in column 6, demonstrates their substantial explanatory power, with an R^2 of 0.31. A third of the variation in conditional yield skewness is explained by the monetary policy cycle. Skewness is high during and soon after monetary easing cycles. Intuitively, these are episodes of upward-tilted interest risk because the Fed has been lowering rates and the next monetary tightening cycle is likely to begin soon. Vice versa, skewness is low during and early after monetary tightening cycles, periods where investors are turning their attention to downward risk to interest rates. Thus, skewness appears to capture the changing balance of interest rate risk over the business cycle.⁴

⁴We have also found evidence for the cyclical behavior of skewness using various business cycle variables, including NBER recession dummies, industrial production *grpwth*, the output gap, the Chicago Fed National

The relationship of skewness with the monetary policy cycle is so strong to even drive out the relationship with the yield curve. When we add these indicators to a regression with RSK and the yield-curve variables, the slope becomes insignificant, as shown in column 7 of Table 3. This finding may be understood in light of the fact that the slope of the yield curve mainly reflects the stance of monetary policy (Rudebusch and Wu, 2008). Our two cyclical indicators apparently provide a more nuanced measure of the monetary policy cycle.

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, we have found that ISK is also closely related to such cyclical indicators. Table 3 shows that changes in the unemployment rate explain about 20 percent of the variation in ISK. The coefficient on the unemployment rate remains statistically significant even after adding yield curve variables.

Overall, we find strong contemporaneous correlations 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 to be high, and vice versa. One take-away from these results is that changes in conditional skewness plausibly represent shifts in perceptions about the future balance of interest rate risk. In addition, the strong contemporaneous correlations with macroeconomic and financial variables make it important to control for these skewness drivers in the following predictive analysis.

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

Activity Index, among others. The two cyclical indicators in Table 3 are so closely related to skewness that when we include them, other macroeconomic variables generally become insignificant. We omit these results for the sake of brevity.

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.⁵ This representation is helpful because it demonstrates that predictability of excess returns can arise from variation in risk premia and from a time-varying bias in beliefs, i.e., from changes in systematic expectational errors. Variation in risk premia has traditionally been the common explanation of empirical results documenting predictability of bond returns or interest rate changes. Recent work, however, has emphasized the possibility that such empirical correlations could be due to the failure of the often implicit assumption of full information rational expectations (FIRE), i.e., to changing biases in beliefs captured by the first two terms in equation (1) (e.g., [Bauer and Swanson, 2021](#), [Bacchetta et al., 2009](#), [Buraschi et al., 2021](#), [Cieslak, 2018](#), [Piazzesi et al., 2015](#)).

This representation guides our empirical work. We first establish whether skewness predicts excess returns. We implement that analysis in two different ways: conventional predictive regressions for excess returns on Treasury bonds, similar to [Cochrane and Piazzesi \(2005\)](#) and many others, and an analysis of high-frequency interest rate changes around FOMC announcements, i.e., monetary policy surprises.⁶ Next, we disentangle the source of predictability by using consensus 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. For the holding period we

⁵Excess 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.

⁶High-frequency interest changes closely correspond to (negative) excess bond returns, since the risk-free return at such frequencies is close to zero.

choose one quarter, because ISK is based on derivative contracts with expiration of about 1-4 months in the future; annual excess returns are more common in empirical work, but we want 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.⁷

A vector of predictors, X_t , is observable at the end of month t . We are most interested in the predictive power of option-implied yield skewness, thus X_t contains ISK. We use the average of ISK over the last five business days of the month to smooth out the high-frequency movements in this series. The predictors in X_t in most cases also include, besides a constant, the level, slope and curvature of the yield curve, which we calculate as the first three principal components of yields with annual maturities from one through ten years. Controlling for the shape of the yield curve is important since a natural null is the hypothesis that it reflects all the information that is relevant for expectations of future interest rates, i.e., the spanning hypothesis investigated by [Duffee \(2011\)](#) and [Bauer and Hamilton \(2018\)](#), among 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 relative to the yields-only specification

⁷As in [Cieslak and Povala \(2015\)](#), we scale excess returns by maturity so that all excess returns have the same duration and similar volatility.

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

Notes: 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 from one to ten years maturity (appropriately scaled); *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, following [Neuberger \(2012\)](#); *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 [Giacoletti 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.

(measured by R^2). The coefficient on ISK is negative and highly significant, indicating that high skewness predicts low bond returns and thus 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 yields—i.e., that the predictive power of ISK is not subsumed by the shape of the yield curve—indicates a violation of the spanning hypothesis.

As a more reliable method of inference that accounts for potential small-sample problems, [Bauer and Hamilton \(2018\)](#) propose a bootstrap method to test the spanning hypothesis in predictive regressions for bond returns. Using their bootstrap procedure leads to a small-sample p -value on the coefficient of ISK that is below 1%. The fact that this small-sample test leads to the same result as the asymptotic standard errors in [Table 4](#) is due to the fact that ISK exhibits only moderate autocorrelation in monthly data (the first-order autocorrelation coefficient is 0.72), which alleviates problems with inference for persistent predictors.

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 substantially 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. Consistent with this view, a regression that includes as its only predictor *residual skewness*, i.e., ISK orthogonalized with respect to the information in the yield curve, yields a highly significant coefficient of -0.34 and R^2 of about five percentage points. This confirms that the new information in conditional skewness is statistically and economically significant for future bond 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. However, 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, in addition to the yield curve predictors, it is only ISK that exhibits significant predictive power, but not RSK.

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.⁸ 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).

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 10-year yield advocated by [Giacoletti et al. \(2021\)](#) (GLS). The sample for this regression is shorter because the GLS variable is currently available only until 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.

⁸We estimate the cycle factor by detrending one- through ten-year yields with a moving average of core CPI inflation, predicting $\bar{r}\bar{x}_{t,t+3}$ using the one-year and the average yield cycles, and calculating the fitted values.

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.

An important question about all of these results is how robust they are across sample periods. To address it we estimate our baseline predictive regression—including yield curve predictors and ISK—for a variety of different sample periods, and the results are shown in Appendix Table B.2. Overall, the predictive power of ISK is robust across different samples. Of particular interest is the role of the first ZLB episode is, since during that time yield skewness was substantially elevated, and Treasury yields increased after the Fed lifted the policy rate off the ZLB, at least for some time. Table B.2 includes results for a pre-ZLB subsample, that is, the sample period ending in November 2008, before the Fed lowered the policy rate to a level near zero. In this subsample the predictive power of conditional skewness was even stronger, indicating that the ZLB episode does not play a unique role in explaining our main results.

In additional, unreported analysis we have investigated further predictive models that control for other variables. First, our earlier estimates in Section 3 documented a strong contemporaneous relationship with implied skewness for a level-slope interaction effect and for indicator variables capturing the state of the monetary policy cycle. Including these variables in the excess return regressions has no material effect on the predictive power of ISK, supporting the view that the variation in ISK *orthogonal* to the yield curve and business cycle indicators—the residual we plotted in Figure 1C—contains relevant information for the future course of interest rates. Second, in recent work Crump and Gospodinov (2019) documented that option-implied skewness in equity markets, measured by the CBOE skew index, predicts Treasury bond returns. We find that in predictive regressions including both the skew index and ISK, the coefficients on both variables are highly significant. Furthermore, in regressions that also include yield-curve factors, ISK is strongly significant while the skew index is insignificant. Lastly, the asset pricing literature has focused on the role option-implied variance (e.g., Choi et al., 2017). Such a measure captures the market uncertainty but does not possess directional information. Indeed, our findings are unchanged when controlling for option-implied variance or volatility using measures calculated from our Treasury options or

the TYVIX index.

Our interest rate data are the smoothed GSW Treasury yields of [Gürkaynak et al. \(2007\)](#), but we have also run predictive regressions using the popular unsmoothed Fama-Bliss Treasury yields. [Cochrane and Piazzesi \(2005\)](#) famously documented in this dataset that a single linear combination of forward rates captures essentially all of the predictive power of the yield curve for future excess returns across bond maturities. Our evidence with the Fama-Bliss data, which we omit for the sake of brevity, also shows that ISK strongly and robustly predicts future bond returns, for both the averaged bond return as well as for individual returns for 2-5 years bond maturities. Importantly, this finding is robust to controlling for the usual yield factors, all five annual forward rates, or the powerful Cochrane-Piazzesi factor.

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 ([Bauer and Swanson, 2021](#); [Cieslak, 2018](#)).⁹

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 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

⁹Monetary 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 \(2021\)](#) supports the view that information effects are likely to be small.

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. For statistical inference, we report White heteroskedasticity-robust standard errors, because ε_t is not serially correlated between FOMC announcements. Our sample contains 213 FOMC announcements from the beginning of 1994 (when the FOMC first started publicly stating a target for the policy rate) 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 results for these regressions. Information in the yield curve alone does not have any predictive power, 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 significant predictive power and explains close to six percent of the variance in the monetary policy surprise.¹⁰ 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 10%, as shown in column (3). The slope's coefficient is statistically negative, while the coefficient of ISK is significantly positive, mirroring the findings for the return regressions in Table 4.

If instead of ISK we include RSK as a predictor, it is also found to exhibit significant predictive power, as shown in column (4). But in this case, the R^2 is only around 3%. Column (5) includes both ISK and RSK with the yield curve factors. The slope and ISK are the most significant predictors in this specification, and the predictive power is similar to the specification without RSK. Similarly to return regressions, the information in implied skewness appears more relevant for predicting future rate changes than the information in

¹⁰ISK is taken as the mean of the daily ISK observations over these 22 days, or roughly one month, before the FOMC meeting. RSK is calculated as in Section 2 but based on sums over the 22 trading days before the FOMC meeting (instead of a calendar month). Moderate changes to these window lengths do not materially affect our results.

Table 5: Predicting FOMC surprises with ISK and RSK

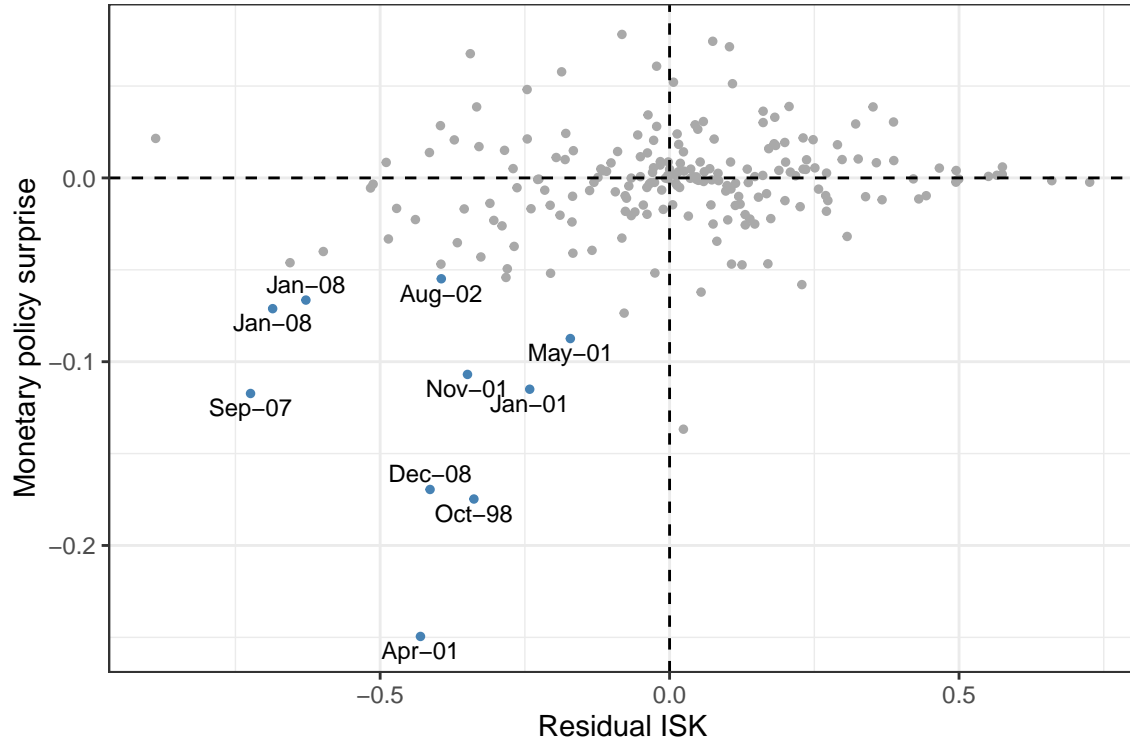
	(1)	(2)	(3)	(4)	(5)
Level	-0.0002 (0.001)		-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Slope	-0.001 (0.002)		-0.007*** (0.003)	-0.003 (0.002)	-0.007*** (0.003)
Curvature	-0.016 (0.013)		-0.025* (0.014)	-0.015 (0.013)	-0.029** (0.014)
ISK		0.030*** (0.009)	0.043*** (0.012)		0.059*** (0.016)
RSK				0.004** (0.002)	-0.005* (0.003)
Constant	0.003 (0.007)	-0.010*** (0.003)	0.018** (0.009)	0.008 (0.008)	0.017* (0.009)
Observations	213	213	213	213	213
R ²	0.009	0.056	0.099	0.027	0.112

Notes: 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 from one to ten years maturity (appropriately scaled) measured on the day before the announcement; *ISK* is option-implied yield skewness, *RSK* is realized yield skewness based on daily changes in futures prices and implied volatilities, following [Neuberger \(2012\)](#), both implied and realized skewness use data over the month (22 trading days) before the FOMC announcement. White heteroskedasticity-robust standard errors are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

realized skewness.

Some previous studies have documented predictive power of macroeconomic and financial variable for FOMC policy surprises. Useful predictors include the federal funds rate and employment growth ([Cieslak, 2018](#)), as well as macroeconomic news, such as the surprise component in the monthly nonfarm payrolls number ([Bauer and Swanson, 2021](#)). Appendix [B.2](#) shows that ISK retains its predictive power even if we control for the other variables.

Figure 4: Skewness and monetary policy surprises



Notes: 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. Top ten most influential observations are shown in blue and labeled.

These results raise the question what type of monetary policy decisions the information in ISK helps predict. Importantly, monetary policy surprises tend to be asymmetric. Whether measured against survey forecasts or the expectations in money market futures, the Fed's policy decisions have more often been easing surprises than tightening surprises, and the largest surprises are typically dovish (Cieslak, 2018; Schmeling et al., 2021). 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 thus captures the information that ISK adds to the information in the yield curve (compare to specification in column 3 in Table 5). A univariate regression for these observations yields an R^2 of nine percent and a t -statistic on the slope coefficient of 3.6 (using White standard errors). The ten most influential observations, as measured by the contribution to the slope coefficient, are shown in blue and labeled with the month containing the FOMC announcement. All ten of these observations are for large dovish surprises, which generally occurred early on in monetary easing cycles, including those starting in January 2001 and September 2007. Essentially all of these surprises were partly anticipated by unusually low levels of skewness ahead of the announcements. An important part of the predictive power of ISK clearly stems from its ability to anticipate dovish surprises during Fed easing cycles.

These turning points for monetary policy are particularly uncertain times. There is ample evidence in the literature that disagreement and uncertainty about the macroeconomic outlook, measured in a number of ways, tend to be high during recessions and early in monetary easing cycles (e.g., [Patton and Timmerman, 2010](#); [Dovern et al., 2012](#); [Schmeling et al., 2021](#)). This raises the possibility that expectational errors may be part of the explanation, which we will investigate next.

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). Under the FIRE hypothesis, the only explanation for these results are changes in (subjective) bond risk premia. However, more generally the established predictability may also arise in part from systematic forecast errors that are related to skewness, as indicated by the right-hand side of equation (1). To investigate this possibility, we use survey forecasts as proxies for subjective expectations.

Specifically, we calculate survey forecast errors for interest rates using the Blue Chip Financial Forecasts (BCFF). This is a monthly survey that contains forecasts for the current quarter (nowcasts) and each of the next five quarters.¹¹ The survey respondents provide

¹¹The surveys conducted before 1997 extend out only four quarters.

their forecasts for the future quarterly averages for a number of different interest rates. We consider both the ten-year Treasury yield, which is a natural starting point since implied skewness is based on options for Treasury securities with maturities between 6.5 and 10 years, and the federal funds rate, motivated by the central role of the Fed’s policy rate for bond markets and influential previous work on these forecast errors by [Cieslak \(2018\)](#). For both interest rates we obtain daily values from FRED and calculate forecast errors as the difference between the quarterly average realized value and the consensus forecast, which is the arithmetic mean of the individual forecasts.

For each forecast horizon from $h = 0$ to 5 quarters we run monthly predictive regressions of the forecast errors on information available at the time of the survey. Specifically, we estimate the regression

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

where t indexes the month of the survey forecast, $y_{q(t,h)}$ 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.¹² 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.¹³ 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, see [Cochrane and Piazzesi \(2005\)](#). We use $3(h + 1)$ lags to estimate the covariance matrix of the parameter estimates.

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

¹²For example, in January, February, and March, forecasts for $h = 1$ are for the average over the second quarter (April to June).

¹³The 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

<i>(A) Ten-year yield</i>	Current	1Q ahead	2Q ahead	3Q ahead	4Q ahead	5Q ahead
ISK	0.13*** (0.04)	0.29*** (0.11)	0.29* (0.17)	0.29 (0.23)	0.38 (0.27)	0.36 (0.29)
Constant	-0.07*** (0.02)	-0.23*** (0.05)	-0.39*** (0.08)	-0.54*** (0.11)	-0.72*** (0.13)	-0.94*** (0.16)
R ²	0.03	0.04	0.02	0.02	0.02	0.03
ISK	0.13*** (0.04)	0.33*** (0.10)	0.37** (0.15)	0.43** (0.17)	0.58*** (0.18)	0.70*** (0.17)
Level	0.004 (0.01)	0.01 (0.02)	0.01 (0.04)	0.02 (0.04)	0.03 (0.04)	-0.02 (0.07)
Slope	-0.005 (0.02)	-0.05 (0.04)	-0.09 (0.06)	-0.15* (0.09)	-0.20* (0.10)	-0.31** (0.13)
Curvature	-0.01 (0.10)	0.14 (0.22)	0.32 (0.27)	0.35 (0.39)	0.39 (0.50)	0.29 (0.48)
Constant	-0.07 (0.05)	-0.21 (0.13)	-0.35* (0.19)	-0.45 (0.28)	-0.56 (0.39)	-0.40 (0.56)
R ²	0.03	0.05	0.06	0.07	0.10	0.18
<i>(B) Federal funds rate</i>	Current	1Q ahead	2Q ahead	3Q ahead	4Q ahead	5Q ahead
ISK	0.07*** (0.03)	0.31*** (0.11)	0.57*** (0.20)	0.85*** (0.30)	1.13** (0.44)	1.36** (0.53)
Constant	-0.05*** (0.01)	-0.18*** (0.05)	-0.37*** (0.11)	-0.58*** (0.19)	-0.82*** (0.27)	-1.02*** (0.36)
R ²	0.03	0.07	0.08	0.09	0.09	0.12
ISK	0.10*** (0.04)	0.41*** (0.11)	0.80*** (0.19)	1.22*** (0.24)	1.57*** (0.29)	1.81*** (0.39)
Level	0.001 (0.003)	-0.02 (0.02)	-0.06 (0.04)	-0.09 (0.07)	-0.13 (0.09)	-0.14 (0.16)
Slope	-0.01 (0.01)	-0.04 (0.03)	-0.10 (0.07)	-0.14 (0.11)	-0.18 (0.15)	-0.24 (0.17)
Curvature	-0.14** (0.07)	-0.42** (0.17)	-0.68* (0.36)	-1.11** (0.52)	-1.50*** (0.58)	-1.62*** (0.63)
Constant	0.02 (0.03)	0.12 (0.10)	0.30 (0.22)	0.48 (0.34)	0.59 (0.48)	0.50 (0.56)
R ²	0.08	0.14	0.16	0.19	0.20	0.22
Observations	372	371	368	365	362	276

Notes: 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. 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.

beyond. The R^2 of these regressions increases with the horizon from three to 18 percent.¹⁴

Previous work has documented that systematic expectational errors are particularly pronounced for short-term yields (Cieslak, 2018; Brooks et al., 2020; Schmeling et al., 2021). The results for the federal funds rate in the bottom panel of Table 6 confirm this finding. They show that conditional skewness is a very powerful predictor of short rate forecast errors. Even univariate regressions have R^2 ranging from seven to twelve 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.

An important robustness check relates to the issue is the role of ex-ante survey forecast revisions and informational rigidities. Coibion and Gorodnichenko (2015) found forecast revisions to be strong predictors of macroeconomic forecast errors, suggesting an important role of informational rigidities for the expectation formation process. In additional, unreported analysis we found that for the ten-year Treasury yield, ex-ante survey revisions do not have any predictive power for ex-post forecast errors. By contrast, in the case of the federal funds rate, survey revisions do contain relevant information for future forecast errors. These findings line up with the suggestion by Coibion and Gorodnichenko that informational rigidities play a larger role for less persistent time series, since short-term rates are less persistent than long-term rates. Importantly, for both the ten-year yield and the federal funds rate, including ex-ante survey revisions as additional predictors leaves our results materially

¹⁴Appendix B.3 present results for a similar analysis using the consensus forecasts in the Survey of Professional Forecasters. There we also find strong predictive power of conditional skewness for survey forecast errors at all forecast horizons.

unchanged.¹⁵

Under the FIRE hypothesis, forecast errors should be unpredictable using information that was publicly available at the time the forecast was made. In line with other recent work, our results indicate that this hypothesis is unlikely to hold for interest rate forecasts. Importantly, our evidence suggests that the correlation of ISK with future interest rates documented in Sections 3.1– 3.2 is unlikely to be due entirely to its correlation with (subjective) risk premia. Instead, it appears that conditional yield skewness is systematically related to the difference between subjective and true/statistical expectations about future yields and bond returns, that is, to persistently biased beliefs of investors. Specifically, we have run predictive regressions for $y_{t+h} - E_t^s(y_{t+h})$, where y_t is a generic interest rate, and the results have shown that $E_t(y_{t+h}) - E_t^s(y_{t+h})$ varies strongly over time. This indicates that variation in the first term on the right-hand side of equation (1), changes in expectational errors, contribute to time variation in statistical bond risk premia and the predictive power of conditional yield skewness.

Another take-away from our results is that improvements in forecasts about the course of monetary policy appear to be the main 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.¹⁶ Consistent with these findings, our predictive regressions for excess bond returns for different maturities, reported in Appendix Table B.1, showed the largest R^2 for the one-year bond, the shortest maturity in that analysis. Finally, the results for monetary policy surprises around FOMC announcements also support the same conclusion. Overall, our evidence suggests that ISK has the strongest predictive power for short-term interest rates, and in particular for declines in short-term rates the resulted from dovish surprises in the Fed’s policy actions during easing cycles.

¹⁵Cieslak (2018) also found that the “predictability of FFR forecast errors cannot be explained with information rigidities such as sticky or noisy information” (fn. 14).

¹⁶We have carried out similar analysis for yields of shorter maturities (1y, 2y, 5y), omitted for the sake of brevity, which showed that the predictive power of ISK systematically increases when the maturity gets shorter.

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 , $E_t r x_{t,t+h} = E_t p_{t+h}^{(n-h)} - p_t^{(n)} - y_t^{(h)}$, can be written as

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

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](#)).¹⁷ This decomposition corresponds to a log version of equation (1). To quantify the importance of these two terms, we use the following variance decomposition:

$$\text{Var}(E_t r x_{t,t+h}^{(n)}) = \text{Cov}(E_t^s r x_{t,t+h}^{(n)}, E_t r x_{t,t+h}^{(n)}) + \text{Cov}(ee_t, E_t r x_{t,t+h}^{(n)}).$$

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. Thus t indexes quarters and $h = 1$. The maturities n we consider are 5, 9, 21 and 41 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 $E_t^s(r x_{t+1}^{(n)}) = -(n-1)E_t^s(y_{t+1}^{(n-1)}) + n y_t^{(n)} - y_t^{(1)}$.¹⁸ To estimate statistical expectations $E_t y_{t+1}^{(n-1)}$,

¹⁷[Schmeling et al. \(2021\)](#) denote as “expectation errors” the difference of expected future interest rates and survey expectations, which explains a large share of realized excess returns (on fed funds futures). This differs from our definition and empirical approach: We call such a difference a “survey forecast error” and reserve the term expectational error as the (scaled) difference between survey expectations and statistical expectations of future interest rates. Since we decompose expectations, the surprise component $y_t - E_t y_{t+h}$ is not part of our calculations.

¹⁸This 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 \(2021\)](#) use a different methodology in which they calculate subjective expected returns from survey-implied zero-coupon yields.

which are used in the calculations both of expectational errors ee_t and of statistical bond risk premia $E_t rx_{t,t+1}^{(n)}$, 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. All observed yields with maturities of one year or longer are calculated from the GSW yield curves, and the one-quarter yield is the three-month T-bill rate.

Table 7: Risk premia and expectational errors

Mat.	Yields only			With ISK		
	$Var(E_t^{(n)} rx_{t,t+1}^{(n)})$	RP (%)	EE (%)	$Var(E_t rx_{t,t+1}^{(n)})$	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

Notes: Variance of statistical bond risk premia, i.e., of expected excess returns $E_t rx_{t,t+1}^{(n)}$, and relative contributions, in percent, of survey-based subjective risk premia (RP), $E_t^s rx_{t,t+1}^{(n)}$, and expectational errors (EE), $ee_t = (n - h) (E_t^s y_{t+h}^{(n-h)} - E_t y_{t+h}^{(n-h)})$. 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, $E_t y_{t+h}^{(n-h)}$, 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 h is one quarter. For details see text.

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.

3.5 Skewness during the COVID pandemic


As we noted earlier, skewness has reached all-time high values in the wake of the COVID 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 period immediately preceding the COVID lockdown and the COVID period itself serve as a nice showcase of our findings.

Figure 5 displays the main actors in the reported evidence: ISK, the ten-year Treasury yield and its survey forecasts from BCFF, and the slope of the yield curve, measured 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 10-year yields during this period? Was it related to expectational errors of forecast surveys?

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 the 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.

Figure 5: Skewness and interest rates since 2019



Notes: 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.

The COVID 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 COVID pandemic.

4 A potential explanation: heterogeneous beliefs

We have documented the following empirical facts. First, Treasury yields exhibit time-varying conditional skewness, which is related to the shape of the yield curve. Second, a simple measure of the skewness, ISK, contains substantial predictive power for future bond returns. Third, this predictive power is particularly strong when both the yield curve

and ISK are included in forecasting regressions, violating the spanning hypothesis. Fourth, ISK also predicts interest rate forecast errors based on surveys of professional forecasters, with expectational errors being quantitatively important for statistical bond risk premia, in violation of the FIRE hypothesis. Fifth, information in skewness is particularly useful for predicting dovish monetary easing surprises early on in recessions, when disagreement about the economic outlook is particularly high.

A natural question is whether all this evidence is consistent with an economic mechanism. We show that this is indeed the case, using a simple two-agent heterogeneous-beliefs (HB) model along the lines of [Basak \(2005\)](#). As a novel result, we show that although the fundamental shocks of the model are Gaussian, time-varying skewness arises endogenously in this framework. Differences in beliefs are directly linked to expectational errors, giving rise to a time-varying wedge between subjective and objective beliefs. We derive the bond pricing implications using tools from [Ehling et al. \(2018\)](#), and show that the model implies that yields are non-linear functions of state variables thereby generating non-normal distribution of yields and breaking the linear spanning of bond excess returns with yields. Time-varying skewness is directly linked to bond risk premia, consistent with our evidence.

In what follows we describe the main assumptions of the model and key results. Details and derivations are provided in [Appendix C](#). The model is set in continuous time. Consumption is exogenous and follows a geometric Brownian motion,

$$dC_t/C_t = \mu dt + \sigma dz_t.$$

There are two agents in the economy who disagree about the true dynamics of consumption growth, specifically about the mean growth rate μ . Agent 1 knows the true value of μ , and thus has rational expectations. Agent 2, on the other hand, does not know this and has to form beliefs μ_t^s . Our interpretation is that agent 1 represents true, statistical expectations, while agent 2 represents the consensus survey forecasts which are inconsistent with FIRE.¹⁹ By assuming that agent 1 has rational expectations, we directly link disagreement to expectational errors, $\mu - \mu_t^s$, i.e., to the difference between (true) statistical expectations and (biased) survey-based expectations about consumption growth.

¹⁹This perspective is consistent with [Reis \(2020\)](#) who posits that bond traders are better informed than the general public, which is proxied by surveys.

We assume that subjective beliefs follow

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

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 of agent 2 is not important for our results, so we do not need to take a stand on which specific mechanism is at play.²⁰ What is important for our model is that the agents disagree about expected consumption growth, which they generally will since FIRE does not hold for agent 2.

Disagreement, which corresponds to the bias in agent 2 beliefs, is measured as

$$\Delta_t \equiv \frac{\mu - \mu_t^s}{\sigma},$$

and becomes the key state variable of our model. Gaussian dynamics of μ_t^s translate into Gaussian dynamics of Δ_t .

We assume that agents have power utility with risk aversion γ . The appendix derives the equilibrium consumption allocations and solves for the real interest rate, bond yields, and risk premia. The expression for the real short rate helps understand the basic intuition of the model:

$$r_t = \rho + \gamma\mu - \frac{1}{2}\gamma(\gamma + 1)\sigma^2 + \frac{1}{2}(1 - \gamma^{-1})f(\lambda_t)(1 - f(\lambda_t))\Delta_t^2.$$

The terms in the expression for r_t are (i) the rate of time preference, (ii) consumption smoothing, (iii) precautionary savings, and (iv) the key “risk-sharing” effect, arising whenever $\gamma \neq 1$. This last term involves the likelihood ratio λ_t and the nonlinear sharing rule $f(\lambda)$, both of which are defined in the appendix. It is this risk-sharing effect that is new in the HB model and generates skewness.

²⁰Our evidence is based on the consensus forecasts, and one would need evidence about both consensus and individual forecasts to distinguish between information rigidities and diagnostic expectations, as argued by [Bordalo et al. \(2020\)](#).

The intuition is as follows: When $\Delta_t \neq 0$, the investor who thinks output will be high sells the bonds short, and an investor with the opposite view matches the position on the other side to clear the markets. Ex ante, each investor expects to capture wealth from the other investor and, hence, both expect future consumption to be higher than without disagreement about output. It is for this reason that the real interest rate and bond yields depend on the dispersion in beliefs between the agents.


Bond yields are generally a quadratic function of Δ_t and are non-normally distributed despite the Gaussian state variables. Thus, yields have time-varying skewness, and conditional skewness is related to expectational errors. Furthermore, the appendix shows that disagreement and expectational errors are connected to bond risk premia. The effect is non-linear because of the non-linear dependence of yields on Δ_t , thus a linear combination of yields cannot span bond excess returns. These features of the model are qualitatively consistent with the empirical evidence summarized above.

To understand how skewness is related to expectational errors, we have to resort to simulations. For each initial value Δ_0 , ranging from zero to two, we simulate 100,000 paths of the bias in beliefs Δ_t and the real interest rate r_t , over a one-quarter horizon that matches maturity of options in our data. We calibrate the parameters as follows: Risk aversion is $\gamma = 5$, annualized volatility of consumption growth is 3.5%, $\sigma = 0.035$, and the persistence of biased beliefs, captured by the mean-reversion coefficient, is $\kappa = 0.1$. We calculate the sample skewness coefficient of the interest rate changes over one quarter conditional on Δ_0 . Figure 6 plots this skewness coefficient against the dispersion in beliefs, Δ_0^2 .

Yield skewness depends on the dispersion in beliefs about mean endowment in a non-linear but monotone way. If Δ_0^2 is small, skewness is positive. Skewness then turns negative for values of Δ_0^2 exceeding 1. For $\Delta_0^2 = 4$, meaning that $\mu_t^s = \mu \pm 2\sigma$, skewness reaches about -1.75. Thus, for reasonable values of disagreement/bias in beliefs about consumption growth, our model is capable of generating interest rate skewness in the range observed in the data.

In light of this tight connection between skewness and disagreement, it becomes clear that the model qualitatively matches the evidence we have documented in this paper. Conditional skewness of interest rates is time-varying as a result of changes in disagreement and biases in beliefs, Δ_t . Bond risk premia are related to Δ_t and skewness, but in a non-linear fashion

Figure 6: Bias in beliefs and skewness



Notes: Sample skewness of interest rates simulated from HB model. For each starting value of disagreement, Δ_0 , 100,000 paths of the real short rate are simulated over a one-quarter horizon, using daily time increments, and the skewness coefficient is calculated for the distribution of the final value of the short rate. Parameter settings are risk aversion $\gamma = 5$, volatility of consumption growth $\sigma = 0.035$, and persistence of biased beliefs captured by the mean-reversion coefficient $\kappa = 0.1$.

so that the spanning hypothesis does not hold and information in both the yield curve and skewness is required to capture changes in risk premia. Finally, there is an important role for expectational errors, which are directly tied to disagreement and risk premia.

The model assumes heterogeneous, biased beliefs about expected consumption growth in order to link expectational errors to interest rate skewness and bond risk premia. Table 8 provides empirical support for this assumption. It documents that survey forecast errors for real GDP growth are predictable based on interest rate skewness, measured by ISK. Fore-

Table 8: Predicting real GDP forecast errors

	(1)	(2)	(3)
Level	0.16 (0.12)	0.15 (0.11)	
Slope	0.13 (0.22)	-0.02 (0.20)	
Curvature	-1.39 (0.94)	-1.94** (0.86)	
ISK		1.77*** (0.60)	1.34* (0.71)
Revision	-0.51 (0.82)	-0.17 (0.69)	-0.34 (0.34)
Observations	123	123	123
R ²	0.06	0.10	0.03

Notes: Regressions of forecast errors for four-quarter real GDP growth (over the current and subsequent three quarters) in the Survey of Professional Forecasters, using surveys from monthly surveys from January 1990 to May 2021. *ISK* is option-implied yield skewness, averaged over the month preceding the survey deadline; *Level*, *Slope* and *Curvature* are the first three principal components of end-of-month Treasury yields from one to ten years maturity, 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.

cast errors are calculated as the difference between realized GDP growth over the current and next three quarters and the consensus (median) expectation in the Survey of Professional Forecasters. As in the earlier results, predictability using *ISK* is stronger when it is used jointly with the yield curve. This evidence suggests that expectational errors about macroeconomic fundamentals, and real growth in particular, are quantitatively important, and supports the model-implied channel connecting biased beliefs about consumption growth with non-normality of bond yields and bond risk premia.

Finally, our model also matches the empirical pattern that the predictive power of skewness is strongest around turning points of monetary policy, when disagreement and uncertainty

about the economic outlook is particularly high.²¹ We documented in Section 3.2 that ISK tends to be unusually low during those times and successfully anticipates Fed easing surprises. According to our model, times of high macroeconomic disagreement indeed coincide with low/negative interest rate skewness, as shown in Figure 6. In other words, this important cyclical empirical pattern in disagreement, skewness, and predictions for interest rates is very much consistent with our model.

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 simple theoretical framework with heterogeneous beliefs.

Our results have implications for asset pricing, macroeconomic forecasting, and investment practice. Forecasters and investment managers in particular are likely to benefit from paying attention to implied yield skewness. Implied rate skewness—which is available in real time, does not require any model or estimation, and can be calculated reliably at a daily frequency—enables forecasters and investors to better gauge the current balance of interest rate risk, improving their forecasts and investment decisions that are predicated on the rates outlook.

²¹A subtle issue is that in our model, disagreement is between the biased belief (consensus survey forecast) and the true expectation, and the latter cannot be measured directly. Our implicit assumption is that high disagreement at policy turning points is associated with elevated expectational errors.

Appendix

A Option-implied moments of Treasury yields

Our Treasury derivatives data are end-of-day prices of Treasury futures and options from CME.²² We focus on the 10-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. In some cases we linearly interpolate option-implied moments to a constant horizon, and then we use 0.2 years as the horizon which is about the average maturity of all option contracts (across all expirations, strikes and put/call prices), and interpolate based on the data for the two expirations surrounding this horizon.

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

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

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 $C(K)$ and $P(K)$.

²²For details see <https://www.cmegroup.com/trading/interest-rates/us-treasury.html>.

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 relative moneyness greater than 15, i.e., $(F_t - K)/\sqrt{(T - t)\sigma^2}$ is at most 15 in absolute value (options that are further out of the money tend to have unreliable/implausible IVs),
- are too far out of the money, with 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, for each trading day t and option expiration T , 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 abstracts 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, since 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 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 maturities ranging from one to ten years, 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 from one to ten years maturity (appropriately scaled). Reverse regression standard errors, using 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.2 reports estimates of our baseline specification—including yield curve predictors and ISK—over a variety of different samples. The predictive power of ISK is generally robust across different samples.

Table B.2: Predicting excess returns: different sample periods

	Full	Pre-2000	Post-2000	Pre-ZLB	Post-crisis	Pre-2018
	(1)	(2)	(3)	(4)	(5)	(6)
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 (0.10)	-1.09* (0.60)	-0.22* (0.12)	-0.27 (0.29)	-0.31** (0.15)	-0.27** (0.12)
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 from one to ten years maturity (appropriately scaled). Reverse regression standard errors, using 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.3: Predicting different monetary policy surprises

	PC1	Target	Path	FF1	FF4	ED4
ISK	0.03*** (0.01)	0.03* (0.02)	0.03*** (0.01)	0.02** (0.01)	0.04*** (0.01)	0.05*** (0.01)
Constant	-0.01*** (0.003)	-0.003 (0.01)	-0.004 (0.004)	-0.01** (0.003)	-0.01*** (0.004)	-0.02*** (0.01)
R ²	0.06	0.03	0.03	0.03	0.05	0.06
ISK	0.04*** (0.01)	0.05** (0.02)	0.05*** (0.02)	0.03** (0.01)	0.05*** (0.02)	0.08*** (0.02)
Level	-0.001 (0.001)	-0.004 (0.002)	0.001 (0.002)	-0.002 (0.001)	-0.002 (0.002)	-0.001 (0.002)
Slope	-0.01*** (0.003)	-0.01 (0.004)	-0.01*** (0.005)	-0.001 (0.003)	-0.01** (0.003)	-0.01*** (0.01)
Curvature	-0.02* (0.01)	-0.04* (0.03)	-0.01 (0.02)	-0.03* (0.02)	-0.03 (0.02)	-0.04* (0.02)
Constant	0.02** (0.01)	0.04** (0.02)	0.02 (0.01)	0.01 (0.01)	0.02* (0.01)	0.03* (0.02)
R ²	0.10	0.06	0.08	0.06	0.08	0.11

Predictive regressions for alternative measures of monetary policy surprises: *PC1* is the same measure used in Tables 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.

Table B.3 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. Overall, the predictive power of ISK for monetary policy surprises is very robust across different measures of these surprises.

Table B.4 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

Table B.4: Predicting FOMC surprises with ISK and macro variables

	(1)	(2)	(3)	(4)
ISK	0.031*** (0.009)	0.028*** (0.009)	0.032*** (0.009)	0.021*** (0.008)
FFR	-0.001 (0.001)			
Annual employment growth	0.455*** (0.169)			
BBK index		0.010** (0.004)		
Change in employment			0.057*** (0.018)	
S&P 500 return				0.135*** (0.049)
Constant	-0.013*** (0.004)	-0.008*** (0.003)	-0.016*** (0.004)	-0.011*** (0.003)
Observations	213	213	213	213
R ²	0.084	0.107	0.130	0.128

Notes: 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 (2021). White heteroskedasticity-robust standard errors are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

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 the macroeconomic variables considered by Bauer and Swanson (2021): 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. In all three

cases, both ISK and the Bauer-Swanson predictor 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, the forecast target is the quarterly average for the constant-maturity 10-year yield from the Fed's H.15 statistical release. Forecasts are reported for the current quarter (nowcasts) and each of the subsequent four quarters. As the SPF consensus forecast we take the median of the individual forecasts.

We run predictive regressions of the form

$$y_{t+h} - \hat{y}_t^{(h)} = \beta' X_t + \varepsilon_{t,t+h}, \quad (\text{B.1})$$

where t indexes the quarterly SPF surveys, y_t is the average 10-year yield in quarter t , $\hat{y}_t^{(h)}$ is the survey consensus forecast made in quarter t for the average 10-year yield in quarter $t+h$, h ranges from 0 to 4, X_t is a vector with predictors, and $\varepsilon_{t,t+h}$ is a forecast error. To ensure that the predictors X_t are observable at the time the forecast is made, we take observations on the day before the response deadline of the survey. Because the forecast errors $\varepsilon_{t,t+h}$ are serially correlated we use Hansen-Hodrick standard errors with h lags.

Table B.5 shows the results. For each forecast horizon, we estimate two specifications, one with ISK only, and one that also includes yield factors. We find that ISK has statistically significant predictive power for all forecast horizons. The specifications that also include yield curve factors show that the slope tends to have additional predictive power for $h > 0$. As before, we find that the slope has a negative coefficient while ISK has a positive coefficient.

C A model of biased beliefs

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

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

The dynamics the expectational error Δ_t are:

$$d\Delta_t = -\kappa \Delta_t dt - \frac{\delta}{\sigma} dz_t.$$

Table B.5: Predicting yield forecast errors in the SPF

	Current		1Q ahead		2Q ahead		3Q ahead		4Q ahead	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
ISK	0.16*** (0.04)	0.19*** (0.04)	0.48*** (0.12)	0.56*** (0.10)	0.57*** (0.18)	0.69*** (0.17)	0.57*** (0.19)	0.75*** (0.26)	0.59* (0.33)	0.83*** (0.29)
Level		0.002 (0.01)		-0.003 (0.03)		-0.01 (0.03)		0.005 (0.06)		0.02 (0.06)
Slope		-0.01 (0.01)		-0.05* (0.03)		-0.10** (0.05)		-0.12 (0.08)		-0.16 (0.10)
Curve		-0.17** (0.09)		-0.13 (0.17)		0.04 (0.26)		-0.04 (0.38)		0.02 (0.46)
Constant	-0.08*** (0.02)	-0.01 (0.04)	-0.25*** (0.06)	-0.09 (0.10)	-0.40*** (0.08)	-0.20 (0.15)	-0.54*** (0.07)	-0.30 (0.27)	-0.69*** (0.14)	-0.45 (0.38)
Observations	118	118	117	117	116	116	115	115	114	114
R ²	0.09	0.14	0.08	0.10	0.07	0.09	0.05	0.08	0.04	0.09

Predictive regressions for forecast errors in the h -quarter ahead median forecast for the ten-year Treasury yield 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. *ISK* is option-implied yield skewness, *Level*, *Slope* and *Curvature* are the first three principal components of Treasury yields from one to ten years maturity (appropriately scaled), measured on the day before the survey deadline. All predictors are measured on the day before 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.

In particular the equation implies that Δ_t is Gaussian variable. In our simulation exercise below, we will assume that $\delta = \kappa\sigma$ (which would hold if there was constant-gain learning).

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

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

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,

$$d\lambda_t = -\Delta_t\lambda_t dz_t.$$

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

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

Then

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

These expressions imply exponentially quadratic form in the state Δ_t . Thus, bond prices are weighted averages of exponentially quadratic functions of a Gaussian state variable. The weights $w_t^{(k)}$, which add up to 1, are affected by the bias in beliefs via λ_t . The real short rate is obtained by applying Ito's lemma to ξ_t and picking out the drift of the result.

Bond risk premia. If $\Delta_t = 0$, the risk premia in the economy 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.1). To

demonstrate the dependence of the bond risk premium on Δ_t , we first write the real bond price as

$$B_{t,T} = \sum_{k=0}^{\gamma} w_t^{(k)} B_{t,T}^{(k)},$$

where $B_{t,T}^{(k)}$ are the artificial (exponential quadratic) bond prices corresponding to the SPD $\xi_t^{(k)}$. Then,

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

The expected bond return in each artificial economy k is going to be the corresponding 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:

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

where α , and β reflect sensitivities of a bond price w.r.t. to its factors (see, e.g., [Ahn et al., 2002](#) for explicit expressions). Thus, the risk premium is:

$$\begin{aligned} \frac{1}{dt} E_t \left(\frac{dB_{t,T}}{B_{t,T}} - r_t dt \right) &= \sum_{k=0}^{\gamma} \frac{w_t^{(k)} B_{t,T}^{(k)}}{B_{t,T}} \times \\ &\quad \left[\tilde{\alpha}_t^{(k)} + \beta_t^{(k)} \gamma^{-1} k \Delta_t + \left(\frac{k}{\gamma} \left(1 - \frac{k}{\gamma} \right) - \left(1 - \frac{1}{\gamma} \right) f(\lambda_t) (1 - f(\lambda_t)) \right) \frac{\Delta_t^2}{2} \right]. \end{aligned}$$

This expression has two implications central to our findings. First, it connects non-normality of bond yields and, in particular, their skewness to the bond risk premiums via Δ_t . Second, the expression explains the evidence about the violations of the bond spanning hypothesis. Indeed, expected excess return is a non-linear function of Δ_t . Thus, it cannot be captured by a linear combination of yields

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.
- Ait-Sahalia, Yacine and Jefferson Duarte (2003) “Nonparametric Option Pricing under Shape Restrictions,” *Journal of Econometrics*, 116 (1–2), 9–47.
- 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.
- Balduzzi, Pierluigi, Giuseppe Bertola, and Silverio Foresi (1997) “A model of target changes and the term structure of interest rates,” *Journal of Monetary Economics*, 39 (2), 223–249.
- 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 (2020) “Interest Rates Under Falling Stars,” *American Economic Review*, 110 (5), 1316–1354.
- Bauer, Michael D. and Eric T. Swanson (2021) “An Alternative Explanation for the ‘Fed Information Effect’,” Working Paper 2020-06, Federal Reserve Bank of San Francisco.

- 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.
- Bondarenko, Oleg (2003) “Estimation of Risk-Neutral Densities Using Positive Convolution Approximation,” *Journal of Econometrics*, 116 (1–2), 85–112.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020) “Over-reaction in macroeconomic expectations,” February, *American Economic Review*, forthcoming.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018) “Diagnostic expectations and credit cycles,” *Journal of Finance*, 73 (1), 199–227.
- Brooks, Jordan, Michael Katz, and Hanno Lustig (2020) “Post-FOMC Announcement Drift in U.S. Bond Markets,” December, Working Paper.
- Buraschi, Andrea, Ilaria Piatti, and Paul Whelan (2021) “Subjective Bond Returns and Belief Aggregation,” *Review of Financial Studies*, forthcoming.
- 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.
- Chernov, Mikhail, Jeremy Graveline, and Irina Zviadadze (2018) “Crash Risk in Currency Returns,” *Journal of Financial and Quantitative Analysis*, 53 (1), 137–170.
- 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. and Nikolay Gospodinov (2019) “Deconstructing the Yield Curve,” Staff Report 884, Federal Reserve Bank of New York.
- Das, Sanjiv (2002) “The surprise element: jumps in interest rates,” *Journal of Econometrics*, 106 (1), 27–65.
- 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 (2021) “Asymmetric Monetary Policy Expectations,” September, working paper.
- Dovern, Jonas, Ulrich Fritsche, and Jiri Slacalek (2012) “Disagreement among forecasters in G7 countries,” *Review of Economics and Statistics*, 94 (4), 1081–1096.
- 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.
- 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.
- 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.
- 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.
- 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 (2021) “Dynamics of Subjective Risk Premia,” working paper, University of Chicago.
- 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.
- Patton, Andrew and Allan Timmerman (2010) “Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion,” *Journal of Monetary Economics*, 57 (8), 803–820.
- 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, Working Paper.
- Reis, Ricardo (2020) “The People versus the Markets: A Parsimonious Model of Inflation Expectations,” December, Working Paper.
- 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 (2021) “Monetary Policy Expectation Errors,” Working Paper.
- Singleton, Kenneth J. (2021) “Presidential Address: How Much ‘Rationality’ Is There in Bond-Market Risk Premiums?” *Journal of Finance*, 76 (4), 1611–1654.
- 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.